



This article appeared in a journal published by Elsevier. The attached copy is furnished to the author for internal non-commercial research and education use, including for instruction at the authors institution and sharing with colleagues.

Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third party websites are prohibited.

In most cases authors are permitted to post their version of the article (e.g. in Word or Tex form) to their personal website or institutional repository. Authors requiring further information regarding Elsevier's archiving and manuscript policies are encouraged to visit:

<http://www.elsevier.com/copyright>



Contents lists available at ScienceDirect

Journal of Experimental Social Psychology

journal homepage: www.elsevier.com/locate/jesp

The value of distrust

Yaacov Schul^{a,*}, Ruth Mayo^a, Eugene Burnstein^{b,c}^a Hebrew University, Jerusalem, Israel^b University of Michigan^c The Warsaw School of Social Psychology

ARTICLE INFO

Article history:

Received 5 July 2007

Revised 9 April 2008

Available online 14 May 2008

Keywords:

Trust

Distrust

Routine action

Decision making

Information processing

ABSTRACT

We assume that a state of distrust is the mental system's signal that the environment is not normal—things may not be as they appear. Hence, individuals sense they should be on guard. In particular, they are likely to avoid routine strategies, ones proven to be optimal and regularly used in normal environments, because these strategies are easily anticipated by whoever may be seeking to deceive them. Conversely, a state of trust is associated with a feeling of safety. The environment is as it normally is and things really are as they appear to be. Thus, individuals see no reason to refrain from doing what they routinely do. Accordingly, we hypothesize that figuring out a new situation depends on the type of environment and the actor's state of mind: in normal environments, where routine strategies are optimal, individuals who trust should outperform those who distrust; however, in unusual environments, where non-routine strategies are optimal, individuals who distrust should outperform those who trust. This paper reports three experiments that manipulate distrust via orienting tasks that participants perform prior to attempting to predict a series of events (Experiments 1 and 2) or solve matchstick arithmetic problems (Experiment 3). Performance success depends on discovering and implementing an appropriate rule. We found that, as predicted, the manipulation of distrust sensitized participants to the existence of non-routine contingencies, that is, contingencies that were not expected.

© 2008 Elsevier Inc. All rights reserved.

Introduction

Distrust is a ubiquitous psychological state that arises whenever we seem unable to take appearances, typically others' declarations and behaviors, at face value. This usually occurs because we recognize that their interests conflict with ours. Somewhat less common are instances in which we are unaware of any explicit conflict of interest but nonetheless feel uneasy, sensing that the situation is not normal, that others may behave unpredictably or cause something unexpected to happen.

The social psychological significance of distrust is best understood in comparison with its opposite, a state of trust. When protagonists trust each other, they may deliberately give each other the benefit of the doubt and assume that their interests are shared, or at least that each has the other's interests in mind. Ordinarily, when the stakes are not high, trust is the default state, so that without thinking much about the other, individuals feel the environment is normal and there is no need to worry. Trust, therefore, entails a belief that the other's actions will benefit the protagonist or that the situation is benign and nothing detrimental to one's interests will occur (Robinson, 1996). On the other hand, distrust denotes a perception of vulnerability due either to fear of the

other's motives, intentions, and prospective actions, or to vague forebodings that things are not as they appear and something unpredictable may occur (Koslow, 2000; Kramer, 1999).

Typically, a state of distrust is focused—it is attached to a specific target, often a person (but also, perhaps through analogy, an organization such as a political party or an inanimate object such as a car). Focused distrust can be triggered either because one knows something about the target (e.g., motivations, intentions, or past behaviors, see review in Kramer, 1999) or because one draws inferences about the other from the nature of the situation (e.g., the presence of temptations to defect, the lack of inducements to reciprocate, or the absence of a binding commitment; e.g., King-Casas et al., 2005; Seabright, 2004; Yamagishi, Cook, & Watabe, 1998). In such cases not only does the protagonist doubt the target's messages and actions, but the target becomes less liked (Winston, Strange, O'Doherty & Dolan, 2002). This can lead to discounting of the target's recommendations or judgments, as well as to a tendency to avoid the target and to be on guard when interacting with him or her.

A state of distrust can also be unfocused in the sense that it is not attached to a particular source. It may reflect residual activation of previous episodes of focused distrust, or the impact of cues that are typically associated with deception (e.g., the amount of details in the message, vocal tension, or fidgeting; see DePaulo et al., 2003 for a meta-analytic review). In such cases, perceivers

* Corresponding author.

E-mail address: yschul@huji.ac.il (Y. Schul).

may not be aware of why they distrust or, even more likely, they may not even consciously experience this state. Still, we propose that states of trust and distrust, regardless of whether they are focused or unfocused, conscious or unconscious, are associated with different patterns of thoughts and actions.

Thoughts and actions under trust and distrust

In order to understand the habitual thought and action processes under trust and distrust we must examine their meaning in everyday functioning. Briefly, trust connotes safety and transparency; individuals believe there is nothing to be feared in transactions between them and others. Distrust, in contrast, is associated with the concealment of truth and a lack of transparency. It is a state of uncertainty, but not the kind associated with outcomes that are inherently probabilistic, as in playing a slot machine or roulette. Rather, distrust reflects the receiver's perception of the source's intention (to mislead) and, potentially, the receiver's theory about the truth (Schul, Mayo, Burnstein, & Yahalom, 2007). This characterization is important because people are not only the targets of misinformation attempts, they are also the source of such attempts. As a result, unlike the possibility of losing to a slot machine, the threat of losing to a person engages recursive reasoning: individuals who suspect they may be deceived assume their potential deceiver will attempt to mask the deception by using knowledge acquired as successful deceivers as well as targets of deception in the past. That is, the deceiver may put on a display of routine actions, the ones each party expects of the other in normal (i.e., truthful) environments, where the problem is typically one of coordinating actions to achieve a shared goal. Therefore, unlike people who trust, those who distrust attempt to ascertain the other's attempts at deception by searching for signs that the other's behavior is departing from what is routine in the situation.

What are the implications of this analysis to the thought processes triggered under diffused trust and distrust? We propose that other things being equal, when a state of trust is active, one tends to believe, to follow the immediate implications of the given information. In contrast, when a state of distrust is active, one tends to search for non-obvious alternative interpretations of the given information, because distrust is associated with concealment of truth (cf., Fein, 1996; Schul, Burnstein, & Bardi, 1996). Thus, in distrust, the mental system become more open to the possibility that the ordinary schema typically used to interpret the ongoing situation may need to be adjusted. We propose that this pattern of thought may occur even in cases of diffuse distrust, where the situation that triggered the distrust bears no phenotypic resemblance to the current situation.

We (Schul, Mayo, & Burnstein, 2004) investigated this conjecture by comparing contexts of trust versus contexts of distrust in respect to the associative links they activated in processing messages. Our reasoning was as follows: when a source is trusted, receivers have no doubt the message is true, and automatically encode it as such. This in turn causes them to spontaneously bring to mind 'routine' concepts which are typically congruent with the message. When a source is distrusted, however, the message is doubted and, as a result, 'non-routine' concepts, typically incongruent with the message, spontaneously come to the receivers' mind. This conjecture was tested using single words as messages and priming facilitation to indicate the associative structure activated in response to a prime word. We found, as predicted, that when a prime word appeared together with a stimulus which signaled distrust it facilitated associations that were incongruent with the meaning of the target word (e.g., *dark* activated *light*). However, when a prime word appeared in the context of trust, the prime activated associations that were congruent with it (e.g., *dark* activated *night*).

The research reported in this paper attempts to extend and generalize these findings by investigating more general and abstract phenomena. Specifically, since the Schul et al. (2004) study involved reactions to verbal information, it could be interpreted in line with either a narrow or a broad perspective. According to the narrow perspective, the facilitation of incongruent cognitions under distrust is interpreted as the tendency of individuals to resist persuasive intent. According to the broader perspective, trust and distrust are associated with different types of thought processes: under trust one turns to routine mental outcomes whereas under distrust one turns to the non-routine. According to this broader interpretation, states of trust (versus distrust) influence how people consider information and generate inferences even when linguistic processes in general, and persuasion processes in particular, are irrelevant.

Routine and non-routine mental actions

By definition, 'routine' strategies are those that have proven most useful in an individual's normal or typical environment and are thus most likely to be activated by default. Consider the inferences one makes about a car whose color is a shiny red. Typically, when there is no reason to distrust, perceivers make *correspondent* inferences, from the way the car looks to other characteristics (e.g., how well it runs). Such routine inference process makes sense under trust, where things tend to be the way they appear. However, consider the reactions of a buyer who distrusts sellers. It is quite possible that the inference might be opposite, namely, good looks might be covering something—a negative relationship between how things look and how they are.

More generally, our analysis suggests that when people attempt to impose a structure on our complex world, they develop many schemata or inference rules¹ that are used to predict one attribute from others. These rules vary in their level of routine—the likelihood that they will be used in a normal environment. We hypothesize that those who trust tend to access inference rules that are more routine and prevalent than those who distrust. As a result, those who trust succeed more in making inferences in environments that are typical, but those who distrust do better in environments that are unusual, unexpected, or non-routine (see below).

To explore this hypothesis we employed experimental paradigms that make minimal use of linguistic inferences. Instead, we compare how people figure out the environment using abstract rules that vary in their level of routine. Specifically, Experiments 1 and 2 utilized the multiple-cue probability-learning paradigm, which is based on Brunswik's lens model and Hammond's Social Judgment Theory (see Doherty & Kurz, 1996, for a review). In this paradigm participants are given several cues (e.g., scores on different tests) on each of many trials, and asked to use them in predicting an outcome (e.g., success on a job). Once a prediction is made, participants are informed about the actual outcome, and a new trial begins. Initially, of course, participants do not know how to use the cues in making their predictions. However, over trials, based on the feedback they get, participants can, and do, discover the rule linking cues to outcome. Since the cues can be combined in different ways to determine the criterion, participants' performance success depends on having the appropriate rule in mind.

To illustrate the nature of the rules used in Experiments 1 and 2, consider the following example. Imagine having to predict applicant's on-the-job success (y) from two predictors: the applicant's level of motivation (x_1), and his or her level of education (x_2). Consider the following two inference rules: according to one, success is

¹ In the decision-making literature, these are termed hypotheses (e.g., McKenzie, 2004).

linked positively to both motivation and education. We call this the positive-linear rule. According to the other rule, on-the-job success is linked positively to motivation, but negatively to education. For this reason we call it the negative-linear rule.

Which of these rules is more routine? Brehmer's research (1974, 1980) offers a well-established method for answering this question with respect to rule-discovery and rule implementation. He compared people's intuitions about the prevalence of positive-linear, negative-linear, and non-linear rules, as well as the ease of learning these rules. He reported that (i) participants estimated that positive-linear rules are more prevalent than negative-linear rules, which in turn were viewed as more prevalent than non-linear rules; and (ii) the ease of learning these rules followed this pattern as well. Consequently, we assume that positive-linear rules tend to be more routine, meaning that they are more likely to be accessed when people are trying to draw inferences from one attribute to another. Note that the above results should not be understood as suggesting that the positive-linear rules are highly routine in an absolute sense, meaning that people consider them frequently in all situations. Rather, Brehmer's findings should be interpreted in a relative sense. Given that people think about linear rules, they are likely to consider hypotheses about positive-linear contingencies more readily than those about negative-linear contingencies. Note that we do not assume that rules are explicit or verbalizable. Research on implicit learning (e.g., Eitam, Hassin, & Schul, 2008) show that people can pick up regularities and act using rules based on these regularities, without being aware of the existence of the rules.

Of course, there are other factors that influence which inference rules are activated and used in any given context. Activation is influenced by factors such as chronic or recent activation and folk theories about applicability (e.g., Sedikides & Skowronski, 1991). However, other things being equal, our formulation predicts that under conditions of trust individuals will behave routinely and use the strategy they normally use. In Brehmer's model, this means they will attempt to utilize the more routine rule. In contrast, when distrust is activated, people depart from the routine and try other (non-dominant) rules to fit the data.

Because understanding environments and making correct predictions is helped by having the right inference rules, our analysis predicts that trusting individuals should succeed more in those environments characterized by dominant rules, while those who distrust should succeed more in understanding unusual environments, which are characterized by non-dominant rules. It is important to note that while our task involves repeated trials—multiple predictions—which allow us to explore different stages of performance (predisposition, learning, and discovery), in real life many tasks involve one-shot attempts. The predisposition for routine or non-routine, may be even more critical in determining people's actions in the single-trial case.

Experiment 1

Participants were asked to make predictions in one of two environments, which were constructed according to a positive-linear rule (the routine) or a negative-linear rule (the non-routine). Based on our theoretical analysis we hypothesize that under trust people tend to activate routine inference rules, while under distrust they activate non-routine rules. Therefore, it follows that distrust is an advantage when the environment is constructed according to the less routine rule (the negative-linear) and, conversely, distrust is a liability when the environment is constructed according to more routine rule (the positive-linear). In statistical terms this implies that the trust/distrust factor should interact with rule type in determining the success of predictions.

Method

Participants

Eighty students at the Hebrew University participated in the experiment. They were assigned randomly to one of four experimental conditions. The data of one participant were eliminated from the analyses due to failure to follow instructions.

Prediction environments

We generated, from a normal distribution, a set of 40 pairs of cues x_1 ($M = 175$, $SD = 62$) and x_2 ($M = 169$, $SD = 45$). x_1 and x_2 were virtually uncorrelated ($r = -.05$). We used them to generate two outcome variables, y_{pos} and y_{neg} , using an error term that was sampled from a normal distribution with $SD = 30$. The regression Eqs. (1a) and (1b) below relate y_{pos} and y_{neg} to x_1 and x_2 . Note that according to Eq. (1a) the criterion (y_{pos}) increases with x_1 and x_2 , while according to Eq. (1b) the criterion (y_{neg}) increases with x_1 but decreases with x_2 . In other words, the former is positive-linear (assumed to be routine) and latter is negative-linear (assumed to be non-routine).

$$\hat{Y}_{pos} = -274 + 2.01 * x_1 + 2.03 * x_2 \quad (R^2 = .98) \quad (1a)$$

$$\hat{Y}_{neg} = 412 + 2.01 * x_1 - 1.94 * x_2 \quad (R^2 = .98) \quad (1b)$$

In addition, we created an equivalent set which included only 10 predictors/outcome triads, to be used for practice.

Distrust manipulation

The manipulation was identical to the one validated in Schul et al. (2004, Experiment 2). Briefly, it involved creating an orienting task that did (or did not) trigger distrust (the control). The orienting task itself is described below. To prepare the stimuli for the orienting task we approached male and female students in the library and asked them to fill out a short questionnaire. Females were asked to answer truthfully. Males were asked to impersonate a female in giving their responses. Specifically, we informed the male students that their answers would be given to other students who would try to separate answers of male imposters from those of female truth-tellers. In the experiment proper we used the answers given to five questions: "What do you like to do in your leisure time?", "What do you think about plastic surgery?", "Explain in a few words how to change a tire on a car", "Describe what's in your handbag right now", and "Describe what kind of gesture you would consider romantic". The responses to each question varied in length from several words to several sentences. We chose four sets of replies (two male, two female) to be used as stimuli in the experiment proper.

Participants in the experiment proper were informed that the experiment consisted of two interwoven tasks, both of which involved predictions. Then they practiced each of these tasks. The person-judgment task was the distrust manipulation. In each person-judgment trial, participants were shown the set of questions and the answers attributed to another person (as described above). They were told that they could read and think about these answers for as long as they needed. They were also informed about the judgment they were asked to make. When they were ready to report their judgment they pressed a key, making the questions and answers disappear. In their place, participants saw a judgment scale with a question. The kind of judgment that the participants made determined the experimental condition (distrust vs. control).

In the distrust condition, participants were informed about the actual procedure for constructing the questions and answers and asked to rate how confident they were that the person who gave the answers was a female (truth-teller) or an imposter (i.e., a male). Ratings were made on a 5-point scale (1 = certainly an imposter, 5 = certainly a female). Unlike the controls (see below),

participants in the distrust condition were confronted with attempts to trick them. We reasoned (Schul et al., 2004) that although the threat of deception is minor, in the sense of having no consequences to the participant's well-being, the effort to ferret out imposters would nonetheless be sufficient to activate the mental structures associated with deception, and therefore to induce a state of diffuse distrust.

In the control condition participants were instructed to rate the spontaneity of the person who provided the answers, without being told how the answers were elicited. Ratings were made on a 5-point scale (1 = not spontaneous, 5 = very spontaneous). We consider it reasonable that this orienting task induced a state of trust because this is the default state when people are trying to be cooperative with no reason for doubt (Gilbert, 1991). Still, we refer to it as a "control" to denote that it is not an explicit manipulation of trust.²

Procedure

After practicing the person-judgment task (see Distrust manipulation, above), participants practiced the job-prediction task. Here, they were asked to play the role of a placement officer who makes judgments about an applicant's potential success on a job on the basis of two important characteristics. However, they were not told either what these characteristics were or what the job was. Rather, they were told that they would be making a series of predictions about the applicants' potential success on the basis of two characteristics, and they were encouraged to improve their prediction success from trial to trial based on the feedback they would get.

The predictors (the two characteristics) were indicated by two bar graphs. Participants were informed that the higher the bar, the more the applicant has of the particular characteristic. The prediction was made by keying in a number from 100 to 800. Following their prediction, participants were given "the actual success of the applicant", a value (ranging from 100 to 800) that was generated on the basis of either the positive Eq. (1a) or the negative Eq. (1b) prediction rule (a between-participants condition). Participants underwent 10 practice trials to familiarize them with the procedure and the functional relationship between the cues and the outcome.

Following the practice trials, participants made two additional person-judgments, reinforcing (or not) the distrust frame of mind that had been activated during practice. Then they made 40 additional predictions.

Design

The experiment consisted of a 2 (control vs. distrust) \times 2 (positive-linear vs. negative-linear rule) between-participants design.

Results

Our analysis has to do with predisposition for activating an abstract general inference rule in the early stages of prediction. We hypothesized that under trust individuals activate routine rules (e.g., positive-linear) and under distrust non-routine ones (e.g., negative-linear). This implies an interaction, namely, under trust participants should succeed more in a routine environment and under distrust more in a non-routine environment. However, since both rules are easy, participants should realize fairly early that they are making systematic mistakes while using the rules.

² Schul et al. (2004) used this manipulation in Experiment 2 and a more direct manipulation of trust and distrust in Experiment 1. The pattern of findings was similar in the two experiments, reinforcing the suggestion that participants in the control condition were actually trusting and that the two distrust manipulations indeed induced distrust.

Having done so, they are likely to search for and switch to an alternative rule. This implies that the predicted interaction should be more likely early in the sequence of trials.

We partitioned the 40 trials into two blocks. For each participant we generated a prediction equation, predicting his or her response in each of the 20 trials within a block from the magnitude of the two cues (i.e., x_1 and x_2). Prediction success was indicated by the proportion of the variance (R^2) accounted for by the two cues. If our analysis is correct, the predicted interaction should be more robust in the first than in the second block.

A 3-way mixed-model ANOVA (control/distrust and positive/negative rule as between factors, block as a within-participant factor) was performed on prediction success (computed within block). Means are presented in Table 1. The analysis of the first block indicated the predicted interaction, $F(1, 75) = 6.61, p < .05$. The analysis of the second block gave no support for the predicted interaction, $F(1, 75) = 0.08, p = .77$. The dissimilarity of the two interaction patterns is revealed in a significant three-way interaction, $F(1, 75) = 6.17, p < .05$.

Interestingly, participants in the distrust condition who had to discover and use the positive-linear rule made the greatest improvement in prediction success, from $M = .58$ (block1) to $M = .74$ (block2), $t(75) = 2.56, p < .05$. Indeed, during the second block distrustful/positive-rule participants were slightly (but non-significantly) more successful than positive-rule controls ($M = .72$). This may suggest that having discovered that the initial rule does not fit, distrustful individuals were readier to shift to an alternative rule and thereby increase their prediction success on later trials.

Finally, the prediction success for the negative-linear rule was not significantly inferior to that for the positive rule, even when we consider only the control participants. This is at odds with the findings of Brehmer (1974). We can only speculate that the participants suspected, prior to the end of the first block or perhaps even during the practice trials, that the rule they were testing (presumably, the positive-linear one) did not fit the data, and they therefore started to search for an alternative rule. Experiment 2 makes use of considerably more difficult rules to avoid this problem.

Discussion

A state of distrust is the mental system's signal that one should be on guard because appearances might be misleading. In these circumstances routine responses are problematic because they can be anticipated by those seeking to mislead or harm the protagonist. Conversely, a state of trust is associated with a feeling of safety and transparency, indicating that the situation is as it appears and strategies that were successful in the past are equally appropriate now, or that there is no reason why they should not continue to be used. The results of Experiment 1 are consistent with this line of reasoning. When the to-be-predicted events were congruent with an inference rule which is more prevalent in normal environments (i.e., positive-linear), participants who trusted

Table 1
Prediction success according to orienting task, type of rule, and block (Experiment 1)

	Positive-linear	Negative-linear
<i>Block 1</i>		
Control	.71	.68
Distrust	.58	.79
<i>Block 2</i>		
Control	.72	.78
Distrust	.74	.82

discovered the rule more readily and so were more successful than those who distrusted; however, when the to-be-predicted events fit the less prevalent (i.e., negative-linear) rule, distrustful participants discovered the rule more readily and were more successful than trustful ones. Moreover, this effect occurred only in the first half of the experiment, suggesting that the different mental states influenced the kind of rule that came to mind early in the series of predictions.

The hyper-sensitivity of distrustful participants to negative-linear rules was explained by assuming that (i) positive-linear rules dominate negative-linear rules in terms of likelihood of activation, and (ii) participants who distrusted were tuned to the less dominant rules. However, it could be argued that the superior performance of distrustful individuals in predicting the negative-linear rule is due to their vigilance to the presence of inverse relationships, rather than the rule's relative prevalence in normal environments. Such an argument resonates with the findings of Schul et al. (2004) if we assume that the tendency to activate incongruent associations, which is what has been observed under distrust, is simply another manifestation of a more general tendency of being tuned to inverse relationships.

In Experiment 2 we employed rules that bear no obvious connection to inverse relationships. This allowed us to test unambiguously whether those who distrust are in fact hyper-sensitive to non-routine operations. Moreover, the greater difficulty associated with the non-linear rules implies that the pattern of prediction success in Experiment 2 should differ from that observed in Experiment 1. Specifically, we hypothesized that due to the difficulty of rules in Experiment 2, participants in both conditions would fail to identify the rule in the early trials. However, because those who distrust are more open to non-routine rules, they should be more likely to pursue alternative avenues in searching for the rule, implying that in the later trials their performance should be superior to that of the controls.

Experiment 2

Experiment 2 replicated the design of Experiment 1, using a different pair of inference rules that allow prediction from a pair of predictors to an outcome. According to the *max* rule, outcomes are affected by the maximum of the two predictors, whereas according to the *min* rule, they are affected by the minimum of the two (see Method section below). Assuming that x_1 and x_2 are positively-scaled mental constructs (e.g., motivation and ability), discovering and learning the *max* rule should be easier when people focus on gains or when they have a positivity bias, since a high value on any of the predictors is associated with a high criterion. Conversely, learning the *min* rule should be easier if people are tuned to negative events—when they have a negativity bias. In this case people pay greater attention to the more negative of the two cues.

Rozin and Royzman (2001) summarized research showing the prevalence of a negativity bias (see Fazio, Eiser, & Shook, 2004 for more recent evidence), which may indicate that the *min* rule is more dominant, that it, it tends to be triggered mentally earlier than the *max* rule. Accordingly, given our argument that distrust tends to activate non-dominant rules while trust activates dominant rules, it follows that those who distrust should make more accurate predictions than those who trust in an environment that implements the *max* rule. We refer to this hypothesis as the strong version to denote that it is based on an assumption about the greater availability or expectedness of the *min* rule.

One could raise two main objections to the strong hypothesis. First, the mechanism through which negativity biases are translated into dominance of inference rules is not well-specified. One

might even argue that having a negativity bias sensitizes individuals to positive (incongruent) information rather than negative information (see the research on the incongruity effect in impression formation; Srull, 1981). Second, the likelihood of focusing on negative information depends on the context. In particular, it has been shown that the prevalence of a negativity bias is influenced by the nature of the judgment (Skowronski & Carlston, 1989), the motivational state of the individual (Higgins, 2002), and the object being evaluated (Ganzach & Czaczkes, 1995; Moon & Conlon, 2002). Therefore, the assumption about the dominance of the *min* rule in normal situations similar to our trust condition might be false.

When it is unknown whether the *min* rule or the *max* rule is more dominant under trust, our analysis makes a weaker prediction. In this case it is predicted that whatever tendency is found under trust will be weakened or reversed under distrust. Thus, both the strong and weak versions lead us to expect an interaction, namely, that prediction accuracy in the environments of *min* and *max* rules depends on whether the individuals are in a state of trust or distrust. The strong version predicts, in addition, a particular type of interaction. Specifically, people who trust should have an advantage over those who distrust in discovering or learning a *min* rule, and this advantage should be attenuated or even reversed in the case of a *max* rule.

Method

Participants and procedure

One hundred and eight students at the Hebrew University participated in the experiment. They were randomly assigned to the four experimental conditions (control/distrust \times min/max rule). Five participants had to be eliminated from the analyses because they failed to follow the instructions or because of technical failure. The experimental procedure was identical to that in Experiment 1.

Prediction environments

We generated, from a normal distribution, a set of 40 pairs of cues: x_1 ($M = 179$, $SD = 69$) and x_2 ($M = 185$, $SD = 70$). x_1 and x_2 were negatively correlated ($r = -.19$). We used them to generate two outcome variables, y_{\max} and y_{\min} , using an error term that was sampled from a normal distribution with $SD = 30$. The following regression equations related the outcomes to the cues:

$$\hat{Y}_{\max} = -15 + 0.09 * x_1 + 0.02 * x_2 + 1.97 * \max(x_1, x_2) \quad (R^2 = .96)$$

$$\hat{Y}_{\min} = 104 + 0.02 * x_1 + 0.03 * x_2 + 2.07 * \min(x_1, x_2) \quad (R^2 = .99)$$

As is apparent from the equations, the minimum/maximum term captures the difference between the two prediction rules. In fact, when this term is not involved in the prediction, the following two prediction equations are found:

$$\hat{Y}_{\max}^* = 75 + 1.02 * x_1 + 1.04 * x_2 \quad (R^2 = .68)$$

$$\hat{Y}_{\min}^* = 10 + 1.11 * x_1 + 1.02 * x_2 \quad (R^2 = .70)$$

Clearly, in the two cases the predicted score is an increasing function of both x_1 and x_2 , and the differences between the two are not readily apparent. This makes successful prediction in Experiment 2 more difficult, as individuals are likely to start with a simple rule that includes only the two predictors (x_1 and x_2) and only at a second stage try configural models that include the minimum or the maximum (e.g., Brehmer, 1974). This implies that unlike Experiment 1, the hypothesized interaction is more likely in the second block of trials.

As in Experiment 1, we also created an equivalent set which included only 10 predictors/outcome triads, to be used for practice.

Results

For each participant we computed two indices of prediction success. First, as in Experiment 1, we used the proportion of variance of the participant's predictions accounted for by the predictors (i.e., R^2). In the *min*-rule condition the predictors were x_1 , x_2 , and $\min(x_1, x_2)$, whereas in the *max*-rule condition they were x_1 , x_2 , and $\max(x_1, x_2)$. Second, we assessed the sensitivity of each participant to the configural term by computing the correlation between the maximum (*max*-rule condition) or the minimum (*min*-rule condition) and the participant's prediction. As might be expected, the two indices were highly correlated ($r(102) = .81$, $p < .05$), and therefore, we report only analyses based on the proportion of variance explained (R^2). As in Experiment 1, we assessed prediction success separately in the first and the second blocks of 20 trials. The means are displayed in Table 2.

A three-way mixed-model ANOVA (control/distrust \times rule \times block) revealed the hypothesized interaction in the second block, $F(1, 100) = 5.87$, $p < .05$. In the control condition, participants were more accurate under the *min*-rule (.70) than the *max*-rule (.59). In the distrust condition, in contrast, participants were more accurate in the environment that implemented the *max*-rule (.66) than in the one with the *min*-rule (.58). The first-block predictions failed to show significant main effects or interaction, $F_s < 1$. The difference between Experiments 1 and 2 in the pattern of the three-way interaction is consistent with the presumed greater difficulty of the non-linear rules used in Experiment 2 (Brehmer, 1974). That is to say, unlike in Experiment 1, participants in the present experiment had great difficulty in discovering the non-linear form of the rules early in the sequence, because they were most probably first trying rules involving the two predictors and omitting the configural term.

Discussion

Experiments 1 and 2 demonstrated that an orienting task which casts doubt on the veridicality of information influenced participants' ability to make accurate predictions even though the predictions were unrelated to the orienting task. We hypothesized, and found, that distrustful individuals detect the less prevalent contingencies more readily than trusting individuals. That is, their prediction performance was more accurate in a less typical environment, and less accurate in a more typical environment.

The findings can rule out several mechanisms that might be used to explain how the manipulation worked. It could be suggested that, unlike the control condition, the distrust manipulation sensitized participants to opposites, because the participants in the distrust condition had to consider two opposite alternatives (i.e., the answers were given by a truth-telling female or a cheating male). However, it is unclear why the awareness of opposites should improve performance with the minimum rule while hindering performance with the maximum rule.

The orienting task may have also influenced people's tendency to elaborate on the information and think about it critically.

Although critical thinking might be part of the story, it cannot be the whole story. In particular, it cannot explain the following two phenomena without evoking auxiliary assumptions: (i) obstruction of performance with the routine rule (i.e., why would those who distrusted perform less well with a routine rule than with a non-routine rule) and (ii) block effects (i.e., why would the rule-by-trust/distrust interaction occur in the first block in Experiment 1 and in the second block in Experiment 2). Moreover, we note in brief that this mechanism can be tested directly by determining whether one of the two orienting tasks was associated with more prolonged cognitive activity than the other. Using prediction latency as an indicator, we find no evidence in our data that the orienting task influenced prediction speed in either Experiment 1 ($F(1, 75) = 0.71$, $p = .40$) or Experiment 2 ($F(1, 100) = 2.68$, $p > .1$), nor did the orienting task interact with the type of rule or the block. The absence of significant interactions using prediction time as the DV contrasts with the presence of significant interaction using prediction accuracy as the DV.

Experiment 3 was designed to replicate the design of Experiments 1 and 2 using a considerably simpler manipulation and performance task.³ Specifically, the distrust manipulation in Experiments 1 and 2 was indirect, and its direct effect on trust or distrust was not assessed. Therefore, it could be argued that the manipulation primed participants' tendency to think innovatively, to come up with non-obvious answers, or to search for novel solutions. Although these might be part of a distrust state of mind, such effects might be also evoked independently.

Experiment 3 attempted to eliminate such alternative explanations. The manipulation of distrust/trust involved identical processing instructions for different stimuli—an untrustworthy or a trustworthy face. We expected that having to think about a face and keep it in memory would induce a state of mind compatible with the face. Moreover, because the processing instructions were identical, any difference between the conditions must reflect the nature of face that the participants viewed—trustworthy in one condition and untrustworthy in the other. Experiment 3 also used a different performance task than was used in the earlier experiments. In addition to being much simpler and having no error component, performance success was based on two rules that could be more readily classified as routine and non-routine.

Experiment 3

Trust and distrust were manipulated in this experiment by having participants process one of two faces, either a face which conveyed trust or one which conveyed distrust. Participants were asked to form an impression of the person whose face they saw, and to remember this impression. We reasoned that participants who processed the trustworthy face would be more likely to be in a trusting mode of thinking, while those who processed the untrustworthy face would be more likely to be in a distrust mode of thinking.

Once the trust/distrust mode had been induced, participants were given a series of matchstick arithmetic problems (see Method section below), all of which could be solved by moving a single matchstick. The problems were classified a priori as routine or non-routine. Analogously to the previous results, we predicted that trusting participants—those who had processed a trustworthy face—would be more successful in solving the routine problems, while distrusting participants—those who had processed the untrustworthy face—would be more successful in solving the non-routine problem.

³ We thank reviewers of an earlier version of this paper for pointing out these issues.

Table 2
Prediction success according to orienting task, type of rule, and block (Experiment 2)

	Max	Min
<i>Block 1</i>		
Control	.58	.64
Distrust	.62	.64
<i>Block 2</i>		
Control	.70	.59
Distrust	.58	.66

Method

Participants

Sixty individuals who were sitting by themselves in the university library were asked to participate in a short study for a small monetary compensation. Those who agreed were assigned randomly to the distrust or the trust condition.

Procedure

After signing the consent form, participants were given a booklet whose first page informed them that the study had several parts, the first of which involved looking at a person's face and forming an impression of him. Participants were also informed that later on they would be asked to recall this person's face and their impression of him. They were told that in order to make this more difficult they would be asked to work on a different task for a while.

After viewing the face for as long as they wanted, participants were given the matchstick task (see below). They had 3 min to solve eight matchstick arithmetic problems. Once the time was up, participants were asked to describe in their own language what they thought when they were looking at the person's face. Then they were given a list of characteristics (wise, warm, swindler, happy, shy, trustworthy, independent, romantic, competitive, friendly, young) and asked to indicate the characteristics that fit the person whose face they saw. Finally, they were asked whether they had solved similar arithmetic problems in the past.

The faces

Participants were shown one of two faces that had been created in an earlier study (Schul et al., 2004) to convey either trustworthiness or untrustworthiness. The faces are reproduced in Fig. 1.

The matchstick task

Participants were given a series of eight arithmetic problems presented in Roman numerals (see Table 3). All the problems appeared on one page, together with a conversion table showing how the Roman numerals I–XII are translated into standard Arabic numerals. The participants were asked to imagine that the Roman numerals were written with matchsticks. To help them in this visualization, we used a font specifically designed to display Roman numerals as if made from matchsticks. The participants were informed that each of the problems could be solved by moving a single matchstick which transforms an incorrect equation into a correct one. To “move” the matchstick, they had to draw a matchstick in the new location and an arrow to connect it to the matchstick's original position. The participants were informed that they had 3 min to solve all eight problems.

Unbeknown to the participants, there were two types of problem. Problems 1–4 involved the transformation of one Roman numeral into another (see Table 3). The transformations required for the solution of problems 5–8, in contrast, involved removing a matchstick from a number and using it to change an arithmetic sign. The former have been shown to be more typical or routine, while the latter are less typical and non-routine (e.g., Ollinger & Knoblich, 2003). Moreover, it is likely that the conversion table which appeared on the same page cued participants to the conversion of Roman to Arabic numerals, thereby reinforcing their natural tendency to think of the routine transformation rule.

Design

The experiment manipulated two factors: the participant's mode of thinking (trust vs. distrust) and the type of problem (routine vs. non-routine). The mode of thinking is a between-participants factor; the problem type is a within-participant factor.

Trustworthy face



Untrustworthy face



Fig. 1. The faces shown to trigger trust (top) or distrust (bottom) in Experiment 3.

Table 3
The matchstick arithmetic problems (Experiment 3)

Original (incorrect) equation	Correct equation
<i>Routine transformation (change of two numerals)</i>	
IX = X – III	IX = XI – II
VI = VIII – IV	IV = VIII – IV
II = III + I	III = II + I
VII = II + III	VI = III + III
<i>Non-routine transformation (change of a numeral and an operator)</i>	
XI = VI – IV	X = VI + IV
VII = III – III	VI = III + III
II = VII + IV	III = VII – IV
IX = X – VI	IV = X – VI

Results

One participant failed to understand the instructions and five participants had solved similar matchstick problems in the last year. The results of these participants were eliminated from the analyses. Next we examined how participants described the face they had seen, that is, which characteristics they selected as fitting the person whose face it was. Our goal was to eliminate those participants who had failed the manipulation, namely, those who interpreted the face unambiguously in the opposite way. Specifically, we eliminated those who had seen the trustworthy face and indicated that it belonged to a swindler (5 out of the 27 participants), and those who had seen the untrustworthy face and indicated that it belonged to a trustworthy person (1 out of 27

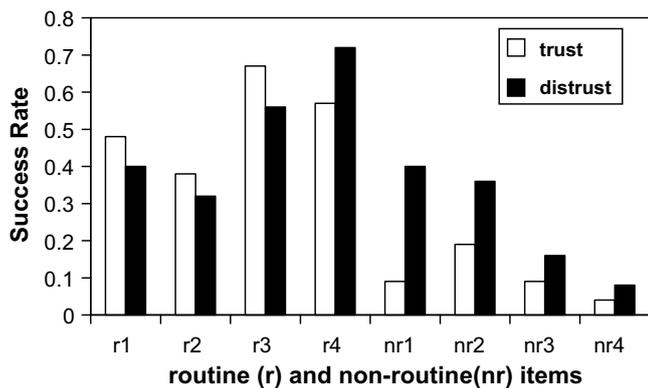


Fig. 2. Rate of correct solution as a function of the induced mode of thinking and type of problem (Experiment 3).

participants).⁴ Thus, the analyses reported below are based on the results of 48 participants.

Fig. 2 presents the proportion of participants who correctly solved each of the eight problems. For three of the four routine problems, participants in the trust condition succeeded more frequently than those in the distrust condition. In contrast, for all four non-routine problems those in the distrust condition succeeded more frequently than those in the trust condition. A two-way ANOVA (routine vs. non-routine problems \times trust vs. distrust manipulation) revealed that the routine problems (mean probability of success = .51) were easier on average than the non-routine ones ($M = .18$), $F(1,44) = 55.12$, $p < .05$. More importantly, the differences between the routine and the non-routine problems varied as a function of the manipulation. The planned interaction contrast was statistically significant, $t(44) = 1.86$, $p < .05$, one tail.

Interestingly, the pattern of success for the routine items is analogous to the one found in Experiment 1. Both experiments involved processing rather easy problems and in both of them participants in the distrust condition realized after a few performance trials that they were on the wrong track and needed to switch to a different rule. Accordingly, unlike the case of the first three routine problems, the success of the distrusting participants in solving the fourth routine problem was superior to that of participants in the trust condition.

In the case of the non-routine problems, the differences between the trust and the distrust conditions were more pronounced in the problems which appeared earlier in the sequence. The third and, even more so, the fourth non-routine problem, yielded very few solutions. This was probably due to the time constraint; participants failed on the last problem because the 3-min interval allowed for responding had ended before they had time to consider it.

In our final analysis, then, we compared the first three routine problems to the first three non-routine problems. A two-way ANOVA revealed that the three routine problems were successfully solved more often than the three non-routine problems, $F(1,44) = 28.12$, $p < .05$. Importantly, the interaction contrast which compared the differences due to problem type (routine vs. non-routine) for the trust and distrust conditions was more pronounced in the current analysis, $t(44) = 2.76$, $p < .05$, than in the analysis based on all eight problems. Specifically, trusting participants per-

formed much better with routine problems ($M = .51$) than with non-routine ones ($M = .13$). Distrusting participants succeeded less often than trusting participants with the routine problems ($M = .43$), but more often than trusting participants with the non-routine ones ($M = .31$).

General discussion

Distrust is typically viewed as a mental state caused by the threat of being deceived. Research shows that people who are in such a state behave differently toward others than do people who trust (Kramer, 1999). Since (i) everyone must deal with potentially untrustworthy individuals many times over their lifetime, (ii) these encounters are often socially or otherwise significant, and (iii) this has probably been the case throughout the history of our species as well as that of other group-living primates, it is likely that the mental representation of distrust has evolved into a unified structure that integrates the cognitive, affective, and action tendencies associated with the possibility of deception (Cosmides & Tooby, 2005). In this respect, distrust resembles other integrated mental states such as our system of emotions (e.g., Lerner, Small, & Loewenstein, 2004; Niedenthal et al., 2005). The existence of an integrated representation implies that action tendencies which are associated with distrust can be triggered even in the case of unfocused distrust, when the cause of the distrust is difficult to specify and completely unrelated to the information being processed.

Indeed, the three experiments reported in this manuscript showed that manipulating participants' distrust through an orienting task influenced their performance on another task unconnected to the former. This prediction was based on the intuition that since distrust is associated with the notion that appearances are false or misleading, it sensitizes individuals to departures from the expected and increases the likelihood that they will search for irregularities and non-routine contingencies. The current findings, together with those reported by Schul et al. (2004), are consistent with this intuition.

Our approach to distrust is likely to recall Zajonc's (1965) classic $D \times H$ analysis of social facilitation, in which he demonstrated that the dominant response in an individual's hierarchy of response tendencies is increasingly likely to be performed as arousal due to the presence of others increases. Or, in our terms, an aroused person is more likely than a non-aroused person to do what he or she routinely does in the situation. Let us assume that distrust causes greater arousal than trust. If so, then our analysis suggests an important qualification to Zajonc's model: when the arousal is due to distrust of others, then aroused individuals are less likely to do what they routinely do than non-aroused individuals.

Could it be that what created the differences between our experimental conditions was participants' mood, being positive (in trust) or negative (in distrust)? We view trust and distrust as affects, with positive and negative valences, respectively. Nonetheless, we believe it unlikely that the pattern of our findings could be predicted a priori from the valence alone. The reason being that the presumed action tendencies associated with negative emotional states (greater caution and more stimulus-based processing) and positive emotional states (greater risk-taking and more heuristic processing, e.g., Clore & Huntsinger, 2007) cannot be mapped easily into the tendency to detect and use routine or non-routine rules. Moreover, there are important distinctions that need to be made within the domain of negative affects, because different negative emotions can be shown to be associated with different action tendencies (e.g., Gilbert & Miles, 2000; Lerner et al., 2004; Zeelenberg, van Dijk, Manstead, & van der Plicht, 2000). Attributing a specific action tendency, such as searching for non-routine solutions, to the valence alone implies

⁴ Twenty of the 27 participants who saw the untrustworthy face reported their impression that the person was a swindler, while 6 out of the 27 participants who saw the trustworthy person indicated that the person was trustworthy. On the one hand, this may mean that the criterion for trustworthiness is stricter than that for a swindler in our sample. On the other, trusting may be the default reaction when one forms an impression, so that it is less likely to be noted than distrust.

that all negative (or positive) affective states should manifest it. We find this unlikely.

Trust and distrust have been used so far as if they were antonyms. We would like to emphasize that this was done for ease of communication. In fact, we believe that the two span a continuum, so that the absence of trust need not activate a state of distrust or vice versa (Ullmann-Margalit, 2004). Of particular interest is the in-between region where individuals are neither fully trusting nor fully distrusting. In any event, situations of pure trust are associated with coordination games and those of pure distrust with zero-sum games. Psychologically, this means that in the former cases individuals believe that the other has only benign intentions, shares their interest totally, and what he or she says is unquestionably valid. In the case of extreme distrust, individuals are equally confident that the other's intentions are totally malign, his or her interests are wholly incompatible with their own, and what he or she asserts is best interpreted according to a theory they have about how, given the situation, others would likely try to dupe them. To illustrate, consider the impact of a witness's testimony on jurors who know that the witness has a strong vested interest. Not only will the jurors doubt the witness' testimony, but they may well think of the alternative position as likely, so that the testimony of a distrusted person can boomerang (Dean, Austin, & Watts, 1971; Jasperson & Fan, 2002). Moreover, untrustworthy individuals seem to be more memorable than trustworthy ones (Mealy, Daood, & Krage, 1996) and perhaps more easily recognized as such by naïve judges (Yamagishi, Tanida, Mashima, Shimona, & Kanazawa, 2003). This is consistent with the conjecture that people's lay theory of deception influences their attention, detection and recall, as well as how they encode what is stated by a distrusted source.

In contrast to the extreme cases, the in-between region of neither-trust-nor-distrust includes cases where receivers are unsure about the veracity of message they receive. Consequently, they are unable either to embrace the messages or completely reject them. The solution in many cases is to accept them and yet to elaborate on them by constructing arguments for and against the messages. Put differently, when receivers are uncertain about the trustworthiness of a source they think about implications that follow the message (i.e., they behave as if it were true) as well as implications consistent with its opposite (i.e., they behave as if it were false). Schul et al. (1996), as well as Fein (1996), referred to this state as suspicion and reported findings that support this reasoning.

Finally, a caveat is in order. Our experiments involved diffuse distrust, a state in which distrust is activated without being attached to a particular person (and, in fact, distrusters may even be unaware of the reason for their concern). We also discussed an alternative state, focused distrust, where the threat of deception is associated with a specific individual. We believe diffuse and focused distrust might have different cognitive consequences: The diffuse state might encourage perceivers, first, to process incongruent or unusual implications of a message; and, second, to avoid taking appearances at face value, but rather to be on the lookout for some (unspecified) unexpected or uncommon event. We think that focused distrust, on the other hand, might lead to a more rigid information-processing strategy. The reason for this is that focused distrust, by definition, is triggered when receivers bring to mind a theory that describes why and how the threatening source would generate a misleading message. Such a theory provides the distrustful person with a recipe for interpreting the message (e.g., recall the ironic interpretation many gave to Nixon's statement "I am not a crook!"). Accordingly, unlike individuals with diffuse distrust, those with focused distrust are tuned to cues relevant to their theory. Their theory may often be wrong, as research on the detection of deception

shows (Anderson, DePaulo, Ansfield, Tickle, & Green, 1999; DePaulo, Charlton, Cooper, Lindsay, & Muhlenbruck, 1997). And, if so, then by elaborating in terms of a wrong theory, receivers are likely to miss the opportunities afforded by the obvious as well as the non-obvious implications of the message.

Acknowledgments

Preparation of this paper was supported by grants from the US-Israel Binational Science Foundation (BSF), the Israel Academy of Science (ISF), and from K Mart Center for Retailing and International Marketing. We wish to thank Shirly Ronen and Maytal Bar-Hai for their help in running the experiments and coding the data, and Gideon Keren and Naomi Yahalom for commenting on an earlier version of this paper.

References

- Anderson, D. E., DePaulo, B. M., Ansfield, M. E., Tickle, J. J., & Green, E. (1999). Beliefs about cues to deception: Mindless stereotypes or untapped wisdom. *Journal of Nonverbal Behavior, 23*, 66–89.
- Brehmer, B. (1974). Hypotheses about relations between scaled variables in the learning of probabilistic inference tasks. *Organizational Behavior and Human Decision Processes, 11*, 1–27.
- Brehmer, B. (1980). Effects of cue validity on learning of complex rules in probabilistic inference tasks. *Acta Psychologica, 44*, 201–210.
- Clare, G. L., & Huntsinger, J. R. (2007). How emotions inform judgment and regulate thought. *Trends in Cognitive Science, 11*, 393–399.
- Cosmides, L., & Tooby, J. (2005). Neurocognitive adaptations designed for social exchange. In D. Buss (Ed.), *The handbook of evolutionary psychology* (pp. 584–627). New York: Wiley.
- Dean, R. B., Austin, J. A., & Watts, W. W. (1971). Forewarning effects in persuasion: Field and classroom experiments. *Journal of Personality and Social Psychology, 18*, 210–221.
- DePaulo, B. M., Charlton, K., Cooper, H., Lindsay, J. J., & Muhlenbruck, L. (1997). The accuracy-confidence correlation in the detection of deception. *Personality and Social Psychology Review, 1*, 346–357.
- DePaulo, B. M., Lindsay, J. J., Malone, B. E., Muhlenbruck, L., Charlton, K., & Cooper, H. (2003). Cues to deception. *Psychological Bulletin, 129*, 74–118.
- Doherty, M. E., & Kurz, E. M. (1996). Social judgment theory. *Thinking and Reasoning, 2*, 109–140.
- Eitam, B., Hassin, R. R., & Schul, Y. (2008). Non-conscious goal pursuit in novel environments: The case of implicit learning. *Psychological Science, 19*, 261–267.
- Fazio, R. H., Eiser, J. R., & Shook, N. J. (2004). Attitude formation through exploration: Valence asymmetries. *Journal of Personality and Social Psychology, 87*, 293–311.
- Fein, S. (1996). Effects of suspicion on attributional thinking and the correspondence bias. *Journal of Personality and Social Psychology, 70*, 1164–1184.
- Ganzach, Y., & Czaczkes, B. (1995). The learning of natural configural strategies. *Organizational Behavior and Human Decision Processes, 63*, 195–206.
- Gilbert, D. T. (1991). How mental systems think. *American Psychologist, 46*, 107–119.
- Gilbert, P., & Miles, J. N. V. (2000). Sensitivity to social put-down: Its relationship to perceptions of social rank, shame, social anxiety, depression, anger and self-other blame. *Personality and Individual Differences, 29*, 757–774.
- Higgins, E. T. (2002). How self-regulation creates distinct values: The case of promotion and prevention decision making. *Journal of Consumer Psychology, 12*, 177–191.
- Jasperson, A. E., & Fan, D. P. (2002). An aggregate examination of the backlash effect in political advertising: The case of the 1996 U.S. Senate race in Minnesota. *Journal of Advertising, 31*, 1–12.
- King-Casas, B., Tomlin, D., Anen, C., Camerer, C. F., Quartz, S. R., & Montague, E. (2005). Getting to know you: Reputation and trust in a two-person economic exchange. *Science, 308*, 78–83.
- Koslow, S. (2000). Can the truth hurt? How honest and persuasive advertising can unintentionally lead to increased consumer skepticism. *The Journal of Consumer Affairs, 4*, 245–268.
- Kramer, R. M. (1999). Trust and distrust in organization: Emerging perspectives, enduring questions. *Annual Review in Psychology, 50*, 60–98.
- Lerner, J. S., Small, D. A., & Loewenstein, G. (2004). Heart strings and purse strings: Carryover effects of emotions on economic decisions. *Psychological Science, 15*, 337–341.
- McKenzie, C. R. M. (2004). Hypothesis testing and evaluation. In D. J. Koehler & N. Harvey (Eds.), *Blackwell handbook of judgment and decision making* (pp. 200–219). Oxford: Blackwell.
- Mealy, L., Daood, C., & Krage, M. (1996). Enhanced memory for faces of cheaters. *Ethology and Sociobiology, 17*, 119–128.
- Moon, H., & Conlon, D. E. (2002). From acclaim to blame: Evidence of a person sensitivity decision bias. *Journal of Applied Psychology, 87*, 33–42.
- Niedenthal, P. M. et al. (2005). Embodiment in attitudes, social perception, and emotion. *Personality and Social Psychology Review, 9*, 284–311.

- Ollinger, M., & Knoblich, G. (2003). Prioritizing the new: Insight creates set for familiar problems. In F. Schmalhofer & R. Young (Eds.), *Proceedings of the first European cognitive science conference* (pp. 253–258). Mahwah, NJ: Erlbaum.
- Robinson, S. L. (1996). Trust and breach of the psychological contract. *Administrative Science Quarterly*, 41, 574–599.
- Rozin, P., & Royzman, E. B. (2001). Negativity bias, negativity dominance, and contagion. *Personality and Social Psychology Review*, 5, 296–320.
- 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., Burnstein, E., & Yahalom, N. (2007). How people cope with uncertainty due to chance or deception. *Journal of Experimental Social Psychology*, 43, 91–103.
- Seabright, P. (2004). *The company of strangers: A natural history of economic life*. Princeton University Press.
- Sedikides, C., & Skowronski, J. J. (1991). The law of cognitive structure activation. *Psychological Inquiry*, 2, 169–184.
- Skowronski, J. J., & Carlston, D. E. (1989). Negativity and extremity biases in impression formation: A review of explanations. *Psychological Bulletin*, 105, 131–142.
- Srull, T. K. (1981). Person memory: Some tests of associative storage and retrieval models. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 7, 440–463.
- Ullmann-Margalit, E. (2004). Trust, distrust, and in between. In R. Hardin (Ed.), *Distrust* (pp. 60–82). New York, NY, USA: Russell Sage.
- Winston, J. S., Strange, B. A., O'Doherty, J., & Dolan, R. J. (2002). Automatic and intentional brain responses during evaluation of trustworthiness of faces. *Nature-Neuroscience*, 5, 277–283.
- Yamagishi, T., Cook, K. S., & Watabe, M. (1998). Uncertainty, trust, and commitment formation in the United States and Japan. *American Journal of Sociology*, 104, 165–194.
- Yamagishi, T., Tanida, S., Mashima, R., Shimona, E., & Kanazawa, S. (2003). You can judge a book by its cover: Evidence that cheaters may look different from cooperators. *Evolution and Human Behavior*, 24, 290–391.
- Zeelenberg, M., van Dijk, W. W., Manstead, A. S. R., & van der Pligt, J. (2000). On bad decisions and disconfirmed expectancies: The psychology of regret and disappointment. *Cognition and Emotion*, 14, 521–541.