On the impact of germs and demons

Taly Bonder¹, Ido Erev¹, Elliot A. Ludvig²

¹The Faculty of Industrial Engineering and Management, Technion - Israel Institute of Technology, Israel

²Department of Psychology, University of Warwick, Coventry, UK

Abstract

Historical and experimental analyses suggest contradictory deviations from efficient reasoning. In some cases, people tend to oversimplify and ignore important factors like germs, while in others they seem to overcomplicate and consider non-existent factors like imaginary demons. The current study shows how this apparent contradiction can be the product of a tendency to rely on small samples of past experience. Simulations demonstrate that reliance on small samples triggers apparent over-simplicity when simple choice rules are counterproductive but better for most sets of samples; the opposite over-complexity bias emerges when the optimal rule is simple but fails in most samples. The descriptive value of this hypothesis is demonstrated in two preregistered studies with 300 Mechanical-Turk participants. Study 1 shows the qualitative pattern predicted by the reliance on the small-samples hypothesis. Study 2 compares alternative formulations of the sampling process and clarifies the importance of a distinction between choice and sampling processes.

Keywords: decisions from experience, base-rate neglect, similarity rules, heuristics, take the best, categorization

Introduction

During the medieval period, most people did not believe that handwashing had a general positive effect (Best & Neuhauser, 2004). On the other hand, at that time many people believed that certain psychological problems resulted from demonic possession of the apparently sick individuals and should be treated by religious ceremonies (Kemp & Williams, 1987). These historical observations seem to reflect contradictory deviations from optimal decision-making. The failure to see the value of hand washing suggests an overly simple cognitive model that overlooked the impact of germs¹. In contrast, the belief in evil spirits reflects the use of an overly complex model that supposed non-existent causes and led to incorrect treatments.

A similar coexistence of contradictory deviations from optimality is suggested by the modern study of multi-dimensional judgment and decision-making. In many settings people pay too much attention to the possibility of complex interactions between different dimensions, but in other settings, people miss important interactions and overuse simple rules. The significance of the tendency to focus on complex interactions is highlighted by Meehl's (1954) classic study of classification decisions based on the x-dimensional MMPI (McKinley et al., 1948). Meehl found that the decisions of expert psychiatrists were less accurate than decisions derived by simple linear rules. Subsequent analysis showed that an important part of the experts' errors reflected a counterproductive effort to capture complex interactions (Dawes et al., 1989). A similar overinterpretation of complexity has been found in human face classification (Reed, 1972), machine autonomous learning (Carpenter & Grossberg, 1988), and sequential learning from experience (Plonsky et al., 2015). The contrasting tendency to miss important interactions is

¹ To clarify the relationship between these historical observations and sensitivity to complex interactions, note that the discovery of the value of washing hands (which facilitated the discovery of the impact of bacteria) by Ignaz Semmelweis was based on the detection of the interaction between the secondary association of medical doctors that worked in the maternity ward and the death rate of new mothers (Best and Neuhauser 2004).

perhaps best illustrated by research that focuses on the descriptive value of simple heuristics. For example, Tversky and Kahneman (1974) showed that when an optimal judgment requires the weighting of base rates and new data using Bayes theorem, many people neglect the base rates and make their judgments solely based on the representativeness of the new data. A similar overuse of simple rules underlies the tendency to accept e-license agreements without considering their content (Bakos et al., 2014).

Previous research has explained the coexistence of these contradictory deviations as arising from disparities in the structure of the situation. For example, Juslin et al. (2003) showed that people favour complex interactions when the rewards are predicted by non-additive rules, especially when the predictive value of each information cue is probabilistic. In particular, the reward structure directs behavior towards complexity if cues are not fully associated with the rewards, or if a level of abstraction is required to achieve the task goal. Later, Pachur and Olsson (2012) showed that when people learn from a comparison between alternatives, they tend to use simple cue-based strategies, and when they learn by estimating a direct criterion for one alternative, they tend to form complex combinations.

The current research proposes and tests the novel hypothesis that these contradictory deviations from optimal complexity arise because people select the choice rule that led to the best outcomes in similar situations in the past. This hypothesis deconstructs the underlying cognitive processes into two parts: one part recalls similar past experiences, and a second part identifies feasible choice rules and chooses according to the best-performing rule in the recalled experiences. This deconstruction suggests two classes of contributors to the apparent contradiction. The first is the possibility that people consider and use choice rules of different complexity in different settings. The explanations described above belong to this class. Our

analysis focuses on the second class: We consider the possibility that the nature of the recall processes implies that the apparent complexity of the underlying processes depends on the incentive structure. Specifically, when the recall processes involve a reliance on small samples (Fiedler 2000; Fiedler and Kareev, 2006; Hertwig et al., 2004; Teodorescu and Erev, 2014; Chater et al., 2020), rules of fixed complexity can trigger behavior that appears to reflect different complexity, depending on the incentive structure of the task.

The first part of the paper presents a numerical example that illustrates our theoretical starting point. It focuses on a 3-dimensional decision task and demonstrates that when the recall process involves a reliance on small samples of past experience, changes in the incentive structure can alter the direction of the expected deviation from optimal complexity.

The second part of our investigation examines human behavior in a pair of 3-dimensional decision tasks. Study 1 shows the qualitative pattern predicted by the reliance on small samples model, but also shows quantitative differences between the predicted and observed results. Study 2 compares alternative explanations for the deviations from the predictions of the basic small-samples model. The results favour a generalization of the basic small-samples model that increases the set of factors that impact how samples are drawn from memory without changing the reliance on the small-samples assumption.

These observations delineate sufficient conditions for the co-existence of both over-simplification, which appears to reflect the reliance on a simple heuristic that impairs performance, and over-complexity, which suggests a failure to use a simple, effective heuristic.

Method

A numerical example

Consider decision-makers that face the two sets of problems presented in Table 1. Both sets include eight basic problems that appear repeatedly in random order. Each problem involves a choice between the Status Quo (SQ, 0 with certainty) and an Action with a negative or positive expected value (EV). In each set, for seven of the problems, the Action yields a deterministic outcome, but for the eighth problem, the Action yields a probabilistic outcome with either a high-valued rare event (Rare Treasure; Set 1) or a low-valued rare event (Rare Disaster; Set 2). As a result, this eighth problem has a positive EV (Set 1) or a negative EV (Set 2). Decision-makers do not receive prior information concerning the incentive structure, but they can rely on two sources of information: (1) a signal which includes three cues that characterise each problem, and (2) full feedback concerning the outcomes for the two options in all previous trials. As the cues distinguish between the different problems, each set can be described as a multi-dimensional choice task.

The optimal, expected-value (EV) maximizing, strategy in Set 1 is simple: “if the first cue is S, select the Action; otherwise, select SQ.” This strategy implies full reliance on the most effective cue. Thus, it is an instance of the “take-the-best” heuristic rule (TTB; Gigerenzer et al., 1991; see also Broder, 2000). The optimal strategy in Set 2 is more complex: to maximize EV, the decision-maker should learn to follow the rule: “if the signal is SSS, select SQ; in other cases, if the first cue is S, select the Action, otherwise, select SQ.”

To clarify the impact of the recall processes, we simulated a set of virtual agents programmed to behave according to a model assuming a simple set of choice rules and restrictive recall processes. The model assumes that the agents consider only two choice rules, Choose Left

or Choose Right, and after gaining experience, the agents select the rule that led to the best average payoff in the set of trials recalled before each choice. The recall process involves sampling with replacement κ times (κ is a free parameter, and the current example uses $\kappa=4$) from the set of previously experienced trials with the exact same signal (same string of cues) as the current trial. In addition, the model assumes perfect memory of past trials and random choice in reaction to new signals (i.e., exemplars; Nosofsky 1986; Rouder & Ratcliff, 2006). To clarify the relationship to the models presented below we call this model “Same-Signal.”

Table 1 shows that the Same-Signal model predicts over-complexity in Set 1: The virtual agents behave as if they fail to master the simple TTB rule that maximizes EV in this set. In particular, they avoid taking the Action for the eighth, probabilistic problem, when the signal is SSS. In that problem, small samples tend to miss the rare treasure, thus, the value of the Action is most often under-estimated in the samples drawn from memory. In Set 2, the virtual agents exhibit over-simplicity: they behave as if they use the TTB rule, even though that behaviour does not maximize EV in this set. Again, the deviation from maximization occurs for the probabilistic problem, when the signal is SSS, and the model selects the Action too frequently. In that problem, small samples tend to miss the rare disaster, thus, the value of the Action is most often over-estimated.

Table 1. The two sets of choice problems considered in the numerical example and in Study 1.

Study 1: Model predictions and results								
Incentive structure			Results		Basic models		Post-hoc models	
Problem	Signal	Payoff from Action	EV	Exp. results	Optimal action rate	Same signal ($\kappa=4$)	Same- or-all ($\varepsilon=.6$, $\kappa=2$)	Same-all- or freq. ($\varepsilon=.5$, $f=.22$, $\kappa=2$)
<i>Set 1: Rare Treasure</i>								
1	CCC	-2	-2	0.33	0	0	0.24	0.20
2	CCS	-2	-2	0.41	0	0	0.24	0.27
3	CSC	-2	-2	0.39	0	0	0.24	0.27
4	CSS	-2	-2	0.4	0	0	0.24	0.35
5	SCC	+2	+2	0.67	1	1	0.64	0.55
6	SCS	+2	+2	0.69	1	1	0.63	0.62
7	SSC	+2	+2	0.69	1	1	0.63	0.62
8	SSS	-2, .9; +20, .1	+0.2	0.43	1	0.34	0.31	0.29
Average Action rate				0.5	0.5	0.42	0.4	0.4
<i>Set 2: Rare Disaster</i>								
1	CCC	-2	-2	0.37	0	0	0.29	0.24
2	CCS	-2	-2	0.4	0	0	0.29	0.31
3	CSC	-2	-2	0.43	0	0	0.29	0.32
4	CSS	-2	-2	0.46	0	0	0.29	0.39
5	SCC	+2	+2	0.7	1	1	0.69	0.59
6	SCS	+2	+2	0.73	1	1	0.69	0.67
7	SSC	+2	+2	0.77	1	1	0.69	0.66
8	SSS	+2, .9; -20, .1	-0.2	0.72	0	0.66	0.62	0.66
Average Action rate				0.57	0.38	0.46	0.48	0.48
MSD (over Studies 1 and S1)						0.051	0.01	0.008

Note: All problems involved a choice between the Status Quo (a payoff of 0) and an Action. The payoff from the Action is described in the table. The right-hand columns present the predictions of three models, and the observed Action rates as described below.

Study 1

Study 1 examined human behavior in the two sets of problems described in Table 1.

Procedure. Each participant was randomly assigned to one of two conditions, Rare Treasure (Set 1) or Rare Disaster (Set 2) and faced each of the eight problems in that condition 20 times (cf. Table 1). The order of the 160 choice tasks was random. Figure 1 presents the experimental design and the provided instructions. The SQ and the Action alternatives were displayed as buttons on a computer screen, with three visual cues presented on each button (Figure 1d). The SQ button always yielded an outcome of zero, while the Action button yielded either positive or negative outcomes, pending the set and problem. The cues on the SQ button were constant, while the binary cues on the Action button, i.e., circles (C) and squares (S), changed from trial to trial, mapped to one of the eight possible choice problems (see Table 1, Signal Column). The mapping was identical for all participants.

Participants. 100 Mechanical Turk workers (52% females, mean age 34.3) participated either in the Rare-Treasure ($n = 50$) or in the Rare-Disaster ($n = 50$) experimental group. Participants received the following instructions: "This is a decision-making task. It includes multiple choice trials and will take you around 15 minutes to complete. Your aim is to collect as many points as possible. Use the 3 visual cues to help you decide. The more points you gain, the higher is your chance to win a bonus of 1\$." The task consisted of 8 blocks of 20 trials; each experimental block was followed by a self-paced break. Responses were not limited in time, and full feedback about both options was provided after each trial for 2 s. The participants received a fee of \$2 for completing the task and, as implied by the instructions, had a chance to win a bonus of \$1 depending on their points gained. All participants provided informed consent before participating, and the study was approved by the Technion – the Israel Institute of Technology ethics committee.

Figure 1. Method - instructions provided to participants (A), timeline of the experimental manipulation (B), a single trial display (C), and possible cue signals in the eight choice problems (D).

Results

The Results column in Table 1 presents the proportion of trials when the Action was selected (i.e., Action rate) over all 20 trials with each of the eight problems. The modal results agree with the predictions of the Same-Signal model: In Set 1 (Rare-Treasure condition), when the simple TTB heuristic is optimal, participants behaved as if they used a more complex choice rule. Both the EV and the TTB heuristic favor taking the Action in Problem 8 (SSS), yet the observed Action rate was only 0.43. In contrast, in Set 2 (Rare-Disaster condition) when the optimal strategy was relatively complex, participants behaved as if they used a simpler rule. The optimal rule (by EV) favors avoiding the Action and going with the SQ choice in Problem 8 (SSS), the observed Action rate was actually 0.72. These choice rates, implying deviations from maximization, were significantly different from random choice ($t(99) = 2.82$, $p = .006$, $d = 0.28$,

95% CI [0.52, 0.63]).² Figure 2 presents the observed Action rates in blocks of 10 trials and shows that most of the learning occurred in the first 5 blocks. The aggregate choice rates in the last 11 blocks were relatively stable.

Figure 2. The mean Action rate in blocks of 10 trials in the Rare Treasure and the Rare Disaster sets. The median Action rate in Problem 8 deviates from maximization in both sets. In Set 1, when the simple TTB rule is optimal, the deviation results in a tendency to favor a counterproductive complex rule. In Set 2, the deviation aligns with the TTB rule when a more complex rule is optimal.

Though the modal results agree with the prediction of the Same-Signal model, the predictions of this model were too extreme when compared to the human behaviour. The following sections present alternative generalizations of this model that provide a better fit.

The Same-or-All model. The first generalization allows for the possibility that decision-makers use two sampling rules. Under the first rule, as in the Same-Signal model, the sample (of size κ) is taken from previous trials with the exact same signal. Under the second rule, the sample is taken from all previous trials. The probability of using the second sampling rule is ϵ , and with probability $1 - \epsilon$, the same-signal rule is used. In all other respects, the Same-or-All model is identical to the same-signal model. To reduce the risk of overfitting the data, the

² A pre-registered paired-samples t-test analyses compared the deviation of the experimental results from the predictions of the Same-Signal model and the deviation of the results from the predictions of the take-the-best rule. The results showed that the Action rate in Problem 8 was predicted better by the Same-Signal model than by the take-the-best rule (Set 1: $t(49) = 15.93$, $p < .001$, $d = 2.25$, 95% CI [0.36, 0.5]; Set 2: $t(49) = 26.9$, $p < .001$, $d = 3.81$, 95% CI [0.66, 0.77]). URL: https://osf.io/h6sfj/?view_only=44bba656d5894a6db6b1af4043972ad4.

parameters were estimated based on the aggregate choice rate observed in Study 1 and based on the aggregate rates observed in a study of Set 3 presented in the Supplementary Materials (Study S1). To facilitate the estimation of ε independently of κ , in Set 3, both options in each of the 8 problems provide a fixed payoff (and κ does not affect the prediction of the model). We used a grid-search procedure (in the range: $\kappa = 1$ to 20; $\varepsilon = 0$ to 1) and selected the parameters that minimized the Mean Squared Deviation (MSD) between the observed and reproduced rates. The best fit was obtained with the parameters $\kappa = 2$ and $\varepsilon = 0.6$. The same-or-all column in Table 1 presents the prediction of the model with these estimated parameters.

The Same-All-or-Frequency model. One apparent shortcoming of the Same-or-All model involves the fact that it cannot capture the difference between signals CCC and CCS in Set 1 (Action rates of .33 and .41 respectively). This difference can be captured by introducing a third sampling rule. Specifically, in certain cases people may sample past experience with the same frequency of the squares/circles as the current trial. This hypothesis leads to a Same-All-or-Frequency model that samples according to the following three rules: all past trials with probability ε , same frequency of squares with probability f , and the same signal otherwise $(1-\varepsilon-f)$. The model best fits the data with the parameters $\varepsilon = 0.5$, $f = 0.2$, and $\kappa = 2$, as estimated by grid search on Sets 1-3³. Note that the Same-or-All model is a special case of the Same-All-or-Frequency model with $f = 0$, and the Same-Signal model is a special case of the Same-or-All model with $\varepsilon = 0$. The final columns in Table 1 present results from all three models.

³ Excluding the first encounter of the participants with each problem resulted in the parameters $\varepsilon = 0.51$, $f = .19$, and $\kappa = 2$. Fitting the Same-or-All and the Same-Signal models provided parameters identical to those reported above.

Study 2

Study 2 compares the predictive value of the models proposed above by focusing on two pre-registered variants of Set 2 (Table 1). Both variants used the same problems as Set 2 and differed with respect to which signal was associated with the Rare-Disaster problem. In Set 4 (Set 3 is in the Supplementary Materials), the signal was CSS (4th Problem), and in Set 5 the signal was SCC (5th Problem). Table 2 shows that this change does not affect the prediction of the Same-Signal and the Same-or-All models, but it affects the predictions of the Same-All-or-Frequency model, which predicts a higher choice rate of the negative EV Rare-Disaster action in Set 4 (when the signal is CSS) than in Set 5 (when the signal is SCC). This prediction is a product of the fact that the frequency-similarity criterion implies different behavior in the two cases. When the signal includes the cue “S” twice (the signal is CSS in Set 5), the frequency-similarity rule leads to the recall of past experiences with the signals SSC and SCS in which the Action led to a deterministic gain. In contrast, when the signal includes the cue “C” twice (the signal is SCC in Set 4), this rule leads to the recall of past experiences with the signals CCS and CSC in which the Action led to a deterministic loss.

Procedure and Participants. The procedure described in Study 1 was used, with the exception that it focused on the two sets of problems presented in Table 2. Two hundred Mechanical Turk workers (61% females, mean age 34) were randomly assigned either to Set 4 ($n = 100$) or to Set 5 ($n = 100$). They read the same instructions specified in Study 1 and performed the task of 160 trials. They received a fee of \$2 for completing the task and had a chance to win a bonus of \$1 depending on their points gained.

Table 2. Model Predictions and Experimental Results in Study 2.

Model predictions for Sets 4 & 5								
Incentive structure			Results			Models		
Problem	Signal	Payoff from Action	EV	Exp. results	Optimal action rate	Same signal ($\kappa=4$)	Same- or-all ($\varepsilon=.6$, $\kappa=2$)	Same-all- or freq. ($\varepsilon=.5$, $f=.22$, $\kappa=2$)
<i>Set 4</i>								
1	CCC	-2	-2	0.42	0	0	0.30	0.25
2	CCS	-2	-2	0.54	0	0	0.30	0.32
3	CSC	-2	-2	0.54	0	0	0.30	0.32
4	CSS	+2, .9; -20, .1	-0.2	0.73	0	0.66	0.63	0.68
5	SCC	+2	+2	0.67	1	1	0.68	0.60
6	SCS	+2	+2	0.73	1	1	0.68	0.72
7	SSC	+2	+2	0.72	1	1	0.68	0.72
8	SSS	-2	-2	0.45	0	0	0.30	0.25
Average Action rate				0.6	0.38	0.46	0.48	0.48
<i>Set 5</i>								
1	CCC	-2	-2	0.49	0	0	0.30	0.25
2	CCS	-2	-2	0.52	0	0	0.30	0.31
3	CSC	-2	-2	0.52	0	0	0.30	0.31
4	CSS	-2	-2	0.56	0	0	0.30	0.39
5	SCC	+2, .9; -20, .1	-0.2	0.63	0	0.66	0.63	0.55
6	SCS	+2	+2	0.69	1	1	0.68	0.66
7	SSC	+2	+2	0.69	1	1	0.68	0.66
8	SSS	+2	+2	0.71	1	1	0.69	0.73
Average Action rate				0.6	0.38	0.46	0.48	0.48
MSD						0.09	0.026	0.024

Results. The Action rates, summarized in the Results column of Table 2, agree with the predictions of the Same-All-or-Frequency model: The deviation from maximization in the Rare-Disaster problem was larger when the signal was CSS (.73 in Set 5) than when the signal was SCC (.63 in Set 4). The difference in Action rate was significant ($t(197) = 3.12, p = .002, d =$

0.44, 95% CI [0.04, 0.18])⁴. Table 2 also shows that the Same-All-or-Frequency model provides more accurate predictions than the competing models.

In addition, the results show that all three models under-predict the Action rate across the problems. The mean Action rate was 0.6 in Set 4 and Set 5, whereas the models predict 0.46-0.48. A similar high rate (.57) was observed in Set 2 of Study 1. Notice that these high rates are inconsistent with the predictions of loss aversion (Kahneman & Tversky, 1979), and the status-quo bias (Samuelson & Zekhauser, 1990). Under one explanation of this deviation, the participants use even more similarity rules than assumed by the current models. Specifically, people may choose to participate in decision-making experiments in Mechanical Turk because they were successful in similar experiments in the past, and these past successes guided part of their choices during the experiment. Appendix 2 presents, and rejects, two further alternative explanations for these high Action rates.

General Discussion

Many deviations from rational choice seem to reflect the neglect of important evidence and the use of heuristics that result in oversimplification, whereas other deviations appear to reflect sensitivity to irrelevant evidence and the use of overly complex rules. The current analysis takes two steps toward clarifying these observations. It starts with a numerical example demonstrating that these contradictory deviations from optimal complexity can be captured without assuming that people consider rules of different complexity in different settings: when people choose between options based on their performance in a small sample of past experiences, changes in the incentive structure can affect the apparent complexity of the rules that best fit the observed behavior. Thus, behaviors that appear to reflect the use of decision rules

⁴ This hypothesis and analyses were pre-registered:
https://osf.io/38dtg/?view_only=4e0dfd08279c4cc1a22c1a83e8f0819a.

of different complexity can be the product of the tendency to choose among simple rules based on a small sample of past experiences.

Second, Study 1 demonstrates the descriptive value of the reliance on small samples hypothesis. A basic model assuming reliance on a sample of 4 past experiences correctly predicted the condition under which the simple take-the-best rule captures the results. Importantly, the model correctly predicted behavior that appeared consistent with take-the-best, when take-the-best impaired maximization (was too simple), and correctly predicted deviations from take-the-best that reflect the use of counterproductive, complex rules, when take-the-best was the optimal strategy.

Our analysis contributes to the study of decision-making in two ways. First, it highlights the descriptive value of the reliance on small-samples hypothesis in the context of multi-dimensional decision tasks. This generalization is potentially important as it suggests that reliance on small-samples hypothesis can shed light on choices between cognitive rules. Though the current models assumed very simple rules, even simpler than take-the-best, the tendency to rely on small samples is also likely to affect choice among more complex rules.

Second, the current analysis demonstrates that generalizing the small-samples model to address a multi-dimensional choice task requires refinement of the similarity rules. The basic model, which assumes sampling of past experiences with the exact same signal, was too limiting. The results of Study 2 highlight the advantage of the model that assumes three distinct similarity rules: same, all, or frequency-matched. Yet, people are likely to consider more rules. For example, the observed general tendency to prefer the Action over the status quo implies a reversed loss-aversion bias, which cannot be explained by the current models, and potentially

suggests a reliance on previous experiences on the Mechanical Turk platform. We hope to address this topic in future research.

The wider implications of the current analysis involve the debate concerning the interaction between the impact of decision rules and similarity (see Hahn & Chater, 1998; Pothos, 2005). The rules considered above can be described as summaries of the cognitive processes that determine perceived similarity. Thus, the current analysis demonstrates that it is difficult to identify the underlying decision rules without a good understanding of perceived similarity.

We believe that the practical implications of the present investigation involve the clarification of the importance of the interaction between cognitive processes and the incentive structure. To illuminate this interaction, it is constructive to recall that some of the fake news concerning Covid19 (e.g., denying it, and the assertion that the vaccination includes a chip that sends information to the government) are similar to the examples that have inspired our analysis. We suggest that the historical hand-washing example (Best & Neuhauser 2004) is similar to Sets 2, 4 and 5: It is natural to assume that certain cues, like sticky hands, triggered effective learning to wash hands; that is, the outcome from washing hands in these conditions was rewarding. Yet, given other cues, like lack of stickiness, the outcomes of washing hands were typically negative: a waste of time without a clear benefit. Thus, medieval physicians learned to behave as if they ignored cues of the type “I have just touched a sick person,” and behave as if they use simple rules of the type “wash hand if they are sticky.”

The example of religious ceremonies (Kemp & Williams, 1987) is more similar to Set 1. While our analysis does not explain why certain religious ceremonies were developed, it can potentially explain why they were overused. It is natural to assume that in most cases, during the medieval

period, trusting local religious leaders and accepting the common norms was reinforcing. Thus, it is possible that the rare cases, which disagreed with this norm, were missed.

References

- Bakos, Y., Marotta-Wurgler, F., & Trossen D. R. (2014). Does anyone read the fine print? Consumer attention to standard-form contracts. *Journal of Legal Studies*, 43 (1), 1–35.
- 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.
- Best, M., & Neuhauser, D. (2004). Ignaz Semmelweis and the birth of infection control. *BMJ Quality & Safety*, 13(3), 233-234.
- 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.
- Carpenter, G. A., & Grossberg, S. (1988). The ART of adaptive pattern recognition by a self organizing neural network. *Computer*, 21(3), 77-88.
- Chater, N., Zhu, J. Q., Spicer, J., Sundh, J., León-Villagrà, P., & Sanborn, A. (2020). Probabilistic biases meet the Bayesian brain. *Current Directions in Psychological Science*, 29(5), 506-512.
- Dawes, R. M., Faust, D., & Meehl, P. E. (1989). Clinical versus actuarial judgment. *Science*, 243(4899), 1668-1674.
- Fiedler, K. (2000). Beware of samples! A cognitive-ecological sampling approach to judgment biases. *Psychological Review*, 107(4), 659-676.
- Fiedler, K., & Kareev, Y. (2006). Does decision quality (always) increase with the size of information samples? Some vicissitudes in applying the law of large numbers. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 32(4), 883-903.
- Gigerenzer, G., Hoffrage, U., & Kleinbölting, H. (1991). Probabilistic mental models: a Brunswikian theory of confidence. *Psychological Review*, 98(4), 506-528.
- Hahn, U., & Chater, N. (1998). Similarity and rules: Distinct? exhaustive? empirically distinguishable? *Cognition*, 65(2-3), 197-230.
- Hertwig, R., Barron, G., Weber, E. U., & Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychological Science*, 15(8), 534-539.

- 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.
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2), 263-291.
- Kemp, S., & Williams, K. (1987). Demonic possession and mental disorder in medieval and early modern Europe. *Psychological Medicine*, 17(1), 21-29.
- McKinley, J. C., Hathaway, S. R., & Meehl, P. E. (1948). The Minnesota Multiphasic Personality Inventory: VI. The K Scale. *Journal of Consulting Psychology*, 12(1), 20-31.
- Meehl, P. E. (1954). *Clinical versus statistical prediction: A theoretical analysis and a review of the evidence*. University of Minnesota Press.
- Pachur, T., & Olsson, H. (2012). Type of learning task impacts performance and strategy selection in decision making. *Cognitive Psychology*, 65(2), 207-240.
- Nosofsky, R. M. (1986). Attention, similarity, and the identification–categorization relationship. *Journal of experimental psychology: General*, 115(1), 39-57.
- 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.
- Pothos, E. M. (2005). The rules versus similarity distinction. *Behavioral and Brain Sciences*, 28(1), 1-14.
- Reed, S. K. (1972). Pattern recognition and categorization. *Cognitive psychology*, 3(3), 382-407.
- Rouder, J. N., & Ratcliff, R. (2006). Comparing exemplar-and rule-based theories of categorization. *Current Directions in Psychological Science*, 15(1), 9-13.
- Samuelson, W., & Zeckhauser, R. (1988). Status quo bias in decision making. *Journal of risk and Uncertainty*, 1(1), 7-59.
- Teodorescu, K., & Erev, I. (2014). Learned helplessness and learned prevalence: Exploring the causal relations among perceived controllability, reward prevalence, and exploration. *Psychological Science*, 25(10), 1861-1869.
- Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157), 1124-1131.

Supplementary Materials

Set 3

Set 3 introduced deterministic payoff rules, providing further data for parameter estimation. In this Set, the SQ option (all pentagons) again always paid out 0, and the payout for the Action option paid out -2 or +2, depending on the number of squares (0 or 1 squares give -2, and 2 or 3 squares give +2). The payout appeared variable to the participant until the pattern was identified. Notice that the deterministic nature of the problems implies that a reliance on small-samples does not impair maximization. Moreover, frequency-based similarity does not lead to overgeneralization, Table 3.

Table 3. Model Predictions and Experimental Results in Set 3.

Model Predictions and Experimental Results in Set 3								
Problem	Incentive structure		EV	<i>Results</i>		<i>Basic models</i>		Post-hoc models
	<i>Signal</i>	<i>Payoff from Action</i>		<i>Exp. results</i>	<i>Optimal action rate</i>	<i>Same signal ($\kappa=4$)</i>	<i>Same-or-all ($\varepsilon=.6, \kappa=2$)</i>	<i>Same-all-or freq. ($\varepsilon=.5, f=.22, \kappa=2$)</i>
1	CCC	-2	-2	0.24	0	0	0.35	0.30
2	CCS	-2	-2	0.26	0	0	0.35	0.30
3	CSC	-2	-2	0.24	0	0	0.35	0.30
4	CSS	+2	+2	0.78	1	1	0.65	0.70
5	SCC	-2	-2	0.24	0	0	0.35	0.30
6	SCS	+2	+2	0.81	1	1	0.65	0.70
7	SSC	+2	+2	0.81	1	1	0.65	0.70
8	SSS	+2	+2	0.79	1	1	0.65	0.70
Average Action rate				0.52	0.5	0.5	0.5	0.5

Procedure and Participants. Set 3 used the same procedure as in Study 1 with the exception that it focused on the set of problems presented in Table . Fifty Mechanical Turk workers read the instructions specified in Study 1 and performed the task of 160 trials. They

received a fee of \$2 for completing the task and had a chance to win a bonus of \$1 depending on their points gained.

Results. The right-hand column in Table 3 presents the mean Action-rate over all 20 trials with each cue combination (problems). As per the pre-registered hypotheses, a one-way analysis of variance showed that the correct response rate in Set 3 was higher than the correct response rate in Sets 1 and 2 in Study 1 ($F_{(2,147)}=8.87$, $p<.001$, partial $\eta^2=.11$). Bonferroni post-hoc tests highlighted the significant differences between Set 3 and Set 1: Rare-Treasure ($p=.001$) and Set 3 and Set 2: Rare-Disaster ($p=.002$). Figure 3 shows that Set 3 group showed a steeper learning rate, compared to the two groups of Study 1. Inspection of the MSD shows that the best predictive model was the Same-All-or-Frequency model, a result consistent with the findings of Study 1.

Figure 3. The mean Risk rate in blocks of 10 trials in Set 3. Dotted gray: Problems 1-3 & 5 (negative deterministic payoff), dashed black: Problems 4 & 6-8 (positive deterministic payoff).

Set 6

Group 6 used the payoff structure and the method described in Set 5 but received modified instructions regarding the bonus payment: "This is a decision-making task. It includes multiple choice trials and will take you around 15 minutes to complete. Your aim is to collect as many points as possible. Use the 3 visual cues to help you decide. You will begin the game with a bonus of 50 cents. For each point you gain in the game you will gain 1 cent, and for each point you lose in the game you will lose 1 cent." The manipulation aimed to clarify the instructions provided in Sets 1-5, inspecting its influence on the risk-seeking tendency showed in Set 5.

Table 4. Model Predictions and Experimental Results in Group 6.

Experimental results in Set 6								
Problem	Incentive structure		EV	Results		Basic models		Post-hoc models
	Signal	Payoff from Action		Exp. results	Optimal action rate	Same signal ($\kappa=4$)	Same- or-all ($\varepsilon=.6$, $\kappa=2$)	Same-all- or freq. ($\varepsilon=.5$, $f=.22$, $\kappa=2$)
1	CCC	-2	-2	0.38	0	0	0.30	0.25
2	CCS	-2	-2	0.44	0	0	0.30	0.31
3	CSC	-2	-2	0.44	0	0	0.30	0.31
4	CSS	-2	-2	0.51	0	0	0.30	0.39
5	SCC	+2, .9; -20, .1	-0.2	0.68	0	0.66	0.63	0.55
6	SCS	+2	+2	0.78	1	1	0.68	0.66
7	SSC	+2	+2	0.79	1	1	0.68	0.66
8	SSS	-2	+2	0.75	1	1	0.69	0.73
Average Action rate				0.6	0.38	0.46	0.48	0.48
MSD						0.119	0.016	0.015

Procedure and Participants. Group 6 used the procedure specified in Study 2 Set 5. Forty Mechanical Turk workers read the instructions and performed the task of 160 trials. They received a fee of \$2 for completing the task and won a bonus of \$0-1 depending on their points gained.

Results. The results replicated the findings of Set 5, with no significant influence of the bonus payment explanation on the Action rate.

Set 7

Group 7 underwent the experimental manipulation identical to that of Group 6 except for visual representation on the SQ button. Instead of including identical pentagons, the three visual cues of the SQ button showed an inverse representation of the cues of the Action button (e.g., if the cues of the Action button were CSC, the cues on the SQ button showed SCS). This manipulation intended to control for the visual novelty of the Action button. Consistent with Sets 1-6, the SQ button was placed on the right side of the display and yielded 0 points in all trials.

Procedure and Participants. Group 7 used the procedure specified in Study 2 Set 5. Forty Mechanical Turk workers read the instructions described in Set 6 and performed the task of 160 trials. They received a fee of \$2 for completing the task and had a chance to win a bonus of \$0-\$1 depending on their points gained.

Results. The results replicated the findings of Set 5, with a small and not statistically significant ($p > .05$) decrease of 0.03 in Action rate.

Table 5 Model Predictions and Experimental Results in Group 7.

Experimental results in Set 7								
Problem	Incentive structure		EV	Results		Basic models		Post-hoc models
	Signal	Payoff from Action		Exp. results	Optimal action rate	Same signal ($\kappa=4$)	Same-or-all ($\varepsilon=.6$, $\kappa=2$)	Same-all-or freq. ($\varepsilon=.5$, $f=.22$, $\kappa=2$)
1	CCC	-2	-2	0.36	0	0	0.30	0.25
2	CCS	-2	-2	0.42	0	0	0.30	0.31
3	CSC	-2	-2	0.44	0	0	0.30	0.31
4	CSS	-2	-2	0.53	0	0	0.30	0.39
5	SCC	+2, .9; -20, .1	-0.2	0.61	0	0.66	0.63	0.55
6	SCS	+2	+2	0.71	1	1	0.68	0.66
7	SSC	+2	+2	0.73	1	1	0.68	0.66
8	SSS	-2	+2	0.75	1	1	0.69	0.73
Average Action rate				0.57	0.38	0.46	0.48	0.48
MSD						0.133	0.0249	0.02