

RESEARCH ARTICLE

WILEY

The common origin of both oversimplified and overly complex decision rules

Taly Bonder¹  | Ido Erev¹ | Elliot A. Ludvig²  | Yefim Roth³ 

¹Data and Decision Sciences, Technion-Israel Institute of Technology, Haifa, Israel

²Psychology, University of Warwick, Coventry, UK

³Human Services, University of Haifa, Haifa, Israel

Correspondence

Ido Erev, Data and Decision Sciences, Technion-Israel Institute of Technology, Haifa, Israel.

Email: erev@technion.ac.il

Abstract

Many deviations from rational choice imply the neglect of important evidence and suggest the use of simple heuristics. In contrast, other deviations imply sensitivity to irrelevant evidence and suggest the use of overly complex rules. The current analysis takes two steps toward identifying the conditions that trigger these contradictory deviations from efficient reasoning. The first step involves a theoretical analysis. It shows that the contradictory deviations can be captured without assuming the use of rules of different complexity in different settings. Both deviations can be the product of a reliance on small samples of similar past experiences. This reliance on small samples triggers apparent overcomplexity when the optimal rule is simple, but more complex rules yield better outcomes in most cases; the opposite tendency, oversimplification, emerges when the optimal rule is complex, and simple rules yield better outcomes in most cases. The second step involves a preregistered experiment with 325 participants (Mechanical Turk workers). The experiment shows that human decision makers exhibit the pattern predicted by the reliance-on-small-samples assumption. In the experiment, participants chose between the status quo and a risky alternative in a multi-attribute decision with three binary cues. They used uninformative cues when this strategy was best in most cases yet ignored two informative cues when this strategy was best in most cases. In addition, describing the cues as recommendations given by three experts increased the tendency to follow the modal recommendation (even when reliance on only one of the experts was optimal), but people still behaved as though they relied on a small sample of past experiences.

KEYWORDS

advice taking, decisions from experience, heuristics, take-the-best, underweighting of rare events

1 | INTRODUCTION

People often behave as if they use simple adaptive rules that save cognitive effort by ignoring part of the relevant information

(Gigerenzer & Goldstein, 1996; Payne et al., 1993; Tversky & Kahneman, 1974). For example, only two of every 1000 retail software shoppers access the software license agreement (Bakos et al., 2014). In certain settings, however, people appear to exhibit the opposite bias and seem to use overly complex rules. One example is suggested by studies of classification decisions based on multiple

The author order was alphabetically determined.

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2023 The Authors. *Journal of Behavioral Decision Making* published by John Wiley & Sons Ltd.

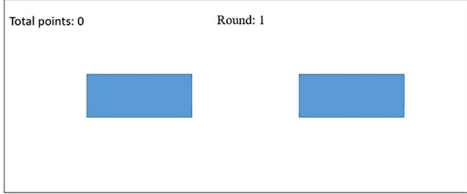
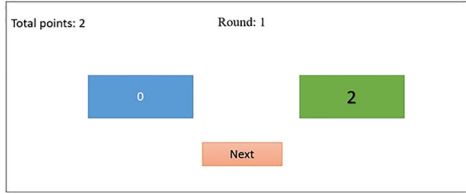
Basic instructions: In each trial of this study you will be asked to choose between two options. Your goal is to maximize your number of points.		
Trial	Pre-choice screen	Feedback screen
1		

FIGURE 1 Typical instructions and screenshot from one trial in the basic clicking paradigm for studying decisions from experience. On each trial, participants select one option and receive feedback about both options.

cues. Meehl (1954) found that the decisions of expert psychiatrists were less accurate than decisions derived from simple linear rules, and Dawes et al. (1989) showed that an important part of the experts' errors reflects a counterproductive effort to capture complex interactions. Another example involves the popularity of conspiracy theories, which are often built on complex explanations (Douglas et al., 2017; Marsh et al., 2022). For instance, over a third of Americans believe that global warming is a hoax (Swift-Hook, 2013), maintaining that mainstream scientific research reflects plots by powerful and malevolent groups.

These contradictory deviations from optimal reasoning can also be found in daily health-related decisions. For example, the belief that health problems reflect the impact of demons (a belief that was more popular in the middle ages) suggests the use of overly complex rules: rules that assume high sensitivity to factors that (at least not based on extant research) do not have a large effect on health. In contrast, the tendency to skip careful hand washing (also more common in the middle ages) suggests overly simplified rules that ignore the impact of germs that can be addressed by washing hands.

The leading explanations for using oversimplified rules focus on the role of cognitive costs and limitations (Payne et al., 1993). For example, the tendency to ignore license agreements can be explained as an effort to minimize cognitive effort. In contrast, the leading explanations for relying on overly complex rules focus on the role of motivated inference (Kunda, 1987). To illustrate, the belief that global warming is a hoax can be explained by assuming that it helps polluters maintain a positive self-image. Similarly, the belief in the existence of nonlinear patterns helps expert psychiatrists explain why their diagnostic decisions should not be overruled by a simple linear algorithm. Together these explanations imply that the direction of the deviations from optimal reasoning is a product of the relative importance of two classes of incentives: cognitive costs and the benefits from self-serving inferences.

One of the main challenges, under this cognitive costs and motivated inference explanation, is the clarification of the way these incentives impact reasoning. The current research addresses this challenge by focusing on simple situations in which behavior is likely to be

affected by objective incentives. Our analysis builds on two observations: First, recent research with simple choice tasks has demonstrated that variations in the objective incentive structure can change the direction of deviations from optimal choice (Erev & Plonsky, 2023). Second, the leading explanations for the impact of incentive structure on decisions from experience assume a reliance on small samples of past experiences.¹ In the context of multicue decision making, this assumption predicts behavior that appears to reflect the use of both oversimplified and overly complex decision rules. These observations are used below to predict sufficient conditions for the co-existence of these contradictory deviations from optimal reasoning. Sections 2 and 3 describe the two motivating observations and the way they can impact the direction of the deviations from optimal decisions in multicue tasks, Section 4 presents an experiment that tests these implied predictions, and Sections 5 and 6 highlight the implications of our results.

2 | CONTRADICTORY DEVIATIONS FROM MAXIMIZATION IN STATIC CHOICE TASKS

Erev & Plonsky, et al. (2023) reviewed research that examines decisions from experience in the basic clicking paradigm presented in Figure 1. In each trial of the reviewed experiments, the participants were asked to click on one of two or three keys. Each key was associated with a static payoff distribution (that did not change during the experiment), and each click was followed by the presentation of a single draw from each of the distributions. Each trial's payoff was the draw from the payoff distribution of the selected key, and the participants' goal was to maximize their overall earnings.

The review showed that certain changes in the incentive structure can reverse the direction of the observed deviations from maximization. Specifically, Erev et al. highlighted six pairs of contradictory

¹The value of the reliance-on-small-samples explanation was demonstrated in four choice-prediction competitions (see a review in Erev & Plonsky, 2023). The high predictive value of this hypothesis can be the product of an ecologically reasonable effort to respond to patterns (Plonsky et al., 2015).

deviations from maximization. Table 1 summarizes four problems (B1–B4) that demonstrate one of these six pairs. In each problem, participants chose between two options with different experienced payoffs. The choice rates in Problems B1 and B2 reveal a deviation from maximization that suggests underweighting of rare events (Barron & Erev, 2003). The participants behaved as if they maximized expected earnings under the belief that the probabilities of the rare outcomes (−10 and +10) are below 10%. In contrast, the choice rates in Problems B3 and B4 reveal a deviation from maximization that suggests overweighting of rare events. In this case, replacing “11 for sure” with a long-shot lottery with the same expected value increased the choice rate of Option Y. Participants now behaved as if they believed that the probability of the rare +20 outcome is above 10%.

In addition, Erev and Plonsky (2023) demonstrated that the direction of the deviations from maximization they documented can be captured with a simple, *sample-of-5* model (Erev & Roth, 2014). This model predicts random choice in the first trial and then a reliance on a sample of five past trials on all subsequent trials (randomly drawn with replacement; see Appendix A). For example, consider the decision in Trial 11 of a study that focuses on Problem B2, assuming that in the first 10 trials Option Y yielded nine “−1” outcomes and one “+10” outcome. Under the sample-of-5 model, the agent chooses Option Y only if the mental sample includes the +10 outcome. The probability that this mental sample of five past outcomes includes the +10 outcome is only $1 - 0.9^5 = 0.41$ (even though the observed average payoff from Y is positive). The sample-of-5 column in Table 1 presents the predictions of this model for Table 1's problems.

3 | RELIANCE ON SMALL SAMPLES IN MULTICUE DECISION TASKS

The observation that a reliance on small samples can reverse the deviations from maximization in basic choice tasks suggests that this tendency can also reverse the direction of the deviation from optimal reasoning in multicue tasks. To illustrate this point, we focus on the multicue choice paradigm presented in Figure 2. In each trial in this paradigm, the decision maker first sees a “signal,” consisting of three cues, and then chooses between the status quo and a risky bet that

can lead to a gain or a loss. Importantly, the signal provides partial information concerning the payoff distribution associated with the risky key. The exact payoff distribution given each signal is presented in Table 2. In this setting, a bias toward oversimplification entails ignoring part of the cues even when this behavior impairs performance, and the opposite bias implies sensitivity to cues (and to their interactions) that cannot help predict the sign of the expected outcomes.

Notice that Figure 2's multicue paradigm is a generalization of Figure 1's basic clicking paradigm. The main generalization is the possibility that the decision maker faces several incentive structures in the same experimental session and receives signals that allow discrimination between these different incentive structures. The basic clicking paradigm is the special case where the cues are always the same, and there is only one incentive structure. The main difference between Figure 2's paradigm and the paradigms used in studies of classification learning in multicue tasks (e.g., Juslin et al., 2003; Pachur & Olsson, 2012) involves the current focus on the incentive structure. In those previous classification studies, the participants' goal was to maximize the proportion of correct responses; instead, we focus on situations in which the decision makers are motivated to maximize the number of points they earn. This difference is likely to be important when the strategy that maximizes the probability of correct responses differs from the strategy that maximizes the expected number of points. For example, in a binary choice between a status quo and a risky choice leading to a gain of 2 points 90% of the time, and −20 otherwise, the strategy that maximizes the percentage of correct answers is “always choose the risky option”, and the strategy that maximizes the expected number of points is “never choose the risky option”.

3.1 | The signal-based-sample-of-5 model

The observation that Figure 2's multicue paradigm can be described as a generalization of Figure 1's basic paradigm suggests that the sample-of-5 model, proposed to capture behavior in the basic clicking paradigm, can also be generalized to capture behavior in the multicue paradigm. The present section evaluates this suggestion by focusing

TABLE 1 Summary of experimental demonstrations of underweighting and overweighting of rare events and the predictions of two reliance-on-small-samples models.

Problem	Option X Payoff	Option Y		Choice rate of option Y		
		Payoff	Expected value (EV)	Observed experimental results	Sample-of-5 model	PAS
B1	0 for sure	+1, 0.9; −10, 0.1	−0.1	58%	59%	59%
B2	0 for sure	−1, 0.9, +10, 0.1	+0.1	27%	41%	41%
B3	12 for sure	11 for sure	+11	12%	0%	11%
B4	12 for sure	10, 0.9, +20, 0.1	+11	22%	41%	24%

Note: The left-hand columns present the payoff distributions for Option X and Option Y in four different problems (B1–B4). The notation “v1, p1; v2, 1−p1” means “win v1 with probability p1, win v2 otherwise.” The right-hand columns present the results over 100 trials of experimental studies conducted by Nevo and Erev (2012) and Teoderescu et al. (2013) and the predictions of two models. The model PAS (partially attentive sampler) is described in Section 4.

<p>Basic instructions (see full instruction in Appendix B):</p> <p>In each trial of this study you will be asked to choose between two options. One option that maintains the status quo, and an option that can lead to a gain or a loss of points.</p> <p>The payoff from the risky option will vary from trial to trial. You will receive three cues that can help predict the exact payoff in each trial. <i>The cues are recommendations of three experts. “T” (take) implies that the expert recommends “Taking the risk”, and “A” (avoid) implies that the expert recommends avoiding the risk. For example, the sequence “TTA” implies that the first two experts recommend taking the risk, and the third expert recommends avoiding it.</i></p> <p>Your goal is to maximize your number of points. Your total number of points at the end of the experiment will determine the probability of earning the \$0.5 bonus in addition to the \$1 show up fee.</p>		
Trial	Pre-choice screen	Feedback screen
1	<p>Total points: 0</p> <p>Round: 1</p> <p>The three cues: TTT</p> <div> <div>Status quo (Zero for sure)</div> <div>Risk (Gain or loss)</div> </div>	<p>Total points: 2</p> <p>Round: 1</p> <p>The three cues: TTT</p> <div> <div>0</div> <div>2</div> </div> <p>Next</p>
2	<p>Total points: 2</p> <p>Round: 2</p> <p>The three cues: ATA</p> <div> <div>Status quo (Zero for sure)</div> <div>Risk (Gain or loss)</div> </div>	<p>Total points: 2</p> <p>Round: 2</p> <p>The three cues: ATA</p> <div> <div>0</div> <div>2</div> </div> <p>Next</p>

FIGURE 2 The basic instructions (used in Condition Experts in the experiment described below) and sample screenshot from two rounds in the multicue clicking paradigm. Participants selected one option on each round and received feedback about both options. In Condition Abstract, the italics were replaced with: *Each cue will be the letter “T” or the letter “A”. For example, the signal “TTA” implies that the first and second cues are “T”, and the third cue is “A”.*

on a simple generalization of the sample-of-5 model. The model, referred to as “signal-based-sample-of-5” (hereafter SB-sample-of-5), is similar to the sample-of-5 model with the additional constraint that only past experiences with the current signal are considered. Specifically, the SB-sample-of-5 model assumes random choice in the first experience with a specific signal and then, in subsequent reactions to that signal, a reliance on a sample of five previous experiences with that exact signal (drawn with replacement).

The SB-sample-of-5 column in Table 2 presents the predictions of this model for an experiment that examines repeated decisions in two sets of eight problems (the right-hand columns in Table 2 present the results of an experiment, described in Section 3, that tests these predictions). Each decision maker in this experiment faces a sequence of 160 choices: 20 rounds with each of the eight problems, in one of the sets, in random order. Each round involves a choice between the status quo (payoff of 0 with certainty) and a risky action that can lead to a gain or a loss. Before each choice, the decision makers receive a signal that includes three cues. As explained in Figure 2, each cue summarizes the recommendation of one of three experts. The cue T stands for the recommendation “take the risk,” and the cue A implies “avoid the risk.” Each choice is followed by the presentation of full feedback showing the payoff from the two possible choices in the last trial.

Set Rare-Treasure in Table 2 creates an environment in which the recommendation of one of the experts (the leftmost cue) always maximizes expected value (EV), and the recommendation of the other experts does not add payoff maximizing information. Thus, the EV-maximizing strategy is to follow the best expert. This strategy implies full reliance on the most effective cue and is thus an instance of the “take-the-best” heuristic rule (Bröder, 2000; Gigerenzer et al., 1991). Yet, given the signal TTT (taking the risk is recommended by all three experts), this optimal strategy leads to a loss in 90% of the cases and to a large gain in the other 10% of the cases. The SB-sample-of-5 model, however, tends to go against the experts, by avoiding this risky option in this one case, predicting behavior that appears to reflect an unduly complex theory. The predicted behavior of most “SB-sample-of-5” agents can be (incorrectly) explained with the assertion that they believe that if all experts recommend taking the risk, this agreement reflects a plot of “corrupted experts.”

Set Rare-Disaster is similar to Set Rare-Treasure with the exception of the outcome distribution given the signal TTT (Problem 8, when all three experts recommend taking the risk). All the experts err in this set (the expected value is negative), even though in 90% of the trials, the payoff is positive. Thus, the EV-maximizing strategy is relatively complex. It requires sensitivity to all three experts and prescribes the following rule: “avoid the risk if all three experts recommend taking it;

TABLE 2 Two incentive structures (sets of problems), predictions and experimental results.

Incentive structure				Predictions		Experimental results	
Problem	Signal	Payoff from risk	EV from risk	Optimal risk-rate	SB-sample-of-5	Condition Abstract	Condition Experts
Set rare-treasure						<i>n</i> = 80	<i>n</i> = 80
1	AAA	−2	−2	0	0.02	0.11	0.10
2	AAT	−2	−2	0	0.02	0.14	0.12
3	ATA	−2	−2	0	0.02	0.14	0.12
4	ATT	−2	−2	0	0.02	0.13	0.36*
5	TAA	+2	+2	1	0.98	0.82	0.57*
6	TAT	+2	+2	1	0.98	0.86	0.85
7	TTA	+2	+2	1	0.98	0.87	0.86
8	TTT	−2, .9; +20, .1	+0.2	1	0.41	0.32	0.59*
Set rare-disaster						<i>n</i> = 82	<i>n</i> = 83
1	AAA	−2	−2	0	0.02	0.15	0.09
2	AAT	−2	−2	0	0.02	0.14	0.09
3	ATA	−2	−2	0	0.02	0.16	0.11
4	ATT	−2	−2	0	0.02	0.22	0.45*
5	TAA	+2	+2	1	0.98	0.82	0.55*
6	TAT	+2	+2	1	0.98	0.91	0.91
7	TTA	+2	+2	1	0.98	0.89	0.92
8	TTT	+2, .9; −20, .1	−0.2	0	0.59	0.68	0.85*

Note: All problems involved a choice between the Status Quo (a payoff of 0) and the Risk. The payoff from the Risk is described in the table (notations as in Table 1). The right-hand columns present the results of the experiment described below. The bold emphasizes the difference between the two sets.

*Significantly different ($p < .001$) from the rate in the abstract condition.

in all other cases follow the best expert.” Notice that while the recommendation of the best expert impairs EV given the signal TTT, it always minimizes the probability of regret (it maximizes the probability of obtaining the best possible outcome in each choice). As in Set Rare-Treasure, most SB-sample-of-5 agents take the action that leads to the best outcomes most of the time. In the current set, this behavior can be (incorrectly) explained with the assertion that these agents learn to follow an oversimplified rule and take the risk if and only if the best expert predicts a positive expected return.

In summary, under the SB-sample-of-5 model, modifications to the incentive structure can change the direction of the deviation from optimal reasoning. In Set Rare-Treasure, agents that behave in accordance with this model choose as if they prefer a complex explanation over the take-the-best rule, even when the take-the-best rule maximizes the expected return. In contrast, in Set Rare-Disaster, these agents behave as if they follow the simple take-the-best rule, even when this rule is suboptimal.

3.2 | The impact of prior information

One shortcoming of the analysis presented above is that it focuses on virtual agents that base their decisions on past experiences in the experimental session and do not consider older past experiences. Human agents likely consider a wider set of past experiences,

including old experiences that resemble the current situation and have occurred before the beginning of the experimental session. Which past experiences human decision makers consider also likely depends on the description of the choice task. Specifically, the information that decision makers receive concerning the meaning of the cues can be important. For example, the cues could be described as recommendations of three experts, as above, or the cues could be described as abstract letters without a particular semantics. Describing the cues as recommendation from experts, as opposed to abstract letters, would be expected to increase the initial tendency to take the risk with an increase in the number of Ts in the signal. The experiment below thus also assessed the impact of prior information on choice in a multicue environment.

4 | EXPERIMENT

The current experiment examined human behavior in the two sets of problems described in Table 2, under two prior information (hereafter “description”) conditions. Condition Experts used the instructions presented in the main part of Figure 2 and described the signal as the recommendation of three experts. Condition Abstract dropped this cover story (see the note in Figure 2). The full instructions also included an answer-until-correct quiz and are presented in Appendix B.

4.1 | Method

4.1.1 | Participants and design

The experiment used a 2×2 between-participant experimental design. Each participant faced one of the sets under one of the two description conditions for 160 trials. The participants were 325 Mechanical Turk workers, recruited through Cloud Research (48% females, mean age 42). They were randomly assigned to one of four groups. Each group included at least 80 participants. This sample size provides 95% power to detect an effect size $d = 0.57$ (the magnitude of the effect of the incentive structure given the signal TTT under the SB-sample-of-5 model) with $\alpha = .05$ in a pairwise test. Only those who completed the experiment were counted as valid participants. On average, the participants completed the task in 12 min (the maximum time allowed was 1 h). The participants received a show-up fee of \$1 for completing the task and, as indicated in the instructions, had a chance to win a bonus of \$0.5 depending on the points gained (63% of the participants earned the bonus). The exact probability of winning the bonus increased linearly with the accumulated number of points; this bonus computation rule induces risk neutrality by rational agents (Roth & Malouf, 1979). All participants provided informed consent before participating, and the study was approved by the Technion-Israel Institute of Technology ethics committee.

4.1.2 | Procedure

After providing consent, each participant received general instructions about the task (see Appendix B) and made a total of 160 choices. On each round, they picked between the status-quo option and the risky option (see Figure 2). These were indicated by two blue boxes on the screen with the corresponding text written on them. The location (left or right) of the status-quo option and the risky option were counterbalanced (but fixed for each participant). Text on top of the screen indicated the round number (out of 160), the total points earned thus far, and the cues for that round (e.g., TTT). Participants indicated their selection by a mouse click. After each selection, the selected key was highlighted and the points earned were displayed, until participants clicked on the “next” button and moved on to the next trial. Payoff from the unselected option was also displayed (i.e., full feedback) on the other key. The next choice was then immediately available.

Each participant faced each of the eight problems 20 times. The problems each had a different three-letter cue and a different underlying pay out rate (see Table 2). The problem presented on each trial was randomly determined (drawn without replacement). That is, the probability of facing the same problem on consecutive trials was around 1/8. On the risky trials with the TTT cue, the rare outcome appeared 10% of the time (truly random). The experiment was programmed using oTree (Chen et al., 2016).

4.2 | Results

The right side of Table 2 presents the main results: the proportions of trials in which the risky option were selected (i.e., risk-rate) over all 20 trials with each of the eight problems. In agreement with the pre-registration (<https://osf.io/n38gm>), the results in Condition Abstract reveal the contradictory deviation from optimal responding to the signal TTT predicted by the SB-sample-of-5 model. When the simple take-the-best heuristic is optimal (Set Rare-Treasure) and predicts a choice of the risky option, the participants behaved as if they used a more complex choice rule: The observed risk-rate was only .32 (95% CI [.26, .38]), which is significantly lower than .5 ($t(79) = -5.75, p < .0001, d = 0.64$). In contrast, when the simple take-the-best heuristic is sub-optimal (Set Rare-Disaster), the participants behaved as if they used that heuristic. The risk-rate in this case was .68 (95% CI [.61, .75]), which is significantly higher than .5 ($t(81) = 5.01, p < .0001, d = .55$).

To clarify the impact of the set and the description condition on the reaction to the signal TTT, we conducted three analyses of variance. The first focuses on the mean choice rates over all 20 reactions to this signal (presented in Table 2). While the interaction between set and description was not significant ($F(1,321) = 2.61, p = .107$, partial $\eta^2 = .008$), the results reveal two main effects: The risk-rate was higher when the Risk option increased the probability of success (Set Rare-Disaster) than when the Risk option maximized EV (Set Rare-Treasure; $F(1,321) = 50.4, p < .001$, partial $\eta^2 = .227$), and was higher in Condition Experts than in Condition Abstract ($F(1,321) = 94.1, p < .001$, partial $\eta^2 = .136$).

The two additional analyses focused on the last 10 responses to the signal TTT and distinguished between the decisions made when the observed payoff from the Risk option (over all previous experiences with the signal TTT) was positive or negative. The mean risk-rates in these cases are presented in Table 3. Notice that the probabilistic payoff rule means that “lucky” participants made all these choices after observing positive average payoffs, and some “unlucky” participants made all these choices after observing negative average payoffs.² The lower row in Table 3 shows that both analyses reveal two main effects and no significant interactions. In all cases, whether participants were lucky or unlucky, or whether they encountered the experts or abstract description, participants took the risky option more often in the rare-disaster set with the TTT cue. This behavior implies deviations from efficient use of the take-the-best rule even after observing that these deviations impair earnings.

4.2.1 | Individual differences

To assess the impact of individual differences, Figure 3 presents the risk-rate of each participant in the last five responses to the signal TTT as a function of the average payoff from the Risk option in the first 15 experiences with TTT. Note that the average payoff is


²For example, participants that faced the first +20 outcome in the 13th experience with signal TTT, in the last 10 trials (given the signal TTT) made 3 decisions with negative running average and 7 decisions with positive running average.

TABLE 3 The risk-rate during the last 10 presentations of the signal TTT, as a function of the sign of the average payoff from the risk option in the previous presentations of this signal.

Description condition	Set	Average payoff from risk over all previous experiences with TTT					
		Positive Risk-rate (n)			Negative Risk-rate (n)		
Abstract	Rare-treasure	0.42 (56)			0.10 (50)		
	Rare-disaster	0.80 (45)			0.52 (60)		
Experts	Rare-treasure	0.64 (64)			0.39 (53)		
	Rare-disaster	0.95 (52)			0.63 (56)		
		$F[1,213]$	p	Partial η^2	$F[1,215]$	p	Partial η^2
Significance test (analysis of variance)	Set	47.76	<.001	.070	41.95	<.001	.068
	Description	15.93	<.001	.191	13.47	<.001	.163
	Interaction	0.50	=.471	.002	3.17	=.076	.014

Note: n is the number of participants that contributed to the computation of each statistic. The sum is larger than the number of participants as some participants experienced both positive and negative average payoffs (in the last 10 TTT trials) and contributed to both measures.

FIGURE 3 Risk-rate in the last five reactions to the signal TTT as a function of the average payoff in the first 15 experiences with TTT.



determined by the number of rare events observed in the first 15 experiences with this signal; for example, observing only one loss of 20 and 14 gains of 2 equates to an average payoff of 0.533.³ When

the average payoff is close to 0 (when the number of rare events in the first 15 experiences is 1 or 2), the results reveal substantial variability. Whereas most participants exhibited the pattern predicted under the reliance-on-small-samples hypothesis (i.e., selected the risky option even after experiencing a negative average payoff), a significant minority exhibited the opposite bias (4 of the 24 participants who observed only one loss of 20 in the first 15 experiences with TTT

³We chose to focus on the last 5 here, rather than the last 10 (as in Table 3), as the current analysis examines individual differences and the focus on the last 5 increases the proportion of the participants that saw at least two rare outcomes (the most interesting case) before the choice we analyze.

of Condition Abstract avoided the risky option in their last 5 choices). Under one explanation of the large individual-difference pattern, this behavior reflects “sampling inertia”: The four participants that behave

as if they overweight the rare events rely on the same small sample in all of the five final trials, and this sample includes the rare event.

TABLE 4 The risk-rate in the first two trials (out of 160) of the experiment.

Condition	Modal cue	Risk-rate	
		First trial	Second trial
Abstract	A	0.82	0.56
Abstract	T	0.81	0.69
Experts	A	0.25	0.24
Experts	T	0.95	0.90

4.2.2 | Additional effects of the description condition

Table 2 shows that in 6 of the 16 problems, the difference between the risk-rates in the two description conditions was significant. In all six cases, the direction of the significant difference can be explained with the assertion that the description of the cues as experts' recommendations increased the tendency to follow the rule “take the risk if the cue includes at least two T's.” In Condition Experts, this rule implies following the modal recommendation and taking the risk when

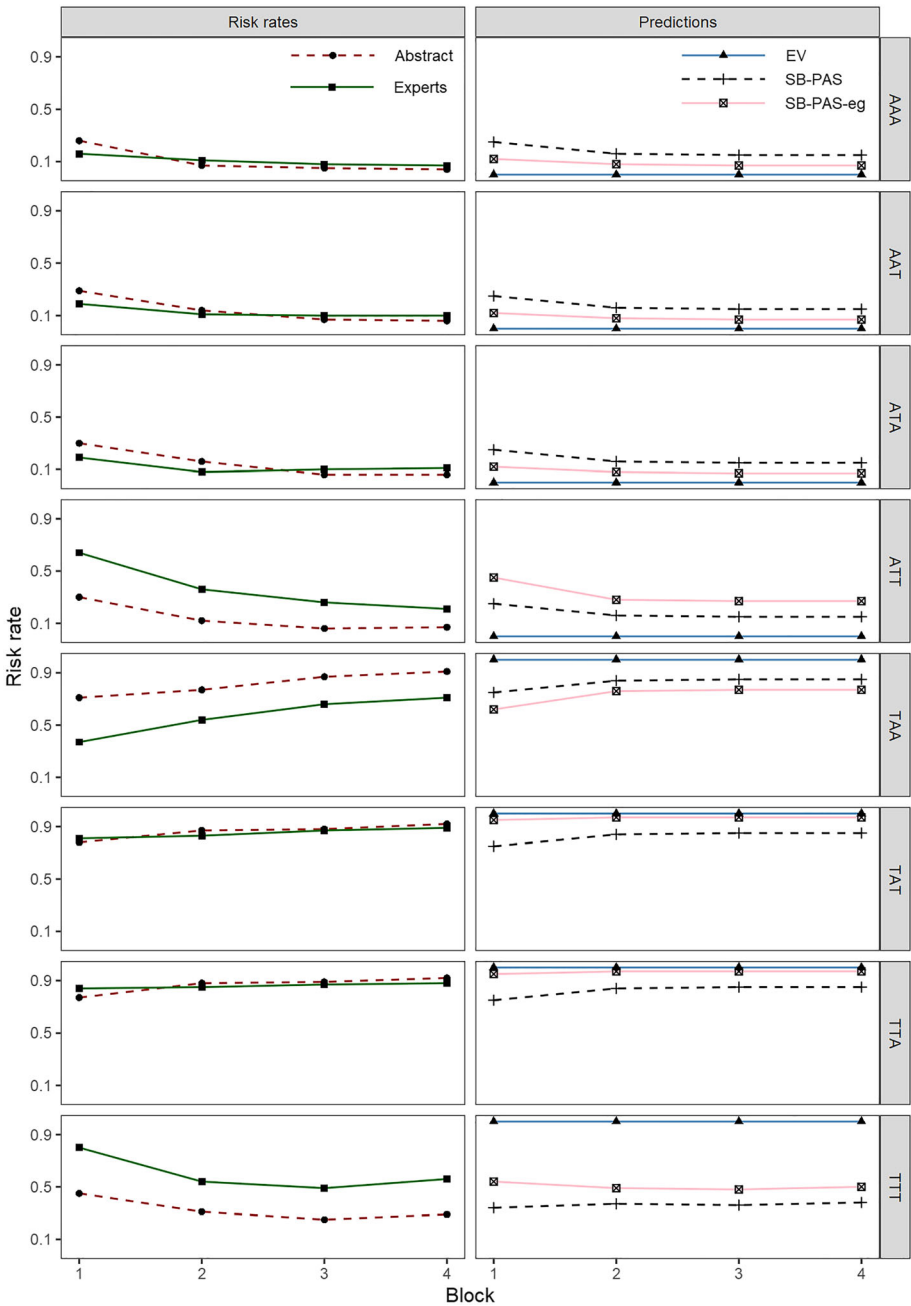



FIGURE 4 Risk-rates in blocks of five trials in each problem of set rare-treasure: Observed rates by condition (left), and the predictions of the models described below (right). EV, expected value maximization; SB-PAS, partial attentive sampler model; SB-PAS-eg, partial attentive sampler with estimated generalization.

FIGURE 5 Risk-rates in blocks of five trials in each problem of set rare-disaster: Observed rates by condition (left), and the predictions the models described below (right). EV, expected value maximization; SB-PAS, partial attentive sampler model; SB-PAS-eg, partial attentive sampler with estimated generalization.



most experts recommend it (see a similar pattern in Erev et al., 2022). Importantly, in five of the six significant cases, following the modal recommendation impaired expected return. Thus, while the additional information could clarify the fact that one of the cues was “best,” it did not increase the tendency to behave in accordance with the take-the-best rule.

Another indication of the difference between the two description conditions is suggested by the risk-rates in the first two trials of the experiment, as presented in Table 4. In Condition Experts, participants start by following the modal recommendation, as one might expect in the absence of other information. In Condition Abstract, participants start with a high risk-rate for all cues, but the first feedback (after the first choice and before the second choice) strongly reduced the risk-rate. This high initial risk-rate in Condition Abstract can be explained by assuming generalization from old experiences (that occurred before the beginning of the current experiment) in situations in which the feedback was limited to the obtained payoff, and initial exploration was effective. The quick decrease of risk-rate after obtaining full feedback supports this explanation.

4.2.3 | The effect of experience

The left side of Figures 4 and 5 presents the risk-rate in four blocks of five trials, split by the eight different signals. The results in Condition Abstract reveal relatively fast learning in the deterministic problems and slow learning in the problems with rare outcomes (signal TTT).⁴ The results in Condition Experts reveal that experience reduced the deviation from maximization that comes from reliance on the experts (in reaction to the signals ATT and TAA) but did not eliminate this tendency. In Condition Experts, the effect of experience on the reaction to the signal TTT was negative (i.e., the maximization rate declined with experience) in Set Rare-Disaster and positive in Set Rare-Treasure.

⁴An additional study of the current sets of problems, summarized in Bonder (2022), reveals that the speed of learning toward maximization in the deterministic problems is sensitive to the nature of the signals. Bonder used graphical signals and found slower learning toward maximization rate in these problems. While the maximization rate with the current letter-based cues reached 90%, the maximization rate in her study did not exceed 75%. Under one explanation of the difference between the two studies, it reflects the ease of discriminating between the different signals.

5 | IMPLICATIONS FOR DESCRIPTIVE MODELS

In order to clarify the implications of the current results to studies of decisions from experience, it is constructive to examine the differences between our predictions and the observed results. While the results support the main qualitative predictions of the SB-sample-of-5 model and reveal sufficient conditions for the emergence of overcomplexity and oversimplification, they also show important deviations from the model's quantitative predictions. The current section presents two model-based analyses that shed light on the differences between human behavior and the predictions of the SB-sample-of-5 model.

The first analysis builds on recent research that compares alternative generalizations of the basic sample-of-5 model based on their ability to predict behavior in simple decisions from experience (Erev & Plonsky, 2023). Across a set of 207 different choice problems, a comparison of 10 alternative models highlighted the value of the Partially Attentive Sampler (PAS) model, which generalizes the sample-of-5 model in two key ways. First, the PAS model allows for variability in the sample size, which can be any positive integer. In addition, PAS allows for reliance on past experiences that have occurred before the beginning of the target experiment. This probability of reliance on older experiences diminished with experience with the target problem. PAS approximates the impact of reliance on old experiences by assuming random choice in these trials. The right-hand column in Table 1 presents the prediction of PAS for four of the problems considered by Erev et al. and highlights its advantage over the basic sample-of-5 model.

To evaluate if similar assumptions can help capture behavior in the current multicue choice tasks, we derived the prediction of a model, referred to as signal-based-PAS (hereafter SB-PAS) that generalizes PAS to the current setting. SB-PAS assumes behavior in accordance to PAS with the additional constraint that only past experiences with the current signal are considered. The dotted curves in the right column in Figures 4 and 5 represent the predictions of SB-PAS to the current conditions with the parameters estimated by Erev et al. (i.e., we did not estimate free parameters to fit the current data; see Appendix C). The results show that SB-PAS model captures the main results observed in Condition Abstract. The mean squared deviation (MSD) between the observed and the reproduced choice rates (in blocks of 5 trials) in this condition is .004. In comparison, the MSD of SB-sample-of-5 model is .016.

In Condition Experts, the MSD between the observed choice rates and the predictions of SB-PAS is .030. While SB-PAS also outperforms the SB-sample-of-5 model in this case ($MSD = .056$), the model does not capture the main results well. One obvious shortcoming of SB-PAS involves the approximation of the impact of old experiences. As noted above, SB-PAS approximates this impact by assuming random choice. This assumption is clearly wrong in Condition Experts as the participants tend (at least initially) to behave as if they rely on past experiences in which following the modal recommendation was effective.

In the second model-based analysis, we considered alternative refinements of SB-PAS that can capture the difference between the two description conditions. Because Figures 4 and 5 show that the difference between the two description conditions diminishes with time, we chose to focus on the abstraction of the impact of generalization from old past experiences that occurred before the beginning of our experiments. As noted above, SB-PAS approximates the impact of reliance on old experiences by assuming random choice in some of the trials. To examine the impact of relaxing this assumption, we considered a variant of SB-PAS that replaces the random choice assumption with the assumption that when agents rely on old experiences, their risk-rates are similar to the risk-rates in the second trial in Condition Experts (presented in Table 4).⁵ The prediction of this model, referred to as "SB-PAS with estimated generalization" (SB-PAS-eg), is presented in the right column of Figures 4 and 5. The results show that this modification to the assumed impact of older experiences captures the main effect of describing the cues as expert recommendations and reduces the MSD score from .030 to .011.

6 | GENERAL DISCUSSION

Many deviations from rational choice seem to reflect the neglect of important evidence and the use of heuristics that imply oversimplification, whereas other deviations appear to reflect sensitivity to irrelevant evidence and the use of overly complex rules. The existing explanations of these contradictory deviations imply that they echo the relative importance of contradictory incentives: cognitive costs and the benefits from self-serving inference. The current analysis takes two steps toward clarifying the impact of incentives on the direction of the apparent deviations from efficient reasoning. The first step shows that the contradictory deviations can be captured without assuming the use of rules of different complexity in different settings. Instead, when the choice between options is based on small samples of similar past experiences, changes in the incentive structure can affect the apparent complexity of the rules that best fit the observed behavior. The modal behavior of agents that rely on small samples agrees with the prescription of overly simple rules when these rules lead to the best outcomes in most cases and suggests the use of overly complex rules when the optimal rule is simple but leads to disappointing outcomes in most cases.

We also showed that human decision makers exhibit the pattern predicted by a reliance on small samples of similar past experiences. This hypothesis correctly predicted behavior consistent with the take-the-best heuristic (Gigerenzer et al., 1991), even when take-the-best impaired maximization (was too simple) and also correctly predicted deviations from take-the-best that reflected the use of counterproductive, complex rules, when take-the-best was the optimal strategy.

The current analysis clarifies the variation in the complexity of rules that best capture behavior in two ways. First, our results

⁵We focus on the second trial to reduce the impact of the apparent exploration in the first trial. Using the first trial does not change the results.

demonstrate that in certain settings, behavior that appears to reflect rules of different complexity can be predicted with models that assume rules of fixed complexity. Thus, in these settings, the existence of contradictory deviations from optimal decisions does not imply the use of rules of different complexity. For example, the apparent contradictory reactions to demons and germs in the middle ages may not reflect rules of different complexity. Rather, the apparent difference can be the product of the incentive structure. In the middle ages, accepting the common belief in demons was rewarding (in most cases, but had extreme negative outcomes in some cases), and considering the impact of small organisms that cannot be observed (but can be addressed by careful hand washing) was not.

The second clarification focuses on tasks that involve sequential decisions like the decisions made while installing a new product and while making clinical recommendations. Our analysis can shed light on the initial decision, and this decision is likely to affect the complexity of the subsequent processes. For instance, the initial decision to skip reading instructions and contracts implies that the subsequent processes will not use this information, and is therefore likely to reflect oversimplification. The current reliance-on-small-samples hypothesis suggests that oversimplification of this type (as demonstrated by Bakos et al., 2014) is likely when reading the relevant information is important in expectation, but in most cases, skipping reading is effective (see Roth et al., 2016).

Another example involves the tendency of experts to trust their intuition rather than using a simple and effective linear regression algorithm (as demonstrated by Meehl, 1954). Our results suggest that preferring intuition that implies counterproductive complexity is more likely when the common outcome of this behavior is rewarding. The reward can be a subjective feeling of competence, but it can also be extrinsic. Similarly, the tendency to read and share fake news is likely in situations in which an incorrect belief in complex interactions is shared by the decision maker's peers and other people that evaluate his or her decisions.⁶

The main shortcoming of the current investigation involves its reliance on models that ignore the possibility of generalization between similar signals. For example, the initial reaction to the signal TAT will most likely be affected by past experiences with the signal TTA. We have considered models that try to capture this likely tendency and found that they do not improve prediction accuracy in a clear way. In natural settings, however, such between-signal generalizations are likely to be important.

Another limitation of our analysis is suggested by gambling phenomena that cannot be easily captured by the current reliance-on-small-samples hypothesis. For example, many gamblers appear to use complex “problem-solving” rules to devise “optimal” strategies to win (Ejova & Ohtsuka, 2020) and invest effort to “change” their luck

(e.g., Ohtsuka & Chan, 2010). We hope to address these and similar shortcomings in future research.

The wider theoretical implications of the current analysis include the generalization of the basic study of decisions from experience to address decisions in multicues tasks. To clarify the interesting implications of this generalization, it is constructive to focus on the relationship between the reliance-on-small-samples hypothesis (supported here and in basic studies of decisions from experience, see review in Erev & Haruvy, 2016) and the hypothesis that people tend to use the take-the-best rule (supported in studies of multidimensional decisions, Gigerenzer & Goldstein, 1996). These two hypotheses are similar in important ways: both assume a focus on only part of the available information, and previous research shows that they can help approximate the optimal strategy. The main difference between these hypotheses involves the distinction between sampling (or cue-based memory search) rules, and choice rules. The reliance-on-small-samples hypothesis concerns the way people collect information (retrieve memories of similar past experiences), while the take-the-best rule concerns with the way people use the available information. Under one interpretation of the predictive advantage of the reliance-on-small-samples hypothesis, in the current setting, the basic properties of the way people collect stored information tend to be more general than the way they use the available information.

ACKNOWLEDGMENTS

Ido Erev acknowledges support from the Israel Science Foundation (grant nos. 535/17 and 861/22). Yefim Roth acknowledges support from the Israel Science Foundation (grant no. 1857/22).

DATA AVAILABILITY STATEMENT

The raw data and pre-registration are available online (<https://osf.io/yna58/>).

ORCID

Taly Bonder  <https://orcid.org/0000-0002-3436-7818>

Elliot A. Ludvig  <https://orcid.org/0000-0002-0031-6713>

Yefim Roth  <https://orcid.org/0000-0003-2732-0496>

REFERENCES

- Bakos, Y., Marotta-Wurgler, F., & Trossen, D. R. (2014). Does anyone read the fine print? Consumer attention to standard-form contracts. *The Journal of Legal Studies*, 43(1), 1–35. <https://doi.org/10.1086/674424>
- Barron, G., & Erev, I. (2003). Small feedback-based decisions and their limited correspondence to description-based decisions. *Journal of Behavioral Decision Making*, 16(3), 215–233. <https://doi.org/10.1002/bdm.443>
- Bonder, T. (2022). *Complexity and simplification in decision from experience*. Dissertation, Technion, Haifa, Israel.
- Bröder, A. (2000). Assessing the empirical validity of the “take-the-best” heuristic as a model of human probabilistic inference. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 26(5), 1332–1346. <https://doi.org/10.1037/0278-7393.26.5.1332>
- Chen, D. L., Schonger, M., & Wickens, C. (2016). oTree—An open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance*, 9, 88–97. <https://doi.org/10.1016/j.jbef.2015.12.001>

⁶Our analysis does not assume, of course, that reading contracts, trusting scientists, or assuming linear models (and ignoring the possibility of complex interactions) is always optimal. Rather, it assumes that there are important situations in which people could benefit from investing more effort in checking and reading contracts (Roth, 2020; Teodorescu & Erev, 2014), and even from trusting careful scientific research (more than trusting other sources).

- Dawes, R. M., Faust, D., & Meehl, P. E. (1989). Clinical versus actuarial judgment. *Science*, 243(4899), 1668–1674. <https://doi.org/10.1126/science.2648573>
- Douglas, K. M., Sutton, R. M., & Cichocka, A. (2017). The psychology of conspiracy theories. *Current Directions in Psychological Science*, 26(6), 538–542. <https://doi.org/10.1177/0963721417718261>
- Ejova, A., & Ohtsuka, K. (2020). Erroneous gambling-related beliefs emerge from broader beliefs during problem-solving: A critical review and classification scheme. *Thinking & Reasoning*, 26(2), 159–187. <https://doi.org/10.1080/13546783.2019.1590233>
- Erev, I., Ert, E., Plonsky, O., & Roth, Y. (2023). Contradictory deviations from maximization: Environment-specific biases, or reflections of basic properties of human learning? *Psychological Review*. *forthcoming*.
- Erev, I., & Haruvy, E. (2016). Learning and the economics of small decisions. In J. H. Kagel & A. E. Roth (Eds.), *The handbook of experimental economics* (Vol. 2, pp. 638–716). Princeton University Press. <https://doi.org/10.1515/9781400883172-011>
- Erev, I., & Plonsky, O. (2023). The J/DM separation paradox and the reliance on small samples hypothesis. In K. Fiedler, P. Juslin, & J. Denrell (Eds.), *To appear in sampling in judgment and decision making*. Cambridge University Press.
- Erev, I., Roth, Y., & Sonsino, D. (2022). Decisions from valuations of unknown payoff distributions. *Decision*, 9, 172–193. <https://doi.org/10.1037/DEC0000172>
- Erev, I., & Roth, A. E. (2014). Maximization, learning, and economic behavior. *Proceedings of the National Academy of Sciences*, 111(supplement_3), 10818–10825. <https://doi.org/10.1073/pnas.1402846111>
- Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, 103(4), 650–669. <https://doi.org/10.1037/0033-295X.103.4.650>
- Gigerenzer, G., Hoffrage, U., & Kleinbolting, H. (1991). Probabilistic mental models: A Brunswikian theory of confidence. *Psychological Review*, 98(4), 506–528. <https://doi.org/10.1037/0033-295X.98.4.506>
- Juslin, P., Olsson, H., & Olsson, A. C. (2003). Exemplar effects in categorization and multiple-cue judgment. *Journal of Experimental Psychology: General*, 132(1), 133–156. <https://doi.org/10.1037/0096-3445.132.1.133>
- Kunda, Z. (1987). Motivated inference: Self-serving generation and evaluation of causal theories. *Journal of Personality and Social Psychology*, 53(4), 636–647. <https://doi.org/10.1037/0022-3514.53.4.636>
- Marsh, J., Coachys, N. C., & Kleinberg, S. (2022). The compelling complexity of conspiracy theories. *Proceedings of the Annual Meeting of the Cognitive Science Society*, 44(44), 2185–2192. Retrieved from <https://escholarship.org/uc/item/0sw059jb>
- Meehl, P. E. (1954). Clinical vs. statistical prediction: A theoretical analysis and a review of the evidence.
- Nevo, I., & Erev, I. (2012). On surprise, change, and the effect of recent outcomes. *Frontiers in Psychology*, 3, 24. <https://doi.org/10.3389/fpsyg.2012.00024>
- Ohtsuka, K., & Chan, C. C. (2010). Donning red underwear to play mah-jong: Superstitious beliefs and problem gambling among Chinese mah-jong players in Macau. *Gambling Research: Journal of the National Association for Gambling Studies (Australia)*, 22(1), 18–33.
- Pachur, T., & Olsson, H. (2012). Type of learning task impacts performance and strategy selection in decision making. *Cognitive Psychology*, 65(2), 207–240. <https://doi.org/10.1016/J.COGLPSYCH.2012.03.003>
- Payne, J. W., Bettman, J. R., & Johnson, E. J. (1993). *The adaptive decision maker*. Cambridge University Press. <https://doi.org/10.1017/CBO9781139173933>
- Plonsky, O., Teodorescu, K., & Erev, I. (2015). Reliance on small samples, the wavy recency effect, and similarity-based learning. *Psychological Review*, 122(4), 621–647. <https://doi.org/10.1037/a0039413>
- Roth, A. E., & Malouf, M. W. (1979). Game-theoretic models and the role of information in bargaining. *Psychological Review*, 86(6), 574–594. <https://doi.org/10.1037/0033-295X.86.6.574>
- Roth, Y. (2020). The decision to check in multialternative choices and limited sensitivity to default. *Journal of Behavioral Decision Making*, 33(5), 643–656. <https://doi.org/10.1002/bdm.2181>
- Roth, Y., Wänke, M., & Erev, I. (2016). Click or skip: The role of experience in easy-click checking decisions. *Journal of Consumer Research*, 43(4), 583–597. <https://doi.org/10.1093/jcr/ucw053>
- Swift-Hook, D. T. (2013). The case for renewables apart from global warming. *Renewable Energy*, 49, 147–150. <https://doi.org/10.1016/J.RENENE.2012.01.043>
- Teodorescu, K., Amir, M., & Erev, I. (2013). The experience–description gap and the role of the inter decision interval. *Progress in Brain Research*, 202, 99–115. <https://doi.org/10.1016/B978-0-444-62604-2.00006-X>
- Teodorescu, K., & Erev, I. (2014). On the decision to explore new alternatives: The coexistence of under- and over-exploration. *Journal of Behavioral Decision Making*, 27(2), 109–123. <https://doi.org/10.1002/bdm.1785>
- Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157), 1124–1131. <https://doi.org/10.1126/science.185.4157.1124>

How to cite this article: Bonder, T., Erev, I., Ludvig, E. A., & Roth, Y. (2023). The common origin of both oversimplified and overly complex decision rules. *Journal of Behavioral Decision Making*, e2321. <https://doi.org/10.1002/bdm.2321>

APPENDIX A: DEMONSTRATION OF THE DERIVATION OF THE PREDICTION OF THE SAMPLE-OF-5 MODEL IN PROBLEM B2 (0 OR −1, .9; +10)

Trial	Payoff from risk	Predicted risk-rate	Explanation
1	−1	.5	Random choice in the first reaction to the signal
2	−1	0	Drawing 5 times from all past experience (with replacement) provides an average payoff of −1
3	−1	0	
4	−1	0	
5	−1	0	
6	−1	0	
7	+10	0	
8	−1	$1-(6/7)^5 = .537$	The probability that the sample (of size 5 with replacement) will not include a + 10 outcome
9	−1	$1-(7/8)^5 = .487$	
10	−1	$1-(8/9)^5 = .445$	
11	−1	$1-(9/10)^5 = .409$	
12	−1	$1-(10/11)^5 = .379$	
13	−1	$1-(11/12)^5 = .353$	
14	+10	$1-(12/13)^5 = .330$	
15	−1	$1-(12/14)^5 = .537$	
16	−1	$1-(13/15)^5 = .511$	
17	−1	$1-(14/16)^5 = .487$	
18	−1	$1-(15/17)^5 = .465$	
19	−1	$1-(16/18)^5 = .445$	
20	−1	$1-(17/19)^5 = .427$	

APPENDIX B: EXPERIMENTAL INSTRUCTIONS AND COMPREHENSION QUESTION

Mandatory: Please read the instructions carefully and enter your details below

In each trial of this study you will be asked to choose between two options. One option that maintains the status quo (provides a payoff of 0 points for sure), and an option that can lead to a gain or a loss of points.

The payoff from the risky option will vary from trial to trial. You will receive three cues that can help predict the exact payoff in each trial. *Each cue will be the letter “T” or the letter “A”. For example, the signal “TTA” implies that the first and second cues are “T”, and the third cue is “A”.*

Please type the word “proceed” in the comments field below. This is to make sure that you read and understand the current instructions.

Your goal is to maximize your number of points. Your total number of points at the end of the experiment will determine the probability of earning the \$0.5 bonus in addition to the \$1 show-up fee.

Please enter the following information:

Gender

Age

Condition Abstract	Condition Experts
<p>To make sure you understood the study instructions, please answer the following question:</p> <p><i>What does the cues ATT implies?</i></p> <p>A. That the last cue is T, and the first two cues are A.</p> <p>B. That the last two cues are T, and the first cue is A.</p> <p>C. That the first cue is T, and the last two cues are A.</p> <p>D. That the first two cues are T, and the last cue is A.</p> <p>Please choose <input type="button" value="A"/> <input type="button" value="B"/> <input type="button" value="C"/> <input type="button" value="D"/></p>	<p>To make sure you understood the study instructions, please answer the following question:</p> <p><i>What does the cues ATT implies?</i></p> <p>A. That the last two experts recommend taking the safe option, and the first expert recommends avoiding it.</p> <p>B. That the last two experts recommend taking the risky option, and the first expert recommends avoiding it.</p> <p>C. That the three experts recommend taking the safe option.</p> <p>D. That the three experts recommend taking the risky option.</p> <p>Please choose <input type="button" value="A"/> <input type="button" value="B"/> <input type="button" value="C"/> <input type="button" value="D"/></p>

Note: The text in italics appears only in Condition Abstract. In Condition Experts this text was replaced with: *The cues are recommendation of three experts. “T” (take) implies that the expert recommends “Taking the risk”, and “A” (avoid) implies that the expert recommends avoiding the risk. For example, the sequence “TTA” implies that the first two experts recommend taking the risk, and the third expert recommends avoiding it.*

APPENDIX C: THE PARTIALLY ATTENTIVE SAMPLER (PAS) MODEL, THE TWO GENERALIZATIONS OF THIS MODEL (SIGNAL-BASED PAS (SB-PAS), AND SB-PAS WITH ESTIMATED GENERALIZATION (SB-PAS-eg))

PAS (Erev & Plonsky, 2023) generalizes the sample-of-5 model in two ways. First, it allows for the possibility that the sample size depends on the task (problem) and on the agent. The sample size used by agent i in task T , denoted as $\kappa_{i,T}$, is a free parameter. Second, it allows for the possibility of reliance on old past experiences (experiences that occur before the beginning of the current study) that diminishes with experience. In the current binary choice task, it is equal to $\delta_{i,T}^{\left[\frac{t-1}{\kappa_{i,T}}\right]}$, where $\delta_{i,T}$ is a second free parameter, and t is the trial number (in the current problem).

Erev et al. estimated (using a constrained maximum likelihood estimation) the distribution of PAS parameters in the population based on the results of 87 distinct binary choice task. Table A1 presents the results.

TABLE A1 The estimated distribution of the PAS parameters.

$\kappa_{i,T}$	$\delta_{i,T}$							
	0.01	0.05	0.1	0.2	0.3	0.4	0.7	1.0
1	9.3	3.4	2.2	2.2	1.7	2.2	2.8	9.5
2	3.3	1.6	0.9	1.1	0.6	1.0	1.3	
3	1.5	1.2	0.8	0.4	0.4	0.8	0.5	
4	0.9	0.4	0.4	0.3	0.2	0.4	0.5	
5	1.0	0.4	0.4	0.5	0.3	0.3	0.5	
6	0.5	0.3	0.3	0.1	0.2	0.3	0.5	
7	0.6	0.4	0.3	0.3	0.2	0.2	0.2	
8	0.7	0.2	0.2	0.2	0.2	0.2	0.4	
9	1.1	0.4	0.4	0.4	0.2	0.4	0.3	
12	0.5	0.2	0.3	0.2	0.2	0.5	0.6	
15	0.6	0.4	0.2	0.1	0.1	0.2	0.2	
20	1.1	0.4	0.2	0.5	0.2	0.5	0.5	
25	1.0	0.5	0.4	0.4	0.3	0.2	0.4	
30	1.9	0.9	0.8	0.7	0.4	0.8	0.7	
VL	3.2	2.4	1.8	2.1	2.9	2.7	3.0	

Note: Cells indicate the percentages of cases. When the value of $\kappa_{i,T}$ is “VL”, all past experiences were equally weighted as expected when the sample size is very large ($\rightarrow \infty$).

The model signal-based PAS (SB-PAS) generalizes PAS to multicue tasks. This model is similar to PAS with the additional constraint that only past experiences with the current signal are considered.

The original version of PAS, and SB-PAS, approximates the impact of reliance on old past experiences by assuming random choice. The model SB-PAS-eg relaxes this assumption. It assumes that the impact of reliance on old past experiences can be approximated by the choice rate in the second trial of the experiment (after the clarification of the nature of the feedback).