Journal of Behavioral Decision Making
J. Behav. Dec. Making, **12**: 55–69 (1999)

How are Base-Rates Used, When They are Used: A Comparison of Additive and Bayesian Models of Base-Rate Use

NATHAN NOVEMSKY* and SHIRIT KRONZON
Princeton University, USA

ABSTRACT

Previous research has uncovered many conditions that encourage base-rate use. The present research investigates how base-rates are used when conditions are manipulated to encourage their use in the lawyer/engineer paradigm. To examine the functional form of the response to base-rate, a factorial design was employed in which both base-rate and the individuating information were varied within-subject. We compared the performance of several models of base-rate use, including a model that allows base-rate and individuating information to be combined in a strictly additive fashion, and a model which presumes that respondents use Bayes' Rule in forming their judgments. Results from 1493 respondents showed that the additive model is a stronger predictor of base-rate use than any other model considered, suggesting that the base-rate and individuating information are processed independently in the lawyer/engineer paradigm. A possible mechanism for this finding is discussed. Copyright © 1999 John Wiley & Sons, Ltd.

KEY WORDS base-rate; additivity; exemplars; Bayes' theorem

Nearly a quarter of a century has passed since Kahneman and Tversky introduced the 'base-rate fallacy'. Their effect can be seen in the following example. Consider a description of Jack:

> Jack is a 45-year-old man. He is married and has four children. He is generally conservative, careful, and ambitious. He shows no interest in political and social issues and spends most of his free time on his many hobbies which include home carpentry, sailing, and mathematical puzzles (Kahneman and Tversky, 1973, p. 241).

Research participants in one condition were told that this description had been randomly drawn from a pool of descriptions that consisted of 70 engineers and 30 lawyers. Participants in another condition were told that the pool consisted of 30 engineers and 70 lawyers. Participants were then asked to assess the probability that Jack is an engineer. Results showed that regardless of the type of base-rate information given, participants produced probability judgments around 90% that Jack is an engineer.

This result produced quite a stir among psychologists. Over the next twenty years, many follow-up studies were conducted to see how robust (or fragile) this result is. Researchers labored to learn when

* Correspondence to: Nathan Novemsky, Department of Psychology, Princeton University, Princeton, NJ 08544, USA. E-mail: novemsky@princeton.edu

CCC 0894–3257/99/010055–15$17.50

Accepted 15 May 1998

and how they could make people pay attention to a simple and relevant piece of information and indeed have uncovered conditions where base-rate use is readily observed. The present research is aimed at the next logical question: Now that we can induce sensitivity to base-rate, how are people using that base-rate information? Before describing the present research we will address some normative and descriptive issues that have arisen in relation to this paradigm.[1]

WHEN ARE BASE-RATES USED?

Since their first lawyer/engineer study, many replications of Kahneman and Tversky's paradigm have been conducted in which participants were asked to assess the probability that an individual belongs to a given category on the basis of a personality description and the relative frequency of that category in the population. These studies uncovered many circumstances that make people more sensitive to the base-rate. Results show an enhanced effect of base-rate when base-rates are varied within-subject (Fischhoff and Bar-Hillel, 1984; see also Fischhoff, Slovic, and Lichtenstein, 1979; Birnbaum and Mellers, 1983); when the base-rate is given a causal interpretation (Ajzen, 1977; Bar-Hillel, 1980; Tversky and Kahneman, 1980; Locksley and Stangor, 1984); when the problem is framed as repetitive rather than unique (Kahneman and Tversky, 1973); when the base-rate is derived from a representative sample (Wells and Harvey, 1977); when the individuating information lacks credibility (Schwarz *et al.*, 1991); when the individuating information lacks diagnosticity (Ginossar and Trope, 1980); when participants bring a scientific orientation to the problem (Zukier and Pepitone, 1984; see also Schwarz *et al.*, 1991); when inferential rules suggesting the use of base-rates have been activated (Ginossar and Trope, 1987); when the base-rate is presented after the individuating information (Krosnick, Li, and Lehman, 1990); when the problem is reframed into frequency instead of probability (Gigerenzer and Hoffrage, 1995); and when participants are shown the random sampling process (Gigerenzer, Hell, and Blank, 1988).

HOW ARE BASE-RATES USED?

Although great progress has been made in illuminating the circumstances under which base-rate information is used, little attention has been paid to the particular form of the response to base-rate in the lawyer/engineer paradigm. In this paper we investigate exactly how people integrate base-rate information into their judgments, but before we do, we will first consider normative guidelines for base-rate usage. We will then examine a descriptive theory.

Normative model of base-rate use

Probability theory offers compelling normative guidelines for how base-rate and individuating information should be combined. Specifically, Bayes' rule in odds form states that the posterior odds are equal to the product of the prior odds and the likelihood ratio. In the present context, the likelihood ratio represents the extent to which the description favors the target profession over the other profession in the pool. The prior odds are constructed from the base-rate. The posterior odds correspond to the judgment made by participants that the target individual belongs to the target profession. In Exhibit 1, the graphs implied by Bayes' rule are plotted for different likelihood ratios

[1] Of course, many studies have looked at other base-rate problems, such as the cab problem (Tversky and Kahneman, 1980) and the medical diagnosis problem (Eddy, 1982). The present paper focuses on the lawyer/engineer problem, where only the base-rate information is presented numerically, while the individuating information is given in the form of a personality description. In other base-rate problems, participants are faced with the task of combining multiple pieces of numerical information. This distinction will become important when we consider possible mechanisms for base-rate use.

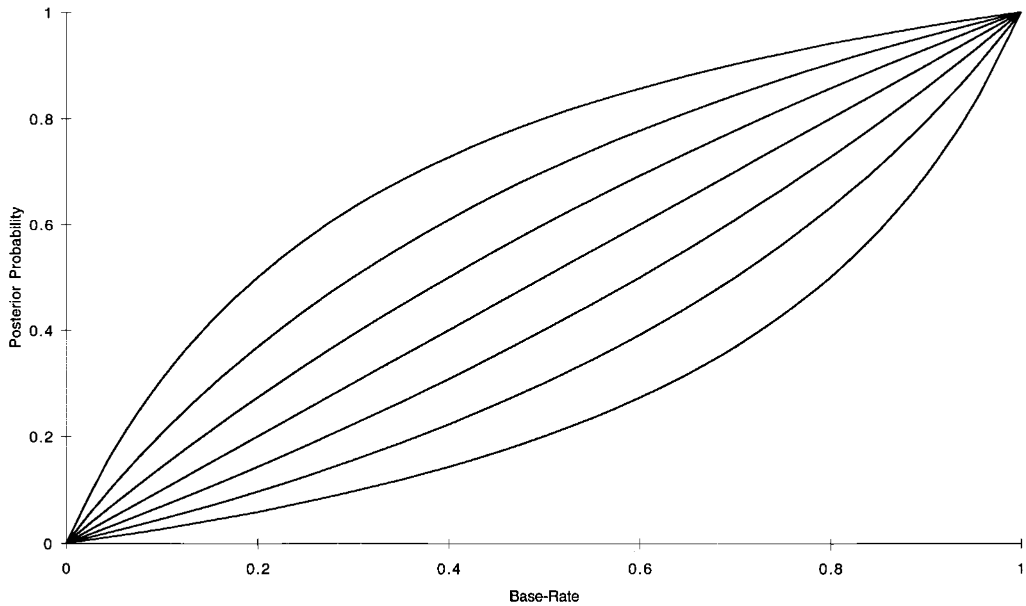

Exhibit 1. Bayesian model. Posterior probabilities are plotted as a function of base-rate. The likelihood ratios from top to bottom are 4, 3/2, 7/3, 1, 3/7, 2/3, and 1/4

using Bayes' rule in probability form, with base-rate on the horizontal axis and posterior probabilities on the vertical axis. In this format, we can see a strong interaction between the base-rate and the likelihood ratio, i.e. the distance between the curves varies as the base-rate changes. If probability theory is descriptively accurate, judgments in the lawyer/engineer paradigm should show a multiplicative interaction between base-rate and individuating information.

Comparing the original Kahneman and Tversky (1973) study to Bayes' rule, it is clear that participants' judgments are not compatible with the normative approach provided by probability theory. According to Bayes' rule, two factors should influence the probability that Jack is an engineer. The first factor is the likelihood ratio, a measure of the probability that Jack's profile is that of an engineer as opposed to a lawyer. Bayes' rule puts no constraints on the value of the likelihood ratio. The second factor is the base-rate. Bayes' rule requires that participants provide a higher probability in the '70 engineer' condition as compared to the '30 engineer' condition. In fact, the odds of Jack being an engineer should be ($\{70/30\}/\{30/70\} =$) 5.44 times higher in the '70 engineer' condition. Since participants gave only slightly different responses in the presence of different base-rates, their judgments are clearly not predicted by Bayes' rule.

Several researchers have questioned the normative role of Bayes' rule in the lawyer/engineer paradigm (see Koehler, 1996). A complete discussion of this issue is beyond the scope of this paper. The present research is not attempting to define what people should do; rather, the goal is to describe the pattern of actual responses in the lawyer/engineer context. To date, when differences in responses between different base-rate conditions have been observed, researchers have suggested that participants' responses are consistent with Bayes' rule (e.g. Gigerenzer, Hell, and Blank, 1988).[2] In this study we examine this assumption.

[2] Gigerenzer, Hell, and Blank (1988) did compare a complete base-rate neglect model to a Bayesian model and found that responses were slightly closer to complete neglect when the sampling of descriptions was stated, as usual, in the questionnaire. Responses shifted slightly closer to the Bayesian model when the sampling was done in front of respondents.

One normative issue that is relevant to this paper involves the prior odds used in the Bayesian analysis. In particular, the application of Bayes' rule that treats the stated base-rate as the prior probability has been contested: Koehler (1996) has claimed that research participants in the lawyer/engineer paradigm may have prior probabilities that do not equal the base-rate. It is important to note that only two pieces of information are presented in the lawyer/engineer paradigm: the description and the base-rate. In the Bayesian framework, people are presumed to have prior odds and a new piece of relevant information with which to update those odds. The new piece of information, in this case the description, is used to calculate the likelihood ratio. All that is left to construct the prior odds is the base-rate. For this reason, the application of Bayes' rule to the lawyer/engineer paradigm treats the base-rate as the prior probability. To be sure, in many real-world judgments there is certainly information other than a base-rate that may influence one's prior probability. However, this paper describes judgments involving base-rates when there is no other 'prior' information present.

Descriptive model of base-rate use

While probability theory offers a compelling normative rule for how people *should* use base-rate information, this rule does not necessarily reflect how people *do* use base-rate information. Gigerenzer, Hell, and Blank (1988) have linked base-rate use with Bayesian reasoning in the lawyer/engineer paradigm. In fact, they have defined an index of base-rate use that is constructed by comparing deviations from complete base-rate neglect with deviations from Bayesian reasoning. In this paper we compare the performance of several models of base-rate use, with the aim of finding the most accurate and parsimonious model available. Bayes' rule is one of several possible models that will be tested.

An important descriptive model of the integration of multiple sources of information stems from work in many areas of psychology. For over twenty years, researchers studying areas ranging from psychophysics to perception to social judgment, have successfully modeled the integration of multiple sources of information in the production of judgments using an additive model. For instance, McBride (1989) showed that people integrate taste and smell additively to produce judgments of intensity. Anderson and Cuneo (1978) showed that children's perceptions of area can be described by the simple rule: area = height + width.[3] Anderson and Butzin (1978) showed that achievement and need are integrated using an additive model to predict judgments of deservingness. In fact, Anderson and his colleagues have developed an influential framework called information integration theory where additivity seems to flourish in many judgment tasks (Anderson, 1980, 1981, 1996).

It is worth noting that Anderson and his colleagues found adding rules in tasks in which a multiplying rule would be expected: the Height + Width rule for the perception of area (Anderson and Cuneo, 1978), the Length + Density rule for judged numerosity of a row of beads (Cuneo, 1982), the Diameter + Height rule for judgments of cone volume (Wilkening, 1980) and the Subjective Probability + Utility rule for judgments in a roulette game (Anderson, 1980). To explain these effects, Anderson suggested that adding rules 'reflect basic and natural modes of information integration' (Anderson, 1981, p. 34). He later wrote that the appearance of an adding rule where multiplication would be expected is taken as evidence that there exists a general purpose integration operator that is additive. This operator is activated when people recognize the relevance of several pieces of information, but do not have a well-developed rule for combining them. Anderson even gave this operator the status of a biological capability, stating that it provides moderately accurate results even in situations where the true integration rule is multiplicative (Anderson, 1996). We can think of Anderson's operator as a mental shortcut or 'additivity heuristic' which provides easier and more natural computations than a multiplicative operation, but yields moderately accurate results.

[3] See Gigerenzer and Richter (1990) for a critique of this study.

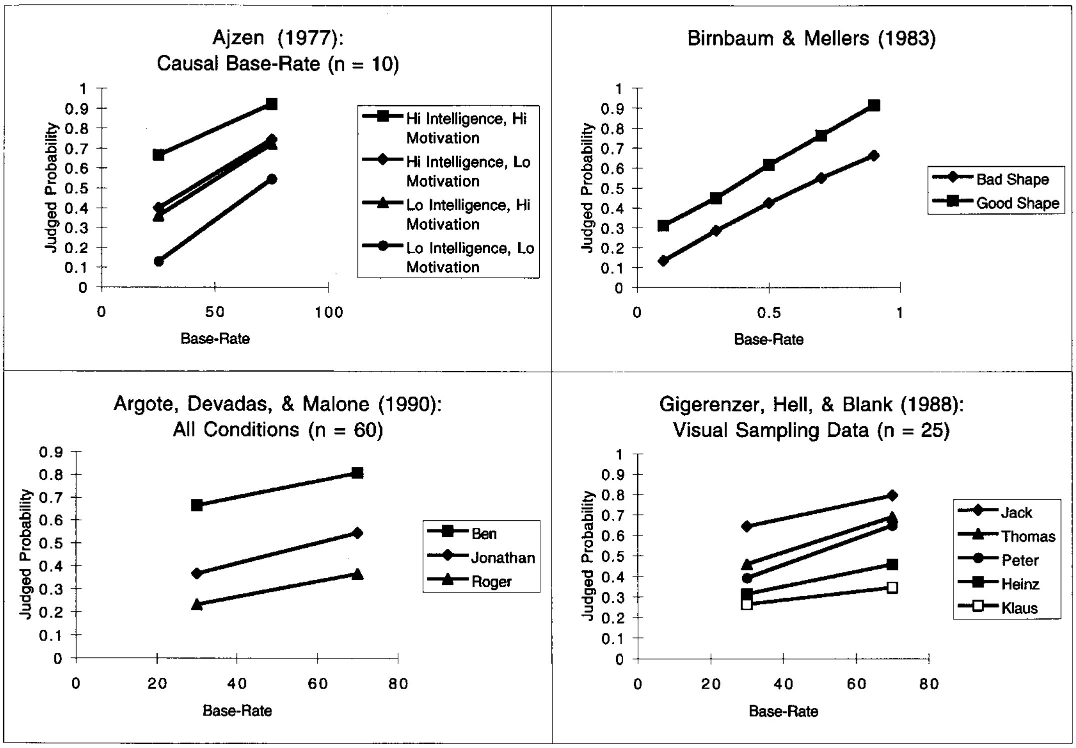

Exhibit 2. Base-rate effects in past studies

In this paper, we will examine additivity in the context of the lawyer/engineer paradigm. It is important to note that when we refer to an 'additive' model, we are refering to a linear model that is the sum of two terms, where each term arises from a distinct attribute of the judgment object (e.g. base-rate or individuating information). This linear combination could be a simple sum or average, or some other linear combination. The relevant aspect of the general additive model to the present research is that it does not contain a multiplicative interaction.

To see if additivity has been observed in the lawyer/engineer paradigm, we looked to past studies that manipulated individuating and base-rate information factorially.[4] We then re-examined these data to see if a clear pattern characterizes the observed judgments. The results from several of these studies are shown in Exhibit 2. Mean probability judgments are plotted along the ordinate, while base-rates are displayed on the abscissa. A distinct pattern emerges in these data: The difference between the mean judgments for different levels of the qualitative variable appears to be the same at both the low base-rates and high base-rates. Thus, the difference in judged probabilities for different descriptions appears to be independent of base-rate. Birnbaum and Mellers' (1983) data shown in the third graph of Exhibit 2 are particularly relevant to the present investigation. They used more than two base-rates which allow us to look for a non-linear response pattern. As can be seen, in Exhibit 2, their results are well approximated by a strictly linear pattern.

This pattern of results suggests that people may integrate the individuating and base-rate information using not a Bayesian process, but an additive operator. An additive model predicts, for example,

[4] This restriction left us with a surprisingly small number of experiments to consider.

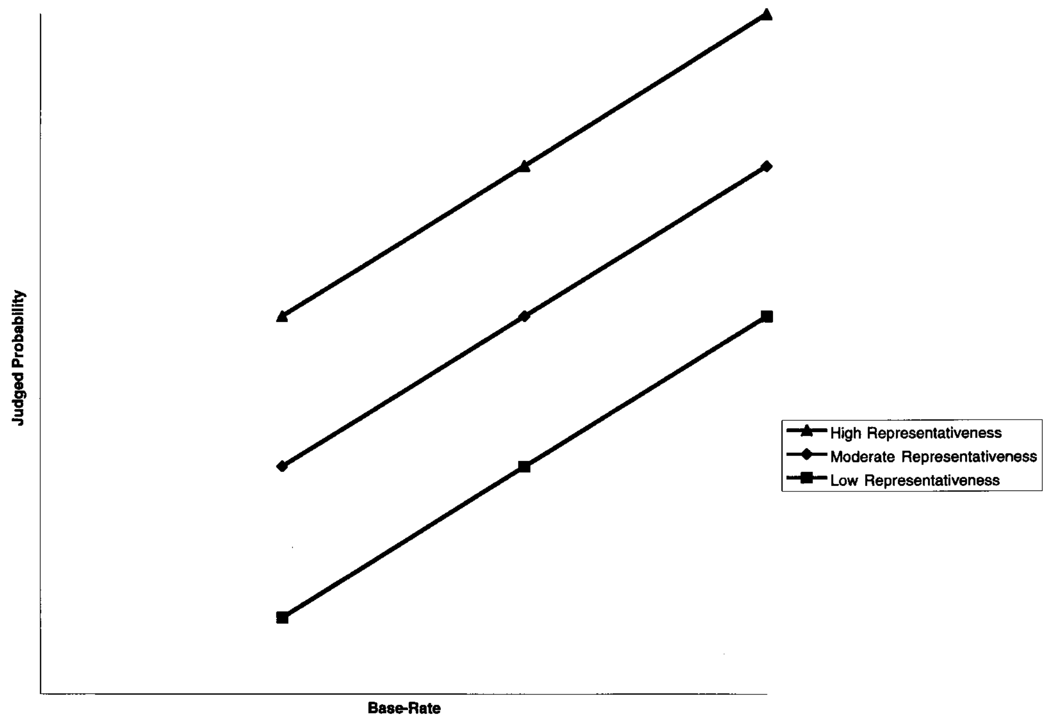

Exhibit 3. Additive model

that a change in base-rate from 10% to 50% will have the same effect on participants' final probability judgments, regardless of the likelihood ratio of the description. In other words, all descriptions should get the same probability boost when the base-rate changes from 10% to 50%. The additive model does not predict the interaction seen in Exhibit 1 but, instead, predicts the parallel lines shown in Exhibit 3.

To summarize, recent research has focused on finding exceptions to complete base-rate neglect. As a result, past studies were often designed to show differences between conditions with different base-rates. These studies do not allow an examination of the quantitative form of the base-rate effect. However, most descriptions of these experiments take the response to base-rate, at least implicitly, as evidence of a normatively appropriate judgment (see Gigerenzer Hell, and Blank, 1988 for an explicit discussion of Bayesian responding). The present research compares several descriptive models to find the model that best characterizes the response to base-rate. We expect the additive model to provide a better description of individuals' response to base-rate in the lawyer/engineer paradigm than a model derived from Bayes' rule.

The present study used a factorial design to examine the functional form of the response to base-rate. Factorial designs have not been widely used in the lawyer/engineer paradigm, possibly because the researchers were more interested in finding conditions that elicit base-rate effects than in determining the exact form of the integration rule. The method of using a factorial design to map out integration rules when multiple sources of information are present was initiated by Anderson (1962), who conducted research to determine the integration rule that governs the synthesis of several adjectives into a likability judgment of a person. More recently, Massaro and Friedman (1990) has stressed the importance of the factorial design in determining how multiple sources of information are combined in the domain of visual perception.

To examine the form of the response to base-rate information, we created a situation that would induce respondents to use that information. Towards this end, the base-rate was varied within-subject (Fischhoff and Bar-Hillel, 1984) and presented after the individuating information (Krosnick, Li, and Lehman, 1990). To reduce the impact of the individuating information, each description consisted of only four adjectives instead of the usual full paragraph description. These four adjectives were not attributed to a professional clinician, but were said to have been checked as self-descriptive by the target individual. With these conditions in place, participants would be encouraged to use the base-rate, allowing the manner of this use to be examined.

METHOD

Participants

Participants were 1493 undergraduates sampled approximately equally from four major universities in southern California and the Northwest. Participants were paid to fill out a one-hour questionnaire, in which the present experiment was embedded.

Materials and procedure

The materials used in this study were modeled after Kahneman and Tversky's (1973) lawyer/engineer study. Research participants were asked to assess the probability that an individual for whom a personality profile was given belonged to a specified profession. All participants were presented with the same six personality profiles. Each profile had supposedly been drawn at random from a pool of profiles containing members of a target profession and one other profession. The profile consisted of four adjectives that the target individual was said to have checked as self-descriptive. After each profile, participants were given base-rate information: they were told that the profile shown was randomly selected from a pool of 100 profiles, a certain number of which were from the target profession. Participants were then asked to judge the probability that the person described is a member of the target profession. A sample question is shown below:

> Accountant or real-estate agent?
> Julie is persuasive, attractive, intelligent, and careful. The above description was drawn at random from a pool of 10 accountants and 90 real-estate agents, all women.
> On a scale from 0% to 100%, what is your probability that Julie is one of the accountants in the pool? ____%.

Independent variables

Representativeness

The first independent variable is the representativeness of the profiles. In this experiment, representativeness was manipulated by varying the fit of the set of four adjectives with the corresponding target professions. Some profiles were more representative of their target profession than others. Each profile was presented beneath a heading that listed two professions, a target profession and an alternative profession. In the example above, the target profession is 'accountant' and the alternative profession is 'real-estate agent'. For each profile, half of the participants were asked about one profession, and the other half were asked about the other profession. Thus, half of the participants stated their probability that Julie is an accountant and half of the participants gave a probability that Julie is a real-estate agent.

Six personality profiles were used: (1) shy, serious, organized, and sarcastic (computer programmer/flight attendant); (2) loud, energetic, colorful, and disorganized (high school coach/dentist); (3) argumentative, flashy, self-confident, and competitive (lawyer/engineer); (4) persuasive, attractive, intelligent, and careful (real-estate agent/accountant); (5) precise, orderly, independent, and tough-minded (surgeon/architect); (6) quiet, prompt, diligent, and conservative (librarian/lab technician).

Base-rate

The second independent variable was the base-rate. Each personality profile was paired with each of three base-rates. For each profile, one third of the participants saw a base-rate of 10 members of the target profession, another third saw a base-rate of 50 members of the target profession, and the last third were given the base-rate of 90 members of the target profession. The base-rate was varied in such a way that every participant saw each base-rate twice.

For the purpose of analysis, the design can be conceptualized as a three-cell between-subject design with 12 replications. The three cells correspond to the three base-rates, 10, 50, and 90. The twelve replications were produced by considering separately each of the six profiles with its two corresponding professions. In other words, there were six profiles, and two questions that could be asked about each profile, giving us a total of twelve replications. Each replication consists of three sets of judgments, one set for each base-rate.

RESULTS

The mean judged probabilities for each base-rate are plotted in Exhibit 4. Each line represents one profession. We used six descriptions, each paired with two professions for a total of twelve professions.

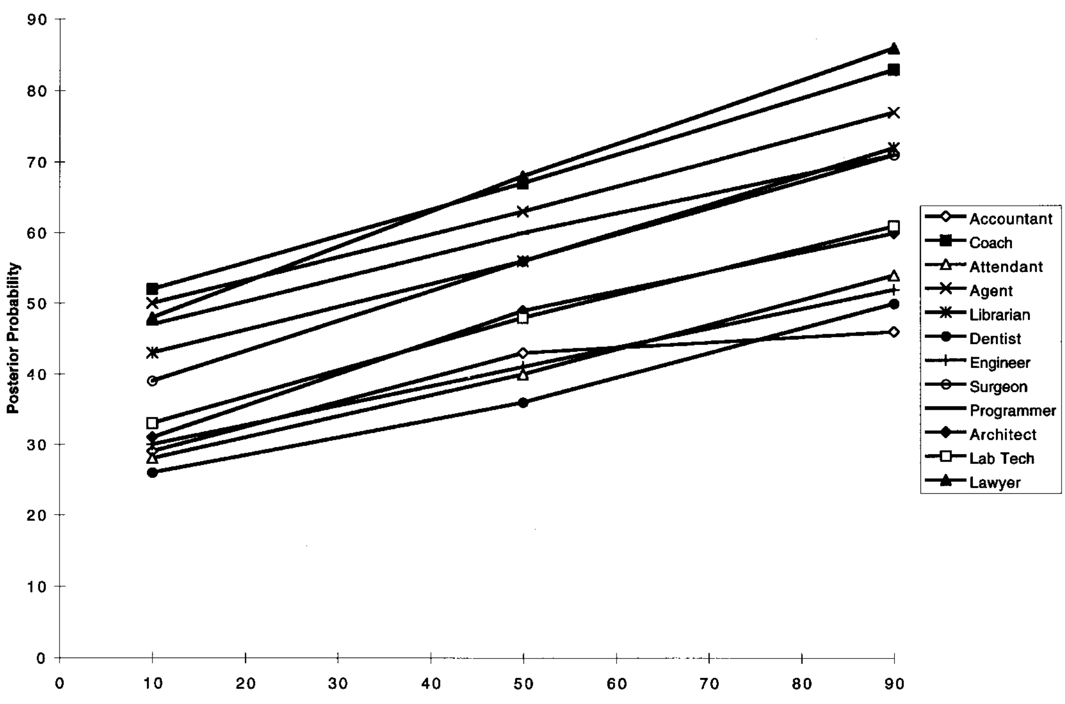

Exhibit 4. Mean judged probability as a function of base-rate

We plotted one line for each profession as a function of the base-rate. Looking at the graph, there seems to be an additive pattern emerging, characterized by parallel lines. As noted in the introduction, this pattern could be caused by any weighted average that combines the individuating information and base-rate. A sum, an average, or another weighting of the two types of information would yield parallel lines. The parallel result is incompatible with any interaction between the individuating information and the base-rate (such as the interaction contained in Bayes' rule). The analyses that follow will quantitatively compare the actual responses to several models described below.

To measure the fit of each model, the root mean square residuals (RMSR) were calculated. The RMSR is the square root of the average squared difference between the predicted and observed values. The larger the RMSR, the worse the fit of the model. The dataset consists of 3 cells (one for each base-rate) times 12 replications (i.e. professions), or 36 cells of probability judgments. To calculate the RMSRs for each model, each replication was treated separately. In other words, the RMSR for a given model was calculated by considering the three cells (10%, 50%, and 90% base-rates) in a replication, and pooling the RMSR across those three cells. For instance, three questions about accountants appeared, one for each of the three base-rates. For each model, a score was generated by taking the average RMSR across the three cells that refer to accountants. Each model had twelve RMSRs, one for each target profession.

Since all responses to a single question (e.g. 'What is the probability that Julie is an accountant when the base-rate is 10?') are not identical, a model that has a fixed prediction for that question cannot predict all responses without error. For this reason, the RMSR cannot be zero in the presence of variability within any single cell in the design. This fact implies that there is a non-zero lower bound to the RMSRs for all models of the present data. The smallest RMSR is produced by a model that predicts exactly the cell mean for every cell in the design. This model will be referred to as the cell means model and its RMSR will provide a yardstick with which to measure the RMSRs of the other models. To facilitate comparison of the models, all twelve RMSRs were placed on a common scale. Specifically, each RMSR for a given replication was divided by the corresponding cell means model RMSR for that replication. Because the cells means model provides the lowest RMSR for any replication, the resulting normalized RMSRs are then greater than or equal to one. A normalized RMSR equal to one indicates the best possible fit a model could achieve. The further from unity the normalized RMSR falls, the worse the fit of the model.

We computed the RMSRs on individual data rather than on group means for two reasons. First, using individual data, we need not assume anything about the form of the distribution of responses. Outliers and non-normality in the set of responses would be masked by just retaining the means. Second, means represent a summary of the data. The models under consideration are expected to provide a parsimonious and accurate summary of the set of individual responses. If we discarded individual data and modeled the means, the models would be trying to provide a parsimonious account, not of the dataset, but of a smoothed summary of the dataset. Now that we have explained how we measured the fit of each model, we describe how we constructed each model of the data.

Bayesian model

To construct the Bayesian model we used individual judgments from the 50% base-rate group. The 50% base-rate condition was chosen because that base-rate does not favor either profession, thus providing a judgment that is least affected by the base-rate, and most representative of the likelihood ratio. Using this condition to construct the model avoids the problem of assuming anything about the form of the base-rate response while testing the form of that response. The judgment of each individual in the 50% base-rate group was used to infer two new judgments; assuming that individuals obeyed Bayes' rule, we generated judgments for the 10% and 90% base-rate. For example, an individual who answered 80% in the 50% base-rate group would have a likelihood ratio of 4. Using this likelihood ratio and the 10% base-rate we can determine what this individual would have answered in the 10% base-rate condition if she were responding in a manner consistent with Bayes' rule. Using this transformation, we can generate two sets of judgments that best represent what the 50% group would have answered had they been assigned to the other two base-rate conditions. For the 'accountants' replication, we transformed all responses to the question that asked about a pool of 50% accountants. This produced a set of inferred judgments for the 10% accountants question and a set of inferred judgments for the 90% accountants question. The mean of the inferred 10% judgments was used to predict actual responses for those who received the 10% accountants question. Similarly, the inferred 90% accountants judgments were used to predict actual responses to the 90% accountants question. The 50% base-rate condition was predicted by the mean of those actual responses. This calculation was performed independently for each of the twelve replications.

Using the prediction curves for all twelve replications, a residual score was calculated for each of the 8958 responses in the experiment. These residuals were then squared, and an average was calculated for each of the twelve replications. The square roots of these averages provide the twelve RMSRs for the Bayesian model. These RMSRs were then normalized, as described above to provide twelve measures of fit displayed in Exhibit 5.

Additive model

Uniquely determining the additive model requires a slope and a level of representativeness. We used a single slope to model all twelve replications; this slope was the overall slope of judged probability as a function of base-rate using all 8958 responses. The level of representativeness was inferred from the 50%

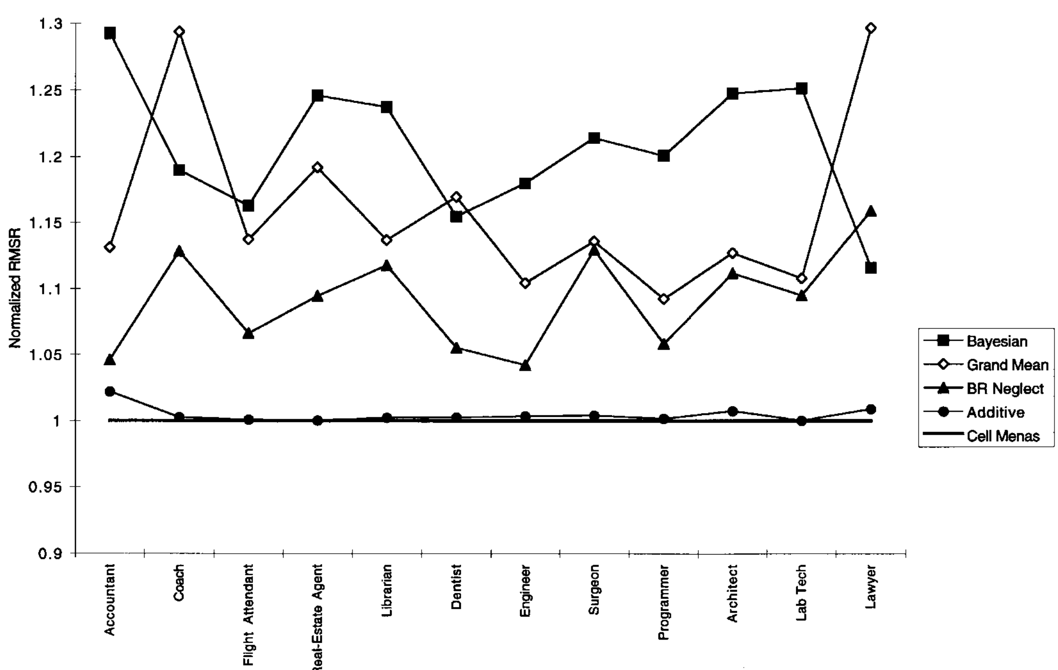

Exhibit 5. Normalized root-mean-squared-residuals

base-rate condition, again, because this condition provides a judgment with the least influence of the base-rate. Using the mean of the 50% base-rate condition and the overall slope, a straight line was constructed for each of the twelve replications. For example, the predictions for the accountants replication were constructed using the mean of the 50% accountants responses and the slope generated by combining all twelve replications; the architect replication used the mean of the 50% architects responses and the same overall slope as the other replications. These lines were then used to generate predictions for the 10%, 50%, and 90% base-rate groups. From these predictions, a residual score was calculated for each response. As with the Bayesian model, these residuals were combined into twelve RMSRs, one for each of the twelve replications. These twelve RMSRs were then normalized (see Exhibit 5).

Comparison of the additive and Bayesian models

To facilitate the comparison of these two models, three other models are also displayed in Exhibit 5. The first model, the cell means model, predicts participants' responses on the basis of the observed mean for each replication and each base-rate. The second model, the grand mean model, makes the simple prediction, 51.9% (the grand mean) for all responses. This model represents the expectation that neither representativeness nor base-rate had any effect on participants' judgments. The final model is the complete base-rate neglect model. For a given profession, the prediction of this model is the mean response for the 50% base-rate condition in that profession. This model assumes an effect of individuating information, but no effect of base-rate on judgment.

In Exhibit 5, along the ordinate, the normalized RMSR is plotted separately for each of the twelve replications. Displayed along the abscissa are the twelve replications used in this study.

Of the four models competing to match the cell means model, the additive model is clearly the closest to the cell means model. In fact, the RMSRs of the additive model differ on average by less than half a percent from the RMSRs of the cell means model. By contrast, the next closest model, the complete base-rate neglect model, has RMSRs that are 9.2% larger than those of the cell means model. This model can be statistically distinguished from the additive model, $t(11) = 8.65$, $p < 0.0001$. This difference between the additive model and the complete base-rate neglect model suggests that the additive model is not simply capturing a failure to draw participants' attention to the base-rate. If participants were not using base-rate, the additive model and the complete base-rate neglect model would yield identical predictions.

The next best-fitting model is the grand mean model. In predicting no differences between any cells in the design, this model clearly loses a large part of the variance captured by the earlier models that allow for effects of the individuating and base-rate information. Because the grand mean model can be generated with no knowledge of the experimental hypotheses or procedures, any model that is claimed to fit the data should provide smaller residuals than this model.

Surprisingly, the Bayesian model provides the worst fit of any of the models shown in Exhibit 5, with RMSRs 20.8% larger than the cell means model. This deviation is nearly 50 times worse than the deviation of the additive model. Clearly, participants were not using Bayes' rule when constructing their judgments of probability, but instead were doing something best described by the additive model.

DISCUSSION

There has been a great deal of research since Kahneman and Tversky wrote in 1973 that base-rates are 'largely ignored' (p. 242). Most of this research focused on conditions and presentational formats that induce participants to be sensitive to the base-rate. Now that there are a few tools at our disposal to induce base-rate use, the question arises of exactly how people use the base-rate. To answer this question, we attempted to provide the optimal conditions for the use of base-rate by within-subject variation of the base-rate, by presenting base-rate information after individuating information, and by using only four adjectives to minimize the influence of the individuating information. The present study provides strong evidence that Bayes' rule does not accurately describe how people use base-rate information in the lawyer/engineer paradigm. Actual base-rate use seems to map best onto an additive model. The remainder of this section describes a mechanism that may explain this additive result.

Additivity has often been taken as evidence that two processes are operating independently. In line with this idea, the present results suggest that the base-rate and individuating information are being processed independently. A specific mechanism that may explain this finding has recently been put forth by Kahneman and his colleagues (Kahneman, 1995; see also Schreiber and Kahneman, in press). Kahneman discusses the evaluation of the properties of sets of objects. He delineates two distinct types of characteristics used to describe sets. Sets have characteristics that can refer to the attributes of individual members of the set. For example, we can say that a particular bird has colorful feathers, has a large wing span, or has a beak. These statements refer to an individual bird. These 'intensional' attributes describe the type of instance that is included in the set. Sets also have characteristics that refer not to the properties of individual members, but to properties of the set as a whole. These are the 'extensional' attributes of the set. For example, birds are more numerous in rural areas than in cities, there are over 10 million birds in the USA, and there are more birds in northern regions in the summer than in the winter, are all examples of extensional attributes. These statements are not about any one bird, but about the set of birds. They are quantitative statements about the group of instances that compose the set.

Research suggesting that representations of many sets and categories are based on exemplars and/or prototypes has a long history. In particular, some researchers have examined the roles of these two representations in the construction of social categories (Smith and Zarate, 1990). For the purposes of this discussion, we only need to assume that either exemplars or prototypes or both play an important role in the representation of categories. The most important property of exemplars and prototypes is that they have the dimensionality of a single instance of the category. In other words, exemplars and prototypes are representations of the intensional information without the extensional information. For simplicity, this discussion will use the term exemplar to refer to the internal representations that may be exemplars, prototypes or some combination of the two.

When individuals are required to make judgments about sets, they can extract their intensional information directly from the exemplar, but they must retrieve extensional information by some other process. The extensional information could be obtained by counting exemplars, judging how easily exemplars come to mind (Tversky and Kahneman, 1973, see also Schwarz *et al.*, 1991), assessing the variation or fuzziness associated with an exemplar, or other processes. Whatever the process, assessing extensional information requires going beyond the natural intensional representation and using a qualitatively distinct process.

Kahneman and his colleagues have found support for the representative exemplar idea in a several judgment domains. They studied hedonic judgments of temporally extended episodes (Varey and Kahneman, 1992: Kahneman *et al.*, 1993; Schreiber and Kahneman, in press). These episodes can be thought of as a set of moments. Summary hedonic judgments of these episodes have been found to neglect the duration of the episode (i.e. the extensional information), while particular moments (the representative exemplar) had a large impact on the summary judgment. When duration information was made salient, it seemed to be used in an additive fashion (Schreiber and Kahneman, in press). In a different domain, Kahneman and Ritov (1998; see also Desvousges, 1992) found that judgments of willingness to pay to save sets of animals threatened by environmental problems neglected the size of the set. When the size of the set was made salient, this information contributed to willingness to pay in an additive manner. As Anderson (1996) suggests, when multiple sources of information are deemed relevant to a judgment but there is no well-developed mechanism for combining this information, an additive judgment is likely to result.

The judgments people made in the present study can be conceived of as judgments about sets of professionals. The extension of the set is the percentage of engineers in the pool of descriptions. The intensional information about the set is contained in participants' own exemplars or prototypes of engineers. Kahneman would assert that these two types of information are used quite differently in assessing the probability at hand. The first component, the intensional component, is derived from a judgment of the fit between exemplars of lawyers and the four adjectives included in the description. When participants are asked for a judgment in the present study, they form a representation based on the four adjectives and compare this representation to their existing exemplars of engineers. The other component of participants' judgments is the extensional component. The set of engineers included in the pool of profiles had a certain size described by the base-rate. Participants do not include this extensional information in their exemplars of engineers, but instead must rely on some other process to assess this information.

We can gain some insight from recasting the base-rate results from the past twenty years in terms of representation by exemplar. In the original Kahneman and Tversky (1973) study, research participants used the intensional information that came naturally from an exemplar representation. Given no other prodding or favorable conditions, they made little attempt to incorporate the extensional information. In the within-subject manipulation of base-rate (Fischhoff and Bar-Hillel, 1984) participants' attention is drawn to base-rate. This induces them to make sense of the base-rate information and incorporate it into their judgments. Similarly, when participants are shown the random sampling process (Gigerenzer, Hell, and Blank, 1988), their attention is drawn to the pool of descriptions. This again induces participants to go the extra mile and process the extensional information. The additive pattern observed can be thought of as the result of two separate processes. The individuating information is incorporated into the judgment using the natural exemplar representation. The base-rate information, however, is not encoded in this representation and must be incorporated through some other process.

It is important to note that the present findings are limited only to the lawyer/engineer paradigm. The exemplar mechanism and the additive results may not necessarily generalize to all base-rate problems. When people are faced with several pieces of numerical information (e.g. hit rates and false alarm rates) and no personality description, it seems unlikely that they would use representative exemplars in the manner described above. Our study suggests that participants form an exemplar when individuating information is given as a personality description and only the base-rate information is given numerically.

While this account closely matches the data, more research is needed to bolster the concept of representative exemplars. While direct support for this mechanism has been found in areas such as willingness to pay judgments (Kahneman and Ritov, 1998), more research is needed to delineate when representation by exemplar is guiding judgment. Exactly how the extensional information is incorporated into judgments is also an important question for future research. Any process that integrates this information requires going beyond the natural intensional representation and using a method of assessment qualitatively distinct from the assessment of intensional information. Since extensional processing is absent or extremely limited under some circumstances, it may require more effort than the process of using an exemplar to make a judgment. It would be interesting to see if manipulations which produce more superficial processing and inhibit effortful processes in other contexts (such as cognitive load) would affect the use of extensional information. Whatever this process may be, the present research demonstrates that its ultimate effect on judgment is best modeled by an additive combination of intensional and extensional information.

ACKNOWLEDGEMENTS

We express our great appreciation to Daniel Kahneman and Rich Gonzalez who provided us with constructive suggestions at various stages of this research.

REFERENCES

Ajzen, I. 'Intuitive theories of events and the effects of base-rate information on prediction', *Journal of Personality and Social Psychology*, **35** (1977), 303–314.

Anderson, N. H. 'Application of an additive model to impression formation', *Science*, **138** (1962), 817–818.

Anderson, N. H. 'Information integration theory in developmental psychology', in Wilkening, F., Becker, J. and Trabasso, T. (eds), *Information Integration by Children* (pp. 1–45), Hillsdale, NJ: Erlbaum, 1980.

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

Anderson, N. H. *Methods of Information Integration Theory*, New York: Academic Press, 1982.

Anderson, N. H. *A Functional Theory of Cognition*, Mahwah, NJ: Erlbaum, 1996.

Anderson, N. H., and Butzin, C. A. 'Integration theory applied to children's judgments of equity', *Developmental Psychology*, **14** (1978), 593–606.

Anderson, N. H. and Cuneo, D. O. 'The height + width rule in children's judgments of quantity'. *Journal of Experimental Psychology: General*, **107** (1978), 335–392.

Argote, L., Devadas, R. and Melone, N. 'The base-rate fallacy: contrasting processes and outcomes of group and individual judgment', *Organizational Behavior and Human Decision Processes*, **46** (1990), 296–310.

Bar-Hillel, M. 'The base-rate fallacy in probability judgments', *Acta Psychologica*, **44** (1980), 211–233.

Birnbaum, M. H. and Mellers, B. A. 'Bayesian inference: combining base rates with opinions of sources who vary in credibility', *Journal of Personality and Social Psychology*, **45** (1983), 792–804.

Cuneo, D. O. 'Children's judgments of numerical quantity: a new view of early quantification', *Cognitive Psychology*, **14** (1982), 13–44.

Desvousges, W. 'Measuring non-use using contingent valuation: an experimental evaluation of accuracy', *Research Triangle Monograph*, **1**, 1992.

Eddy, D. 'Probabilistic reasoning in clinical medicine: problems and opportunities', in Kahneman, D., Slovic, P. and Tversky, A. (eds), *Judgment Under Uncertainty: Heuristics and Biases* (pp. 249–267), New York: Cambridge University Press, 1982.

Fischhoff, B. and Bar-Hillel, M. 'Focusing techniques: a shortcut to improving probability judgments?', *Organizational Behavior and Human Performance*, **34** (1984), 175–194.

Fischhoff, B., Slovic, P. and Lichtenstein, S. 'Improving intuitive judgment by subjective sensitivity analysis', *Organizational Behavior and Human Performance*, **231** (1979), 339–359.

Gigerenzer, G., Hell, W. and Blank, H. 'Presentation and content: the use of base rates as a continuous variable', *Journal of Experimental Psychology: Human Perception and Performance*, **14** (1988), 513–525.

Gigerenzer, G. and Hoffrage, U. 'How to improve Bayesian reasoning without instruction: frequency formats', *Psychological Review*, **10** (1995), 684–704.

Gigerenzer, G. and Richter, H. R. 'Context effects and their interaction with development: area judgments', *Cognitive Development*, **5** (1990), 235–264.

Ginossar, Z. and Trope, Y. 'The effects of base rates and individuating information on judgments about another person', *Journal of Experimental Social Psychology*, **16** (1980), 228–242.

Ginossar, Z. and Trope, Y. 'Problem solving in judgment under uncertainty', *Journal of Personality and Social Psychology*, **52** (1987), 464–474.

Kahneman, D. 'Bartlett Lecture', July 1995, Experimental Psychology Society Birmingham, UK.

Kahneman, D., Fredrickson, B., Schreiber, C. and Redelmeier, D. 'When more pain is preferred to less', *Psychological Science*, **4** (1993), 401–405.

Kahneman, D. and Ritov, I. 'Values, attitudes, and dollars', manuscript in preparation, 1998.

Kahneman, D. and Tversky, A. 'On the psychology of prediction', *Psychological Review*, **80** (1973), 237–251.

Koehler, J. J. 'The base rate fallacy reconsidered: descriptive, normative, and methodological challenges', *Behavioral and Brain Science,* **19** (1996), 1–53.

Krosnick, J. A., Li, F. and Lehman, D. R. 'Conversational conventions, order of information acquisition, and the effect of base rates and individuating information on social judgments', *Journal of Personality and Social Psychology*, **59** (1990), 1140–1152.

Locksley, A. and Stangor, C. 'Why versus how often: causal reasoning and the incidence of judgmental bias', *Journal of Experimental Social Psychology*, **20** (1984), 470–483.
Massaro, D. W. and Friedman, D. 'Models of integration given multiple sources of information', *Psychological Review*, **97** (1990), 225–252.
McBride, R. L. 'Psychophysics — a link between R&D and marketing', *Food Technology in New Zealand*, **24** (1989), 25–27.
Schreiber, C. and Kahneman, D. 'Determinants of the remembered utility of aversive sounds', *Journal of Behavioral Decision Making* (in press).
Schwarz, N., Strack, F., Bless, H., Klumpp, G., Rittenauer-Schatka, H. and Simons, A. 'Ease of retrieval as information: another look at the availability heuristic', *Journal of Personality and Social Psychology*, **61** (1991), 195–202.
Schwarz, N., Strack, F., Hilton, D. and Naderer, G. 'Base rates, representativeness, and the logic of conversation: the contextual relevance of "irrelevant" information', *Social Cognition*, **9** (1991), 67–83.
Smith, E. and Zarate, M. 'Exemplar and prototype use in social categorization', *Social Cognition*, **8** (1990), 243–262.
Tversky, A. and Kahneman, D. 'Availability: A heuristic for judging frequency and probability', *Cognitive Psychology*, **5** (1973), 207–232.
Tversky, A. and Kahneman, D. 'Causal schemas in judgment under uncertainty', in Fishbein, M. (ed.), *Progress in Social Psychology* (pp. 81–103), Hillsdale, NJ: Erlbaum, 1980.
Varey, C. and Kahneman, D. 'Experiences extended across time: evaluation of moments and episodes', *Journal of Behavioral Decision Making*, **5** (1992), 169–185.
Wells, G. L. and Harvey, J. H. 'Do people use consensus information in making causal attributions?' *Journal of Personality and Social Psychology*, **35** (1977), 279–293.
Wilkening, F. 'Development of dimensional integration in children's perceptual judgment: experiments with area, volume, and velocity', in Wilkening, F., Becker, J. and Trabasso, T. (eds), *Information Integration by Children* (pp. 47–69), Hillsdale, NJ: Erlbaum, 1980.
Zukier, H. and Pepitone, A. 'Social roles and strategies in prediction: some determinants of the use of base-rate information', *Journal of Personality and Social Psychology*, **47** (1984), 349–360.

Authors' biographies:
Nathan Novemsky is a doctoral candidate in the Department of Psychology at Princeton University. His research focuses on decision-making and hedonic judgments.

Shirit Kronzon is a doctoral candidate in the Department of Psychology at Princeton University. Her research interests include negotiation and decision-making.

Authors' address:
Nathan Novemsky and **Shirit Kronzon**, Department of Psychology, Princeton University, Princeton, NJ 08544, USA.

