ISSN 1471-0498

DEPARTMENT OF ECONOMICS

DISCUSSION PAPER SERIES

CHOICES BETWEEN SIMPLE AND COMPOUND LOTTERIES: EXPERIMENTAL EVIDENCE AND NEURAL NETWORK MODELLING

Daniel John Zizzo

Number 57

January 2001

Manor Road Building, Oxford OX1 3UQ

Choices between Simple and Compound Lotteries: Experimental Evidence and Neural Network Modelling

Daniel John Zizzo*
Christ Church, Oxford

January 5, 2001

Abstract

An experiment on choices between single and compound lotteries is presented, and results are calibrated with neural network models. Many subjects tend to average out probabilities, though behavior becomes more rational with more exposure to compound lotteries in the practice stage. The Prior Knowledge Model hypothesizes that subjects categorize stimuli according to the prior knowledge acquired in their long-run learning history; practice stage cues help them referring to the relevant learning history. The trained networks predict the behavior of about 3/4 of the subjects with transitive preferences; the model can explain where we would expect the trained networks to fail.

JEL Classification Numbers: C91, D81
Keywords: conjunction fallacy, neural networks, heuristics, probability compounding

---

* Thanks to James Engle-Warnick and participants to the Experimental Economics Workshop at Oxford University for helpful advice, and Edmund Fantino because, without him, this research would have never been started. The usual disclaimer applies.

1 Introduction

You are getting information on a potentially attractive high-tech low-priced stock, whose main product is a new material to isolate radioactivity emissions from nuclear waste. The company is currently losing money, which may explain the low price. You learn two things about this material. First, the news appears promising that the Department of Energy will certify this product for commercial use. Second, the material is likely to have some advantages relative to its competitors in the market. Both are good things, so on average the stock looks good and you purchase some shares.

We may simplify our analysis of this decision by considering that there are only two states of the world, one where the product makes a big revenue to the company ('winning' state), and one where it does not ('losing' state). We can ignore the case where only some revenue is made on the product because the company would still go bankrupt. The decision can then be simplified to a binary choice compound lottery where you win if the Department of Energy certifies the product *and* if the material is actually better in some respects than that of its competitors. By the monotonicity axiom of probability (axiom 3 in Hoel, 1984), it then follows that $Prob(certification\&betterproduct) \leq min[(Prob(certification), Prob(betterproduct)]$: assuming that the two events are independent (obviously, they need not be), the probability of winning the compound lottery would be equal to just the product of 'winning' the simple lotteries. However, this is not what the economic agent in our example did. What he did is to average out the probabilities of the two events to assess how good the lottery looks like. This is an example of what Tversky and Kahneman (1982) called the 'conjunction fallacy' using verbal tasks: the fallacy occurs whenever a subject judges the conjunction of two events to be more rather than less (or at least equally) likely to occur than one of the events alone.

This paper has two objectives. First, we present the results of a behavioral (rather than verbal) experiment that analyzes how subjects choose between simple and compound lotteries, after having been exposed to different fractions of compound lotteries in the practice stage. The conjunction fallacy is committed between about 20 and 50% of the times, according to the experimental condition. Many subjects tend to average out probabilities, but behavior becomes more rational the more the exposure to compound lotteries in the practice stage. There is no serious problem of transferability of knowledge from old to new compound lottery tasks.

Second, we present two neural network models relating specific cognitive hypotheses to specific quantitative predictions on probability compounding. Neural networks are artificial intelligence models inspired by analogy with the brain and realizable in computer programs. They have been used to model a variety of cognitive processes involving learning by example, such as pattern recognition and categorization (Taraban and Palacios 1994), arithmetic learning (Anderson 1998) and bounded-rational play in normal form games (Zizzo and Sgroi 2000). Our interest is exclusively in neural networks as models of decision-making rather than as econometric tools: hence, the focus will not be on learnability but rather on what is learnt *with a specific training history,* reflecting certain psychological hypotheses.

According to the Tabula Rasa Model, agents enter the laboratory virtually naïve about probability compounding, and behave in the testing stage simply according to the reinforcement learning in the practice stage. The population of trained tabula rasa networks does not have sufficient variance to explain the behavioral variance across subjects: we conclude that the model is falsified.

According to the Prior Knowledge Model (PKM), subjects categorize stimuli according to the prior knowledge acquired in their learning history, not just the immediate history. The long-run learning history is formalized as parsimoniously as possible, by differences by 6% or less in the composition of the training set relative to the benchmark of no compound lotteries faced in one's own learning history. Further, in the PKM, the role of the practice stage is that of providing cues to the subjects about the nature of the task, and so to what aspect of their learning history they should be using. If there are no or insufficient cues on compound lotteries during the practice stage, subjects are more likely to misperceive the probability mapping. This may lead to more additive stimuli as subjects may think of compound stimuli as additive signals of the quality of a good. It may also lead to more idiosyncratic and less consistent behavior. The PKM-trained networks can explain about three fourths of the choices of subjects with consistent (transitive) preferences; further, the PKM can explain where we would expect the PKM-trained networks to fail in predictive power. The PKM gives quantitative indications of what kind of learning environment may produce which cognitive compounding capabilities in real-world economic contexts.

The paper is structured as follows. Section 2 reviews the literature. Section 3 describes the experimental design. Section 4 presents the basic results, and how describable they are by rules or convex combinations of rules. Section 5 introduces the neural network and sets the scene for the neural network models, presented in sections 6 and 7 respectively. Section 8 concludes.

2 Review

The original formulation of the conjunction fallacy task was based on a verbal description of a character (e.g., Bill, described as 'intelligent but unimaginative, compulsive and generally lifeless' and so on) and then on judgements on the 'probability' of certain simple and compound events (e.g., 'Bill is an accountant', 'Bill is an accountant who plays jazz for a hobby'). The verbal nature of the task has induced the belief that it is due just to a misunderstanding of the word 'probability' (Hertwig and Gigerenzer 1999). Presentations of the task in terms of 'frequency' reduces fallacy committal (Gigerenzer 1996), though Mulford and Dawes (1999) found a significant committal in the recollection of the frequency for personal events. Moreover, strong hints that should induce probabilistic reasoning, and training in formal logic, did not reduce fallacy committal (Zizzo, Stolarz-Fantino, Wen, and Fantino 2000), nor did having a short training period (Benassi and Knoth 1993). Another worry is that the verbal description may induce representativeness-based reasoning (Tversky and Kahneman 1982), that would disappear in a non-verbal choice, and the latter may be more of interest to the economist. A quarter of the subjects still committed the fallacy when the framing description was eliminated in Stolarz-Fantino, Fantino, and Kulik (1996). Zizzo, Stolarz-Fantino, Wen, and Fantino (2000) found the fallacy quite robust to monetary incentives, some repetition, and the usage of the word 'likelihood' rather than probability; a reduction in the complexity of the problem did significantly reduce the fallacy committal, though.

The at least partial variability of the conjunction fallacy rate according to the verbal formulation of the problem strongly suggests the need for a behavioral test that may help assessing whether economists should or should not be concerned about probability compounding, independently of issues of representativeness or of linguistic ambiguities.

Fantino and Savastano (1996) first trained subjects on simple and possibly also compound lotteries. The lotteries had two possible outcomes, either a win or a loss. In case of the simple lotteries, subjects were presented with a circle of a simple color, with the colors being associated with different winning probabilities. In case of the compound lotteries, half of the circle was of one color and the other half of another: each color corresponded to a simple lottery. The circle then disappeared and subjects were asked to choose between two letters, one of which was always associated with a winning probability depending on the circle color(s). In the second part of the experiments, in some rounds subjects directly chose between two lotteries. Fantino and Savastano found that subjects tend to add up probabilities when they are untrained on compound lotteries, while they tended to play more rationally on compound lotteries they had been trained on.

Unfortunately, these results are in conflict with the evidence from verbal tasks, which suggests that subjects average out rather than add probabilities (Zizzo, Stolarz-Fantino, Wen, and Fantino 2000). Moreover, in Fantino and Savastano monetary incentives were very low for a rather boring task. There was a prize of U.S. $5 for having the highest score 'so far', for an experiment lasting 2 hours in 3 different sessions, to which subjects had to participate for course credit: the incentive to do well on any individual task - including the choices in the second part of the experiments - was minimal. Further, subjects never chose between lotteries during the practice stage: this would arguably be a more natural and possibly ecologically valid way of familiarizing subjects with the task. Finally, there was never a choice between a simple lottery and a new compound lottery in Fantino and Savastano's design: this raises the objection that, even when trained, subjects simply chose compounds on the basis of the winning frequencies they had detected in the lengthy training. If this were the case, it would not be appropriate to consider Fantino and Savastano a test of some cognitive process on how people compound probabilities if they are exposed to some probability compounding during training. Rather, we would have the more modest finding that, having faced specific lotteries during training, subjects tend to learn to respond to them according to their observed frequencies. The interpretation of the results is made more difficult by the small sample size (24 subjects divided between 2 conditions).

3 Experimental Design

The experiment presented in this paper has several distinctive features. Like Fantino and Savastano (1996), it is a behavioral task where subjects have to choose between colors associated with probabilities, rather than evaluate the probability of some verbal statement. However, subjects have to make choices between lotteries both in the practice and in the testing stage. In the testing stage, there are significant monetary incentives, and subjects who had practice on compounds face some lotteries never encountered before, not only 'old' compounds. There was also a larger sample size (64 subjects).

Subjects were not verbally told that compound stimuli were compound lotteries, but were able to see the colors associated to the simple lotteries stacked together in correspondence to the compounds. In the condition with no compound lotteries in the practice stage, this was all they had, though the instructions were sufficient to infer that it was a matter of combining the probabilities of the simple lotteries (this claim will be tested later). This is, after all, the situation many decision-makers are likely to find themselves in in real world environments, including that of the investor at the start of this paper: it is up to the decision-maker to recognize that simple lotteries should be combined together in a compound lottery. The case could still be made that stacking colors together may spuriously prime additive probabilities, but Fantino and Savastano (1996) used the 'half circle' system to represent compound probabilities, and as it will become apparent later responses were more and not less additive in their data relative to ours. We also know, from the experiments with verbal data, that explicit verbal statements suggesting the need for probability compounding ('X and Y') do not help. This paper has a different and arguably cleaner strategy to 'tell' subjects they were dealing with probability compounding: in the conditions with compound lotteries in the practice stage, although not informed verbally, subjects had reliable reinforcement cues, increasing in the number of compound lotteries faced during practice.

3.1 Design

The experiment was performed in Oxford in the Fall of 1999. 64 subjects participated to the experiment, one at a time, in a room where only the experimenter was otherwise present. Most subjects were males

(45 out of 64). Subjects were either undergraduate or postgraduate students at Oxford or Oxford Brookes University; they were recruited by email and printed adverts, and some of them had participated to previous unrelated experiments. The experiment was computerized, except for the instructions, which were presented on a printed sheet of paper; subjects were also provided with blank paper and a pen, to take whatever notes they wished during the experiment. The instructions are presented in Appendix A. Average payment was 6.11 U.K. pounds for about 20-25 minutes of work. Subjects were paid immediately at the end, upon signing a receipt and pledge of confidence on the content of the experiment.

The experiment was divided in two stages, the practice stage and the testing stage. In the practice stage, subjects faced 150 choices among three lotteries, each of which was associated with a different color (blue, red and yellow). The colors were univocally mapped into 3 simple lotteries, a High (H_A), a Medium (M_A) and a Low (L_A) lottery. These lotteries corresponded to a winning probability equal to 75, 50 and 25%, respectively. To choose a lottery, subjects had to click a button placed immediately below the color. When a subject chose a lottery, the computer randomly determined whether a subject had won or not by using these probabilities. In all cases, winning a lottery increased the score of a subject by one point. Each point earned in the practice stage was worth 2 U.K. pence.

The computer determined randomly at the start of each session what colors corresponded to H_A, M_A and L_A. The position of the lotteries was determined randomly on he screen each round.

The experiment had three conditions, according to the mix (the "learning mix") of lotteries offered in the practice stage. In the zero compound lotteries condition, subjects faced simple lotteries only. In the CP10% condition, 10% of the choices (i.e., 15 out of 150) that the subjects had to make were choices among two simple lotteries and a compound lottery ('compound lottery tasks'). All compound lotteries presented in the practice stage were made of two colors ('2-compound lotteries'), at least one of which was also simultaneously presented as a simple lottery. Let us label HM_A as the compound lottery made of a H_A and an M_A lottery, and similarly for HL_A and ML_A. The probability of winning a compound lottery was determined as the product of the winning probabilities associated to the simple lotteries (colors) making up the compound. For example, in the case of HL_A, the probability was $0.75 \times 0.5 = 0.375$.

The compound lotteries presented in the practice stage were HL_A and ML_A: HM_A was never presented during practice. Compound lottery tasks were interspersed among the other practice stage choices, according to an order randomly determined by the computer at the start of each session. The computer also randomly determined whether in each compound lottery task subjects faced a HL_A or an ML_A task: for each compound lottery task, there was a 50% chance of encountering each of the two. Figure 1 depicts a practice stage computer screen with two simple lotteries and a compound lottery.

[Insert Figure 1 about here]

Subjects had no initial information about which lottery was good or bad, but: a) the instructions stressed that "whether a bet is good or not depends on the color(s) that you see above the button". After each of the 150 choices, they received feedback on the computer screen on whether they had won or lost the lottery they had chosen: hence, they could infer information from repeatedly playing each type of lottery. The actual winning probability for the compound lotteries was determined according to the product rule for independent events, i.e. as the product of the winning probabilities of the simple lotteries the compound lottery was made of. Hence, in the case of HL, the winning probability was determined as $\frac{3}{4}\times \frac{1}{2}= 3/8$; in the case of ML, as $\frac{1}{2}\times \frac{1}{4}= 1/8$. If they wanted, subjects could take notes on paper on the outcome of their choices. They had to wait at least two seconds between decisions. In the instructions they were told that, to make good decisions in the testing stage, they "might get useful information from the practice stage", and that therefore they should be careful about the decisions they made, and not hurry.

The instructions highlighted that the testing stage had high monetary incentives relative to the practice stage: there were only 10 choices to make, and each point earned in the testing stage was worth 1 U.K. pound at the end of the experiment. The testing stage presented choices between two lotteries at a time. Tasks 1-3 required a choice between simple lotteries, in all possible combinations. Tasks 4-10 were compound lottery tasks. More specifically, tasks 4-7 involved a choice between a simple lottery and a 2-compound lottery of which the simple lottery was a component (e.g., H_A v. HL_A). All possible 2-compound lotteries could be present, including HM_A. Tasks 8-10 involved a choice between each simple lottery and the 3-compound lottery combining H_A, M_A and L_A (HML_A). Figure 2 shows a computer screen from the testing stage.

[Insert Figure 2 about here]

Subjects learnt how many lotteries they had won in the testing stage when they got paid at the end of the experiment. However, and importantly, no feedback was provided after each individual choice in the testing stage. This is because the testing stage was meant to get a snapshot of each subject's preference ordering over simple and compound lotteries, as it emerged after the practice stage: we wanted to minimize the likelihood of dynamic changes in the preference ordering during the testing stage.

4 Results

4.1 Some Basic Results

55.1% of the choices exhibited the conjunction fallacy in the CP0% condition. The fraction decreased to 31.97% and 22.38% in the CP10% and CP20% conditions, respectively. On the one hand, the Spearman correlation between conjunction fallacy committal and the fraction CP of compound lotteries is significant (ρ=-0.544, P<0.01). On the other hand, the results are comparable to the values found with verbal description experiments, and there is still a 20% lower bound on conjunction fallacy committal.

It may be useful to recover the subjects' preference orderings among simple and compound lotteries. The choices made in tasks 1-3 can be used to determine what simple lottery the subject believed as most probable, which one was second most probable and which one was least probable: we shall label these lotteries as H, M and L. H, M and L corresponded to H_A, M_A and L_A for a majority of the choices made by the subjects (71.35%, not significantly different across conditions). Even when this was not true, however, it was possible to infer the subject's preference ordering in terms of H, M and L, as long as choices in tasks 1-3 were not intransitive. For example, if $H_A > L_A$ and $H_A > M_A$ but $L_A > M_A$, one could map $H_A \rightarrow H$, $M_A \rightarrow L$ and $L_A \rightarrow M$. For 6 subjects out of 64, preferences were intransitive even just looking at tasks 1-3, an example of what we might call a 'basic intransitivity' (e.g., $H_A > L_A$, $L_A > M_A$ and $M_A > H_A$). In other cases, a preference ordering could be reconstructed over the simple lottery, but this contradicted with transitivity when the analysis was extended to 2-compound or 3-compound lottery tasks. For example, a contradiction emerged if, in the 3-compound lottery tasks, the choices made were such that $L > HML$ and $HML > H$, hence by transitivity $L > H$. Overall, 27 out of 64 subjects (42.19%) displayed intransitive preferences.

In what follows we shall consider three samples to ensure the robustness of our analysis. The first, the full sample, includes all observations. The second, the pure sample, includes only the subjects with transitive preference orderings (37 out of 64). The third, the robust sample, is equivalent to the pure sample plus all choices with two properties: a) they are made by the subjects with non-basic contradictions; b) they do not lead, by transitivity, to contradictions with the basic preference ordering. For example, a subject for whom $L > HML$ and $HML > H$, but whose other choices were consistent with her basic preference ordering over the simple lotteries H, M and L, would have $L > HML$ and $HML > H$ removed from the robust sample, but all other observations included.

Performance was slightly better on compound lotteries used in the practice stage than on new compounds, i.e. HM_A and HML_A (39.5 vs. 31.74% of conjunction fallacy committal in the full sample, respectively). In the lack of transfer of knowledge from playing "old" to new compounds, we would expect a signficant interaction effect, leading to less fallacy committal with the "old" lotteries than with the new lotteries as the exposure to compounds in the practice stage increases. However, an F test on lottery novelty and CP was not able to detect a significant main or interaction effect in the full, robust or pure samples. We can conclude that transferability of knowledge was not a serious problem in our experiment.

4.2 Rules Description Framework

We could describe behavior in terms of 3 rules, namely product, averaging and summation. Product corresponds to the application of rational compounding (e.g., in line with the monotonicity axiom of probability and, as a consequence, with Bayes' Rule) for the case of independent events.

All the tasks with compounds in the testing stage have two features: a) they involve a choice between a simple and a compound lottery; b) the compound lottery has the simple lottery as a component. Taken together, these features imply that in our setting rational choice corresponds to a very simple rule-of-thumb: always go for the simple lottery. Similarly, summation corresponds to the very simple rule-of-thumb of always going for the compound. In both cases, the rule is insensitive to whether the subject correct detects the winning probability of the simple lottery.

Averaging consists in averaging out the probabilities of the simple lotteries. This means that in some cases simple lotteries should be preferred to compounds ($H > HM$, $H > HL$, $M > ML$) but in other cases compounds should be preferred ($HM > M$, $HL > L$, $HML > L$). Hence, the rule is sensitive to whether a subject has a clear view of winning probabilities. This implies a potential downward bias in the success of averaging in the robust and full samples, relative to the pure sample. In our setting, averaging also means that no a priori prediction can be made in tasks involving choices between M and HML, since they correspond to probabilities of 0.5. We shall tackle this indeterminacy by allowing both for $M > HML$ and $M < HML$: this extra degree of freedom implies an upward bias in the success rate of the rule - which effectively summarizes two closely related rules, one with $M > HML$ and one with $M < HML$.

Figure 3 graphs the fraction of compound lottery tasks that can be rationalized according to one of the three rules in the full sample.

[Insert Figure 3 about here]

Product can explain 63.39% of the choices, with a summation/conjunction fallacy rate of 36.61%. Averaging can explain 64.29% of the choices, a fraction that is insignificantly different from that (68.3%) if averaging were assumed to occur over H_A, M_A and L_A rather than over H, M, L. Product performs better - people are 'more rational' - the greater the fraction of compounds, while summation does worse. Averaging is the simple best rule in the CP0% condition, and is stable for higher CP values. These results do not change if we consider the robust or the pure sample, though averaging tends to perform slightly worse at CP10%. They also do not change if one considers the rules as applying over H_A, M_A and L_A rather than over H, M and L.

Figure 3 can only be considered a preliminary step towards describing what people are doing. This is because subjects make seven choices each in compound lottery tasks, and their behavior is describable according, say, to the product rule only if all or at least a vast majority of these choices can be described by the product rule. If three choices are rationalizable according to the product rule and four according to summation, say, we would not be able to say that the subject is following either rule, even though Figure 3 would tally 3 choices for the product rule and 4 for summation. This intuition requires us to check how the revealed preferences by each subject are consistent with the predictions made by each rule. Since averaging's predictions depend on the preference ordering, and a preference ordering can be constructed only on the basis of non contradictory choices, we shall be dealing with the robust sample and, still better, the pure sample. To allow for some error, a subject's preference ordering will be considered as describable by a rule if it fits at least all the compound lottery tasks minus one. This normally implies a fit of 6 out of 7 choices, but, if a task is M v. HML, it effectively requires only a fit of 5 out of 6; the numbers are lower if two or three choices by a subject are excluded because leading to intransitivity, as it is possible in the case of the robust sample. This increases the chances of mistaken fits, and is an additional reason to assign more weight on the pure sample: in the pure sample the bias does not occur, since only choices by subjects with transitive preference orderings are included.

Table 1 describes the overall fit of what we might call the Rule Description Framework (RDF) in both the pure and the robust sample.

[Insert Table 1 about here]

Overall, 75.68% of the subjects' preference orderings are describable by at least one of the rules in the pure sample (the fraction is slightly higher in the robust sample). The CP0% success rate hovers around 50%, while that in the CP10% and CP20% conditions is around 90%. There is clearly more behavioral dispersion in the CP0% condition, with both averaging and product being more successful in the CP10% and CP20% conditions. Summation disappears in the pure sample and falls in the robust sample for non-zero CP values. An F test on rule correctness using rules and CP as factors shows a significant rules main and interaction effects (e.g., in the pure sample, F=10.599, P<0.001 for the main effect; F=4.274, P<0.05 for the interaction effect); however, the post hoc tests (Tukey and Scheffé) show that the difference between averaging and product is never significant. Unlike what found by Fantino and Savastano (1996), summation appears to underperform relative to the other rules.

4.3 The Novelty Hypothesis

Consider a subject facing a stack of two or three colors in the testing stage of the CP0% condition. She has never faced a stack before, although she has encountered the simple lotteries whose colors are present in the stack. She can infer from the instructions that the stack is 'a bet', and that she should be using the information acquired in the practice stage 'to make a good decision'. The information she has is on the simple lotteries, and so she can infer that what is required is some combination of the probabilities of the simple lotteries.

Assume, however, that she actually perceives the compound lottery as an entirely new lottery. This would be equivalent to facing a new color in the testing stage, e.g. white. We may label this as the Novelty Hypothesis. It implies that the subject will discount whatever information she has on the simple lotteries, and will assign an average probability of about 50% to the 'new' lottery, since she believes she has no prior information about it. Obviously, the more the exposure to compound lotteries in the practice stage, the less likely is this to occur.

The Novelty Hypothesis places testable, if somehow weak, boundaries on observable behavior. H should be preferred to any compound, and any compound should be preferred to L: intuitively, a lottery which you know is 'good' should be preferred to an unknown lottery, and vice versa in the case of a 'bad' lottery. Averaging is consistent with these constraints on 'novelty-rationalizable' behavior, but so are choices inconsistent with averaging, such as ML over M. Notwithstanding the indeterminacy on the relationship between M and any compound, the Novelty Hypothesis is testable because it suggests that the boundaries it places on behavior should be observed by a significantly larger proportion of subject in the CP0% than in the CP10% or the CP20% condition: more precisely, we would expect a negative correlation between CP values and novelty-rationalizable behavior.

68.42, 73.61 and 72.73% of the choices made in tasks involving H or L are compatible with the boundaries placed by the Novelty Hypothesis in the CP0%, CP10% and CP20% conditions, respectively. The correlation between CP and novely-rationalizable behavior is insignificant and positive in either the pure or the robust sample (e.g., in the robust sample, Spearman's ρ=0.038, P=0.303, two-tailed). The result is robust to including all observations (involving H or L) except those by subjects with basic contradictions and hence non-recoverable preference orderings. It is also robust to evaluating the fit of the Novelty Hypothesis in the full sample, using the actual simple lottery probabilities (H_A and L_A) rather than those recovered from Tasks 1-3 (H and L). Finally, the same conclusion emerges if, rather than taking one observation at a time, we consider the fit of the Novelty Hypothesis predictions on the choices overall made by each subject. We conclude that the evidence appears against any significant role for the Novelty Hypothesis.

4.4 The p-Rules Framework

The empirical failure of the Novelty Hypothesis suggests that subjects were able to bring their priors on the simple lotteries to bear on the probability of the compound lottery. Product, summation and averaging are three ways of doing so. A related way to try to describe the data is to allow subjects to be uncertain on whether product or summation is applicable: in the p-Rules Framework (pRF) the subject assigns a probability weight p to the applicability of the product rule, and a weight (1-p) to that of summation. Consider three lotteries with probabilities u, v and z; denominating uv and uvz the compound lottery probabilities, we have

$$uv = p(u+v) + (1-p)(u \times v) \tag{1}$$

$$uvz = p(u+v+z) + (1-p)(u \times v \times z) \tag{2}$$

under the obvious upper constraint that $uv, uvz \leq 1$. These formula allow the computation of preference orderings according to the p value. Table 2 depicts how, according to p, ten different preference orderings emerge in relation to the compound lottery tasks faced in the testing stage. It also shows why one needs not consider averaging separately: P5 and P6, the rules for p values between 0.429 and 0.2, correspond to the averaging rule (P5 with $HML > M$, P6 with $HML < M$).

[Insert Table 2 about here]

When extended to allow for covariation between the simple lotteries outcomes, the pRF can be used to analyze to what degree even apparently transitive preference orderings are compatible with the usage of either a simple rule or a convex combination of simple rules. In section 7.2 and in appendix B, this will lead us to the concept of quasi-intransitivity.

Once again, we consider for what fraction of subjects the behavioral rules fit the choices made in at least all compound lottery tasks except one (see Table 2). Table 2 displays a story similar to the one we saw with the RDF: as CP increases, there is a movement from summative-like to averaging to product-like compounding. In the pRF, this movement could be interpreted as a migration towards higher p values.

Unfortunately, although we are using ten rules rather than the three/four of the RDF, in both the pure and the robust sample the pRF is able to explain the behavior of only two extra subjects relative to the RDF (81.08 vs. 75.68% of the subjects in the pure sample; 87.93 vs. 84.48% of the subjects in the robust sample). The difference is insignificant in a t test (e.g., for the pure sample, t=0.558, d.f.=72, P=0.578, two-tailed). Hence, there is not much value added in using the pRF relative to the simpler RDF.

4.5 Frameworks and Models

The RDF and the pRF are descriptions of behavior that is observed in the experiment. They are not explanations of why people are behaving the way they do. For that we would need a cognitive model of how people take compounding decisions. Neither the RDF nor the pRF explain why subjects may be adding, averaging or multiplying the winning probabilities of the simple lotteries. The pRF is more sophisticated in suggesting one psychological way in which uncertainty about the decision problem at hand translates itself into a behavioral rule: however, nothing in the pRF explains what determines the key p variable. Moreover, introducing p does not shed significant additional light to that provided by the purely descriptive RDF. Because of these difficulties we referred to both the RDF and the pRF as 'frameworks' rather than as 'models'.

Probability axioms provide one principled cognitive model of how rational agents should handle probabilities: however, it is apparent from our analysis so far that a considerable amount of behavior can be better explained by averaging and, in the CP0% condition, even summation. Moreover, the incidence of rational or less rational behavior varies across conditions in meaningful ways that require explanation: behavioral patterns tend to be less dispersed, and more rational behavior occurs, for higher CP values. Similarly, the information averaging model by Anderson (1981) is limited insofar as it can only explain one particular behavioral pattern, namely averaging.

Support theory (Tversky and Koehler 1994) is a descriptive account for dealing with uncertain disjunctive events: under some auxiliary assumption, it can justify the presence of the conjunction fallacy. The assumption is that, in comparing a simple lottery X and a compound lottery XY, subjects fail to recognize that (a) $P(X) = P(XY) + P(X\overline{Y})$, and rather think that (b) $P(X) = P(X\overline{Y})$. Subjects may be sensitive to cues leading to perceiving compound lotteries subadditively (according to (a)) or otherwise (according to (b)). Even so, one would need to move from this reinterpretation of the fallacy to specific predictions - what exactly determines a specific degree of subadditivity, and why -, and this has yet to be done. The pRF might be reinterpreted as a way of balancing out different hypotheses about the addditive or subadditive nature of the fallacy, but this would not solve the quantitative indeterminacy of the pRF. Another possible specification based on support theory might be that there should be greater fallacy committal with conjunctions of similar than of dissimilar events.[1] However, if one interprets similarity in probability assessment as one possible attribute of the similarity of two events, the reverse is actually true (Zizzo, Stolarz-Fantino, Wen, and Fantino 2000).

In the next sections we shall discuss two cognitive models linking specific psychological hypotheses and learning histories to behavioral predictions. Both models use neural networks, and more specifically the same neural network architecture. One of the two models captures the broad psychological intuition from support theory of the conjunction fallacy as a misperception sensitive to the cues of the decision problem.

5 Neural Network Modelling

5.1 A Brief Introduction

In this section I provide a very compact introduction to neural network modelling. Neural networks are used in this paper as psychological models embodying specific hypotheses on cognitive processes and allowing quantitative predictions on the behavior of decision-makers. Hence the focus is *not* on learnability but rather on how *particular training histories* affect probability compounding choices in a suitably parsimonious neural network architecture. The interested reader should refer to Zizzo and Sgroi (2000) and Zizzo (2000) for a

---

[1] For a parallel, see Rottenstreich and Tversky (1994).

more extended treatment and some motivation.

Networks can be thought of as agents that receive external stimuli, process them, and produce an output. A typical network has an input layer of nodes receiving stimuli from the outside (as real numbers). This input is then transmitted to the nodes the input layer is connected to, multiplied by the respective connection weights. Each node on the downstream layer receives input from many nodes. The sum is then transformed according to the activation function and the result is transmitted to the nodes in the further downstream layer. In such a way, the network processes the input until it reaches the output nodes (the output layer), in the form of new real-valued numbers. The activation level of the output nodes expresses the outcome of network processing, i.e., the network's decision. A feature of biological brains is that the connections between neurons are of different strengths, and that they can either increase or decrease the firing rate of the receiving neuron. In neural networks, this is modelled by associating a connection weight to each connection. Training a network to perform a task means to have the network adjust connection weights in such a way as to perform the task.

Neural network simulations usually are divided into two parts, a *training stage* and a *testing stage.* In the training stage, the network adjusts connection weights following a basic learning algorithm (effectively a numerical optimization technique) called backpropagation developed in Rumelhart, Hinton, and Williams (1986). Performance is then tested in the testing stage.

Consider a neural network, or more simply A, to be a machine capable of taking on a number of states, each representing some computable functions mapping from input space to output space, with one *hidden* layer of further computation between input and output. The following definition formalizes this:

Definition 1 *Define the neural network as* $C = \langle \Omega, X, Y, F \rangle$ *where* Ω *is a set of states,* $X \subseteq \mathbb{R}^n$ *is a set of inputs,* Y *is a set of outputs and* $F : \Omega \times X \mapsto Y$ *is a parameterized function. For any* ω *the function represented by state* ω *is* $H_\omega : X \mapsto Y$ *given by* $H_\omega(x) = F(\omega, x)$ *for an input* $x \in X$*. The set of functions computable by* C *is* $\{H_\omega : \omega \in \Omega\}$*.*

Put simply, when the network, A, is in state ω it computes the function H_ω. First we will consider the form of the network in practice, and start by defining an activation function.

Definition 2 *An activation function for node i of layer k in the neural network A is of the logistic form*

$$a_i^k = \frac{1}{1 - \exp\left(-\sum_j w_{ij}^k u_{ij}^{k-1}\right)} \tag{3}$$

where u_{ij}^k is the output of node j in layer $k-1$ sent to node j in layer k, and w_{ij}^k is the weight attached to this by node i in layer k. The total activation flowing into node i, $\sum_j w_{ij}^k u_{ij}^{k-1}$, can be simply defined as t_i.

During training A receives a sequence of inputs p until some some round T, that is always fixed in this paper. The *training sample* consists of T binomial choice lotteries V, the nature of which depends on the specification of the model. The *correct choice* is considered to be a Boolean variable, v, whose values of 1 and 0 correspond respectively to having won and having lost the lottery. The occurrence of a win or a loss is determined randomly with the probability P_V associated to V. We now need to define the method by which during training, when A receives an input and produces an output, its connection weights get adjusted to reflect whether the network faces a win or a loss. The method used is backpropagation. We restrict ourselves to one output node, and this implies that the error function that is commonly used (root mean square error) here takes the simple form of absolute difference between output and correct choice: that is, given an output y_p and a correct choice u_p for an input pattern p, $\varepsilon = \mid u_p - y_p \mid$.

Backpropagation aims to minimize the error function by altering the set of weights w_{ij} of the connections between a typical node j (the sender) and node i (the receiver) in different layers. These weights can be adjusted to raise or lower the importance attached to certain inputs in the activation function of a particular node. Backpropagation is a form of numerical analysis akin to gradient descent search in the space of possible weights. Following Rumelhart, Hinton, and Williams (1986) we use a function of the form:

$$\triangle w_{ij} = -\eta \frac{\partial \varepsilon}{\partial w_{ij}} = \eta k_{ip} o_j \tag{4}$$

where w_{ij} is simply the weight of the connection between the sending node j and receiving node i. As ε is the neural network's error, $\partial \varepsilon / \partial w_{ij}$ measures the sensitivity of the neural network's error to the changes in the weight between i and j. There is also a *learning rate* given by $\eta \in (0, 1]$, that may make learning slower or quicker. Define $\partial\varepsilon/\partial w_{ij} = -k_{ip}o_{jp}$ where o_{jp} is the degree of activation of the sender node o_{jp}. The higher o_{jp} is, the more the sending node is at fault for the erroneous output, so it is this node we wish to correct more. k_{ip} is the error on unit i for a given input pattern p, multiplied by the derivative of the output node's activation function given its input. Calling g_{ip} the goal activation level of node i for a given input pattern p, in the case of the output nodes k_{ip} can be computed as:

$$k_{ip} = (g_{ip} - o_{ip})f'(t_{ip}) = (g_{ip} - o_{ip})o_{ip}(1 - o_{ip}) \tag{5}$$

since the first derivative $f'(t_{ip})$ of the receiving node i in response to the input pattern p is equal to $o_{ip}(1-o_{ip})$ for a logistic activation function. Now assume that a network has N layers, for $N \geq 2$. As above, we call layer 1 the input layer, 2 the layer which layer 1 activates (the first hidden layer), and so on, until layer N the output layer which layer $N-1$ activates. In the case of the network presented in this paper, we shall have $N = 3$.

We can now define the backpropagation learning process.

Definition 3 *Using backpropagation, we first compute the error of the output layer (layer N) using equation 5, and update the weights of the connections between layer N and $N-1$, using equation 4. We then compute the error to be assigned to each node of layer $N-1$ as a function of the sum of the errors of the nodes of layer N that it activates. Calling i the hidden node, p the current pattern and β an index for each node of layer N (activated by i), we can use:*

$$k_{ip} = f'(t_{ip}) \sum_{\beta} k_{\beta p} w_{\beta i} \tag{6}$$

to update the weights between layer $N-1$ and $N-2$, together with equation 4. We follow this procedure backwards iteratively, one layer at a time, until we get to layer 1, the input layer. A variation on standard *backpropagation would involve replacing equation 4 with a momentum function of the form:*

$$\triangle w_{ij}^{t} = -\eta \frac{\partial \varepsilon^{t}}{\partial w_{ij}^{t}} + \mu \triangle w_{ij}^{t-1} \tag{7}$$

where $\mu \in [0,1)$ *and* $t \in \mathbb{N}^{++}$ *denotes the time index (an example game, vector* x*, is presented in each* t *during training).*

Momentum makes connection changes smoother by introducing positive autocorrelation in the adjustment of connection weights in consecutive periods. The connection weights of the network are updated using backpropagation until round T, i.e. the end of the training stage. Backpropagation is not used in the testing stage, which implies that the network has a fixed set of connection weights during testing.

5.2 The neural network architecture

We now describe the neural network architecture that we shall be using for both models. The network is pictured in Figure 4: it has three input nodes, one for each simple lottery, fully interconnected with a hidden layer of 3 nodes, fully interconnected downstream with a simple output node.

[Insert Figure 4 about here]

Input can be represented as a vector of three numbers, whose k-th element is applied to the k-th input node. Inputs take a value of either 0 or 1. If a node receives a value of 1, we can say that the node is activated. A simple node is activated when the network faces a simple lottery: if the lottery is H, the first node is activated and the input vector is [1 0 0]; if the lottery is M, the second node is activated and the input vector is [0 1 0]; finally, if the lottery is L, the third node is activated and the input vector is [0 0 1]. A compound lottery is operationalized simply by activating more than one input node at once, namely the input nodes corresponding to the simple lotteries the compound lottery is made of. Hence, the input vector will be [1 1 0], [0 1 1], [1 0 1] and [1 1 1] in the case of HM, ML, HL and HML, respectively.

Activation gets processed downstream according to the connection weights (positive or negative) of the network for each connection: the details of this processing are discussed in the previous section.

The algorithm produces an output, the activation level of the output node. This is a number between

0 and 1, and can be thought as the winning probability assigned by the network to the lottery at hand. The network is never explicitly taught this probability, even for the lotteries encountered during training. Rather, more realistically, what the network observes in the training stage is simply whether the lottery led to a win or to a loss: a win corresponds to a 'correct answer' of 1, and a loss to a 'correct answer' of 0. This correct answer is then used to update the network's connection weights using the backpropagation algorithm, to minimize the distance between observed output and correct answer. Since wins and losses occur probabilistically, according to the winning probability of each lottery, the best the network can do is to produce an activation corresponding to the winning probability of the lottery. Hence, as in many realistic settings (including that of our experiment), frequency detection occurs purely by repeated positive and negative experiences about the occurrence of one (win) or another (loss) state of the world.

In the testing stage, the network is faced with each of the lotteries (H, M, L, HM, ML, HL, HML) and determines probabilities for each of them. One can then recover the preference ordering of the network from the probabilities assigned by the network to each lottery. One implication, and possibly limitation, of this architecture as a modelling tool is that one can only infer a transitive preference ordering, although predictions can be made (and shall be made) on what the architecture cannot explain.

Another important feature of the network architecture is an application of the parsimony principle: input nodes are provided only for the simple lotteries, rather than separately also for the compound lotteries. There are three related advantages from this restriction. First, it allows us not to need to presume that the cognizing agent has an inbuilt capability of dealing with compound lotteries, as it would be the case if there were dedicated input nodes. Second, it genuinely structures the lottery compounding problem as one of compounding the stimuli associated with the simple lotteries. Third, the network design implies that the number of required nodes increases linearly, rather than exponentially, with the number of simple lotteries.

6 The Tabula Rasa Model

The Tabula Rasa Model (TRM) assumes that agents enter the lab virtually naïve about probability compounding, e.g. as "tabula rasae" or, put it differently, as zero rational automata. Anything they learn they learn by reinforcement learning from the practice stage. Hence, in the TRM the practice stage is for the subjects what the training stage is for the network. This implies that the training stage should be modelled on the basis of the practice stage. The lotteries chosen in the practice stage and the actual feedback received from the computer were used to train individually each one of sixty-four networks, one per subject. We employed ten sequential repetitions of the practice stage, for a total of 1500 rounds. This was considered as a tentative compromise between two opposite needs: that of allowing the networks a fair chance to diversify their behavior by giving sufficient practice, and at the same time that of modelling an actual practice stage of 150 rounds and not, say, of 15000 or 150000 rounds. The learning rate was fixed at 0.6, and the momentum at 0.3: other combinations were used, but the network's performance was found robust to the specific choice, an unsurprising result (for the same result in a different context, see Zizzo and Sgroi 2000). The usage of a positive momentum rate and sequential training were motivated by the desire to model the fact that lotteries were encountered in the practice stage in a particular order, and this might matter for learning. In other words, it was a way to try to incorporate in the learning process any eventual history-dependence from facing a particular sequence of lotteries and lottery outcomes.

Figure 5 compares the performance of the TRM relative to the RDF in the pure sample, both in terms of perfect fit and of all-but-one-choice fit: the TRM performs rather poorly across CP conditions, with an average success rate of just 32.43% in the pure sample and 39.66% in the robust sample.

[Insert Figure 5 about here]

If one correlates the TRM predictions with those of averaging, Bayes and summation, one finds insignificant correlations with product and summation, but a large significant correlation with averaging predictions (e.g., in the pure sample, $\rho = 0.552$, P<0.001, two-tailed). This is still clearer if one decomposes the TRM in the 'tabula rasa rules' that the individual networks have learnt: one can rank the preference orderings over lotteries by these rules and by averaging, product and summation. 60 out of 64 trained networks have one out of just five preference orderings. 63 out of 64 have $\rho \geq 0.927$ if one compares it to that envisaged by the averaging rule (with much lower correlations in the case of product and summation). Hence, the poor performance of the TRM is due to the insufficient behavioral variance that is produced by the practice stage training, even if one allows training over 1500 rather than 150 rounds as it actually was in the experiment.

Put it simply, the networks tend to average too much relative to the data, leading to a poor performance in the light of the only partial success of averaging as a descriptive rule.

The choice of 1500 rounds was deemed a compromise between realism (the subjects faced the decision task only 150 rounds in the practice stage) and giving the model a fair chance. If instead of 1500 rounds we were to use 15000 rounds - the number of rounds of the model presented in the next section -, performance improves to 43.24% in the pure sample: this result is not surprising since we allow for more training on subject-specific training sets. However, the fit is still considerably poorer than either the RDF or the pRF, and this can be shown to be, once again, because of the excessive reliance by the TRM-trained networks on averaging-related heuristics.

By their own nature, computer simulations cannot obviously show that the tabula rasa idea is wrong. Nevertheless, the model's performance does not warrant optimism: the heterogeneity in population cannot be easily explained on the basis of practice stage reinforcement alone. At least in its present embodiment, the TRM cannot be claimed to have support in our data: agents were probably heterogenous already in some important respect when they entered the laboratory.

7 The Prior Knowledge Model

7.1 Design

The Prior Knowledge Model (PKM) is based on two hypotheses: 1) agents categorize stimuli according to the prior knowledge acquired in their long-run learning history, not just the immediate history; 2) the role of the practice stage is that of providing cues to the subjects about the nature of the task, and so to what aspect of their learning history they should be using. As in support theory, cues are here interpreted as pointers about the nature of the task.

The first hypothesis can be formalized by modelling the training stage of the network as the long-run learning history that the agent to brings to bear on the decision task. In what follows we shall try to model long-run learning histories as parsimoniously as possible. To do this, we shall use the case in which the agent has never faced compound lotteries as the benchmark, and we shall bring deviations as small as possible from this benchmark to see what its impact is on the final knowledge of the agent. In the benchmark case, the network faces only simple lotteries in the training stage and yet is asked to give probability assessments of compound lotteries in the testing stage. The deviations from the benchmark will be in terms of changes in the composition of the training set by 6% or less. To simulate the long-run horizon, we use training spells of 15000 rounds. No assumption is made on the relevance of the sequential order in which the lotteries are presented, and therefore training is in random order, and no momentum is used.[2]

The second hypothesis suggests that, if there are no or insufficient cues on compound lotteries during the practice stage, subjects are more likely to misperceive the probability mapping: this is why we might label this hypothesis as the 'misperception hypothesis'. On the one side, it may lead to more additivity than would otherwise be expected from exposure to any kind of compound lottery in one's own long-run learning history. Assume that, in seeing compound stimuli, subjects think in terms of them as signals of quality of a good: they are more likely to think additively, since, say, some nice-smelling and well-colored fish is considered better than an only nice-smelling or an only well-colored fish. This is more likely if subjects are not explicitly aware of dealing with probability compounding, and therefore the less the cues they receive in terms of compound lotteries in the practice stage. On the other side, the misperception hypothesis suggests that behavior will be more idiosyncratic the less the cues received on compound lotteries, with greater chances of behavior (intransitive or otherwise) that the PKM-trained neural networks will not be able to describe. A testable prediction of the PKM then follows: we would expect a positive correlation between CP values and what the PKM-trained networks can explain.

We label CP0-2% the benchmark case with a training set with no compound lotteries. The label reflects the fact that small perturbations relative to the 0% benchmark are unable to change the emergent preference ordering of the trained network. In particular, such preference ordering is robust to the introduction of a 2% fraction of 2-compound lottery tasks whose winning probabilities of the component simple lotteries are independent ('independent' compound lottery tasks in what follows). This means that, when the network is faced with 2-compound lottery tasks (randomly chosen among HM, HL and ML), it receives the feedback

---

[2] Results are robust to the usage of a positive momentum rate. The learning rate is 0.5, but it could change to 0.6 (as in the TRM) or other value without making any difference to the results, except that too low a value slows learning.

that it has won with a probability determined according to the product rule for independent events (H x M, H x L or M x L, according to the case). The CP0-2% preference ordering is also robust to the introduction of a 2% fraction of 2-compound lottery tasks whose winning probabilities of the component lotteries are perfectly correlated ('perfectly correlated' compound lottery tasks in what follows). For these lotteries, the winning probabilities are determined as the lower of the two simple component lotteries: hence, HM=M, HL=L and ML=L. Finally, the CP0-2% preference ordering is also robust to the introduction of unbalanced samples of independent 2-compound lottery tasks: specifically, to the introduction of 2, 4 and 6% random mixes of HM and ML lotteries, and to that of 2 or 4% random mixes of HL and ML lotteries.

We shall show shortly that the CP0-2% benchmark is an averaging-related algorithm. HLML6% is a 'more rational' preference ordering corresponding to the introduction of an unbalanced sample of independent 2-compound lottery tasks as before, specifically a 6% random mix of HL and ML lotteries. With a fraction of 4% or above of independent 2-compound lottery tasks - randomly chosen among HM, HL and M -, we obtain the same preference ordering as the product rule (CP4% case).

Let us define additive stimuli as stimuli whose winning probability is equal to the sum of the probabilities of the simple lotteries (with an upper ceiling of one): introducing additive stimuli in the training set is a way of formalizing an aspect of the misperception hypothesis, as discussed above. Add1%, Add2%, Add3% and Add4% correspond to training samples with a fraction respectively of 1, 2, 3 and 4% of additive stimuli relative to the CP0-2% benchmark. The preference ordering from training the network on Add1% is obtainable also by introducing a 2% of perfectly correlated compound lottery tasks, including both 2-compound and HML lotteries. The same outcome as Add2% is obtainable also with 4% or more of perfectly correlated compound lottery tasks (2-compounds are sufficient). Conversely, Add3% and Add4% cannot be easily justified by means other than the misperception hypothesis.

Overall, the PKM model tries to retain as much parsimony as feasible by having the smallest possible deviations from the 0% compounds benchmark. The emergence of different preference orderings is explained on the basis of the salience of slightly different long-run learning histories.

7.2 Performance

Figure 6 compares the performance of the PKM relative to the RDF in the pure sample, both in terms of subjects whose choices in the compound lottery tasks are fit perfectly, and in terms of subjects for whom at least all choices but one are fit.

[Insert Figure 6 about here]

In t tests, the basic performance of the PKM is statistically indistinguishable from either the RDF or the pRF, both in the pure and in the robust sample: the PKM predicts the behavior of 75.68% of the subjects in the pure sample, and 81.03% in the robust sample. This is an encouraging result, because now we have a model capable of explaining why people might behave the way they do, and not simply a framework describing what they do.

Table 3 graphs the Spearman correlation matrix between RDF rules and PKM networks' preference orderings: CP4-6% and HLML6% are identical and closely related, respectively, to the product rule; CP0-2%, Add1% and Add2% are most closely correlated with averaging; Add3% and Add4% are most closely related to summation.

[Insert Table 3 about here]

Figure 7 graphs the fit of the preference orderings in the different experimental conditions.

[Insert Figure 7 about here]

Two features emerge as CP values increase, mirroring previous results in discussing the RDF and the pRF: a) the distribution shifts upwards: for example, in the pure sample, the overall fit goes from 53.85% in the CP0% condition to 88.89 and 86.67% in the CP10% and CP20% condition, respectively); b) the distribution shifts to the left. We can measure the leftward shift by constructing an ordinal index with a value of 1 assigned to CP4-6%, 2 to HLML6%, 3 to CP0-2%, 4 to Add1% and so on. The average index value is 4.286 (4.807), 3.188 (2.964) and 2.115 (2.614) in the CP0%, CP10% and CP20% condition, respectively: in the robust sample, Spearman's ρ=-0.414; in the pure sample, Spearman's ρ=-0.502. Both features are in agreement with the misperception hypothesis.

On the basis of the misperception hypothesis, we hypothesized a positive correlation between CP values and what the PKM-trained networks can explain. We can operationalize the latter as a dummy variable, which we label LackOfClarity, equal to 1 in correspondence to those subjects whose preference orderings display intransitivities and those whose choices, even if transitive, PKM-trained networks are unable to fit. We can also use another dummy variable, Contradiction, equal to 1 whenever a subject has intransitive and *quasi-intransitive* preferences. Quasi-intransitive preferences are preferences that do not satisfy the relationship $HM > HL > ML$: an appendix contains a discussion of this and a proof of how these preferences cannot, in general, be rationalized with convex combinations of simple rules, even allowing for beliefs in a positive or negative correlation between the winning probabilities of the simple events. Four subjects display quasi-intransitive (though not fully intransitive) preferences, three in the CP0% condition and one in the CP10% condition.

The mean Contradiction values are 52.38, 61.9 and 36.36% in the CP0%, CP10% and CP20%, respectively; the corresponding values for LackOfClarity are 66.67, 61.9 and 40.91%. While the relationship between CP0% and CP10% is ambiguous, it is difficult to deny a significant decrease under the CP20% condition. Spearman's ρ=-0.135 (P=0.145, one-tailed) for Contradiction, which is correctly signed though insignificant; Spearman's ρ=-0.214 (P<0.05, one-tailed) for LackOfCorrectly, which is significant. In t tests comparing the Contradiction and LackOfClarity mean values in the CP20% condition vs. those in the other two CP conditions, t=1.585 (d.f.=62, P<0.07, one-tailed) for Contradiction and t=1.808 (d.f.=62, P<0.05, one-tailed) for LackOfClarity. These t tests results can be replicated with non-parametric Mann-Whitney tests.

An inspection of the preference orderings of the nine subjects in the pure sample that PKM-trained networks cannot explain - six of which are in the CP0% condition - shows that only two of them are compatible with one another. This suggests the difficulty of capturing this idiosyncratic data in a parsimonious model. Nevertheless, as we discussed, the PKM model is able to make predictions on what is less good at predicting, and the evidence is broadly in agreement with these predictions: this holds particularly when one compares the CP20% condition with the other two CP conditions.

8 Conclusions

This paper presents the results from a purely behavioral experiment on choices between simple and compound lotteries, under differential exposure to compound lotteries in the practice stage. About 50% of the subjects commit the conjunction fallacy if they are not implicitly told they are dealing with compound lotteries. There is a robust 20% lower bound on conjunction fallacy committal, even for significant exposure to compounds. Many subjects tend to average out probabilities, but choices of heuristics change with the experienced learning mix.

We develop two neural network models to explain how economic agents deal with probability compounding. We find support for one of the models, the Prior Knowledge Model: subjects categorize stimuli according to the prior knowledge acquired in their long-run learning history, and cues received in the practice stage help them referring to the relevant learning history. The emergence of different preference orderings is explained on the basis of parsimoniously different salient learning histories. Misperception of the problem at hand may lead to stimuli additivity, intransitivities and quasi-intransitivities. The Prior Knowledge Model gives quantitative indications of what kind of learning environment may produce which cognitive compounding capabilities in real-world economic contexts.

Appendix A: Experiment Instructions

Welcome to the experiment!

The experiment is divided into two parts, a practice stage and a testing stage. You are playing the practice stage first.

In the practice stage you have to choose 150 bets. You choose bets by clicking one of the buttons. *Please do not start clicking buttons until having finished reading the instructions.*

When you click a button, your score on the screen will be updated immediately. For each bet that you win in the practice stage, you get 1 point, which is worth 2 pence at the end of the experiment. After having clicked a button, you need to wait at least two seconds before being able to move to the next decision.

In the testing stage you only have to play 10 bets. Again, you choose bets by clicking one of the buttons below. Now you have to wait at least three seconds between decisions. Unlike the practice stage, you will not know the outcome of your bets after each choice. However, at the end of the testing stage, *you will be paid 1 pound for each bet won in the testing stage.* It is up to you, and to your good luck, whether you will be able to earn a full 10 pounds in the few minutes of the testing stage.

How to make good decisions? You might get useful information from the practice stage, information that you might want to use in the testing stage. This is why you should be careful about the decisions you take not only in the testing stage, but also in the practice stage. There is no need to hurry.

Whether a bet is good or not depends on the colour(s) that you see above each button.

Payment both for the practice and the testing stage will be at the end of the experiment.

If you have any questions or doubts, this may be a good moment to raise them: the experimenter will be happy to help.

Good luck!

Appendix B: Quasi-Intransitive Preferences

In this appendix we shall use and extend the pRF to show in what cases we can talk about *quasi-intransitive preferences.* The intuition is that the revealed preference ordering in our experimental setting should be compatible, if not with a rule, with a simple convex combination of rules. The pRF deals with convex combinations of rules and so is is a natural choice in this context. In particular, we know from section 4.4 that neglecting averaging is without loss of generality since averaging itself can be considered a convex combination of summation and product (see Table 2 above).

We extend the pRF by considering the covariance $\gamma(i,j) \in [-1,1]$ between any two lotteries i and j, and thus not restricting ourselves to the product rule as a form of rational compounding.

Proposition 4 *If the probability of a compound lottery is determined according to equation 1 and* $\gamma(H,M) = \gamma(M,L) = \gamma(H,L) = \gamma$*, then* $HM > HL > ML\ \forall\ p \in [0,1]$ *and* $\gamma \in [-1,1]$*, except if simultaneously* $p = 0$ *and* $\gamma = -1$*, in which case HM=HL=ML=0.*

Proof. We can re-write equation 1 to take into account the possibility of positive or negative covariance, and apply this extended formulation to HM, HL and ML:

$HM = p(H + M) + (1 - p)[(H \times M)(1 + \gamma)]$

$HL = p(H + L) + (1 - p)[(H \times L)(1 + \gamma)]$

$ML = p(M + L) + (1 - p)[(M \times L)(1 + \gamma)]$

We now consider all possible cases:

a. $p = 0$ and $\gamma = -1$: then both the additive and the multiplicative terms vanish, so trivially $HM = HL = ML = 0$

b. $p = 0$ and $\gamma > -1$: then $HM = (H \times M)(1 + \gamma) > HL = (H \times L)(1 + \gamma) > ML = (M \times L)(1 + \gamma)$

c. $p = 1$ or $\gamma = -1$: then $HM = H + M > HL = H + L > ML = M + L$

d. $p \in (-1, 1)$ and $\gamma > -1$: $H + M > H + L > M + L$ and $H \times M > H \times L > M \times L$, so trivially $HM > HL > ML$ because of both the additive and the multiplicative term. ■

Proposition 4 tells us that, under the most plausible scenarios, convex combinations of summation and rational compounding agree in predicting that $HM > HL > ML$. There are two exceptions to this. The first one occurs in the limit case where agents believed to have a zero winning probability for all 2-compound lotteries. The second one may occur for idiosyncratic covariances: for example, if the agent believed that $\gamma(H, M) = -1$ while $\gamma(H, L) = \gamma(M, L) = 0$, we would have $HL > ML > HM$. Neither exception seems realistic in our experimental setting. We can now define quasi-intransitive preferences as preferences that are incompatible with the relationship $HM > HL > ML$.

Definition 5 *The preference ordering of a subject i is quasi-intransitive if for i* $\urcorner(HM > HL > ML)$.

References

ANDERSON, J. A. (1998): "Learning Arithmetic with a Neural Network: Seven Times Seven is about Fifty," in *Methods, Models and Conceptual Issues: An Invitation to Cognitive Science, vol. 4*, ed. by D. Scarnborough, and S. Sternberg, pp. 255–299. MIT Press, Cambridge, Mass. and London.

ANDERSON, N. H. (1981): *Foundations of Information Integration Theory.* Academic Press, New York.

Benassi, V. A., and R. L. Knoth (1993): "The Intractable Conjunction Fallacy: Statistical Sophistication, Instructional Set and Training," *Journal of Social Behavior and Personality*, 8, 83–96.

Fantino, E., and H. Savastano (1996): "Humans' Responses to Novel Stimulus Compounds and the Effects of Training," *Psychonomic Bulletin and Review*, 3, 204–207.

Gigerenzer, G. (1996): "On Narrow Norms and Vague Heuristics: A Reply to Kahneman and Tversky," *Psychological Review*, 103, 592–596.

Hertwig, R., and G. Gigerenzer (1999): "The 'Conjunction Fallacy' Revisited: How Intelligent Inferences Look Like Reasoning Errors," *Journal of Behavioral Decision Making*, 12, 275–305.

Hoel, P. G. (1984): *Introduction to Mathematical Statistics*. Wiley, New York.

Mulford, M., and R. M. Dawes (1999): "Subadditivity in Memory for Personal Events," *Psychological Science*, 10, 47–51.

Rottenstreich, Y., and A. Tversky (1994): "Unpacking, Repacking, and Anchoring: Advances in Support Theory," *Psychological Review*, 104, 406–415.

Rumelhart, D. E., R. J. Hinton, and R. J. Williams (1986): "Learning Representations by Back-Propagating Error," *Nature*, 323, 533–536.

Stolarz-Fantino, S., E. Fantino, and J. Kulik (1996): "The Conjunction Fallacy: Differential Incidence as a Function of Descriptive Frames and Educational Context," *Contemporary Educational Psychology*, 21, 208–218.

Taraban, R., and J. M. Palacios (1994): "Exemplar Models and Weighted Cue Models in Category Learning," in *Categorization by Humans and Machines*, ed. by G. V. Nakamura, R. Taraban, and D. L. Medin, pp. 91–127. Academic Press, San Diego.

Tversky, A., and D. Kahneman (1982): "Judgements of and by Representativeness," in *Judgement under Uncertainty: Heuristics and Biases*, ed. by D. Kahneman, P. Slovic, and A. Tversky, pp. 84–98. Cambridge University Press, New York.

Tversky, A., and D. J. Koehler (1994): "Support Theory: A Nonextensional Representation of Subjective Probability," *Psychological Review*, 101, 547–567.

Zizzo, D. J. (2000): *Relativity-Sensitive Behaviour in Economics.* D.Phil. Thesis, University of Oxford.

Zizzo, D. J., and D. Sgroi (2000): *Bounded-Rational Behavior by Neural Networks in Normal Form Games.* Nuffield College Discussion Paper No. 2000-W30, Oxford.

Zizzo, D. J., S. Stolarz-Fantino, J. Wen, and E. Fantino (2000): "A Violation of the Monotonicity Axiom: Experimental Evidence on the Conjunction Fallacy," *Journal of Economic Behavior and Organization*, 41, 263–276.

Address For Correspondence: Daniel John Zizzo, Christ Church College, Oxford OX1 1DP, United Kingdom.

Email: daniel.zizzo@economics.ox.ac.uk

Figure 1. A screen from the practice stage

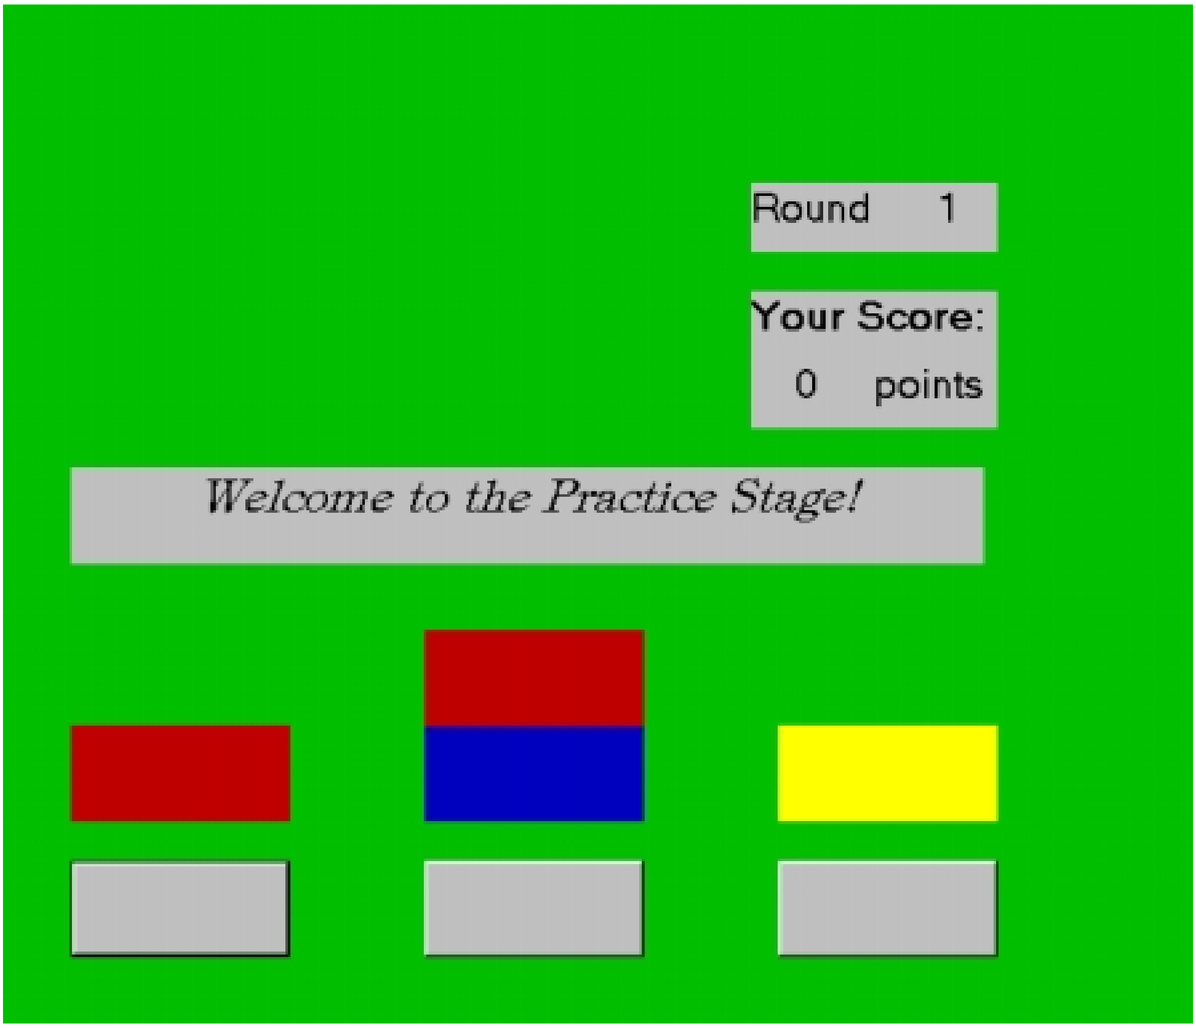

Figure 2. A screen from the testing stage

Remember! Each bet you win in this stage is worth 1 pound, and there are only ten decisions to take!

Figure 3. Fraction of individual choices describable by simple rules (full sample)

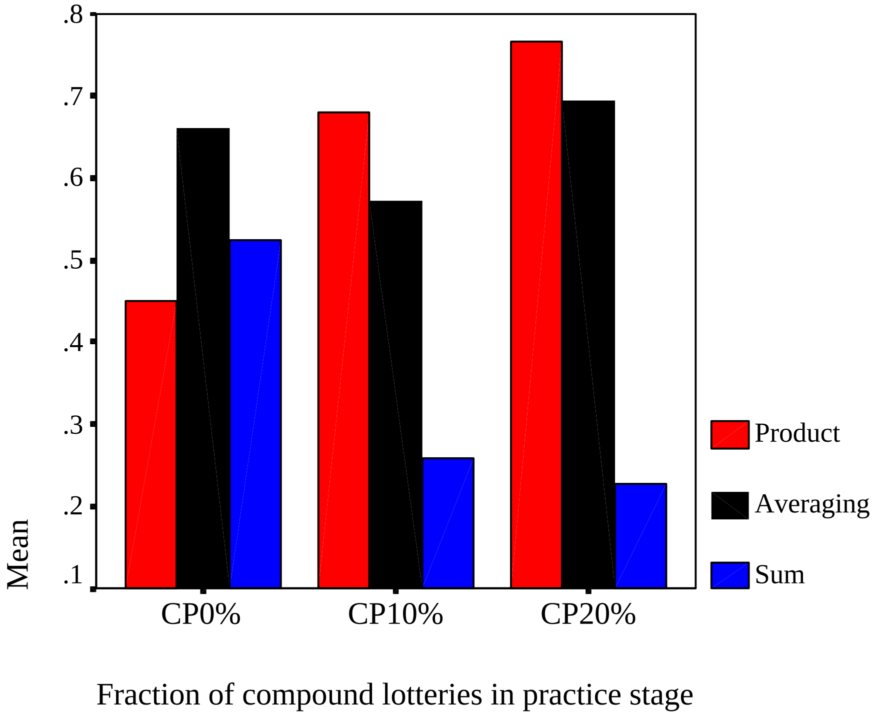

CP0%, CP15% and CP30% correspond to a 0, 15 and 20% fraction of compound lotteries in the practice stage.

Figure 4. The neural network architecture

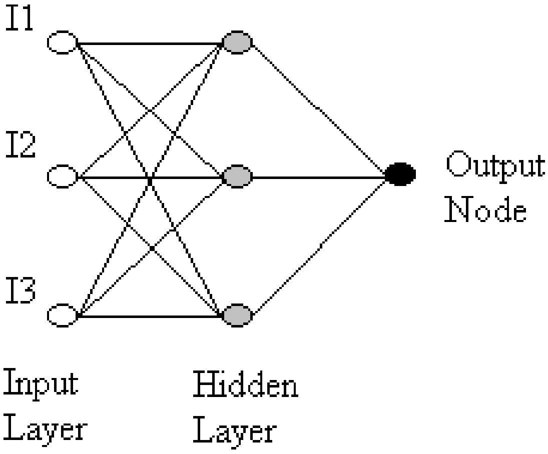

The network receives input at the input nodes, as a Boolean vector of numbers that can be mapped into lotteries: for example, the vector [1 0 0] corresponds to the H lottery, and entails applying a value of 1 to input node 1 (I1) and 0 to input nodes 2 and 3 (I2 and I3). The input gets processed downstream by the neural network, until it produces an output, equivalent to the activation level (between 0 and 1) of the output node. This activation can be interpreted as the probability attached by the economic agent to the lottery.

Figure 5. Performance of rules vs. Tabula Rasa Model (TRM) trained networks

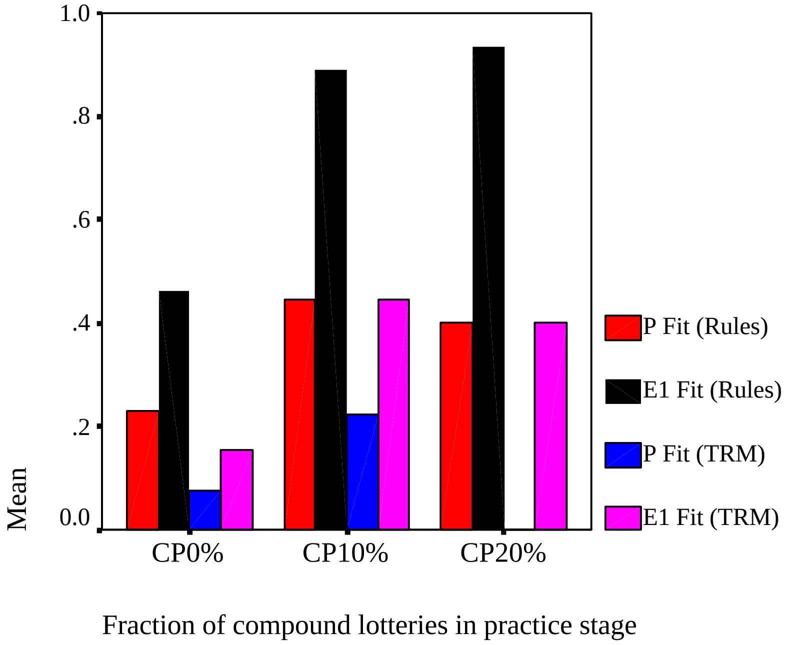

Figure 6. Performance of rules vs. Prior Knowledge Model (PKM) trained networks

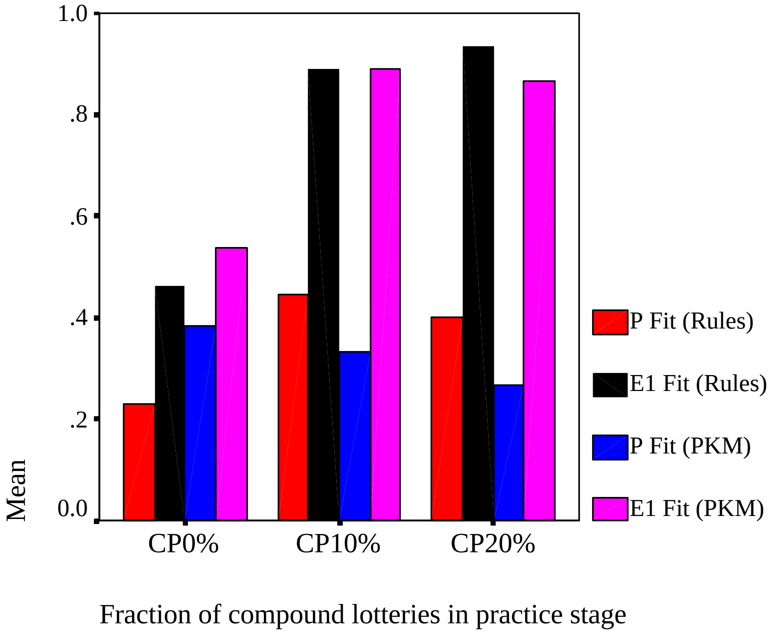

Notes for both figures:

P Fit (Rules): fraction of subjects with perfect fit of RDF rules; E1 Fit (Rules): fraction of subjects for whom the RDF rules fit of at least all but one choice (E1 stands for allowance of 1 error). P Fit (PKM): fraction of subjects with perfect fit by PKM-trained networks; E1 Fit (PKM): fraction of subjects for whom the PKM-trained networks fit at least all but one choice.

Figure 7. Detailed performance of Prior Knowledge Model (PKM) trained networks

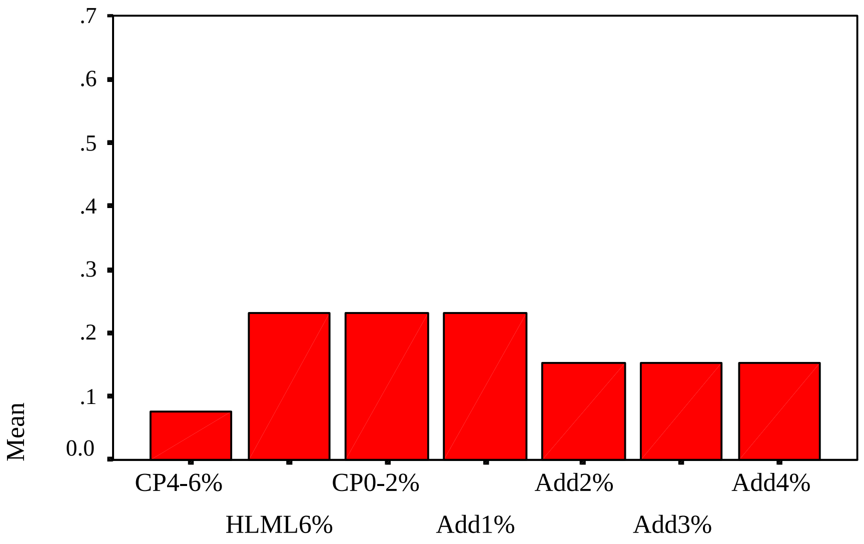

CP0% Condition

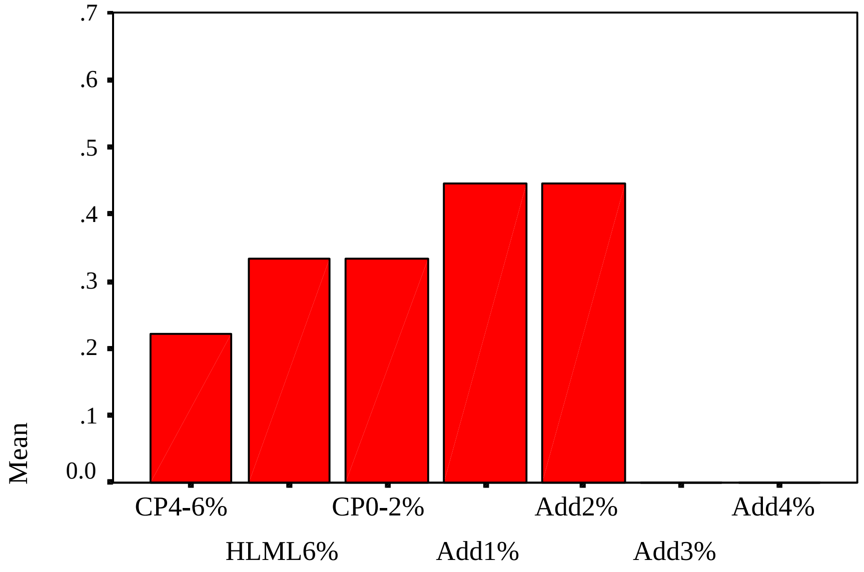

CP10% Condition

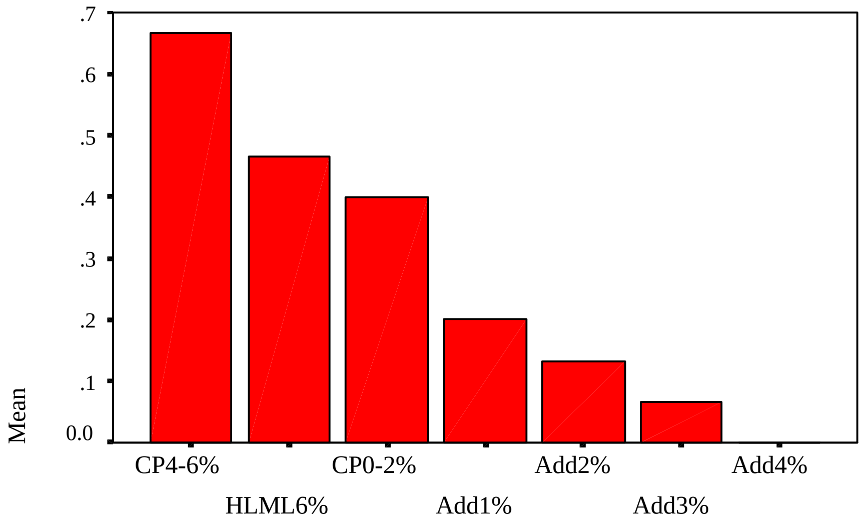

CP20% Condition

Table 1. Fraction of subjects describable by rules in the pure and robust samples

Experimental	Pure Sample				Robust Sample			
Condition	Product	Averaging	Sum	Overall	Product	Averaging	Sum	Overall
CP0%	0.077	0.308	0.154	0.462	0.25	0.45	0.3	0.65
CP10%	0.222	0.667	0	0.889	0.529	0.706	0.059	0.941
CP20%	0.667	0.467	0	0.933	0.714	0.571	0.143	0.952
Total	0.351	0.459	0.054	0.757	0.5	0.569	0.172	0.845

Performance of the Rule Description Framework. CP0%, CP15% and CP30% correspond to a 0, 15 and 20% fraction of compound lotteries in the practice stage. The pure sample includes all subjects with transitive preferences (n=37). The robust sample includes all subjects with transitive preferences over the simple lotteries (n=58).

Table 2. p-Rules Framework

	p value range		Probability	Fraction of subjects fitting p-rule in pure sample			
p-Rule	min	max	ordering restrictions	CP0%	CP10%	CP20%	Total
P1 (sum)	0.693	1	HM>H, M; HL>H, L; ML>M, L; HML>H, M, L	0.154	0	0	0.054
P2	0.6	0.693	HM>H, M; **H>HL**>L; ML>M, L; HML>H, M, L	0.231	0	0	0.081
P3	0.4667	0.6	HM>H, M; H>HL>L; **M>ML**>L; HML>H, M, L	0.154	0	0.067	0.081
P4	0.4286	0.4667	HM>H, M; H>HL>L; M>ML>L; **H>HML**>M, L	0.154	0.222	0.133	0.162
P5 (averaging)	0.2889	0.4286	**H>HM**>M; H>HL>L; M>ML>L; H>HML>M, L	0.154	0.444	0.133	0.216
P6 (averaging)	0.2	0.2889	H>HM>M; H>HL>L; M>ML>L; H, **M>HML**>L	0.308	0.444	0.467	0.405
P7	0.1428	0.2	H>HM>M; H>HL>L; M, **L>ML**; H, M>HML>L	0.231	0.333	0.4	0.324
P8	0.1111	0.1428	H, **M>HM**; H>HL>L; M, L>ML; H, M>HML>L	0.077	0.222	0.533	0.297
P9	0.0769	0.1111	H, M>HM; H>HL>L; M, L>ML; H, M, **L>HML**	0.077	0.333	0.667	0.378
P10 (product)	0	0.0769	H, M>HM; H, **L>HL**; M, L>ML; H, M, L>HML	0.077	0.222	0.667	0.351

p and (1-p) are the weights given to the additive and multiplicative components, respectively, of the compound lottery probability assessment. Any p-Rule *Pw+1* is the same as its adjacent p-Rule *Pw* except for one reversion in the preference ordering: label the lotteries whose probability ranking gets reverted as *pivotal lotteries*. The pivotal lotteries are marked as bold in the table, in correspondence each case of the *Pw+1* lottery. The boundary p-value separating the ranges of any two lotteries always corresponds to the point where the pivotal lotteries are equally probable. Example: P1 and P2 are identical except that in one case HL>H and in the other H>HL; at p=0.693, H=HL.

Table 3. Spearman correlation matrix between rules and Prior Knowledge Model's trained networks' preference orderings

	CP4-6%	HLML6%	CP0-2%	Add1%	Add2%	Add3%	Add4%
Product	**1**	**0.893**	0.679	0.5	0.286	-0.071	0.071
Averaging	0.556	0.815	**0.927**	**0.927**	**0.927**	0.741	0.815
Sum	-0.45	-0.036	0.288	0.468	0.721	**0.883**	**0.847**

The bold numbers correspond to the best describing rule for each network training