Available online at www.sciencedirect.com
ScienceDirect
Journal of Experimental Social Psychology 43 (2007) 91–103
Journal of Experimental Social Psychology
www.elsevier.com/locate/jesp

How people cope with uncertainty due to chance or deception ☆

Yaacov Schul [a,*], Ruth Mayo [b], Eugene Burnstein [b,c], Naomi Yahalom [a]

[a] *Hebrew University, Jerusalem, Israel*
[b] *University of Michigan*
[c] *Warsaw School of Social Psychology*

Received 28 August 2005; revised 19 December 2005
Available online 24 May 2006
Communicated by Fiedler

Abstract

In making social judgments people process effects caused by humans differently from effects caused by non-human agencies. We assume that when they have to predict outcomes that are attributed to non-human causes, people acknowledge their ignorance and try to focus on what is most diagnostic. However, when events are attributed to human agency, they believe that nothing is arbitrary and that one can understand the decision situation well enough to eliminate error. If so, then people should behave differently when an uncertainty is attributed to chance (a non-human agency) or to deception (a human agency). We tested this prediction using the probability-matching paradigm and found reasonable support for our analysis in four experiments. Individuals who attributed uncertainty to deception were less likely to adopt the optimal rule-based strategy than those who attributed it to chance. Indeed, only when the former were prevented from thinking about and elaborating the outcomes (the high-interference condition in Experiment 3) was their performance comparable to the level of individuals in the chance condition.

Keywords: Uncertainty; Distrust; Decision-making; Prediction-strategy; Deception

Man's mind is so formed that it is far more susceptible to falsehood than to truth—Desiderius Erasmus.

To a supreme knower all events might be completely deterministic. However, humans are condemned to survive in a probabilistic world since, whether due to ignorance or inability, we cannot comprehend its complexity. Perhaps the most extreme example of accepting uncertainty in making predictions occurs when we say an event is due to chance. For example, we believe the number of factors determining the behavior of a fair coin tossed up in the air exceeds our knowledge and ability to calculate. Hence, we forgo attempting to do so; instead we attribute the outcome to "chance" or some unpredictable fleeting property called "luck". Our research investigates the implications of this by studying how people cope with different types of uncertainty, due to chance and deception.

Given the pervasiveness of uncertainty, as well as its influence on life, it seems likely that humans have evolved ways to cope with it. And indeed, there is a vast literature on decision-making under uncertainty (see recent reviews in Gilovich, Griffin, & Kahneman, 2002) that describes (and prescribes) a variety of strategies to this end. It appears, however, that uncertainty is not a unitary mental state (Kahneman & Tversky, 1982). This is important because the same level of uncertainty may activate different strategies depending on how it is interpreted. For example, it matters whether uncertainty is quantified or not (e.g. Tversky & Fox, 1995), is about the magnitude of the probability or the outcome (Harel & Segal, 1999; Keren & Gerritsen, 1999), is attributed to chance or luck (Friedland, 1998), and whether it is about

☆ Preparation of this paper was supported by grants from the US-Israel Binational Science Foundation (BSF) and from the Israel Academy of Science (ISF). We thank Orly Carmi, Hila Goldberg, Shirly Ronen, and Alon Epstein for their help in running the experiments and coding the data, and Gideon Keren and Iris Rooij for commenting on an earlier version of this paper.
* Corresponding author.
E-mail address: yschul@huji.ac.il (Y. Schul).

0022-1031/$ - see front matter
doi:10.1016/j.jesp.2006.02.015

an event or about one's beliefs about the event (Teigen, 1994; Wells, 1992). The research reported here explores another characteristic of uncertainty that has received little, if any, attention—namely, its source. Does it matter whether uncertainty is due to the complexity of the world or to the complexity of human psychology? To explore this question, our research compares peoples' reactions to uncertainty when it is attributed to chance or deception.

Briefly, we propose that in dealing with chance people are more likely to use a statistical mode of thinking than when they have to deal with deception. We shall discuss this prediction in more detail after considering the differences between statistical and clinical decision processes.

Statistical versus clinical orientations to decision making under uncertainty

Einhorn (1986) distinguished between two orientations to forecasting, clinical and statistical (see Grove & Meehl, 1996, for a review of the debate about the effectiveness of statistical versus clinical predictions). In clinical forecasting the all-important goal is a sense of understanding that ensues when error is eliminated. Zero error means the decision-makers account for every detail, large and small. To do so they attempt to construct a coherent story integrating everything unselectively. Likelihood estimates, therefore, depend on the "goodness" of the story, so that a detailed scenario consisting of what appears to be a causally linked chain of events can give rise to a higher likelihood estimate than a scenario containing a subset of these events (Kahneman & Lovallo, 1993; Tversky & Kahneman, 1983). This kind of thinking is reminiscent of explanation-based decision-making models (Pennington & Hastie, 1992; Wells, 1992) in the sense of emphasizing individuals' attempts to *understand* the situation by imposing a coherent narrative on it (Bruner, 1986).

Statistical prediction, in contrast, acknowledges the forecaster's ignorance. Rather than trying to fit everything one knows about the event into a coherent story, a forecaster with a statistical orientation is willing to be selective, to consider only a subset of the features of the to-be-predicted event, treating them as essential or defining characteristics, and thus implying that the other features are non-essential or accidental. This allows the forecaster to treat a large number of events as equivalent (cf., the external view in Kahneman & Lovallo, 1993). In so doing, the forecaster is willing to tolerate error (i.e., treat non-identical events as equivalent) to reduce error (Einhorn, 1986).

Coping with the uncertainty associated with deception and chance

We propose that the likelihood of engaging in clinical prediction is greater when uncertainty is attributed to deception rather than to chance. There are several reasons for this hypothesis. To begin with, in making statistical predictions forecasters necessarily simplify—namely, they ignore features of the particular event that they regard as accidental or irrelevant, and consider only those deemed essential. There is evidence, however, that simplification is contrary to the typical processing tendency exhibited when individuals suspect that information is misleading.

Schul, Burnstein, and Bardi (1996), for example, found that when participants read a set of messages describing a person, those who suspected that one or more of the messages might be invalid needed more time to read and integrate the entire set than those who were not suspicious (see also Chiappe et al., 2004). This was interpreted to mean that suspicion increased the complexity of encoding by recruiting multiple scenarios to encode the information (see also Schul, Mayo, & Burnstein, 2004). A similar conclusion has been suggested by Fein and coworkers (Fein, Hilton, & Miller, 1990; Fein, McCloskey, & Tomlinson, 1997; Hilton, Fein, & Miller, 1993).

A second factor that may make individuals reluctant to use statistical thinking in cases of deception has to do with giving up control. The statistical orientation implies that decisions are based on rules which are applied consistently in a mechanized fashion, effectively barring the decision-maker from calling the shots (Einhorn, 1986). Giving up control over the conclusion drawn from the information can be aversive (Langer, 1975), especially when there is the possibility of deception. Tailoring the decision according to the unique aspects of the particular case gives decision-makers a sense of control (Dawes, 1994) and thereby helps them cope with the emotional aspects of the uncertainty.

The two types of uncertainty also differ in terms of people's understanding of how the uncertainty has been generated. People are not only the targets of deception attempts, they are often the sources of such attempts (see DePaulo et al., 2003, for a brief review). The experience of being the deceiver is likely to induce a sense of competence: One believes not only that he or she is able to deceive others, but also that the other's attempts at deception can be unmasked. There is no comparable experience, however, if the uncertainty is non-human in origin. As a result we feel it normal that the factors generating the uncertainty remain either unknown or poorly understood, and thus beyond our control.

This should remind us that deception and other norm violations become concerns only to the extent the other is seen as a free agent, someone whose actions can be predicted and, if necessary, influenced, once his or her intentions are known. There is considerable evidence from experiments with the Prisoner's Dilemma and coordination games that individuals adopt a different strategy when playing with a programmed agent such as a computer than when playing with a person using the identical strategy (e.g. Abric, Faucheux, & Moscovici, 1967; Kelley, Thibaut, Radloff, & Mundy, 1962; Rabinowitz, Kelley, & Rosenblatt, 1966). Sometimes the difference is intriguingly 'irrational'. For instance, in an ultimatum game individuals often reject a lopsided offer from another person but hardly ever when the same amount is offered by a computer (Blount, 1995).

The above analysis suggests that when uncertainty is attributed to chance, and perhaps to any non-human agent like ‘mother nature’, not only do people have a weaker sense of confidence in their ability to tailor a particular event to a decision, but they also have less urgency to do. This may create a greater willingness to adopt a statistical orientation in making decisions compared to occasions when the uncertainty is attributed to the possibility of deception.

The current study

We (Schul & Mayo, 2003) recently investigated the contrast between clinical and statistical modes of forecasting using a variant of the Probability-Matching paradigm (for reviews of Probability-Matching research see Fantino, 1998; Millward, 1971; Myers, 1976; Shanks, Tunney, & McCarthy, 2002). Participants in Schul and Mayo (2003) were shown a bag containing matchboxes. They were informed that: (i) half of these boxes contained a blue token and half a yellow token; (ii) half of the boxes also had a blue sticker on them and half a yellow sticker; and (iii) in two-thirds of the cases the sticker’s color matched the token’s. Their task involved repeatedly sampling a box and predicting the color of the token. Participants were told they would receive money for each correct prediction. In addition they were explicitly instructed to maximize the amount of money they won in the study.

This task allowed us to explore whether individuals can resist the attraction implicit in the clinical mode of thinking and apply a general rule in making predictions. That is, while it is very tempting to take into account contingencies peculiar to a trial (e.g., its position in the series, the outcomes on trials preceding it, etc.) in making a prediction, the optimal prediction strategy uses a very simple decision rule, called *maximization*, whereby one consistently predicts the color of the token according to the color of the sticker. Because the optimal performance strategy requires forecasters to use a single informational cue while ignoring all other available cues, it was expected that optimal performance would typically be rare, which was indeed the case. We believe that one of the impediments for optimal performance is the forecasters’ attempt to consider and make sense of everything they know about the event. Our conjecture is that it was difficult for participants in Schul and Mayo (2003) to abstract a simple rule and use only those cues that are specified by it, while completely ignoring other possible cues.

This analysis implies that if attention is directed to general features of the task, rather than to trial-by-trial contingencies, use of the maximization rule should increase. Indeed, when participants were instructed to *plan* an overall strategy that would allow them to predict correctly as often as possible, they were more likely to hit on maximization than when they were merely asked to explain their performance (Schul & Mayo, 2003). In other words, if individuals think about the decision task abstractly, or as a whole, and use this analysis in planning how to make future predictions, they are able to focus on diagnostic features and ignore non-diagnostic ones (cf., Trope & Liberman, 2003). This allows them to make forecasts while accepting some error, perhaps understanding that the strategy will help them make fewer errors in the long run (Einhorn, 1986).

Schul and Mayo’s (2003) findings suggest that the use of a maximization strategy can serve as a marker for individuals’ willingness and ability to depart from a clinical perspective and engage in rule-based thinking when making forecasts. The current study uses this marker to compare people’s performance in situations involving chance or deception. As our earlier discussion indicates, we hypothesize that participants who view the uncertainty as based on the possibility of deception will take a more clinical approach to prediction. If so, then such participants should be less likely to maximize, that is, to use a single cue in making their forecasts, compared to participants who attribute the uncertainty to chance.

Below we describe three studies investigating this hypothesis. Study 1 explores the influence of deception on predictions in face-to-face interactions. Studies 2 and 3, which involve computer-assisted interaction, test alternative interpretations of the findings in Study 1.

Study 1

Method

Participants

Forty-five students participated in the study. They were allocated randomly to chance ($N = 15$) and deception ($N = 30$, forming 15 pairs) conditions. The only restriction in randomization was that the pairs of participants in the deception condition were of the same sex.

Procedure

After signing up to participate in the study, individuals were contacted by telephone and scheduled to come to the lab. Telephone scheduling permitted us to assure that participants in the same session did not know each other prior to the study.

In the *chance* condition we used a variant of the task employed by Schul and Mayo (2003). Participants were shown a bag containing 60 white matchboxes. They were informed that 30 of them contained a blue token and 30 a yellow token, and that 30 of these boxes had a blue sticker and 30 a yellow sticker. Participants were further told that in 2/3 of the cases the sticker matched the token in color, and that this was true for boxes containing either blue or yellow tokens. On each trial the participant drew one box from the bag (without looking inside) and predicted its content by placing it in one of two locations, marked “blue” and “yellow”. The box was then opened so that the experimenter and the participant could see whether the prediction was accurate. The box was then placed back into the bag, which was reshuffled before a new box was drawn. Each participant made 60 predictions during the experimental session. Participants were informed that if their prediction accuracy was above the average accuracy of other participants in the study, they would receive a monetary bonus.

Participants in the *deception* condition were informed that the experimental task involved two roles (a sender and a receiver) and that each of them would participate as a sender and as a receiver at different phases of the experiment. It was explained that on each trial the sender would place a blue or a yellow token in a box with a blue or a yellow sticker and hand it over to the receiver. The receiver's task was to predict the color of the token. Furthermore, they were informed that the sender was permitted to deceive the receiver (i.e., to place a token with a different color than the sticker in the box) in 1/3 of the cases at the most.[1] Assignment to the roles of sender and receiver was determined by a draw at the onset of the study. Participants were also told that they would switch roles after 30 rounds and then continue for 30 additional rounds. Finally, they were told that the receiver whose predictions were more accurate would receive a monetary bonus. [2]

Of course, in competitive games of this sort the receiver's behaviors could be affected not only by wanting to uncover the deception but also by a desire to influence the sender. That is, participants in a two-person repeated game serve simultaneously in two roles: As receivers they respond to potential bluffs transmitted by the other while also 'sending' information via their predictions that could influence the sender's performance in later trials. Because our analysis assumes the former rather than the latter condition, we instructed receivers to conceal their responses from the senders. Therefore, the sender's behavior could not have been influenced by the receiver's responses.

The prediction task was followed by a structured interview in which participants were asked to explain the rationale behind their placement of the boxes and were probed for their understanding of the normatively optimal strategy.

Results

Preliminary analyses

Senders in the deception conditions were instructed to deceive in no more than 1/3 of the cases. That is, they were allowed to send no more than 10 boxes with stickers that did not match the tokens. Most senders utilized all the opportunities they had to deceive. Fifteen senders sent 10 boxes with a sticker–token mismatch, 13 sent 9 such boxes, and 2 sent 8 such boxes, yielding an overall rate of 31% of the trials with a sticker–token color-mismatch. The proportion of trials in which the color of the token and that of the sticker did not match was not significantly different from that in the chance condition ($M = 33\%$, $t(43) = 1.67$, $p > .10$). This pattern was replicated in the remaining experiments as well. Therefore it will not be discussed further.

Next we examined whether there is a difference between participants who started in the role of senders and those who started as receivers. We compared them with respect to three measures: the probability of correct prediction, the probability of predicting according to the sticker (i.e., maximization), and the probability of deception. We found no evidence for differences as a function of the order in the three measures. Therefore, in subsequent analyses we ignore whether the participants were initially receivers or senders.

Likelihood of making maximization-consistent predictions

Participants in the chance condition made 60 predictions while those in the deception condition made only 30 predictions. Because participants' performance might be affected by factors such as fatigue, all the analyses reported below utilize only the first 30 predictions of the participants in the chance condition.[3] We hypothesized that when uncertainty was associated with the motivation of another person to deceive them, participants would be less likely to use the optimal rule-based prediction than when it was attributed to chance. In line with this hypothesis, participants in the deception condition made maximization-consistent predictions in 70% of the trials, while those in the chance condition did so in 86% of the trials, $t(43) = 3.39$, $p < .01$, $d = 1.03$. Perhaps more striking is the analysis of the number of participants who used the optimal strategy consistently in the 30 predictions. Whereas 5 of the 15 participants in the chance condition used the maximizing strategy throughout, none of the 30 subjects in the deception condition used this rule-based prediction strategy. This difference is highly significant according to Fisher's exact test, $p < .01$. We shall return to this issue in more detail in General discussion.

Participants' reasoning about their actions

During the post-experimental interview, participants were probed for the strategies they used in making the predictions and the reasons behind these strategies. Participants' responses were transcribed, and two coders rated the extent to which each participant (1) searched for regularities in the sender's actions (regularity search, hereafter); and (2) focused on the association between the sticker and the token or used probability terms in justifying his or her

[1] Whereas the probability of a token-sticker mismatch was 33% in the chance condition, it was 33% or less in the deception condition. However, this difference stacks the odds against the experimental hypothesis, since the likelihood of making a maximization-consistent response increases as the likelihood of mismatch decreases.

[2] Obviously, our experimental situation provides a rather weak version of real-life deception. Not only it was done within the context of a scientific investigation, but the cost associated with being deceived were small (cf., Frank & Ekman, 1977). Still, the deception manipulation maintains important characteristics of deception—the conflict of interests between the sender and receiver and the ability of the sender to advance his or her interests through the information sent to the receiver. Thus, such manipulation can put participants in the deception condition in a different frame of mind than those in the chance condition and thereby induce different prediction strategies.

[3] The level of maximization in the first half of the experiment was not significantly different from the second half ($M = .86$ versus .88, $t(15) = 0.59$, $p > .5$).

Table 1
Classification of participants as a function of regularities search and probability focus (Experiment 1)

	Search for regularities		Probability focus	
	Chance (%)	Deception (%)	Chance (%)	Deception (%)
Absence	93	23	7	66
Secondary importance	7	30	7	17
Primary importance	0	47	86	17

Note: Classification is based on the coding of the post-experimental interviews.

action (probability focus, hereafter). Ratings were done on a 3-point scale (0 = absence, 1 = secondary importance, 2 = primary importance). For example, statements like "I was trying to put myself in her position" or "I was trying to discover repeated patterns in the sequence" were coded as 2 in regularity search. Statements like "I made my prediction according to the sticker because this offers a 2/3 chance to win" or "congruent choices were more probable" were coded as 2 in probability focus. Considering the interview transcripts of all 4 experiments reported below, the two coders agreed in 92% of the cases, and disagreements were resolved in discussion.

Table 1 depicts the classification of participants according to the two measures. Whereas participants in the chance condition were primarily focused on the probability with which the sticker and the token co-occurred and were not attempting to ferret out regularities in the sequences, those in the deception condition were much more concerned with the regularities and much less concerned with the probability. Statistical analyses of the pattern of classification within each measure indicate that the chance and the deception participants were statistically different from each other, $\chi^2(2)>15$, p's < .01. It may come as no surprise, therefore, that the two measures were associated with the probability of using maximization response. The Spearman rank correlation indicated that the extent of maximization was positively associated with the probability focus, $r(43) = .68$, and negatively with regularity search, $r(43) = -.50$, both p's < .05.

Discussion

Study 1 shows that individuals who made predictions in situations involving uncertainty associated with chance behaved differently than individuals who had to deal with uncertainty associated with deception attempts. In the latter case, receivers showed a greater degree of reluctance to use maximization. The behavioral difference was associated with differences in the way participants described their performance in the post-experimental interview. Participants in the deception condition were more concerned with finding regularities in the sticker–token sequences than those who performed under chance. Our conjecture was that under fear of deception individuals attempt to understand the other so that they can decipher the other's strategy. Such an attempt at understanding makes it difficult to focus on a particular cue and apply it consistently, and to accept erring as inevitable. Even so, the design of Study 1 permits several alternative explanations of these results.

To begin with, whereas sessions in the chance condition involved a single participant, those in the deception condition included two participants. It could be that the presence of the other participant interfered with the discovery and application of maximization, either because of the presence of another person in the room or because the competition between the two participants created a distracting environment that prevented the receiver from achieving the optimal performance. Study 2 tests this possibility by including a *competition* condition in which two participants were run simultaneously in the same room, each competed against Nature, and the better performer was rewarded with a bonus.

Another mechanism that might be responsible for the difference between the deception and the chance conditions has to do with the sequence of boxes receivers used in the prediction task. In the chance condition this sequence was generated by a random drawing. In the deception condition it was generated by the sender. The different generating mechanisms may lead to different sequences of colors as well as the matches–mismatches pairings (Rapoport & Budescu, 1997). Consequently, differences in predictions by the receivers might reflect the different input sequences in the two conditions. To control for this possibility, Study 2 had two control conditions. In one, receivers who thought they were in the deception condition were actually given random sequences. In the other, receivers who thought they were in the chance condition were actually given sequences generated by a deceiver in one of the deception runs.

Finally, the differences between the responses in the deception and chance conditions could occur because receivers in the deception condition tried to utilize other cues. Faces, in particular, are believed to leak information about deceit (e.g. Akehurst, Koehnken, Vrij, & Bull, 1996). Since receivers could observe the face of the sender, they may have tried to extract facial cues for deceit, which might have interfered with the use of statistical reasoning. To examine this issue, Study 2 included a condition in which a divider separated the sender and receiver so that receivers had no visual access to the senders' face.

Still, before reporting on Study 2 we would like to comment on two directions in which the main finding of Study 1 could be further examined. First, one could ask whether the nature of uncertainty (i.e., chance versus deception) influences people's understanding of the maximization strategy, and/or the use they make of this strategy. Understanding and use are clearly not identical. For example, forecasters may understand that maximization is optimal, yet choose not to maximize because they want to experience the thrill of gambling. In General discussion, we discuss this issue in light of the findings from the set of all experiments.

Second, as noted above, although maximization is known to be the optimal prediction strategy when one predicts from completely random sequences (i.e., under chance), it may not be optimal in the case of the deception, because human senders might be unable to generate random sequences. Consequently, one can ask whether receivers who maximize in the deception condition are better off than those who fail to maximize, and more generally, whether the maximization strategy is effective in enhancing accuracy under chance and deception. This issue will also be discussed in General discussion.

Study 2

One of the strengths of Study 1 is its transparency. Specifically, participants in the chance condition experienced the sampling, so that the mechanism that generated randomness was apparent. Participants in the deception condition experienced a very different source of uncertainty. The sender and the receiver sat face to face, and the sender handed a box to the receiver after placing a token in it, with no ambiguity about the conflict of interests between them or about the sender's wish to deceive. Moreover, all senders played the roles of receivers (and vice versa), so that they were well aware of the possibility of deception. However, all these advantages with regard to the realities of chance or deception came at the expense of the possibility of several alternative explanations discussed earlier. Study 2 therefore sacrifices some of the direct contact between the participants and the mechanisms generating the uncertainty to allow us to control the sequences of sticker–token pairs the participants received.

The major difference between Studies 1 and 2 has to do with the introduction of the computer as a vehicle that either sampled the boxes "randomly" (chance condition) or mediated between the choices of the sender and the responses of the receiver (deception condition). In the latter condition, both participants sat facing each other in the same room. Although they could see each other, they were instructed to limit their interaction to the computer-assisted mode. Each participant had his or her own screen and keyboard, which was hidden from the other participant. The study was run by software that was specifically programmed to display the boxes and tokens graphically, to allow sampling a box, placing a token inside a box, predicting the color of the token, and examining the contents of the box. Importantly, the software allowed us to manipulate the sequence of sticker–token pairs independently of the actual choice of the sender, and to monitor and control the number of deception attempts.

Study 2 consists of two experiments, 2a and 2b. All receivers in the deception conditions in Experiment 2a were informed that they would have to make 30 predictions, and that the senders were allowed to make 10 deception attempts at the most. This leaves open the possibility that receivers were counting the number of deception attempts. Experiment 2b was identical to Experiment 2a, except that receivers were not informed about the number of predictions they were about to make, or the number of deception attempts. Rather, receivers were informed that they would be making a series of predictions and that the senders were allowed to deceive in no more than 1/3 of the trials. Therefore, receivers in Experiment 2b had no incentive to count the number of deception attempts.

Experiment 2a

Method

The experiment included six conditions. Participants in the *chance-baseline* (C1) condition were given a computer-assisted version of the chance condition in Study 1. Briefly, participants were informed about the number of boxes stored in the computer memory, the colors of the stickers and the tokens, and the proportion of boxes with sticker–token mismatch. Then, on each trial the participant pressed a button to select a box randomly. The participant predicted its content using the mouse to drag the box to one of two locations, a blue circle or a yellow circle. Once inside the circle, the box was opened and the participant could see whether or not the prediction was accurate. Each participant made 60 predictions during the experimental session. Participants were informed that if their prediction accuracy would be above the average accuracy of other participants in the study, they would receive a monetary bonus.

The experiment began with a short demonstration of how to make a prediction and how to open the box to see whether the prediction was correct. Then the experimental session started, with the participant proceeding at his or her own pace. Following the prediction phase, participants were given the post-experimental questionnaire similar to the one employed in Study 1.

Whereas participants in the chance-baseline condition were run individually, those in the *chance-competition* (C2) condition were run in pairs. However, each performed individually, with the instructions and procedure being identical to the chance-baseline condition except in two respects. First, the participants were together in the same room rather than in different rooms. Second, the participants were told that they were competing against each other and that the one whose predictions were more accurate would get an extra monetary bonus.

Participants in the *chance-with-deception-sequences* (C3) condition were given the same instructions and procedure as those in the chance-baseline condition. However, the sequences they received were not generated by a random-number function. Rather, participants in this condition were matched with participants in the deception-baseline condition (see below), so that the sequence of boxes presented to each receiver corresponded to a sequence actually generated by a sender in the deception condition.

Participants in the *deception-baseline* (D1) condition were given similar instructions to those in the deception condition in Study 1, with a few changes required by the computer-assisted procedure. As in Study 1, assignment of the participants to the roles of sender and receiver was determined by a draw at the onset of the study. Participants were informed that they would switch roles after 30 rounds and continue thereafter for 30 additional rounds. Finally, they were told that the receiver whose prediction accuracy was higher would receive a monetary bonus.

After the introduction, the experimenter demonstrated how the interaction between the two players was to be achieved using the computer setup. In particular, during the demonstration the sender had to choose a token, place it in a box, and send it. Then the receiver had to make a prediction, open the box, and see whether the prediction was accurate. The sender and receiver were informed that the sender would not have access to the receiver's predictions. Following the demonstration the experimenter repeated the highlights of the experimental situation and answered questions. It was stressed at this point that no verbal interaction would be allowed between the sender and receiver. Moreover, receivers were warned not to provide non-verbal cues because this might permit the sender to take advantage of them. After 30 trials, participants switched roles, and the study continued for 30 additional trials.

Participants in the *deception-divider* (D2) condition were given the same instructions and procedure as those in the deception-baseline condition except that a divider was placed between the two participants. The divider separated the two participants from each other visually so that neither the sender nor the receiver could monitor each other's facial expression.

Finally, participants in the *deception-with-chance-sequences* (D3) condition were given the same instructions and procedure as those in the deception-baseline condition, with one significant difference. Unbeknown to the senders and receivers, rather than being given the sequences constructed and sent by the sender, receivers saw and responded to sequences identical to those in the chance-baseline condition.

Experiment 2b

Method

Experiment 2b included only the chance- and deception-baseline manipulations. These were identical to those in Experiment 2a, except that receivers were not informed about the number of predictions or the number of deception attempts. Instead, they were merely told the senders could deceive them in one-third of the trials at most, thus eliminating the incentive to count. The experiment employed two sequence conditions. In the 30/45 condition the first sender delivered 30 boxes (and was allowed to deceive in 10 or less), whereas the second sender delivered 45 boxes (and was allowed to deceive in 15 or less). In the 45/30 condition these parameters were reversed. The chance condition also included either 30 or 45 prediction attempts, and participants were not informed how many prediction attempts they were to make, only that the sticker and the token would not match in color in one-third of the cases.

Participants

One hundred and seventy students of the Hebrew University participated in Experiment 2a, and 80 students participated in Experiment 2b. They were paid the equivalent of four dollars for their participation. They could also earn an additional two dollars as a bonus.

Results

Preliminary analyses

As in Study 1, the set of receivers who started the experiment in this role were not different in the proportion of maximization-consistent predictions from the set of receivers who started the experiment as senders. Also, the two sets of receivers did not vary in the type of sequences they received. Therefore, the task-order factor was ignored in subsequent analyses.

Comparing chance and deception

For each participant we computed the percent of trials in which predictions were consistent with the sticker, that is, were consistent with maximization. Let us first consider the findings of Experiment 2a. The means are presented in Fig. 1. We tested the critical predictions with a series of planned contrasts.

Responses in the chance-baseline condition (C1, see Fig. 1) were similar to those in the chance-competition condition (C2, $t(52) = 0.24$, $p = .8$), suggesting that the presence of another player in the room, or the introduction of competition between the two participants, did not decrease the rate of maximization. A comparison between the deception-baseline condition (D1) and the deception/divider

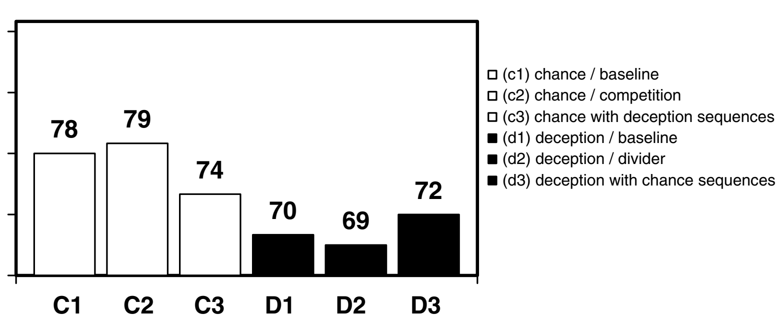

Fig. 1. Proportion of maximization-consistent predictions (Experiment 2a). Note: Bars are plotted relative to the rate of maximization expected from probability-matching.

condition (D2) also found no significant differences in performance ($t(56) = 0.06$, $p = .95$). It appears that the receiver's opportunity to see the sender did not influence maximization. Critically, the mean maximization in the two chance conditions ($M = .79$) was significantly higher than that in the two deception conditions ($M = .70$), $t(110) = 3.32$, $p < .01$, $d = .63$. Thus, Experiment 2a replicated the main results of Study 1 in spite of the more controlled environment, which minimized the number of cues available.

Experiment 2b showed a similar difference between chance and deception. That is, participants in the chance condition made predictions that were in line with the maximization rule more often than those in the deception condition. Considering only the first 30 predictions, receivers in the chance condition predicted according to the color of the sticker in 81% of their predictions, while those in the deception condition did so in 74% of their predictions, $t(76) = 2.12$, $p < .05$, $d = .49$. Importantly, the proportion of maximization responses did not differ between the 45/30 and the 30/45 sequence conditions, $t(76) = .32$, $p > .1$, nor did this factor interact with the chance/deception factor, $t(76) = 1.06$, $p > .1$. Taken together, the two experiments suggest that when uncertainty was due to the possibility of deception participants were more reluctant to maximize, that is, use the color of the sticker as the sole cue in making predictions, than when uncertainty was due to chance.

Participants' reasoning about their action

Table 2a presents the classification of participants in both experiments (excluding conditions C3 and D3 in Experiment 2a) according to the extent of their search for regularities and probability focus, both of which were coded from the post-experimental interviews. As in Experiment 1, the majority of participants in the chance conditions used probabilistic terms to justify their prediction strategy and/or were concerned with the association between the sticker and the token. These participants were not concerned with searching for regularities in the sequences. Participants in the deception conditions showed the opposite patterns, as revealed by significant tests of association, $\chi^2(2) > 35$, p's $< .01$

Table 2a
Classification of participants (N of respondents) and mean proportion of maximization-consistent predictions (in parentheses) as a function of regularities search and probability focus (Experiments 2a and 2b)

	Search for regularities		Probability focus	
	Chance	Deception	Chance	Deception
Absence	64 (.82)	34 (.80)	10 (.68)	60 (.67)
Secondary importance	19 (.76)	20 (.73)	10 (.73)	11 (.71)
Primary importance	6 (.73)	41 (.64)	69 (.83)	24 (.84)
Overall mean	**89 (.80)**	**95 (.72)**	**89 (.80)**	**95 (.72)**

Notes: Classification is based on the coding of the post-experimental interviews. Eleven participants in the chance condition failed to return their questionnaires.

Table 2a also shows the proportion of maximization-consistent responses as a function of participants' classifications. The entries in the bottom row show that when we collapsed over the levels of classification there was a sizable difference between chance and deception in the rate of maximization. However, the difference between chance and deception disappears when we consider each level of classification by itself. This indicates that the chance–deception difference in maximization is highly correlated with the way participants conceived their actions during the post-experimental interview. We caution the reader against interpreting this correlation causally. Because participants were interrogated only after the experiment, they may have responded on the basis of their actions during the experiment proper. Still, it might also be the case that the predictions individuals made during the experiment proper reflected the way they framed the situation, that is, in terms of probabilities, or in terms of searching for and responding to regularities. Clearly, however, participants' actions and reasoning are in line with the hypothesized difference between the handling of uncertainties associated with chance and deception.

The importance of the sequence

Conditions D3 and C3 allowed us to explore the importance of the nature of sequence in determining the response. Receivers in D3 believed they were seeing a sequence of boxes sent by a person who was trying to deceive them. In fact, the sequences were identical to those in the baseline chance condition. The statistical comparison reveals that the rate of maximization in D3 was significantly lower than that in C1 and C2, ($t(82) = 2.13$, $p < .05$, $d = .47$), but not significantly different from that in D1 and D2 ($t(82) = .60$, $p = .55$). This indicates that the smaller likelihood of maximization in the deception conditions could not be attributed to differences in the sequences used under chance and deception.

Receivers in C3 believed the sequence of boxes they were seeing was random. In fact, the sequences had been generated by senders in the deception-baseline condition. The results revealed that the average rate of maximization in C3 was between the chance-baseline and the deception-baseline conditions. The statistical analyses indicated that the rate of maximization in C3 was not significantly different either from the average performance in the two chance conditions, ($t(84) = 1.36$, $p = .18$), or from that in the two deception conditions ($t(84) = 1.42$, $p = .15$).

Table 2b presents the proportions of maximization-consistent responses and the number of participants according to the classification of participants on the post-interview measures. The pattern of classification was very similar to that in the "pure" chance and deception conditions (compare to Table 2a). However, inspection of the means of the maximization responses reveals that the within-row discrepancies between chance and deception do not disappear, as they did in Table 2a. Although this could reflect the small sample size (and consequently, the large variance of

Table 2b
Classification of participants (*N* of respondents) and mean proportion of maximization-consistent predictions (in parentheses) as a function of regularities search and probability focus (Conditions C3 and D3 in Experiment 2a)

	Search for regularities		Probability focus	
	Chance	Deception	Chance	Deception
Absence	16 (.77)	6 (.80)	8 (.63)	21 (.69)
Secondary importance	12 (.70)	12 (.71)	5 (.73)	1 (.83)
Primary importance	2 (.75)	12 (.68)	17 (.80)	6 (.79)
Overall mean	**30 (.74)**	**28 (.72)**	**30 (.74)**	**28 (.72)**

Note: Classification is based on the coding of the post-experimental interviews.

the sampling distribution), it may also indicate that the participant's frame of mind interacted with the nature of the sequence to determine his or her actions. Thus, on the basis of these comparisons we cannot rule out the possibility that the sequence influenced the proportion of maximization under chance (but see the analysis in footnote 5).

Discussion

Research on probability matching (see reviews in Schul & Mayo, 2003; Shanks et al., 2002) shows that people do not make optimal predictions when they are asked to predict a binary outcome from a binary cue. Instead of using a single cue, people use variety of cues, thereby lowering their prediction success. Our results in the chance condition are consistent with this finding. The findings in Studies 1 and 2 indicate that when uncertainty is perceived as due to chance participants maximized in about 80% of their predictions. While less than optimal, this performance is superior to that of participants whose uncertainty stemmed from the threat of deception, who maximized in about 71% of their predictions. We attribute this behavioral effect to the ways people handle uncertainty. Under chance, participants were more willing to forgo elaborate causal analyses and settle for using a single attribute—the color of the sticker—for prediction. This interpretation is supported by the analyses of the participants' justifications for their actions: compared to those in the chance conditions, participants faced with the possibility of deception were more concerned with understanding the pattern of regularities in the sticker–token sequences and less with the probabilistic nature of the situation.

We conjecture that receivers facing deception rejected maximization—the mechanized prediction strategy—and adopted one they believed might allow them to outsmart the sender. With this end in mind they were not merely trying to predict but also to understand the sender. This, in turn, involved constructing a narrative that places each prediction within a "theory of deception", including (but not limited to) the pattern of sticker-chip matches and mismatches on past trials, success or failure on the last trial, and various recursive intuitions regarding the sender's beliefs about the receiver.

Constructing such narratives presumably consumes time and cognitive resources. If so, then it follows that preventing receivers from elaborating in this fashion—for example, by increasing the amount of cognitive interference—should, paradoxically, improve performance. Study 3 manipulates the amount of interference to test this proposition.

Study 3

In this study we compare two levels of interference and a baseline/no-interference condition. The *baseline* condition is similar to the 30/45 condition of Experiment 2b. In particular, receivers were not restricted in the amount of time they could use for making predictions. In the *low-interference* condition participants were required to make a speeded prediction. Although the speeded prediction did not permit lengthy deliberations, receivers still had ample time to plan their response strategy while the sender was preparing the box for the next trial. To impede such planning, participants in the *high-interference* condition were given a large set of simple arithmetic problems and were asked to solve as many as they could in the interval between the prediction response and the time they received a new box.

Method

Participants

One hundred students were recruited to participate in the same way as those in the previous experiments. They were randomly assigned to 5 experimental conditions (see below).

Procedure

The *chance/no-interference* and *deception/no-interference* conditions were virtually identical to the 30/45 chance and deception conditions in Experiment 2b differing only in that participants in Experiment 2b used the computer mouse to indicate their response, whereas those in Study 3 used the keyboard to do so. This was done to make their mode of response identical to that of participants in the interference conditions (see below).

Participants in the *chance/low-interference* and *deception/low-interference* conditions were informed that they had only 2 s to make their prediction. One second prior to the appearance of the box, participants were warned by a buzzer as well as the message "GET READY", which appeared on the screen. After one second, the message was replaced by the box. Participants were informed that predictions made after the 2-s interval was up would be counted as errors.

Receivers in the high-interference condition had to make speeded predictions (like those in the low-interference condition) and also to solve math problems when they were not contemplating a prediction. The experimental procedure allowed receivers in the deception condition an interval of

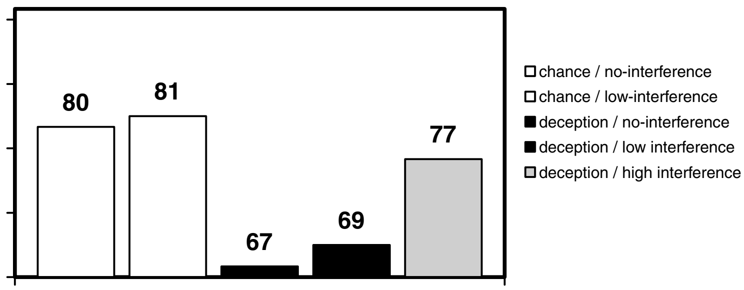

Fig. 2. Proportion of maximization-consistent predictions (Study 3). Note: Bars are plotted relative to the rate of maximization expected from probability-matching.

several seconds, when the senders made up their mind which sticker–token combination they would use in the next trial. It was during this interval that they could think about the sender's strategy using their 'theory of deception' to construct a narrative incorporating such reasoning. To impede their ability to do this, receivers in the *high-interference* condition were given a large set of arithmetic problems at the beginning of the experiment and asked to solve as many as they could in the intervals between each prediction and the appearance of a new box. They were also told that whoever solved more problems correctly would receive an extra monetary bonus.

Participants in the chance conditions received a new box (sent by the computer) immediately after their prediction response, and hence they were not placed under high interference.

Results

Preliminary analyses

Did the interference manipulation actually interfere with making predictions in time? Participants in the low-interference conditions missed about 2% of their predictions (with no difference between chance and deception), whereas those the high-interference condition missed about 8% of their predictions, $\chi^2(2) = 43.01$, $p < .01$. This indicates the manipulation was effective in consuming the participants' processing resources.

As in the previous studies, we examined whether there was a difference between receivers whose first role in the experiment was receiver (first receivers) and those whose first role in the experiment was sender (second receivers). For each participant we computed the proportion of maximization-consistent predictions, that is, predictions that were identical with the color of the sticker. Proportions were computed only from prediction responses made within the response window. A 2-way ANOVA (role order × interference) indicated that unlike Studies 1 and 2, there was a significant difference between the two sets of receivers. Second receivers tended to use the optimal response more often than first receivers (.78 versus .71, $F(1,51) = 4.31$, $p < .05$, $d = .58$). This effect was not moderated by the interference condition, as indicated by a non-significant interaction, $F(1,51) = 1.75$, $p > .1$. Because the difference between the two sets of receivers

Table 3
Classification of participants as a function of regularities search and probability focus (Experiment 3, no- and low-interference conditions)

	Search for regularities		Probability focus	
	Chance (%)	Deception (%)	Chance (%)	Deception (%)
Absence	65	40	37	70
Secondary importance	35	30	15	10
Primary importance	0	30	47	20

Note: Classification is based on the coding of the post-experimental interviews.

indicates that second receivers may have learned something from their experience as senders, our primary analysis compared participants in the chance conditions to first receivers in the deception conditions.

Comparing the effect of interference on performance under chance and deception

Fig. 2 presents the mean proportion of maximization-consistent predictions. Overall, participants in the chance conditions ($M = .81$) used the color of the sticker to predict the color of the token (i.e., maximized) more often than those in the deception conditions ($M = .71$, $t(65) = 3.30$, $p < .01$, $d = .82$). As the figure indicates, performance in the no-interference condition was similar to that for low interference, both under chance ($t(65) = .19$, $p > .1$) and under deception ($t(65) = .40, p > .1$).

Table 3 depicts the classification of participants in the low interference or no interference conditions, according to their responses in the post-experimental interviews. We see trends similar to those observed in Experiments 1 and 2. Participants in the chance condition were more concerned with probabilities and less concerned with finding regularities than participants in the deception condition[4].

Notwithstanding the similarity of the low-interference to the no-interference condition, the level of interference did matter. Participants in the high-interference condition made a greater number of optimal predictions than those in the deception/no-interference and deception/low-interfer-

[4] Because of the small N in the individual cells (3 cells with 2 observations or less), ANOVA on the proportion of maximization-consistent predictions was inappropriate.

ence conditions (.77 versus .68, $t(65) = 1.94$, $p < .06$, $d = .48$). Moreover, participants in the deception/no-interference and the deception/low-interference conditions made significantly fewer maximization-consistent predictions than those in the chance conditions (t's > 2.5, $p < .01$); and the null hypothesis of no difference between the extent of maximization under chance and that under deception with high interference could not be rejected ($t(65) = 0.84$, $p > .4$). Thus, when participants who attributed uncertainty to deception experienced high interference their pattern of prediction responses was similar to that of participants under chance, and dissimilar to participants under deception (with little or no interference).

Analysis of the post-experimental interviews revealed that 80% of the participants in the high-interference condition did not report thinking about probabilities and 70% did not report searching for regularities. This tallies nicely with the suggestion that the high interference condition disturbed the participants' ability to look for and construct models about the reality. Rather, the participants seemed to have used the optimal strategy because of its simplicity, and, as is often the case, simplicity was associated with optimality.

General discussion

In making social judgments people pay special attention to whether effects are caused by human or non-human agencies. We assume that when they have to predict outcomes that are attributed to non-human causes, people acknowledge their ignorance and try to focus on what is most diagnostic. However, when events are attributed to a human agency, they believe that little or nothing is arbitrary and that it is possible to eliminate all error if one understands the decision situation sufficiently. Accordingly, the individual strives for 'no loose ends' by incorporating as many features of the situation as possible into a coherent causal narrative. On the other hand, in attributing uncertainty to chance the person implicitly recognizes that the outcome in question has nothing to do with human intentions or agency and instead is indeterminate—a state of mind which activates statistical thinking.

Since optimal performance in the probability-matching paradigm can be achieved by attending to a single feature and ignoring all else, it is well suited to distinguish those who use statistical thinking from those who do not. Accordingly, it was predicted that individuals under the threat of deception would be disadvantaged compared to those who viewed the uncertainty as stemming from chance.

We found reasonable support for our analysis in four experiments using probability matching: individuals who attributed uncertainty to deception were less likely to make their predictions on the basis of a single cue (i.e., to maximize) than those who attributed it to chance. Indeed, only when the former were prevented from engaging in elaborative thinking (the high-interference condition in Experiment 3) did their performance approach the level of individuals in the chance condition.

Understanding maximization

The analyses presented so far involved the proportion of maximization-consistent predictions, which as hypothesized, was higher when uncertainty was considered due to chance rather than to deception. Obviously, the use of maximization-consistent prediction depends on the person's knowledge that the strategy is optimal. Even with such knowledge, however, the person's choice may still be influenced by other concerns, say, a desire to appear daring or a need for variety (Gal & Baron, 1996; Nies, 1962; Schul & Mayo, 2003).

In the analyses below we focus more directly on participants' understanding of the optimality of maximization. This is done in the "pure" chance and deception conditions (i.e., not including C3 and D3 in Experiment 2a), using two different markers. First, understanding was indicated by the use of maximization-consistent predictions over all 30 trials. Such an index imposes the most stringent criterion for 'knowing', so as to exclude individuals who may have been tempted by goals other than optimal performance. Only 3% of the participants (6/184) in the deception conditions maximized throughout, whereas 13% of the participants (20/151) under chance did so, $\chi^2(1) = 10.19$, $p < .01$. There is, however, a shortcoming inherent to this index—namely, it almost certainly underestimates people's actual understanding of maximization. To maximize throughout, participants would have had to grasp at the very beginning of the experiment that maximization is the optimal prediction strategy, and participants who did so only after making one or more predictions would not have been counted. Moreover, participants are likely to have had other goals, playful or serious, besides optimal performance, in mind while predicting (e.g., a desire to appear risky or curiosity about a different decision rule).

Be that as it may, the second index of participants' understanding is therefore based on data collected during the post-experimental interview. After they were questioned about the reasons for their strategy of prediction, participants were presented with a series of boxes, each with either a blue or a yellow sticker. They were asked to predict the contents of each box (i.e., a blue or a yellow token) assuming that each correct response would earn them $100. Understanding of maximization was indicated by a consistent prediction of the token's color according to that of the sticker. In these circumstance 19% of the participants in the deception condition and 33% of the participants in the chance condition showed that they understood maximization, $\chi^2(1) = 7.85$, $p < .01$.[5] Taken together, the pattern of findings suggests that individuals who attributed their uncertainty to chance were more likely to understand the maximization

[5] Interestingly, when we examined the understanding of participants under deception with random sequences (D3) and under chance with deception sequences (C3) we found that while only 14% of participants in D3 showed understanding, 47% of participants under C3 did so. This provides further support for the suggestion that the nature of the sequence did not play a major role in mediating the difference between chance and deception.

strategy and to use it in their predictions than those who attributed their uncertainty to the possibility of deception.

The maximization response need not reflect understanding, however. Rather, the sticker's color might be used as a signal to predict the token's color mindlessly, even without thinking about rules or making inferences about optimal behaviors. This possibility is consistent with the seemingly paradoxical finding that animals show superior performance to human forecasters in probability matching tasks (see Fantino, 1998; Mackintosh, 1974). Thus, there might be two routes to maximization, either through understanding the rule, or through not elaborating on the information but responding on the basis of the most salient local cue (Schul & Mayo, 2003). A threat of deception draws attention to other, non-salient, cues in the environment (Schul et al., 2004). It, therefore reduces both the likelihood of understanding and the likelihood of responding only on the basis of the salient cue.

Optimality of the prediction strategy

By definition, the sequences of boxes under chance were random. Therefore, participants could have achieved, on the average, no better than 67% accuracy, the level expected if maximization is used. However, because the sequences under deception were generated by senders, they were not random, and as a result were potentially more informative than those in the chance conditions. It is therefore possible that receivers in the deception condition may have had an advantage in terms of predictability. Moreover, because the sequences were not random, there is no simple best rule that maximizes performance accuracy, and in principle receivers could have attained a perfect prediction score by outguessing the sender. In fact, our hypothesis about the reluctance of receivers who suspect they are being deceived to engage in statistical prediction rests on the assumption that they believe they can figure out the sender's strategy and thereby do better than chance. This makes it interesting to examine what would have happened had receivers in the deception condition used maximization.

For each receiver we computed an accuracy score based on perfect maximizing. Analysis reveals that *had* receivers in the deception condition predicted according to the maximization rule their accuracy would have increased on average by 9%, which is significantly larger than the hypothetical gain from perfect maximizing in the chance condition ($M = 6\%$, $t(296) = 2.91$, $p < .05$, $d = .34$). Moreover, despite the disadvantage of receiving relatively uninformative sequences, participants in the chance condition were more accurate ($M = 61\%$) than those in the deception condition ($M = 59\%$, $t(296) = 2.40$, $p < .05$, $d = .28$). Thus, although receivers in the deception condition got non-random sequences, they were unable to take advantage of this fact, and if anything, they would have benefited more than participants in the chance condition from mechanizing predictions using the optimal statistical rule.

Our analysis highlights the cost associated with attempting to understand trial-by-trial contingencies (clinical reasoning), rather than using a simple rule (statistical reasoning) when making predictions. How can the prevalence of such a costly processing mode be explained? And if it is indeed a natural mode of thinking, what is its adaptive value? Our speculation is that the benefits of such processing are due in part to the absence of immediate reliable feedback in the social world which makes it difficult to learn more useful strategies for dealing with uncertainty (Dawes, 1994; Einhorn, 1986). It is also quite likely that finding out the truth or achieving high prediction accuracy is neither the only nor the most important goal in social interactions, and that attempts to understand the trial-by-trial contingencies are meant to serve other ends. For example, the sense of understanding that good stories provide may enhance a person's sense of control. According to Langer (1975), people are motivated to control their own destiny as well as their physical and social environment because they have a need for competence, mastery, or agency.

Moreover, we do not believe that performance under the threat of deception is necessarily less optimal than that under chance. We could have utilized other tasks, such as memorizing events or making attributions (Fein, 1996; Schul et al., 1996), where performance under the threat of deception would have led to a more optimal performance than in its absence (e.g., when receivers and senders trust each other). However, in examining how people handle deception when making predictions may help us begin to understand the riddle implicit in the quote attributed to Erasmus, namely, why the mind is more susceptible to falsehood than to truth.

References

Abric, J. C., Faucheux, C., & Moscovici, S. (1967). Role of the image of the partner on cooperation in a game situation. *Psychologie Francaise, 12*, 267–275.

Akehurst, L., Koehnken, G., Vrij, A., & Bull, R. (1996). Lay persons' and police officers' beliefs regarding deceptive behavior. *Applied Cognitive Psychology, 10*, 461–471.

Blount, S. (1995). When social outcomes aren't fair: the effect of causal attributions on preferences. *Organizational Behavior and Human Decision Processes, 63*, 131–144.

Bruner, J. (1986). *Actual minds, possible worlds*. Cambridge, MA: Harvard University Press.

Chiappe, D., Brown, A., Dow, B., Koonz, J., Rodriguez, M., & McCulloch, K. (2004). Cheaters are looked at longer and remembered better than cooperators in social exchange situations. *Evolutionary Psychology, 2*, 108–120.

Dawes, R. M. (1994). *House of cards*. New York: The Free Press.

DePaulo, B. M., Lindsay, J. J., Malone, B. E., Muhlenbruck, L., Charlton, K., & Cooper, H. (2003). Cues to deception. *Psychological Bulletin, 129*, 74–118.

Einhorn, H. J. (1986). Accepting error to make less error. *Journal of Personality Assessment, 50*, 387–395.

Fantino, E. (1998). Behavior analysis and decision making. *Journal of the Experimental Analysis of Behavior, 69*, 355–364.

Fein, S. (1996). Effects of suspicion on attributional thinking and the correspondence bias. *Journal of Personality and Social Psychology, 70*, 1164–1184.

Fein, S., Hilton, J. L., & Miller, D. T. (1990). Suspicion of ulterior motivation and correspondence bias. *Journal of Personality and Social Psychology, 58*, 753–764.

Fein, S., McCloskey, A. L., & Tomlinson, T. M. (1997). Can the jury disregard that information? The use of suspicion to reduce the prejudicial effects of pretrial and inadmissible testimony. *Personality and Social Psychology Bulletin, 23*, 1215–1226.

Frank, M. G., & Ekman, P. (1977). The ability to detect deceit generalizes across different types of high-stake lies. *Journal of Personality and Social Psychology, 72*, 1429–1439.

Friedland, N. (1998). Games of luck and games of chance: the effect of luck- versus chance-orientation on gambling decisions. *Journal of Behavioral Decision Making, 11*, 161–179.

Gilovich, T., Griffin, D., & Kahneman, D. (2002). *Heuristics and biases: The psychology of intuitive judgments*. New York: Cambridge University Press.

Gal, I., & Baron, J. (1996). Understanding repeated simple choices. *Thinking and Reasoning, 2*, 81–98.

Grove, W. M., & Meehl, P. E. (1996). Comparative efficiency of informal (subjective, impressionistic) and formal (mechanical, algorithmic) prediction procedures: the clinical-statistical controversy. *Psychology, Public Policy, and Law, 2*, 293–323.

Harel, A., & Segal, U. (1999). Criminal law and behavioral law and economics: observations on the neglected role of uncertainty in deterring crime. *American Law and Economic Behavior, 6*, 276–312.

Hilton, J. L., Fein, S., & Miller, D. T. (1993). Suspicion and dispositional inference. *Personality and Social Psychology Bulletin, 19*, 501–512.

Kahneman, D., & Lovallo, D. (1993). Timid choices and bold forecasts: a cognitive perspective on risk taking. *Management Science, 39*, 17–31.

Kahneman, D., & Tversky, A. (1982). Variants of uncertainty. *Cognition, 11*, 143–157.

Kelley, H. H., Thibaut, J. W., Radloff, R., & Mundy, D. (1962). The development of cooperation in the "minimal social situation". *Psychological-Monographs, 76* (Whole No. 538).

Keren, G., & Gerritsen, L. E. M. (1999). On the robustness and possible accounts of ambiguity aversion. *Acta Psychologica, 103*, 149–172.

Langer, E. J. (1975). The illusion of control. *Journal of Personality and Social Psychology, 32*, 311–328.

Mackintosh, N. J. (1974). *The psychology of animal learning*. New York: Academic Press.

Millward, R. B. (1971). Theoretical and experimental approaches to human learning. In J. W. Kling & L. A. Riggs (Eds.), *Experimental psychology* (3rd ed., pp. 905–1017). New York: Holt.

Myers, J. L. (1976). Probability learning and sequence learning. In W. K. Estes (Ed.), *Handbook of learning and motivation* (Vol. 3, pp. 171–205). New Jersey: Erlbaum.

Nies, R. C. (1962). Effects of probable outcome information on two-choice learning. *Journal of Experimental Psychology, 64*, 430–433.

Pennington, N., & Hastie, R. (1992). Explaining the evidence: tests of the story model for juror decision making. *Journal of Personality and Social Psychology, 62*, 189–206.

Rabinowitz, L., Kelley, H. H., & Rosenblatt, R. M. (1966). Effects of different types of interdependence and response conditions in the minimal social situation. *Journal of Experimental Social Psychology, 2*, 169–197.

Rapoport, A., & Budescu, D. V. (1997). Randomization in individual choice behavior. *Psychological Review, 104*, 595–602.

Schul, Y., Burnstein, E., & Bardi, A. (1996). Dealing with deceptions that are difficult to detect: encoding and judgment as a function of preparing to receive invalid information. *Journal of Experimental Social Psychology, 32*, 228–253.

Schul, Y., Mayo, R., & Burnstein, E. (2004). Encoding under trust and distrust: the spontaneous activation of incongruent cognitions. *Journal of Personality and Social Psychology, 86*, 668–679.

Schul, Y., & Mayo, R. (2003). Searching for certainty in an uncertain world: the difficulty of giving up the experiential for the rational mode of thinking. *Journal of Behavioral Decision Making, 16*, 93–106.

Shanks, D. R., Tunney, R. J., & McCarthy, J. D. (2002). A re-examination of probability matching and rational choice. *Journal of Behavioral Decision Making, 15*, 233–250.

Teigen, K. H. (1994). Variants of subjective probabilities: concepts, norms, and biases. In G. Wright & P. Ayton (Eds.), *Subjective probability* (pp. 212–238). New York: Wiley.

Trope, Y., & Liberman, N. (2003). Temporal construal. *Psychological Review, 110*, 403–421.

Tversky, A., & Fox, C. R. (1995). Weighing risk and uncertainty. *Psychological Review, 102*, 269–283.

Tversky, A., & Kahneman, D. (1983). Extensional versus intuitive reasoning: the conjunction fallacy in probability judgment. *Psychological Review, 90*, 293–315.

Wells, G. L. (1992). Naked statistical evidence of liability: is subjective probability enough? *Journal of Personality and Social Psychology, 62*, 739–752.