

On the descriptive value of the reliance on small-samples assumption

Ido Erev* Doron Cohen† Ofir Yakobi‡

Abstract

Experience is the best teacher. Yet, in the context of repeated decisions, experience was found to trigger deviations from maximization in the direction of underweighting of rare events. Evaluations of alternative explanations for this bias led to contradicting conclusions. Studies that focused on the aggregate choice rates, including a series of choice prediction competitions, favored the assumption that this bias reflects reliance on small samples. In contrast, studies that focused on individual decisions suggest that the bias reflects a strong myopic tendency by a significant minority of participants. The current analysis clarifies the apparent inconsistency by reanalyzing a data set that previously led to contradicting conclusions. Our analysis suggests that the apparent inconsistency reflects the differing focus of the cognitive models. Specifically, sequential adjustment models (that assume sensitivity to the payoffs' weighted averages) tend to find support for the hypothesis that the deviations from maximization are a product of strong positive recency (a form of myopia). Conversely, models assuming random sampling of past experiences tend to find support to the hypothesis that the deviations reflect reliance on small samples. We propose that the debate should be resolved by using the assumptions that provide better predictions. Applying this solution to the data set we analyzed shows that the random sampling assumption outperforms the weighted average assumption both when predicting the aggregate choice rates and when predicting the individual decisions.

Keywords: inertia, recency effect, noisy-sampler, noisy-adjuster

*Faculty of Industrial Engineering and Management, Technion, Israel. <https://orcid.org/0000-0001-9889-4070>. Email: erev@technion.ac.il.

†Economic Psychology, Department of Psychology, University of Basel, Switzerland. <https://orcid.org/0000-0001-9888-1560>.

‡Department of Psychology, University of Haifa, Israel. <https://orcid.org/0000-0002-9253-7483>.

This paper was supported by a grant from the Israel Science Foundation (grant 535/17). We thank Dirk Wulff, David Azriel, Paul Feigin, Ori Plonsky, and careful reviewers and editor for their useful comments.

Data are at https://osf.io/4jsuh/?view_only=e002039be04b40dfbd8089341e2b0980.

Copyright: © 2022. The authors license this article under the terms of the [Creative Commons Attribution 4.0 License](https://creativecommons.org/licenses/by/4.0/).

1 Introduction

Early studies of learning in repeated choice tasks highlight the value of simple models that quantify Thorndike's (1898) law of effect. The law of effect states that positive reinforcements increase the propensity of selecting the reinforced actions. The simplest quantifications of this law assume a sequential adjustment process to experienced reinforcements (e.g., see the "noisy-adjuster" model described below in Section 2). Such sequential adjustment models have five important and attractive features: First, they can capture a wide set of behavioral phenomena. For example, Erev and Roth (1998) demonstrate how a 3-parameter sequential adjustment model provides useful predictions of behavior in simple games. Second, they entail a highly efficient process: The decision maker needs to remember only one value per option — the updated subjective value. A third attractive feature is that in static settings, these simple models can approximate optimal choice (Sutton & Barto, 1998). A fourth attractive feature is that the computations these models denote are correlated with well-documented brain activity (Schultz et al., 1997). Finally, several studies have shown that the estimated parameters for models of this type can capture interesting individual differences (e.g., Yechiam et al., 2005).

Given the evidence in support of simple "sequential adjustment" models, the results of a series of choice prediction competitions (Erev et al., 2010a, 2010b; Erev et al., 2017; Plonsky et al., 2019) come as a surprise: While these choice prediction competitions were originally designed to compare alternative sequential adjustment models, these models did not perform well. Instead, the best performing models in these competitions relied on the assumption that people remember many past experiences (see related idea in Gonzalez et al., 2003), but base each choice on a small sample of these memories.

The apparent inconsistency between the evidence in favor of sequential adjustment models and the superiority of sampling models in the competitions has been previously explained in two different ways. The first explanation rests on the fact that because of the competitions' focus on predicting aggregate choice rates, the underlying processes that produce the choice rates can be misrepresented (Birnbaum, 2011; Regenwetter & Robinson, 2017; Spektor & Wulff, 2021; Wulff & van den Bos, 2018; Chen et al., 2021). Thus, it is possible that while individuals actually rely on an efficient sequential adjustment process with an individual-specific adjustment speed, on the aggregate this process is obscured. That is, while aggregate measures may best be captured by models that assume costly memory storage and sampling-based valuation, this in fact misrepresents the underlying processes. The feasibility of this explanation was recently demonstrated by Spektor and Wulff's (2021, hereafter, SW) reanalysis of the data collected by Yakobi, Cohen, Naveh and Erev (2020, hereafter YCNE).

The second explanation assumes that the apparent inconsistency reflects the reliance on different working assumptions that led to different comparisons (Erev, 2020). In accordance with this explanation, the clearest evidence in favor of sequential adjustment models come from studies that do not include a systematic comparison of the assumptions that distinguish

these models from sampling models. For example, Erev and Barron (2005) considered a simple sampling model, and then show how the data can be captured with a more complex sequential adjustment model. The current paper examines this explanation by building on SW's analysis of YCNE's data.

YCNE's original analysis focused on aggregate choice rates. Their results highlight the predictive value of models that assume sampling-based decisions and imply that the main driver for deviation from maximization (of expected payoff) is a tendency to rely on small samples.¹ Conversely, SW's analysis suggests that a simple model that assumes sequential adjustment to the payoffs' weighted average can capture the data better than the models considered by YCNE. In support, SW demonstrate that their model predicts the aggregate choice rates as well as the random sampling from experience models (as used by YCNE), and highlights an interesting pattern of individual differences that imply a new interpretation of YCNE's results. This interpretation suggests that extreme myopia (by 32% of participants), rather than reliance on small samples, is the main driver of the deviations from maximization documented by YCNE.

The current paper extends SW's analysis by considering two differences between models that assume sampling-based decisions (as in YCNE), and sequential adjustment models as considered by SW. The first difference is in the assumptions dictating how the option's subjective valuation process is carried out (i.e., by relying on random sampling or on weighted average). The second difference is in the assumptions dictating how choice is derived from those valuations (i.e., the choice rule). While YCNE limited their analysis to models that assume a deterministic choice rule (i.e., choice of the option with a higher sampled mean), SW's model uses a stochastic (noisy) choice rule. Our analysis clarifies the importance of the different assumptions regarding each process of valuation. We do so by comparing the descriptive value of the random sampling and the weighted average assumptions, while using the same stochastic choice rule as SW use (i.e., keeping the choice rule fixed).

Our results validate the existence of large individual differences, as suggested by SW, but favor a different interpretation of these differences. Our analysis shows that the random sampling assumption provides better predictions (both qualitatively and quantitatively) than the weighted average assumption, even when predicting individual decisions.

2 The Data

YCNE's analysis starts with the observation that the reliance on small samples hypothesis can be used to shed light on the conditions under which high taxation, designed to reduce reckless behavior, is likely to backfire. In certain settings, it predicts a backfiring effect even

¹One example of a deviation from maximization is provided by Group 0.6 in Figure 1. When the Tax was 0.8, most participants in this group preferred the risky option that yields "1.5, .94, -20" (expected payoff of 0.21) over a safe option that provides "0.6 with certainty."

when the tax is carefully designed to ensure that the desired behavior (i.e., safer decisions) maximizes expected return.

To test this prediction, each of 246 participants (Mturk workers) in YCNE studies was assigned to one of the three groups described in Figure 1 (one group in Study 1, and two in Study 2), and faced either three or two tasks (in a within-subject design). Each task included 100 trials, and in each trial the participant was asked to choose between three keys marked as A, B or C. The participants did not receive a description of the incentive structure and had to base their decisions on feedback that was provided after each choice. As demonstrated in Figure 1, the feedback described the obtained and the forgone payoffs. The participants' final compensation was determined by their accumulated payoffs gained during the experiment.

The middle panel in Figure 1 shows that all the tasks involved a choice between a safe option, a moderate risky option, and a counter-productive (low expected return) risky option. The groups differed with respect to the payoff from the safe option (as reflected by the group's names). The tasks faced by each group differed with respect to the magnitude of the variable "Tax" that reduces the payoff from the moderate risky option. This Tax variable simulates the adoption of a policy that tries to reduce accidents (abstracted by the loss of 20 points) by imposing a cost on the most attractive reckless behavior. The results (bottom panel of Figure 1) show that high taxation moved many participants to choose the counterproductive risky option. As a result, accident rates significantly increased.

3 Comparison of the Weighted Average and the Random Sampling Assumptions

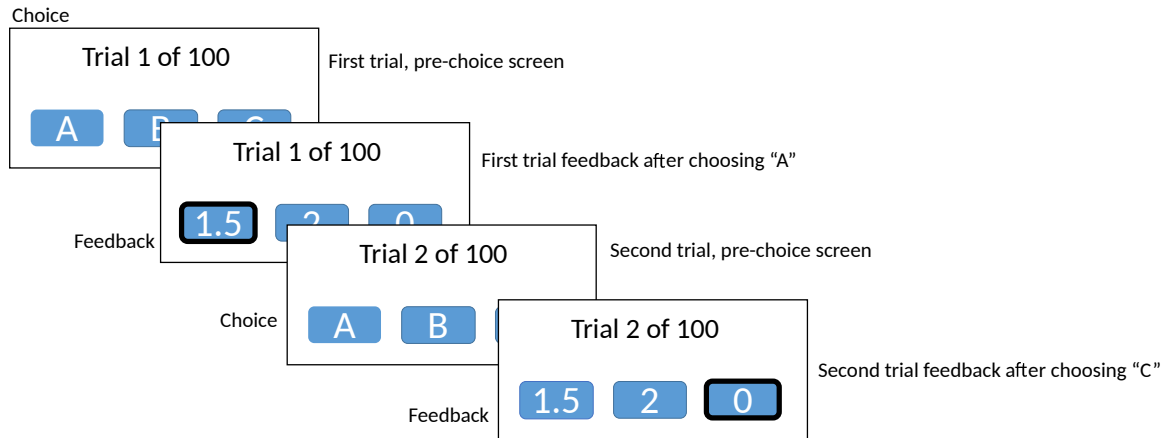
As noted above, SW show that YCNE's main results can be captured with a simple sequential adjustment model that does not include an explicit "reliance on small samples" hypothesis. Their model assumes that the subjective value of Option j for agent i in trial $t + 1$, after observing the payoff $R_{t,j,i}$ (from Option j in trial t) is:

$$Q_{t+1,j,i} = (1 - \alpha_i)Q_{t,j,i} + \alpha_i R_{t,j,i} \quad (1)$$

The initial subjective value is assumed to equal $Q_{t,j,i} = 0$, and α_i is a parameter that captures Agent i 's learning rate. Thus, the subjective value is the weighted average of the observed payoffs, and recent observations receive more weight than older observations. Besides, the model also assumes a noisy ε -greedy response rule. The model, referred to here as the "noisy-adjuster," chooses an option randomly with probability ε_i (Agent i 's error rate parameter), and the option with the highest $Q_{t,j}$ value otherwise.

In addition, SW note that YCNE's analysis ignores the existence of large between-individual differences. SW's highlighted the significance of the individual differences by estimating the parameters of their model for each individual. Their analysis, relying on maximum likelihood estimation (MLE) and shown on Figure 2a, suggests a bi-modal

The experimental screens



Procedure. The participants were assigned to one of three groups. Each participant faced two or three conditions, for a block of 100 trials. The notation “x, p; y” implies “x with probability p, y otherwise. The conditions differ with respect to the value of the variable “Tax”. The term “(2-Tax)” implies that the payoff (with $p = .97$) was 2 minus the tax (the value of the tax was 0, 0.4, or 0.8 as noted in the figure).

	Group 3or0 (Study 1)	Group 1.35 (Study 2)	Group 0.6 (Study 2)
Option			
Safe	3, .45, 0	1.35 for sure	0.6 for sure
Moderate risk	(2-Tax), .97, -20	(2-Tax), .97, -20	(2-Tax), .97, -20
High risk	1.5, .94, -20	1.5, .94, -20	1.5, .94, -20

Main results

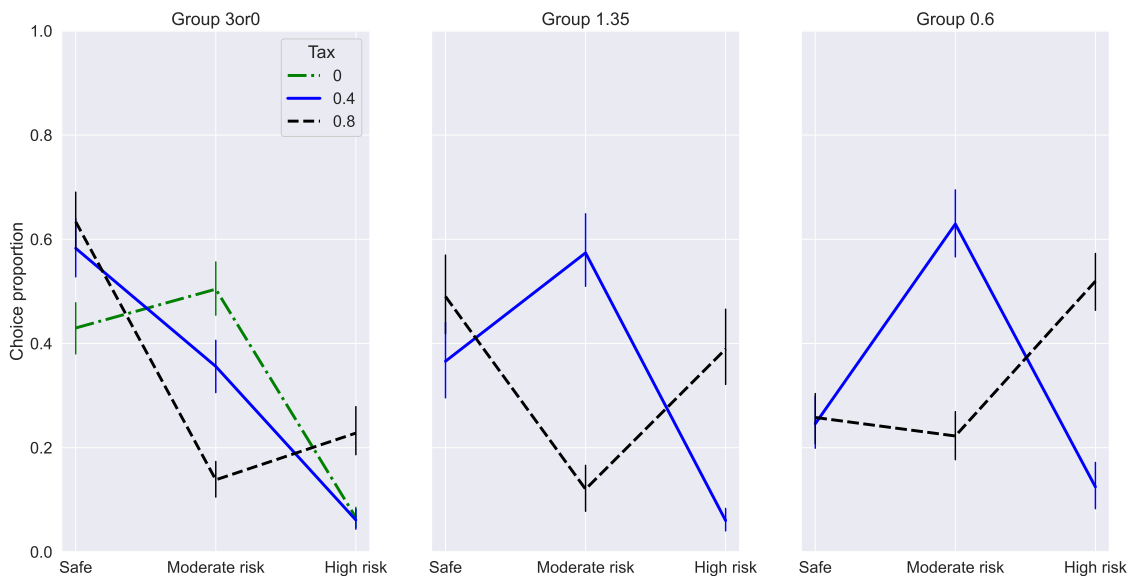


FIGURE 1: The experimental screens (top), the procedure (center), and the main results (bottom) of YCNE.

distribution: About 32% of the decision-makers appear to be “myopic” (their estimated α_i is in the range [.85, 1], suggesting extremely strong positive recency bias), and the rest appear to be “emmetropic” (their estimated α_i is positive and close to 0, suggesting weak positive recency bias).




FIGURE 2: The estimated individual parameters under the two models. Each dot summarizes the two parameters estimated based on all the decisions made by one of the 246 participants. The estimation used standard MLE criterion with a grid search procedure. The grid search for the noisy-adjuster model considers the ranges ε [.01, 1] and α [.01, .99] with steps of 0.01. The grid search for the noisy-sampler model considered the following values for ε : .01, .1, .2, .3, .4, .5, .6, .7, .8, .9, 1, and the following values of κ : 1, 2, 3, 4, 5, 6, 8, 10, 12, 15, 20, 30, 40, 100. The coarser noisy-sampler grid reflects the fact that this estimation was based on simulations. The values of κ are presented on a log scale (e.g., $\kappa = 10^2$ implies $\kappa = 100$). A small error term was added to the data for the sake of visualization.

SW’s analysis demonstrates that a simple sequential adjustment model can provide an elegant and insightful explanation of YCNE’s results. Yet, their analysis *does not* imply that their noisy-adjuster model outperforms sampling-based models. To facilitate a clear comparison of the weighted average and the random sampling assumptions, we chose to compare them while keeping the ε -greedy response rule assumed by SW. Specifically, we compare the predictive value of the noisy-adjuster model with a variant of the same, changing only the computation of the subjective values (Equation 1). The new “noisy-sampler” model assumes that the subjective values in trial $t > 1$ are determined by the average payoffs of each option in a sample of κ_i randomly selected (with replacement) previous trials (where κ_i is a parameter that captures the sample size taken by agent i).²

²To reduce computation time, we also assumed that in the case of $\kappa_i = 100$ the sampling is drawn without replacement, and all previous trials are equally weighted.

Figure 2b presents the estimated parameters (with the MLE procedure used above)³ for the noisy-sampler model. It shows that the change to the computation of the subjective value did not eliminate the variability in the estimated parameters. Yet, the distribution under the noisy-sampler model is more uniform.⁴

To evaluate the predictive value of the two models we build on the fact that each participant in YCNE's study faced at least two conditions (i.e., Tax levels). Our analysis focuses on predicting each of the 100 choices, made by each participant in each condition, based on the parameters estimated on the same participant's decisions in the other condition (or conditions) they faced. For example, the predictions of Condition Tax = 0.8 in Group 3or0 were derived with the parameters estimated based on the participant's (200) decisions in Conditions Tax = 0 and Tax = 0.4.

The accuracy of the predictions was evaluated using a log likelihood criterion. The results, summarized in Table 1, reveal a clear advantage of the random sampling assumption. The random sampling assumption fits the data better (higher log-likelihood score), and more importantly, provides better prediction. The significance of this advantage is reflected by the fact that the noisy-sampler model provided a better prediction of the impact of higher taxation for 157 (64%) of the 246 participants ($p < .001$ in a sign test, compared to the noisy-adjuster model).⁵

Figure 3 presents the log likelihood prediction scores of each participant (average over the two or three conditions faced by the participants) under the two models. Each dot in this figure presents one of the 246 participants. The results show large individual differences, and also show that most dots fall around the 45 degrees line.

The lower rows in Table 1 shows that both models capture the most interesting pattern documented by YCNE's study: The observation that taxation designed to increase the relative attractiveness of a promoted safe behavior can backfire (i.e., increasing the tax from 0.4 to 0.8 increases the accident rate). These results show that the noisy-sampler model also provides a better prediction on this measure of aggregate choice rates.

³Since the estimation of the random sampling model used simulation, it focused on the finite set of parameter combinations as explained in Figure 2's note.

⁴A related difference between the two models is suggested by analysis of the pooled parameters (when estimated using the MLE procedure to fit each of the three groups). The pooled parameters under the noisy-adjuster model imply large differences between the adjustment speed in the three groups: The estimated values of α_g are .01, .02 and .99 in Groups 3or0, 1.35, and 0.6, respectively (the estimated values of the error parameter are relatively stable at .56, .59, and .62). The estimated parameters under the noisy-sampler model are more stable: The estimated values of the sample size, κ_g , are 20, 20, and 10 in Groups 3or0, 1.35, and 0.6, respectively. The estimated value of the error parameter under the noisy-sampler model is .3 in all three groups.

⁵In an additional analysis we evaluated the two model's within-task prediction: We first estimated (using MLE) the parameters for each participant based on her choices in trials 2-51. These parameters were then used to predict the same participant's remaining 49 choices in the same task. The log likelihood scores in this analysis are -46.8 and -37.5 for the noisy-adjuster and the noisy-sampler model. The noisy-sampler model provides a better prediction (i.e., lower log likelihood) for 378 (66%) of the 577 observed sequences of choices.

TABLE 1: Model comparison.

Statistics	The noisy-adjuster model	The noisy-sampler model			
Log likelihood of the best fitting parameter of each participant in each condition (577 sequences of choices, each of 100 trials)	(−39,807)	(−36,038)			
Log likelihood of the prediction of each condition with the parameters that best fit the target participant’s decisions in each of the other conditions (577 sequences of 100 trials)	−58,326	−51,961			
Proportion of participants for which the model provides better predictions (N = 246)	0.34	0.64			
Main aggregate results					
Group	Safe	Tax	Accident Rate		
3or0	3, .45; 0	0	.019	.023	.022
		.4	.014	.020	.019
		.8	.018	.022	.022
1.35	1.35	.4	.021	.022	.021
		.8	.027	.032	.031
0.6	0.6	.4	.026	.028	.026
		.8	.038	.038	.039
MSDx100				.0016	.0009

Note. The accident rate is estimated as 0.03(Moderate risk rate) + 0.06(High risk rate). MSD is mean squared deviation.

4 From predictions to understanding

The advantage of the random sampling assumption over the weighted average assumption of course does not imply that the noisy-sampler model provides an accurate description of the underlying processes. It suggests only that the random sampling assumption provides a better approximation of the data than the weighted average assumption. To clarify the advantage of the random sampling assumption, we compared the two models on how well they predict the observed sequential dependencies found in YCNE’s data. We focus on Study 2 from YCNE (Groups 1.35 and 0.6, 161 participants)⁶, which showed the clearest

⁶See <https://journal.sjdm.org/18/18722/data-and-code.zip> for the data.




FIGURE 3: The LL (log likelihood) scores of the two models with the parameters that best fit each of the 246 participants. Each dot describes one participant. Participants below the 45-degree line are better fitted by the noisy-sampler model.

individual differences in SW's study. Table 2 and Figure 4 summarize the results of a sequential dependency analysis on trials 2 to 100, for each of the four conditions. Table 2 reveals that both assumptions (and the implied models) under-predict the participant's tendency to repeat their previous choice (i.e., the rate of inertia). The main difference between the two models involves the predicted recency effect (estimated by the difference between the choice rates after trials in which the option did or did not lead to the best payoff, see middle rows of Table 2). The median over the 9 recency scores in the data is .16. This value is similar to the median recency score under the noisy-sampler model, and much lower than the median recency score under the noisy-adjuster model (.41).

The top panel in Figure 4 presents the observed recency score as a function of the observed inertia rate for each of the 161 participants. The lower panels in Figure 4 present the predicted rates for each participant, based on the best fitting parameters over the two tax levels. In agreement with SW's analysis, the human data plot (top panel) reveals large individual differences. However, the results do not show the bimodal recency pattern predicted by the noisy-adjuster model (middle plot).

The results summarized in Table 2 and Figure 4 highlight two contributors to the advantage of the random sampling assumption. First, this assumption can capture the detrimental effects of high taxation (that implies underweighting of rare events) without over-predicting the magnitude of the recency effect. Second, the random sampling assumption's predictions are less sensitive to the fact that it ignores the tendency to repeat the last choice (inertia), compared to the weighted average assumption. To illustrate this point, consider the simulated example presented in Table 3. It focuses on the behavior of virtual agents

TABLE 2: Predicted and observed sequential dependencies in Study 2 of YCNE.

Statistic	Group, Tax	Observed			Noisy-adjuster			Noisy-sampler		
		Option X:	Safe	Med Risk	High Risk	Safe	Med Risk	High Risk	Safe	Med Risk
Inertia rate (repeating the last choice)	0.6, .4		.80			.53			.63	
	0.6, .8		.78			.58			.58	
	1.35, .4		.83			.65			.68	
	1.35, .8		.81			.66			.71	
P(X) after a trial in which Option X gave the best payoff	0.6, .4	.41	.64	.	.63	.63	.	.37	.71	.
	0.6, .8	.45	.34	.54	.64	.52	.54	.44	.34	.55
	1.35, .4	.57	.57	.	.77	.52	.	.56	.53	.
	1.35, .8	.62	.	.40	.78	.	.50	.69	.	.49
P(X) after a trial in which Option X did not give the best payoff	0.6, .4	.24	.48	.12	.21	.21	.17	.21	.56	.09
	0.6, .8	.25	.20	.38	.23	.23	.20	.24	.21	.38
	1.35, .4	.37	.38	.06	.35	.10	.12	.38	.36	.08
	1.35, .8	.48	.12	.25	.39	.10	.12	.44	.06	.35
Recency scores	0.6, .4	.17	.16	.	.42	.42	.	.16	.15	.
	0.6, .8	.20	.14	.16	.41	.29	.34	.20	.13	.13
	1.35, .4	.20	.19	.	.42	.42	.	.18	.17	.
	1.35, .8	.14	.	.15	.39	.	.38	.16	.	.

Note. Inertia rate is calculated as the rate in which choice of an option in trial $t-1$ is repeated in trial t . Recency score is calculated as the difference between the choice rate of each option after trials in which that option led to the best payoff, and after all other trials. Missing values appear for conditions in which one of the two risky options could not lead to the best payoff. Median recency scores are highlighted in Bold.

that face 100 repeated choices between “0 for sure” and an attractive risky prospect that provides “+5, .3; -1” (i.e., get +5 with probability .3, -1 otherwise).

The top row in Table 3 focuses on virtual agents that choose in accordance with the noisy-adjuster model (with the parameters $\alpha_i = .99$ and $\varepsilon_i = .01$) in some of the trials, and repeat their last choice in the other trials. The probability of repeating the last choice was 0 in the first two trials and P_{rep} thereafter (with $P_{rep} = 0, .5$ or $.9$). The bottom row presents virtual agents that behave in accordance with the noisy-sampler model (with $k_i = 1$ and $\varepsilon_i = .01$) under the same repetition conditions.⁷ The results reveal that ignoring the

⁷Note that with these parameters, the two models are equal in the amount of information they consider prior to a choice: Both models rely only on one previous outcome. The difference is that while the noisy-adjuster

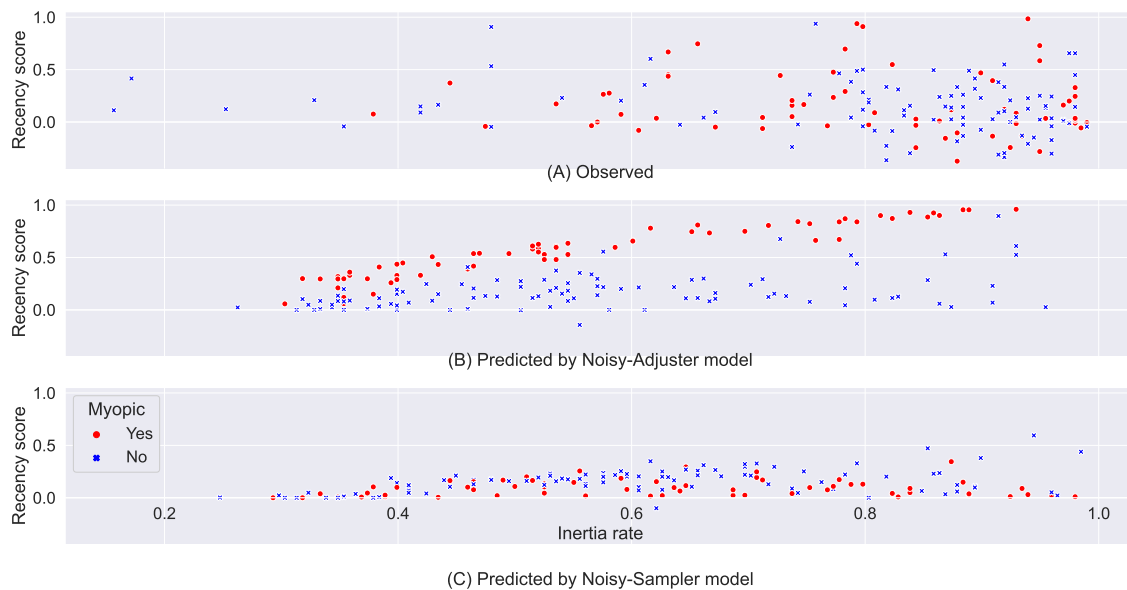


FIGURE 4: The observed and predicted recency scores (Y-axis) as a function of the repetition rate (X-axis) of each of the 161 participants in Group 1.35 and Group 0.6 (Study 2). The participants defined as myopic by SW are marked by Red dots.

rate of inertia had a limited effect on the estimated parameters of the random sampling model (bottom row): The estimation of the true parameters is robust to the level of inertia. Conversely, changes in inertia rates have a much larger effect on the estimated parameters of the noisy-adjuster model: The difference between the generating and estimated parameters increases with inertia, leading to a bias in the predicted R-rates. Increase in the level of inertia increases the estimated error parameter (ϵ_i) of the noisy-adjuster model from 0.01 to 0.70. This overestimation suggests the noisy-adjuster model cannot reproduce the generated choice rates in the presence of inertia, let alone provide useful predictions based on its estimated parameters.

5 Summary

Previous studies of decisions from experience highlight an apparent inconsistency between the results of choice prediction competitions focused on aggregate choice rates, and results of studies that focus on individual decisions. While the competitions favor models assuming reliance on small samples of randomly selected past experiences, many analyses of individual decisions favor models that assume sequential adjustment of choice propensities. The difference between these two classes of models is important as they imply very different cognitive processes. While the reliance on small samples models assume the storage and the

model relies on the last observed outcome, the noisy-sampler model relies on one outcome randomly selected from the set of previous trials.

TABLE 3: Demonstration of the impact of inertia on the estimated parameters, for models that ignore the possibility of inertia.

Models	Statistics	P _{rep} (inertia)		
		0	.50	.90
Generated Risk-rates (observed)		.30	.30	.30
Noisy-adjuster model, true parameters: $\alpha_i = .99, \varepsilon_i = .01$	Reproduced risk rates	.31	.40	.48
	Median estimated parameters	α_i	.90	.90
		ε_i	.01	.40
Noisy-sampler model, true parameters: $\kappa_i = 1, \varepsilon_i = .01$	Reproduced risk rates	.32	.33	.37
	Median estimated parameters:	κ_i	1	1
		ε_i	.03	.03

Note. The generated Risk-rates (observed) row was estimated for each model separately. Median estimated parameters' rows are estimated on those generated Risk rates of each model. Reproduced risk rates' rows are the Risk-rates reproduced with the estimated parameters of each separate model.

use of many past experiences, the sequential adjustment models assume efficient processes that require the storage of only one value (weighted payoff average) per option.

The present research clarifies this debate by highlighting the importance of the distinction between elegant explanations, and prediction-based model comparisons. Specifically, we propose that sequential adjustment models provide more elegant explanations of specific experimental results, but random sampling models tend to perform better in prediction tasks. For example, the weighted average assumption used by SW implies a simple and cognitively efficient process that fits the data we analyzed, but prediction-based model comparison highlights a clear (both qualitative and quantitative) advantage of the random sampling assumption. Our analysis shows that this advantage of the random sampling assumption is not limited to predictions of aggregate choice rates. We find that the random sampling assumption outperforms the sequential adjustment (weighted average) assumption even when the analysis focuses on individual decisions and sequential dependencies.

To understand clearly the implications of the current results, remember that there are important boundaries to the descriptive value of the noisy-sampler model supported here. The clearest boundary, in the context of pure decisions from experience, involves environments with easy-to-detect dynamic structures as illustrated by the thought experiment described in Table 4 (following Plonsky et al., 2015). While the noisy-sampler model (with the parameters estimated above based on YCNE's data) predicts a Top-rate (at trial 100) of only 29%, it is natural to assume most human subjects will quickly learn to select Top after a sequence of four losses (see related observation in Cohen & Teodorescu, 2021). We believe

that the observation that the noisy-sampler model provides useful description of YCNE results, but fails to describe the likely behavior in Table 4's thought experiment, can shed light on the underlying processes. Under one explanation of this pattern, people always try to rely on a small sample of their most similar past experiences. When it is easy to discover the most similar past experiences (in terms of the expected payoff), as in Table 4's thought experiment, choice behavior is likely to deviate from the predictions of the noisy-sampler model. Yet, when it is difficult, or impossible, to detect the most similar past experience (as in the current static setting) the effort to rely on the most similar past experiences leads to behavior that can be approximated with the current noisy-sampler model.

TABLE 4: A thought experiment.

Task: In each trial of this thought experiment, decision makers choose between “Top” and “Bottom”, earning the payoff that appears on the selected key. The payoff from Bottom is always 0, while the payoff from Top is +4 at every fifth trial (i.e., trials 5, 10, 15. . . 100) and -1 otherwise. **What will be the Top rate at Trial 100?**

Interpretation: It is natural to assume that in Trial 100 most human subjects will choose Top. This intuition suggests that the decision is based on the most similar previous trials (e.g., trials divisible by 5, or trials after four consecutive losses from Top). In contrast, the predictions of the noisy-sampler and the noisy-adjuster models (with the parameters that best fit YCNE's data) are 0.29, and 0.15 respectively.

The current similarity-based explanation suggests that the reliance on small random samples assumption can be used to shed light on natural environments in which the payoff distributions are relatively stable. While this set of situations has clear boundaries, it contains many important members. Examples include settings in which safety devices increase accidents (Cohen & Erev, 2018), taxation backfires (YCNE), people over and under-commit to a course of action (Cohen & Erev, 2021), experience reduces the tendency to trust well-calibrated experts (Erev et al., 2022) and it is necessary to enforce rules (Plonsky et al., 2021). Yet, more insight into how people respond to different dimensions of similarity, and how these similarities interact, is necessary to predict behavior when dynamic regularities are easily detectable.

References

- Birnbaum, M. H. (2011). Testing mixture models of transitive preference: comment on Regenwetter, Dana, and Davis-Stober (2011). *Psychological Review*, 118(4), 675–83.
- Chen, M., Regenwetter, M., & Davis-Stober, C. P. (2021). Collective choice may tell nothing about anyone's individual preferences. *Decision Analysis*, 18(1), 1–24. <https://doi.org/10.1287/deca.2020.0417>

- Cohen, D., & Erev, I. (2018). On safety, protection, and underweighting of rare events. *Safety Science, 109*, 377–381.
- Cohen, D., & Erev, I. (2021). Over and under commitment to a course of action in decisions from experience. *Journal of Experimental Psychology: General, 150*(12), 2455–2471.
- Cohen, D., & Teodorescu, K. (2021). On the effect of perceived patterns in decisions from sampling. *Decision* <http://dx.doi.org/10.1037/dec0000159>.
- Erev, I. (2020). Money makes the world go round, and basic research can help. *Judgment & Decision Making, 15*(3), 304–313.
- Erev, I., & Barron, G. (2005). On adaptation, maximization, and reinforcement learning among cognitive strategies. *Psychological review, 112*(4), 912.
- Erev, I., Ert, E., Plonsky, O., Cohen, D., & Cohen, O. (2017). From anomalies to forecasts: Toward a descriptive model of decisions under risk, under ambiguity, and from experience. *Psychological Review, 124*(4), 369–409. <https://doi.org/10.1037/rev0000062>.
- Erev, I., Ert, E., & Roth, A. E. (2010). A choice prediction competition for market entry games: An introduction. *Games, 1*(2), 117–136.
- Erev, I., Ert, E., Roth, A. E., Haruvy, E., Herzog, S. M., Hau, R., Hertwig, R., Stewart, T., West, R., & Lebiere, C. (2010). A choice prediction competition: Choices from experience and from description. *Journal of Behavioral Decision Making, 23*(1), 15–47.
- Erev, I., & Haruvy, E. (2005). Generality, repetition, and the role of descriptive learning models. *Journal of Mathematical Psychology, 49*(5), 357–371.
- Ert, E., Erev, I., & Roth, A. E. (2011). A choice prediction competition for social preferences in simple extensive form games: An introduction. *Games, 2*(3), 257–276.
- Erev, I., & Roth, A. E. (1998). Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *American Economic Review, 88*, 848–881.
- Erev, I., Roth, Y., & Sonsino, D. (2022). Decisions from valuations of unknown payoff distributions. *Decision, 9*(2), 172–193. <https://doi.org/10.1037/dec0000172>.
- Gonzalez, C., Lerch, J. F., & Lebiere, C. (2003). Instance-based learning in dynamic decision making. *Cognitive Science, 27*(4), 591–635.
- Plonsky, O., Apel, R., Ert, E., Tennenholtz, M., Bourgin, D., Peterson, J. C., Reichman, D., Griffiths, T. L., Russell, S. J., & Carter, E. C. (2019). Predicting human decisions with behavioral theories and machine learning. *ArXiv Preprint ArXiv:1904.06866*.
- Plonsky, O., Roth, Y., & Erev, I. (2021). Underweighting of rare events in social interactions and its implications to the design of voluntary health applications. *Judgment and Decision Making, 16*(2), 267.
- Plonsky, O., Teodorescu, K., & Erev, I. (2015). Reliance on small samples, the wavy recency effect, and similarity-based learning. *Psychological Review, 122*(4), 621.
- Regenwetter, M., & Robinson, M. M. (2017). The construct–behavior gap in behavioral decision research: A challenge beyond replicability. *Psychological Review, 124*(5), 533–550.

- Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, 275(5306), 1593–1599.
- Spektor, M. S., & Wulff, D. U. (2021). Myopia drives reckless behavior in response to over-taxation. *Judgment and Decision Making*, 16(1), 114–130.
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction*. MIT press.
- Thorndike, E. L. (1898). Animal intelligence: An experimental study of the associative processes in animals. *The Psychological Review: Monograph Supplements*, 2(4), i.
- Wulff, D. U., & van den Bos, W. (2018). Modeling choices in delay discounting. *Psychological Science*, 29(11), 1890–1894.
- Yakobi, O., Cohen, D., Naveh, E., & Erev, I. (2020). Reliance on small samples and the value of taxing reckless behaviors. *Judgment & Decision Making*, 15(2), 266–281.
- Yechiam, E., Busemeyer, J. R., Stout, J. C., & Bechara, A. (2005). Using cognitive models to map relations between neuropsychological disorders and human decision-making deficits. *Psychological science*, 16(12), 973–978.