

Reliance on small samples and the value of taxing reckless behaviors

Ofir Yakobi*

Doron Cohen[†]

Eitan Naveh[†]

Ido Erev[†]

Abstract

New technology can be used to enhance safety by imposing costs, or taxes, on certain reckless behaviors. The current paper presents two pre-registered experiments that clarify the impact of taxation of this type on decisions from experience between three alternatives. Experiment 1 focuses on an environment in which safe choices maximize expected returns and examines the impact of taxing the more attractive of two risky options. The results reveal a U-shaped effect of taxation: some taxation improves safety, but too much taxation impairs safety. Experiment 2 shows a clear negative effect of high taxation even when the taxation eliminates the expected benefit from risk-taking. Comparison of alternative models suggests that taxing reckless behaviors backfires when it significantly increases the proportion of experiences in which a more dangerous behavior leads to better outcomes than the taxed behavior. Qualitative hypotheses derived from naïve sampling models assuming small samples were only partially supported by the data.

Keywords: safety-enhancement, reliance on small samples, negative effect of taxation, reversed payoff variability effect

1 Introduction

It is often possible to enhance safety by taxing (i.e. imposing extra costs on) specific reckless behaviors. For example, many modern cars implement seat-belt reminders, wherein a loud chime sounds if occupants fail to buckle up. The seat-belt chime enhances safety by imposing costs (the unpleasant noise) on an undesired risky behavior (Lie, Krafft, Kullgren & Tingvall, 2008).

New technological tools that rely on big data and machine learning naturally extend the set of situations in which it is possible and even necessary to tax reckless behaviors. For instance, if autonomous vehicles reliably stop when they perceive pedestrians in the road ahead, people could learn to cross the road even when crossing is illegal, possibly dangerous, and will likely cause traffic jams. Thus, effective designs for autonomous transportation systems should tax this and similar behaviors.

However, previous research suggests that the impact of taxing reckless behaviors is not always positive. In certain settings, imposing a moderate cost on the target behavior tends to be effective, but imposing higher costs can backfire. For example, Katz-Navon, Naveh and Stern (2005) document a U-shaped relationship between the (self-reported) level of standardization in a particular medical unit and the reported error rate in that unit. They find that past some “optimal”

point, more standardization implies more errors, which can be a product of reckless behavior. Similarly, research on automobile safety suggests that extreme regulations, such as banning all use of mobile phones while driving, could shift users to other risky behaviors (Jacobson & Gostin, 2010), and so may have a limited effect on overall safety (McCartt, Kidd & Teoh, 2014).

Under rational economic theory, high taxation backfires when the imposed costs reduce the expected utility from the taxed behavior below the expected utility of a more dangerous alternatives (Allingham & Sandmo, 1972). The current paper extends this analysis by considering situations in which decision-makers have to base their decisions on past experiences (and cannot rely on an accurate description of the payoff distributions). Specifically, we consider decisions from experience between three alternatives by focusing on the full-feedback clicking paradigm presented in Figure 1.

Our analysis starts with the identification of a class of situations in which high taxation is expected to backfire if decision-makers rely on small samples of past experiences, but not if the samples are large enough. To clarify the behavioral implications of these situations, our experimental investigation compares alternative abstractions of decisions from experience between three alternatives. The models that best fit the results suggest that taxing reckless behaviors backfires when it significantly reduces the proportion of experiences in which the target reckless behavior leads to better outcomes than more dangerous behaviors. This effect emerges even when the taxation is optimally designed to enhance safety under the assumption that the agents use all the available data. Thus, it suggests that the reliance-on-small-sample hypothesis can be of practical value.

This research was supported by two grants from the Israel Science Foundation: I-CORE grant no. 1821/12, and grant no. 535/17.

Copyright: © 2020. The authors license this article under the terms of the Creative Commons Attribution 3.0 License.

*William Davidson Faculty of Industrial Engineering and Management, Technion—Israel Institute of Technology. Email: ofiryakobi@gmail.com.

[†]William Davidson Faculty of Industrial Engineering and Management, Technion—Israel Institute of Technology.

(A)


This is a decision-making experiment, consisting of 3 games. In each game you will play 100 rounds, and finishing the whole experiment is expected to take 15–20 minutes. In each round you will have to choose between three options: A, B or C (the three options remain the same throughout each game).

Following each choice, you will be presented with the outcome from the alternative you chose, along with the outcomes you would have gotten had you chosen the other alternatives. Your payment for participating in this game depends solely on the outcome of your choices in these rounds: if you will not pay attention — you may end up with no payment at all. Your payoff will be presented in experimental points. You will be paid using a conversion rate of 1000 points = 1\$ (USD).

You will start the experiment with an initial endowment of 3000 points. Your final payoff in this experiment will be the accumulated points you have received after playing all the experimental rounds (remember, payoffs can be negative as well!).

(B)

A trial starts by asking the participant to make a choice



After making a choice, the choice buttons are disabled and the participant is presented with feedback (the choice made and the resulting payoffs). After 1.2s, the feedback disappears and the buttons are enabled and ready for the next choice

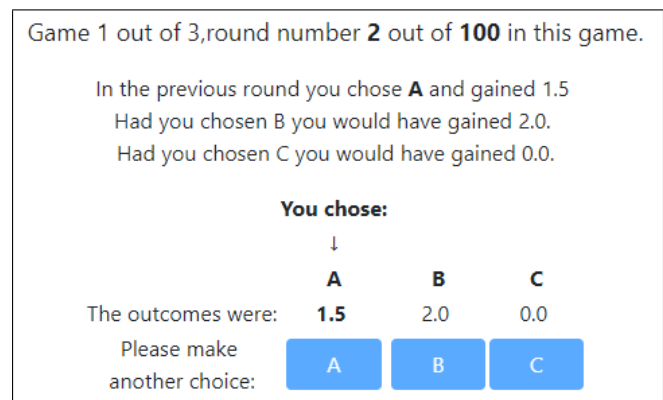


FIGURE 1: The instructions page (A) and screenshots of the main task (B) in the full feedback clicking paradigm.

2 Experiment 1

The current experiment focuses on the choice-tasks presented in Table 1, using the experimental paradigm described in Figure 1. In each trial, the decision-maker chooses between three alternatives: Safe, Low-Risk, and High-Risk, under one of three taxation conditions. The two risky alternatives can lead to an “accident” that costs 20 points, and option Low-Risk is much more attractive than High-Risk (it implies higher EV and lower accident rate) from the decision maker’s point of view.

Assume that your goal is to minimize the accident rate in a 100 trial play of the task described in the top row of Table 1. What is the optimal tax level if the possible options are 0, 0.4 and 0.8?

Rows 2–4 in Table 1 suggest that the optimal level depends on the amount of information used by the decision maker. Row 2 presents the predicted behavior under the assumption that the decision maker tries to maximize her earnings by using all the available information (i.e., feedback from all past experiences). The predictions of this “full data” model were

derived by running a computer simulation. Virtual agents were programmed to select (starting at trial 2) the option that led to the highest average payoff over all the previous trials. For example, in trial 100, these agents selected the option that led to the best average payoff over all previous 99 trials. Notice that this process is consistent with the prediction of the fictitious play model (Brown, 1951), yet in the current context implied beliefs are not fictitious, as using the full data is rational under the accurate belief that the payoff distributions are static. The results reveal that under this assumption high taxation ($Tax = 0.8$) minimizes the expected accident rate.¹

¹Predictions of the full-data model in the long-term (after many trials) converge to the predictions of rational choice theory, assuming risk neutrality. This implies choice of Safe in the current setting, independent of the magnitude of the tax. One example of an environment in which rational choice theory predicts a U-shaped effect of taxation, assuming risk neutrality, involves two types of agents, and the following inequalities: For agents of Type 1, $EV(\text{LowRisk noTax}) > EV(\text{Safe}) > EV(\text{LowRisk Moderate Tax}) > EV(\text{HighRisk}) > EV(\text{LowRisk HeavyTax})$, while for agents of Type 2, $EV(\text{LowRisk noTax}) > EV(\text{LowRisk Moderate Tax}) > EV(\text{HighRisk}) > EV(\text{LowRisk HeavyTax}) > EV(\text{Safe})$. That is, any tax moves Type 1 from Low risk to Safe; Type 2 selects LowRisk if the tax is not heavy; Heavy tax

TABLE 1: Experiment 1 — The choice tasks, predictions, and main results.

	Tax	Alternatives			Accident rate (loss of 20)
		Safe	Low-Risk	High-Risk	
(1) Choice problems		Win 3 with $p=.45$; 0 otherwise, $EV=1.35$	Win 2–Tax if notE1; lose 20 if E1, $P(E1)=0.03$, $EV=1.34-Tax(.97)$	Win 1.5 if notE1 & notE2, lose 20 if E1 or E2, $P(E1 \text{ or } E2)=0.06$, $EV=0.21$	
(2) Full data hypothesis	0	45%	55%	0%	1.6%
	0.4	72%	28%	0%	0.8%
	0.8	88%	3%	9%	0.6%
(3) The small samples hypothesis 1 (The basic Naïve sampler, $\kappa=9$)	0	33%	67%	0%	2.0%
	0.4	43%	57%	0%	1.7%
	0.8	51%	5%	44%	2.8%
(4) The small samples hypothesis 2 (2-stage Naïve sampler, $\kappa=9$)	0	39%	61%	0%	1.8%
	0.4	49%	51%	0%	1.5%
	0.8	60%	4%	36%	2.3%
(5) Results	0	43%	50%	7%	1.9%
	0.4	58%	36%	6%	1.4%
	0.8	63%	14%	23%	1.8%

Note: Row 1 presents the basic choice problem. This problem was examined under 3 tax conditions. Rows 2–4 present the predictions of three models, and Row 5 summarizes the main results. The outcomes of the risky prospects are correlated. Loss from Low-Risk occurs only when E1 occurs, and in this case High-Risk also leads to a loss.

Rows 3 and 4 in Table 1 present the predictions of two models that assume reliance on small samples of past experiences. Our interest in the reliance-on-small-samples hypothesis stems from the observation that this hypothesis captures the basic properties of binary decisions from experience (Hertwig et al., 2004; and see reviews in Rakow & Newell, 2010, and Erev & Haruvy, 2016). Row 3 presents the predictions of the basic naïve sampler model, used by Erev and Roth (2014) to quantify the small-samples hypothesis. This basic model assumes random choice in the first trial, and starting at trial 2, reliance on a random draw (sample) with replacement of k_i past trials. The value of k_i is a property of agent i , drawn from the uniform set $\{1, 2, \dots, \kappa\}$, where κ is the sole free parameter of the model. All previous trials are equally likely to be drawn, and the decision maker selects the alternative with the highest average payoff in that sample. The predictions presented in Row 2 of Table 1 are the average choice-rates in 100,000 computer simulations in which virtual agents (programmed to behave in accordance with the naïve sampler model) faced each of the three problems for 100 trials. The simulations used the parameter $\kappa=9$ that best fitted Erev and Roth’s (2014) results. The model

moves Type 2 from LowRisk to HighRisk.

predicts a clear U-shaped effect of taxation: Increasing the tax from 0.4 to 0.8 increased the expected accident rates from 1.7% to 2.8%. Thus, moderate taxation ($Tax = 0.4$) is optimal.

Roth, Wänke and Erev (2016; and Erev, Gilboa, Freedman & Roth, 2019) highlight an important shortcoming of the basic naïve sampler model: It over-predicts the impact of adding counterproductive risky options to the set of available alternatives. For example, consider the choice between “1 with certainty” and m independent options that provide “equal chance to win 10 or lose 10.” The basic model implies that when m increases, choice-rates over the m risky options will approach 100%. To address this shortcoming, Roth et al. proposed the “two-stage naïve sampler” model. This model assumes that decisions from experience among three alternatives reflect a two-stage process. At the first stage of the current implementation, the decision-makers select one of the riskiest alternatives (one of the two options with the highest observed payoff standard deviations). As in the basic model, choice at $t > 1$ is made based on a random draw with replacement of k_i past trials. The prospect with the highest mean payoff in the sample is the first-stage “tentative” choice. In the second stage the decision maker draws with replace-

ment a second sample of k_2 past experiences, to compare the first-stage tentative choice with the safest option. The value of the tentative choice is estimated by the mean of the two samples ($2k_1$ draws drawn at the first and second stages), and the value of the third alternative is evaluated only based on the second sample (k_2 draws). Row 4 of Table 1 presents the predictions of the two-stage model for the current settings. The predictions were derived using the simulation procedure described above with the parameter $\kappa=9$ that best fitted Roth, Wänke, and Erev's (2016) results. Under the two-stage assumption, the predicted accident-rates are slightly reduced (as the two-stage assumption implies reliance on a larger sample from the risky options), yet it does not change the predicted U-shape effect of taxation. The optimal tax level is still 0.4.

Experiment 1 tests Table 1's predictions. We pre-registered this experiment (<https://osf.io/4haqe>).² Our pre-registered qualitative hypotheses were:

H1: Risk rates (choice rates for the R alternatives low risk and high risk) and expected accident rates (the rate of receiving a negative payoff of 20 expected in the population) will be lower in condition Moderate-regulation (Tax = .4) compared with condition Baseline (Tax = 0).

H2: Risk rates (choice rates for the R alternatives low risk and high risk) and expected accident rates (the rate of receiving a negative payoff of 20 expected in the population) will be higher in condition Over-regulation (Tax = .8) compared with condition Baseline (Tax = 0) and Moderate-regulation (Tax = .4).

H3: Behavior in the current experiment will be captured with a behavioral model that assumes reliance on a small sample of past experiences, the two-stage naïve sampler model used in Roth, Wänke & Erev (2016).

2.1 Method

2.1.1 Participants

Eighty-five participants took part in the study (48 female; $M_{\text{age}} = 35$ years, $SD = 11.55$). They were recruited via Amazon's Mechanical Turk, and were given monetary compensation based on their performance ($M_{\text{pay}} = \$3.35$, $SD = 0.06$).

²Note that there are two minor errors in the pre-registration form: (a) In H1 and H2, "alternative" should be plural (i.e., alternatives) in "...choice rates for the R alternative". (b) We found a typo in the computer code used for generating the predictions of the Naïve sampler models. This typo led to small errors in the quantitative predictions, but did not change the qualitative prediction (in the current context, reliance on small samples implies U-shape effect of taxation, and the basic one-stage model implies larger cost of high taxation than the two-stage model). For convenience and transparency, the two sets of predictions (pre-registered and current) are included in Appendix A.

2.1.2 Procedure

We used a within-subject design. Each participant faced all three experimental conditions, which differed only with respect to the value of Tax (0, 0.4 or 0.8). Participants were informed that their payment would be based on the amount of points they gained in the experiment, with a conversion rate of 1000 points = \$1, and that they would start the experiment with an initial endowment of 3000 points. The instructions and the main experimental screens are presented in Figure 1.

All participants started with condition Tax = 0 as a baseline, and then completed the other two conditions in a randomly assigned order. Each condition included 100 trials. In each round, participants were presented with a choice between three buttons representing the Safe, Low-Risk and High-Risk alternatives (as presented in Table 1). After making their choice, participants received immediate feedback (presented for 1.2 seconds) of both the outcome from the chosen alternative, and the outcomes that the other two alternatives would have yielded had they been chosen. The on-screen position of the buttons was counterbalanced between participants and conditions. The experiment was conducted online using the oTree platform (Chen, Schonger & Wickens, 2016).

2.2 Results and Discussion

Row 5 of Table 1 and Figure 2 present the choice rates and the accident rates (estimated as $.03(\text{Low Risk}) + .06(\text{High Risk})$) as a function of the manipulation of Tax. The results reveal the U-shaped pattern predicted by the reliance-on-small-samples hypothesis, but not the increased estimated accident rate in condition Tax = 0.8 compared to Tax = 0. The risky choices rate (low- plus high-risk rates) in the baseline condition was significantly higher than in Tax = 0.4 ($t(84) = 4.9$, $p < .001$, 95% $CI_{\text{diff}} = [8\%, 23\%]$) and in Tax = 0.8 ($t(84) = 6.6$, $p < .001$, 95% $CI_{\text{diff}} = [13\%, 28\%]$).

The estimated accidents rate in the baseline condition ($Mean = 1.9\%$, $SD = 0.8\%$) was significantly higher than the estimated accident rate in the Tax = 0.4 condition ($Mean = 1.4\%$, $SD = 0.9\%$; $t(84) = 4.8$, $p < .001$, 95% $CI_{\text{diff}} = [0.3\%, 0.7\%]$), but not significantly different from accident rates in the Tax = 0.8 condition ($Mean = 1.8\%$, $SD = 0.8\%$; $t(84) = 0.9$, $p = .369$, 95% $CI_{\text{diff}} = [-1.5\%, 0.4\%]$). The negative effect of high taxation is significant too: The estimated accident rate in the Tax = 0.8 condition is higher than the rate with Tax = 0.4 condition ($t(84) = 2.566$, $p = 0.012$, 95% $CI_{\text{diff}} = [0.1\%, 0.6\%]$), but this pattern was not significant for the risky choices rate ($t(84) = 1.776$, $p = 0.079$, 95% $CI_{\text{diff}} = [-2\%, 13\%]$).

Hence, concerning accident rates, there is support for H1 but only partial support for H2 in that the accident rate in high taxation was not higher than in baseline. Concerning risky choices rate (i.e., choices for low- and high-risk options), we




FIGURE 2: The aggregated predicted and observed (estimated) accident rates in Experiment 1.

observe higher rates for baseline (Tax = 0) than moderate taxation (Tax = .4) supporting H1, but no significant difference between high and moderate taxation as predicted by H2.

Figure 3 presents the predicted and observed choice and accident rates as a function of 25-trial blocks. The results reveal a slow decrease in risk taking with experience, and that the negative impact of high taxation decreases with experience. Comparison of the three models shows that the 2-stage naïve sampler model outperforms the 1-stage model, but the full data model provides better predictions for the impact of experience. The mean square deviation (MSD) scores between the predicted and observed choice rates in Figure 3 are 0.023, 0.014, and 0.011 for the 1-stage, 2-stage and the full data models. Thus, H3 was partially supported by the data in that the U-shape effect of taxation is predicted by the reliance on small samples hypothesis and not the full-data model, but the quantitative performance of the full-data model is better. Experiment 2 was designed to improve our understanding of these observations.

3 Experiment 2

Experiment 1 documents the U-shaped effect of taxation predicted by the reliance on small samples hypothesis (Figure 2), but the results (Figure 3) also suggest that the negative effect of high taxation decreases with experience. Experiment 2 was designed to clarify the conditions under which the negative effect of high taxation is likely to be robust. Our analysis builds on the model I-SAW2 (2-stage inertia sampling and weighting, see Appendix B) that can be described as a joint generalization of the models that best predict the results of Experiment 1: the 2-stage Naïve Sampler (that predicts the U-shape effect), and the full-data model (that best predicts the choice rates). I-SAW2 implies a choice of

the option with the highest “estimated value”, where the estimated value of each option is a weighted average of the mean payoff over all previous trials (as in the full-data model) and the mean payoff in a small sample of past trials (as in the 2-stage naïve sampler model, the sample size is determined by the free parameter κ). The exact weight of the mean payoff over all trials is captured with a free parameter (ω). In addition, I-SAW2 assumes some probability of inertia (simply repeating the last choice, the probability of making a new decision is capture with the parameter ν estimated based on the observed repetition rate). This model, with the parameters estimated based on previous research ($\kappa = 8, \omega = .5, \nu = .9$), predicts that the negative effect of high taxation increases with a decrease in the attractiveness of the safe option.

In order to clarify this prediction, we chose to focus on the four problems presented in Table 2. These problems differ along two dimensions: The payoff from the safe choice (M : 0.6 or 1.35), and the Tax level (Tax : 0.4 or 0.8). Note that when $M = 0.6$, high taxation changes the EV-maximizing choice. The EV-maximizing choice when $Tax = 0.4$ is Low-Risk ($EV(\text{Low Risk}) = 0.952 > EV(\text{Safe}) = 0.6 > EV(\text{High Risk}) = 0.21$), and Safe when $Tax = 0.8$ ($EV(\text{Safe}) = 0.6 > EV(\text{Low Risk}) = 0.564 > EV(\text{High Risk}) = 0.21$).

Row 3 in Table 2 presents the prediction of I-SAW2 (with the parameters described above). It predicts a strong negative effect of taxation when the payoff from Safe is low ($M = 0.6$), and an experience-sensitive effect for taxation when the payoff from Safe is high ($M = 1.35$). When Safe gives the higher payoff, I-SAW2 predicts a large initial (for the first 25 trials) negative effect for the high taxation, but this effect is attenuated by experience – mainly due to the sensitivity to the full experienced sample.

Row 2 in Table 2 presents the prediction of the full-data model. It predicts a positive long-term effect for high taxation in both M levels. In order to clarify the relationship of our analysis to recent research, we chose to compare I-SAW2 to the Accentuation-of-Difference (AOD) model proposed by Spektor et al. (2019) to address decisions from experience among three alternatives. This model is a 5-parameter generalization of the full-data model. It assumes that the subjective value of each option is a running average of a subjective value of the observed payoff. The exact predictions of this model, explained in Appendix C, depend on five parameters: Adjustment speed (α , large values imply strong positive recency), decisiveness ($\theta > 0$, small values imply indecisiveness or random choice), power subjective utility ($\gamma > 0$, values below 1 imply diminishing sensitivity), and two parameters (ψ , and η) that capture a decrease in the subjective value of the target option with an increase in its similarity to the other options. The predictions with the parameters estimated by Spektor et al. (2019; we used the average estimates reported in each of the five experiments in Spektor et al. and assumed symmetric uniform distributions around these values with the minimum value at 0) are

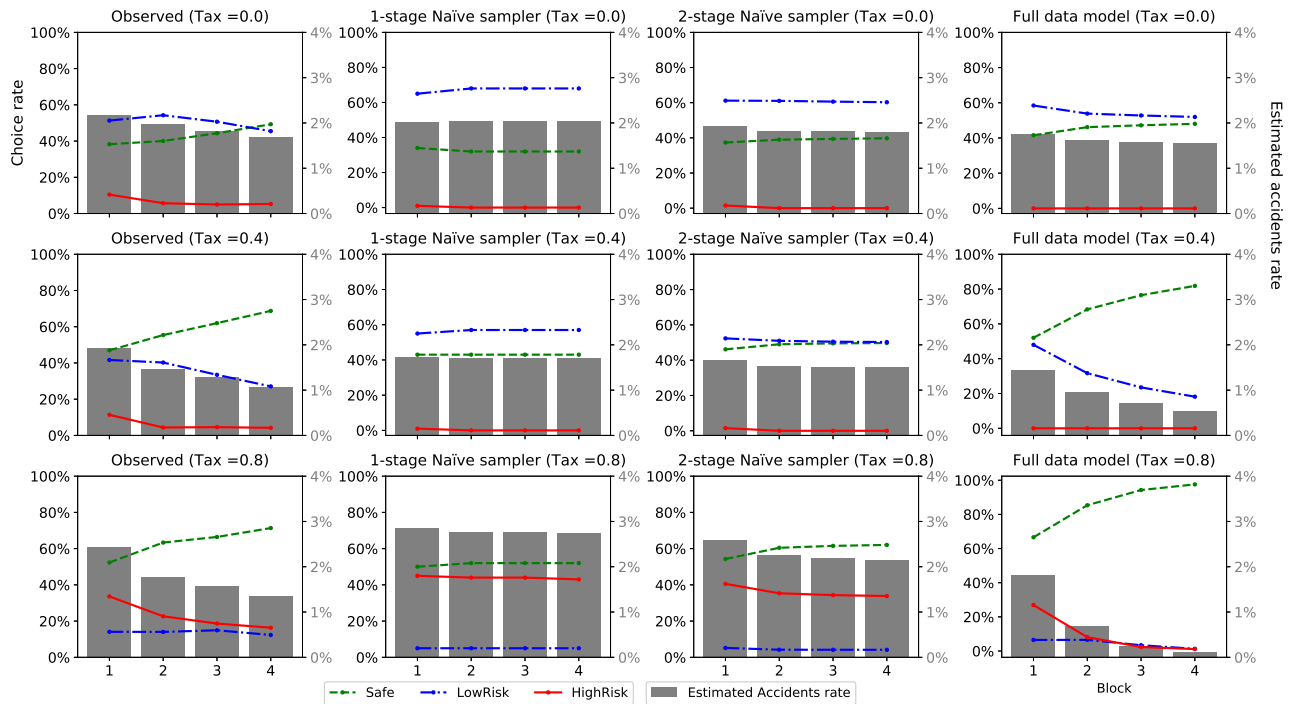


FIGURE 3: The predicted and observed choice and estimated accident rates in Experiment 1 by blocks of 25 trials.

presented in Row 4 of Table 2. They imply limited sensitivity to taxation. We pre-registered these predictions and experimental design (<https://osf.io/a9qfs>) as follows:

In the long term (after the first 25 trials of experience), in the present experiment, the following will hold:

H1: When $M = 0.6$, high taxation (Tax = 0.8) of the attractive risky option (Low-Risk) is counterproductive: It increases the accident rate (the expected rate of losing 20 points).

H2: When $M = 1.35$, high taxation of the attractive risky option (Low-Risk) is effective: It reduces the accident rate.

H3: Both effects of taxation (the negative effect implied by H1 and the positive effect implied by H2) are predicted by the I-SAW2 model. The accentuation-of-differences (AOD) model does not predict the positive effect, and the full-data model does not predict the negative effect. We predict the I-SAW2 model to show the best model fit.

3.1 Method

3.1.1 Participants

One-hundred and sixty-one mTurkers participated in the study (61 female, 98 male, two chose not to disclose; $M_{age} = 36.6$ years, $SD = 10.42$). They were recruited via Amazon’s

Mechanical Turk, and were given a monetary compensation based on their performance ($M_{pay} = \$2.88$, $SD = 0.72$).

3.1.2 Procedure

The design was similar to Experiment 1, with the addition of a between-subject factor, M . Participants were randomly assigned to one of the M conditions (0.6 or 1.35), and played two different gambling games (one for each tax level, 0.4 and 0.8), 100 trials per game, with the choice alternatives as described in Table 2. The order of the two games was randomized for each participant, and the on-screen location of the three choice buttons was randomly determined per participant per game. In addition to a \$0.7 show-up fee, participants were informed that their bonus would be based on the points they gained in the experiment, with a conversion rate of 80 points = \$1. We added an attention check procedure to the first instructions page, where participants had to write a specific word in one of the fields before clicking “next”. Participants who failed to do so were referred to a page stating that they failed to read the instructions and cannot participate further in the experiment.

3.2 Results

Row 5 of Table 2 presents the main experimental results. We conducted a repeated-measures ANOVA with block and tax as within-subject factors, and M as a between-subject factor. Degrees-of-freedom were corrected using Greenhouse-

TABLE 2: The choice tasks, predictions, and main results in Experiment 2.

(1) Choice problem	M	Tax	Trials	Alternatives			Accident rate (loss of 20)
				Safe	Low-Risk	High-Risk	
				M with certainty, EV=M	Win 2–Tax if notE1; lose 20 if E1, P(E1) = 0.03, EV = 1.34–Tax(.97)	Win 1.5 if notE1 & notE2, lose 20 if E1 or E2, P(E1 or E2)=0.06, EV = 0.21	
(2) Full-data	0.6	0.4	1–25	26%	72%	1%	2.3%
			26–100	21%	79%	0%	2.4%
	0.8	1–25	29%	20%	50%	3.6%	
		26–100	44%	39%	18%	2.2%	
	1.35	0.4	1–25	30%	68%	1%	2.1%
			26–100	78%	22%	0%	0.7%
	0.8	1–25	49%	1%	50%	3%	
		26–100	96%	0%	4%	0.3%	
(3) I-SAW2	0.6	0.4	1–25	8%	76%	6%	2.6%
			26–100	22%	78%	0.8%	2.4%
	0.8	1–25	23%	14%	63%	4.2%	
		26–100	30%	10%	60%	3.9%	
	1.35	0.4	1–25	26%	69%	5%	2.4%
			26–100	37%	62%	1%	1.9%
	0.8	1–25	37%	10%	53%	3.5%	
		26–100	70%	3%	27%	1.7%	
(4) AOD	0.6	0.4	1–25	35%	35%	30%	2.9%
			26–100	37%	37%	26%	2.7%
	0.8	1–25	35%	32%	33%	2.9%	
		26–100	38%	31%	31%	2.8%	
	1.35	0.4	1–25	38%	33%	29%	2.7%
			26–100	44%	32%	24%	2.4%
	0.8	1–25	39%	31%	30%	2.7%	
		26–100	46%	28%	26%	2.4%	
(5) Results	0.6	0.4	1–25	15%	67%	18%	3.1%
			26–100	28%	62%	10%	2.5%
	0.8	1–25	13%	21%	66%	4.6%	
		26–100	30%	21%	49%	3.6%	
	1.35	0.4	1–25	21%	69%	10%	2.7%
			26–100	42%	54%	4%	1.9%
	0.8	1–25	28%	15%	57%	3.9%	
		26–100	56%	11%	33%	2.3%	

Note: Row 1 presents the basic choice problem. This problem was examined under 4 (M by Tax) conditions. Row 2–4 present the predictions of three models, and Row 5 presents the main results. The outcomes of the risky prospects are correlated. Loss from Low-Risk occurs only when E1 occurs, and in this case High-Risk also leads to a loss.




FIGURE 4: The predicted and observed (estimated) accident rates by condition (*M* and Tax) in the last 75 trials.

Geisser estimates of sphericity. The analysis of the estimated accident rates reveals three main effects: The accident-rates increased with higher taxation ($F(1, 159) = 57, p < .001$), decreased with the benefit from the safe choice, *M* ($F(1, 159)=21.69, p<.001$) and decreased with trial number ($F(2, 318)=99.54, p<.001$). The three-way interaction predicted by I-SAW2 is insignificant ($F<1$). In the last 75 trials, high taxation increases the estimated accident rates from 2.48% to 3.56% when $M = 0.6$ ($t(85)=6.14, p<0.001, d=0.663$) in support of H1. The estimated accident rates increased from 1.86% to 2.33% when $M = 1.35$ ($t(74)=2.2, p=0.03, d=0.255$), in contrast to H2. Thus, the results confirm our first, but not the second hypothesis: the negative effect of taxation found in Experiment 1 was replicated, but in contrast to our prediction – it persisted even when the certain safe outcome was high ($M = 1.35$). The similarity between the two *M* conditions is easily seen in Figure 4 (solid black line).

Comparison of the pre-registered predictions of the three models to the behavioral results (see Figure 5) shows that the predictions of I-SAW2 are the most accurate supporting our third hypothesis. The mean squared deviations (MSD) of the last three blocks are 0.0084 (I-SAW2), 0.0231 (AOD) and 0.0568 (Full-data model). Beyond the analysis of MSD which we pre-registered, we move to explore two other measures of model accuracy. The correlations between the pre-registered predictions and the observed accident rates (presented in Table 2) are .90, .78 and .78 for I-SAW2, AOD and Full-data, respectively. A similar conclusion is reached by computing the Equivalent Number of Observations (ENO, Erev et al., 2007).³ The ENOs of the pre-registered predictions for the accident rates are 17.1, 3.8, and 2.3 for I-SAW2, AOD and Full-data, respectively.

³In the current context, the ENO of each model is the estimated number of subjects that are needed so that the average accident-rate over these subjects will provide a better prediction for the accident rate of the next (new) subject.

The reversed payoff variability effect: According to the payoff variability effect (Busemeyer & Townsend, 1991, Erev & Haruvy, 2016, and see examples in Table 4 below) an increase in payoff variability moves choice behavior toward random choice. Comparison of the high taxation condition in Experiment 1, and Condition $M=1.35, Tax=.8$ in Experiment 2 reveals a reversal of this effect. The sole difference between these conditions is the variability associated with option Safe, the expected value (EV) maximizing choice. Safe provides 1.35 with certainty in Experiment 2, and variable payoff with the same EV (3 with $p = .45$; 0 otherwise) in Experiment 1. The results show higher maximization rate with variable Safe (63% over all blocks in Experiment 1) than with certain Safe (only 49% in Experiment 2).

The reversed payoff variability effect is the main reason to the lower accuracy (higher MSD score) of the full-data model in Experiment 2. This model predicts similarly high maximization rate (about 84%) in the two “high taxation and attractive safe” conditions. The reliance on small samples models capture the reversed payoff variability effect because they imply high sensitivity to the probability that the maximizing option provides the best payoff. In the current context, payoff variability increases this probability: It is only 6% without payoff variability (in Condition $M=1.35, Tax=.8$ Safe provides the best payoff only when High risk leads to a loss), and near 50% with payoff variability (in the high taxation condition of Experiment 1, Safe provided the highest payoff if it pays 3 or if the high risk option leads to a loss).

3.3 Post hoc analyses

While the pre-registered predictions of I-SAW2 capture the main results, they predict a positive effect of taxation in Condition $M=1.35$, although a negative effect was observed. The post hoc analyses summarized in Table 3 attempt to clarify this failure by comparing the fit of all the models presented above. Table 3 presents the observed and predicted/fitted accident rates in the last 75 trials in the seven 3-alternatives conditions studied here. The left hand columns present the predictions with the pre-registered parameters. Comparison of these models shows that only the naïve sampler models capture the negative effect of taxation in Condition $M=1.35$, but I-SAW2 has the best MSD scores. The bottom rows in Table 3 presents three MSD scores: The first is the MSD between the observed and predicted error rate in last 75 trials. To facilitate comparison with the MSD scores used above, the accident rates were normalized to be between 0 and 1 by dividing them by 0.06, the maximal accident rate. The second score is the MSD between the observed and predicted choice rates in 5 blocks of 20 trials used above, and the third is the mean of the first two.

The right-hand side of Table 3 presents the fitted values that minimize the mean MSD scores. The 1-stage naïve

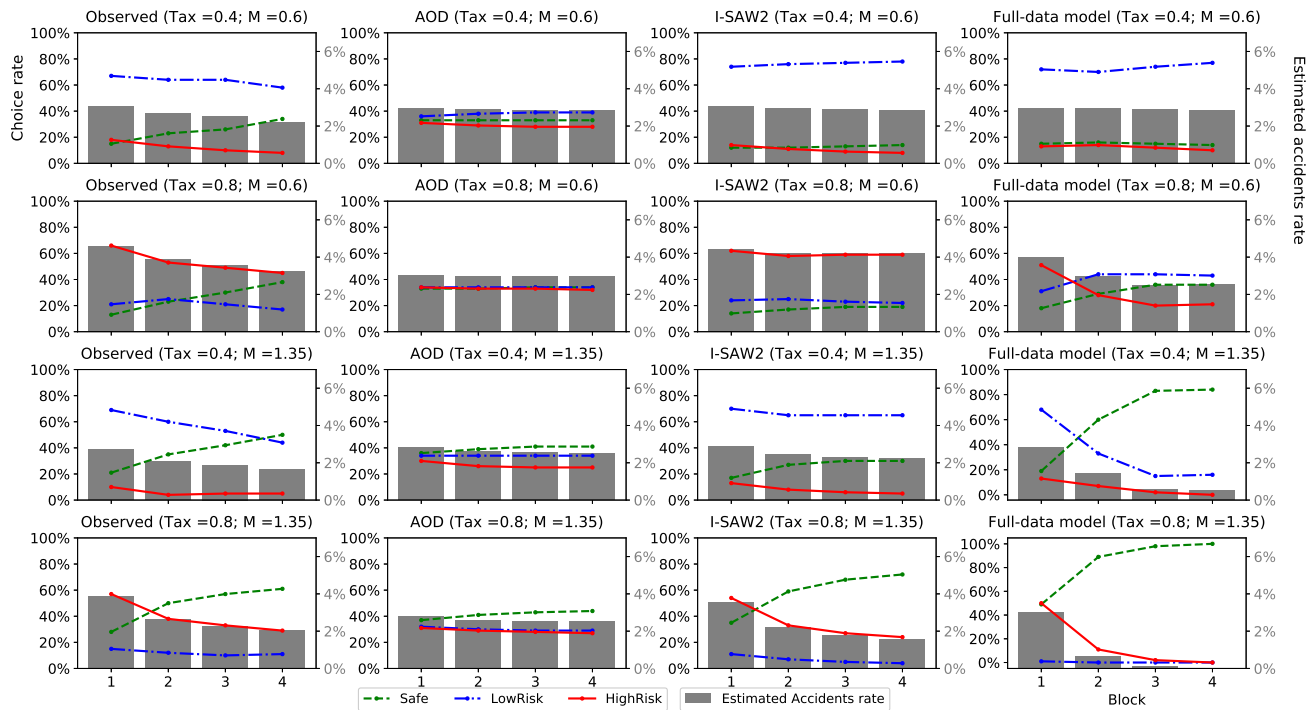


FIGURE 5: The predicted and observed choice and accident rates in Experiment 2 by blocks of 25 trials

sampler (Naïve1) model best fits the results with $\kappa=47$ with a mean MSD score of 0.0033. The two-stage naïve sampler (Naïve2) best fits the results with $\kappa=24$, with a mean MSD score of 0.0031. I-SAW2 best fits the data with $\kappa=25$ and $\omega=0$, with a mean MSD score of 0.0025. (Note that the value of ν , .9, which determines the repetition, was estimated in previous studies.) Thus, the advantage of I-SAW2 over the naïve sampler models is small, and it is not the result of weighting the mean payoff: The estimated weight of the mean payoff (ω) is zero.

AOD best fits the data with the parameters $\alpha=.36$, $\theta=20$, $\gamma=1.3$, $\psi=.6$, $\eta=.7$, and the mean MSD is 0.0028. Additional analysis shows that in the current context AOD’s unique assumption (a decrease in the attractiveness of the target option with its similarity to the other options) has limited effect: The MSD with the constraint “no sensitivity to the similarity between the options” (setting $\eta=0$ and $\psi=0$) is 0.0029 (with the parameter $\alpha=.32$, $\gamma=1.25$, and $\theta=30$). Importantly, both sets of estimates imply strong positive recency ($\alpha=.32$ or $.36$) and this value prescribes strong sensitivity to the small set of most recent outcomes. For example, with $\alpha_i = .36$ (the average value given the estimate) the last 11 outcomes receive more than 99% of the weight $(1-(.64)^{11} = 0.992)$.

In summary, all four fitted models best capture the results under the assumption that the decision makers tend to rely on a subsample of their past experiences, but the size of the sample appear to be larger than the size suggested by the previous research that underlies our pre-registrations. The

other assumptions introduced by these models do not appear to improve the fit of the current conditions in a meaningful way. These assumptions include: sensitivity to the mean payoff over all trials, sensitivity to the similarity between the options, inertia, and a two-stage decision process.

The four fitted models differ with respect to the experiences predicted to be most influential. In order to compare these models, we computed the choice rate of the counter-productive High-Risk option as a function of the number of trials since the last “accident” (loss of 20) from this choice, conditioned on observing no more than one loss in the last 10 trials. Figure 6 presents the observed and predicted rates (after trial 25) by the three best fitted models (the predictions of Naïve1 are almost identical to the predictions of Naïve2 and were thus omitted from the figure). It shows that the AOD model predicts a positive recency, which is not apparent in the behavioral data. Rather, the observed curves, like the curves predicted by the reliance on small samples models, are relatively flat.⁴

⁴Yet, the observed curves are not completely flat: Over all seven conditions, the choice rate of the high risk option is 15.5% immediately after a loss, 21% 2 to 5 trials after the last loss, and only 18% 6 to 10 trials after the last loss. Using subject as a unit of analysis, the middle rate (2-5 trials after a loss) is significantly larger than the other two rates ($t[245]=5.64$, $p<.0001$, and $t[245]= 4.46$, $p<.0001$). This pattern is consistent with the wavy recency effect documented by Plonsky et al., (2015, and see Szollosi, 2019), and suggests one way in which the current models can be improved.

TABLE 3: The observed, predicted, and fitted accident rates in the last 75 trials in all seven conditions (top), and summary of the MSD scores (bottom).

			Pre-registered before Experiment 1			Pre-registered before Experiment 2		Fitted post hoc				
Exp.	M	Tax	Observed	Full	1-stage Naïve Sampler, $\kappa=9$	2-stage Naïve Sampler, $\kappa=9$	I-SAW2 ($\kappa=8, \nu=.9, \omega=.5$)	AOD ($\alpha=.05, \theta=2.7, \gamma=.61, \psi=.63, \eta=1.2$)	1-stage Naïve Sampler, $\kappa=47$	2-stage Naïve Sampler, $\kappa=24$	I-SAW2 ($\kappa=25, \nu=.9, \omega=0$)	AOD ($\alpha=.36, \theta=20, \gamma=1.3, \psi=.6, \eta=.7$)
1	(3, .45,0)	0	1.8%	1.6%	2.0%	1.9%	1.9%	2.8%	1.8%	1.8%	1.8%	1.7%
1	(3, .45,0)	.4	1.3%	0.7%	1.7%	1.6%	1.5%	2.8%	1.3%	1.3%	1.4%	1.5%
1	(3, .45,0)	.8	1.6%	0.3%	2.8%	2.2%	1.8%	2.8%	1.7%	1.6%	1.6%	1.9%
2	0.6	.4	2.5%	2.4%	2.6%	2.3%	2.4%	2.7%	2.2%	2.3%	2.3%	2.5%
2	0.6	.8	3.6%	2.2%	4.9%	3.9%	3.9%	2.8%	3.4%	3.3%	3.3%	3.6%
2	1.35	.4	1.9%	0.7%	2.6%	2.3%	1.9%	2.4%	1.8%	1.8%	1.8%	2.1%
2	1.35	.8	2.3%	0.3%	4.5%	3.6%	1.7%	2.4%	2.3%	2.2%	2.2%	2.4%
MSD accident rates last 75 trials			0.0320	0.0344	0.0097	0.0022	0.0021	0.0006	0.0006	0.0006	0.0007	0.0007
MSD choice rates in 4 blocks			0.0380	0.0378	0.0160	0.0089	0.0314	0.0060	0.0056	0.0045	0.0048	0.0048
Mean MSD			0.0354	0.0361	0.0130	0.0055	0.0267	0.0033	0.0031	0.0025	0.0028	0.0028

4 General Discussion

Previous research suggests that the effort to increase safety by taxing reckless behaviors can backfire. Under certain conditions, the taxation of risky decisions impairs safety. The current analysis clarifies these conditions. It shows that this effect emerges even when the tax is carefully designed to ensure that the expected outcome of the undesirable reckless behaviors is, from the agents’ point of view, lower than the expected benefit from the desirable safe behavior. Experiment 2 shows that the negative effect of high taxation is not eliminated by experience (of 100 trials) with full feedback in which the choice of the problematic risky prospect reduces the expected return from 1.35 to 0.21. Experiment 1 presents a counterexample to the assertion that high incentives are more effective than low incentives (Gneezy & Rustichini, 2000). The results reflect a U-shaped effect of taxation: low taxation is effective, but high taxation backfires.

Overall, the hypotheses derived from naïve sampler models were partially supported by the data. Yet, the pre-registered hypotheses over-predict the negative effect of high taxation in Experiment 1, and under-predict this effect in Experiment 2. Post hoc comparison of alternative models suggests that the main driver of the observed negative impact of high taxation might be a tendency to rely on a subsample of past experiences. The results are best captured with models assuming that on average the decision makers rely

on sample of 13 to 24 past experiences. While this sample size is larger than suggested by the previous research that underlies our pre-registered predictions, it is small enough to trigger the observed U-shaped effect. To understand why reliance on small sample can trigger a negative effect of high taxation, note that reliance on small samples implies high sensitivity to the frequent outcomes. The negative effect of high taxation occurs, under the current hypothesis, when the taxation targets a moderately risky behavior, and reduces the frequent outcome of this behavior below the frequent outcome of a counterproductive riskier option. Thus, high taxation of a moderate risky option, can lead agents that rely on small samples to prefer a riskier option even when this behavior significantly impairs their expected return.

In order to clarify the implications of the current analysis to the basic study of decisions from experience, it is constructive to consider two possible explanations to the relatively large subsample size estimated in the present work (average of 13 to 24 past experiences) compared with previous research (reviewed in Erev & Roth, 2014). The first explanation assumes that with more alternatives (three in the present work and two in Erev and Roth’s review) decision makers rely on larger sample size. The second explanation states that the difference reflects a “relatively flat MSD (error) function”. That is, in the current context the prediction error of the reliance on small samples models is relatively

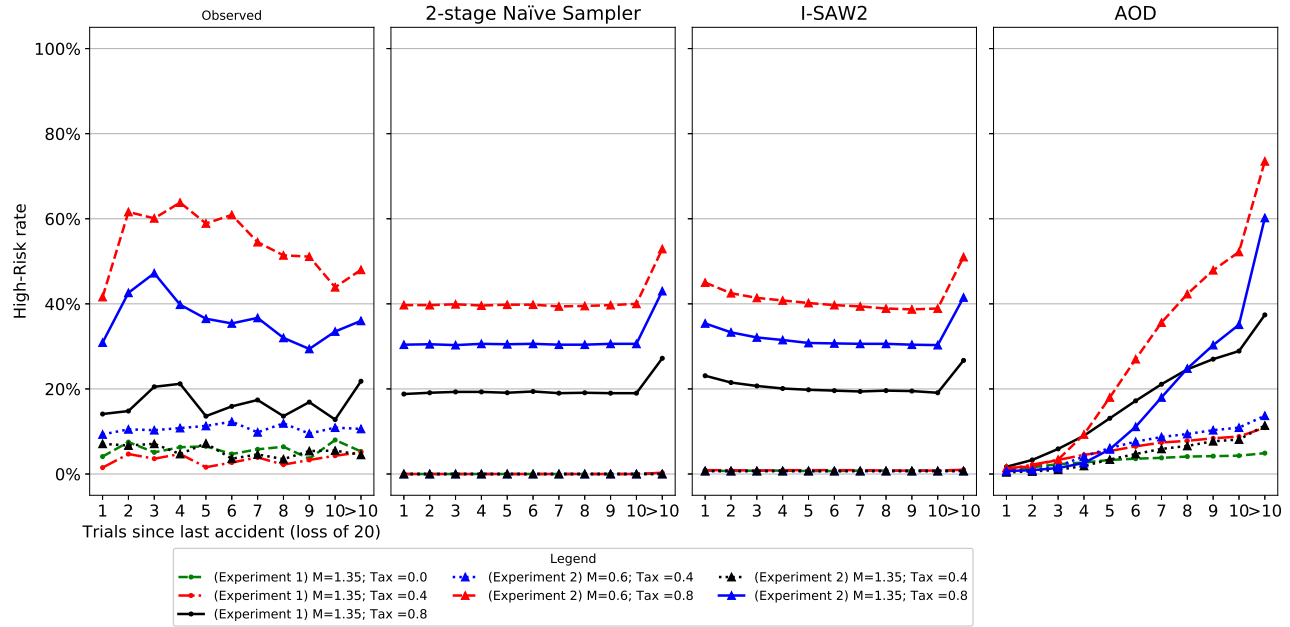


FIGURE 6: Observed and predicted High-Risk rates (after trial 25) as a function of the number of trials since the last accident

insensitive to value of κ in the ranges we examined. While the first explanation questions the practical value of the reliance on small samples hypothesis (it implies that the exact parameter changes with the number of alternatives), the second highlights its potential (it implies relative robustness to the exact parameter).

To compare these explanations, we derived the predictions of the models presented above (with the pre-registered and the re-estimated parameters) for trials 75 to 100 in the six conditions used by Erev and Roth (2014), which allow us to demonstrate the main properties of binary decisions from experience. The left-hand side of Table 4 presents the experimental conditions and main results. The first four problems document the original payoff variability effect (e.g., lower maximization rate in the problem with payoff variability), and the two lower problems demonstrate underweighting of rare events (risk seeking in Problem ER5, and risk aversion in Problem ER6). The right hand columns in Table 4 show that all four fitted models capture these two phenomena. In addition, the results show that the MSD of the models estimated above is only slightly higher than the MSD of the $\kappa = 9$ models. Over the two sets of conditions (Tables 3 and 4), the models with the parameters estimated above outperform the models with the pre-registered parameters. We feel that these observations support the “relatively flat error function” explanation. The results demonstrate that it is possible to capture Erev and Roth’s and the current results with a single one-parameter model (the 2-stage naïve sampler captures the results the best), and the estimated pa-

rameter implies reliance on an average sample size of more than 5 observations. More importantly, the results suggest that accurate estimation of models that assumes reliance on small samples requires large set of experiments, and it is natural to assume that the current estimations can be improved. We hope to address this task in future research.

At first glance the success of the $\kappa=24$ Naïve2 model (and the $\kappa=47$ Naïve1 model) appears to be inconsistent with the observation that the “reliance on small samples” implies underweighting of a 10% event only when the sample size is smaller than 7 (Teodorescu et al., 2013; Shteingart & Loewenstein, 2015).⁵ More careful analysis reveals that the “sample smaller than 7” observation was derived under the assumptions that (1) all choices are made with the same sample size, and (2) sampling without replacement. The current sampling models assume that (1) the actual sample size is uniformly distributed between 1 and κ , and (2) sampling with replacement. Under the current naïve sampler models, a 10% event (e.g., in trials 76 to 100 while facing Problems ER5 or ER6) is underweighted as long as $\kappa < 70$ (average sample size smaller than 35.5).

A second contribution of the current analysis to basic decision-making research involves the clarification of the difference between the repeated decisions considered here and the one-shot decisions from sampling reviewed by Wulff,

⁵Let p be the occurrence rate of rare event. Sample of size k will include the rare event with probability below 0.5 when the following inequality holds $0.5 < P(\text{no rare}) = (1-p)^k$. This inequality implies that $k < \text{Log}(0.5)/\text{Log}(1-p)$. For example, when $p = 0.1$, $k < 6.57$. That is, when k is 6 or lower, most samples do not include the rare events.

TABLE 4: The observed and predicted choice rates of option Action in trials 76 to 100 of the six conditions analyzed by Erev and Roth (2014) under the models considered above.

		Pre-registered before Experiment 1				Pre-registered before Experiment 2		Fitted based on Exp 1 and 2			
Problem	Action (the alternative was “0 with certainty.”)	Observed	Full	AOD		I-SAW2 ($\kappa=8, \nu=.9, \omega=.5$)	($\alpha=.05, \theta=2.7, \gamma=.61, \psi=.63, \eta=1.2$)	AOD		I-SAW2 ($\kappa=25, \nu=.9, \omega=0$)	($\alpha=.36, \theta=20, \gamma=1.3, \psi=.6, \eta=.7$)
				1-stage Naïve Sampler, $\kappa=9$	2-stage Naïve Sampler, $\kappa=9$			1-stage Naïve Sampler, $\kappa=47$	2-stage Naïve Sampler, $\kappa=24$		
ER1	1 with certainty	96%	100%	100%	100%	99%	74%	100%	100%	99%	98%
ER2	11, .5; -9	55%	83%	58%	58%	60%	59%	66%	63%	65%	61%
ER3	-1 with certainty	6%	0%	0%	0%	1%	26%	0%	0%	1%	2%
ER4	9, .5; -11	43%	18%	42%	42%	40%	42%	35%	38%	35%	39%
ER5	1, .9; -10	59%	40%	62%	62%	65%	63%	52%	56%	54%	57%
ER6	10, .1; -1	27%	59%	38%	38%	35%	37%	47%	44%	46%	43%
MSD over 100 trials			0.0436	0.0027	0.0027	0.0032	0.0199	0.0095	0.0059	0.0081	0.0044

Mergenthaler-Canseco and Hertwig (2018). Wulff et al. show that the main properties of decisions from sampling can be captured with the assumption that decision makers behave as if they equally weight all the observed experiences. In contrast, the current results suggest that in repeated decisions, decision makers behave as if they rely on a sub-sample of their past experiences. The difference can be explained with the hypothesis that experiences with repeated decisions increase the effort to response to patterns by selecting the option that led to the best outcome in the most similar past experiences (Plonsky et al., 2015).

A third contribution of the current analysis to basic research involves the documentation of a reversed payoff variability effect. While the original effect implies that payoff variability increases random choice (as demonstrated in Table 4), in the current high taxation conditions, an increase in the payoff variability of the EV maximizing option, increased maximization. Importantly, the reliance on small samples hypothesis correctly predicts the direction of the payoff variability effect.

Wider practical implications of the current results stem from the minimalistic nature of the current experimental paradigm relative to settings examined in previous demonstrations of the negative effects of economic incentives. All the previous demonstrations of the negative impact of economic incentives we are familiar with (see review in Gneezy, Meier & Rey-Biel, 2011) emerge in situations in which the incentives are explicitly described (i.e., decision from de-

scription). The leading explanations of these demonstrations focus on the description. They state that the description provides a signal that changes the subjective utilities, and the exact effect appears to be situation specific. While the current results do not question the validity of these explanations, they demonstrate the existence of a very different contributor to these effects. The negative effect of incentives demonstrated here, is not a result of a specific signal; it appears to reflect a general property of human learning that emerges in choice between more than two alternatives. The current results can be captured with simple quantitative models with clear implications. Specifically, the reliance on small sample models suggest that monetary incentives backfire when they are designed to reduce the choice rates of a moderately undesirable behavior, but increase the proportion of experiences in which an even less desirable behavior leads to the best outcomes.

5 References

Allingham, M. G., & Sandmo, A. (1972). Income tax evasion: A theoretical analysis. *Journal of Public Economics*, 1(3), 323–338, [http://dx.doi.org/10.1016/0047-2727\(72\)90010-2](http://dx.doi.org/10.1016/0047-2727(72)90010-2).

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, <http://dx.doi.org/10.1002/bdm.443>.
- 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, <http://dx.doi.org/10.1016/j.jbef.2015.12.001>.
- Fox, C. R., & Hadar, L. (2006). “Decisions from experience” = sampling error + prospect theory: Reconsidering Hertwig, Barron, Weber & Erev (2004) *Judgment and Decision Making*, 1, 159–161.
- Doron, C. (2016). Risk compensation, over-regulation and the effect of experience (unpublished master’s thesis). Technion, Haifa, Israel.
- 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, <http://dx.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, <http://dx.doi.org/10.3390/g1020117>.
- Erev, I., Ert, E., Roth, A. E., Haruvy, E., Herzog, S. M., Hau, R., et al. (2010). A choice prediction competition: Choices from experience and from description. *Journal of Behavioral Decision Making*, 23(1), 15–47, <http://dx.doi.org/10.1002/bdm.683>.
- Erev, I., Gilboa-Freedman, Gail, & Roth, Y. (2019). The impact of rewarding medium effort and the role of sample size. *Journal of Behavioral Decision Making*, 32(5), 507–520, <http://dx.doi.org/10.1002/bdm.2125>.
- Erev, I., & Haruvy, E. (2016). Learning and the economics of small decisions. *The handbook of Experimental Economics*, 2, 638–700, <http://dx.doi.org/10.1515/9781400883172-011>.
- Erev, I., & Roth, A. E. (2014). Maximization, learning, and economic behavior. *Proceedings of the National Academy of Sciences*, 111(Supplement 3), 10818–10825.
- Erev, I., Roth, A. E., Slonim, R. L., & Barron, G. (2007). Learning and equilibrium as useful approximations: Accuracy of prediction on randomly selected constant sum games. *Economic Theory*, 33(1), 29–51.
- Glöckner, A., Hilbig, B. E., Henninger, F., & Fiedler, S. (2016). The reversed description-experience gap: Disentangling sources of presentation format effects in risky choice. *Journal of Experimental Psychology: General*, 145(4), 486–508, <http://dx.doi.org/10.1037/a0040103>.
- Gneezy, U., Meier, S., & Rey-Biel, P. (2011). When and why incentives (don’t) work to modify behavior. *Journal of Economic Perspectives*, 25(4), 191–210. <https://dx.doi.org/10.1257/jep.25.4.191>.
- Gneezy, U., & Rustichini, A. (2000). A fine is a price. *Journal of Legal Studies*, 29(1), 1–17, <http://dx.doi.org/10.1086/468061>.
- Gneezy, U., & Rustichini, A. (2000). Pay enough or don’t pay at all. *Quarterly Journal of Economics*, 115(3), 791–810 <http://dx.doi.org/10.1162/003355300554917>.
- 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, <http://dx.doi.org/10.1111/j.0956-7976.2004.00715.x>.
- Hertwig, R., & Erev, I. (2009). The description–experience gap in risky choice. *Trends in Cognitive Sciences*, 13(12), 517–523.
- Ho, C., & Spence, C. (2005). Assessing the effectiveness of various auditory cues in capturing a driver’s visual attention. *Journal of Experimental Psychology: Applied*, 11(3), 157, <http://dx.doi.org/10.1037/1076-898X.11.3.157>.
- Jacobson, P. D., & Gostin, L. O. (2010). Reducing distracted driving: regulation and education to avert traffic injuries and fatalities. *JAMA*, 303(14), 1419–1420.
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2), 363–391.
- Katz-Navon, T. A. L., Naveh, E., & Stern, Z. (2005). Safety climate in health care organizations: A multidimensional approach. *Academy of Management Journal*, 48(6), 1075–1089.
- Lie, A., Krafft, M., Kullgren, A., & Tingvall, C. (2008). Intelligent seat belt reminders—Do they change driver seat belt use in Europe?. *Traffic injury prevention*, 9(5), 446–449.
- McCartt, A. T., Kidd, D. G., & Teoh, E. R. (2014). Driver cellphone and texting bans in the united states: Evidence of effectiveness. *Annals of Advances in Automotive Medicine*, 58, 99–114.
- 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.
- Rakow, T., & Newell, B. R. (2010). Degrees of uncertainty: An overview and framework for future research on experience-based choice. *Journal of Behavioral Decision Making*, 23(1), 1–14.
- 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.
- Shteingart, H., & Loewenstein, Y. (2015). The effect of sample size and cognitive strategy on probability estimation bias. *Decision*, 2(2), 107–117.
- Spektor, M. S., Gluth, S., Fontanesi, L., & Rieskamp, J. (2019). How similarity between choice options affects decisions from experience: The accentuation-of-differences model. *Psychological Review*, 126(1), 52–88.
- Szollosi, A., Liang, G., Konstantinidis, E., Donkin, C., & Newell, B. R. (2019). Simultaneous underweighting and overestimation of rare events: Unpacking a paradox. *Journal of Experimental Psychology: General*, 148(12), 2207–2217.

Teoderescu, K., Amir, M., & Erev, I. (2013). The experience-description gap and the role of the inter decision interval. *Progress in brain research* (1st ed., pp. 99 –115). Amsterdam, the Netherlands: Elsevier. <http://dx.doi.org/10.1016/B978-0-444-62604-2.00006-X>.

Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, 5(4), 297–323.

Wulff, D. U., Mergenthaler-Canseco, M., & Hertwig, R. (2018). A meta-analytic review of two modes of learning and the description-experience gap. *Psychological Bulletin*, 144(2), 140–176.

Appendix A

TABLE A1 – Description of alternatives and taxation conditions examined in Experiment 1. The original predictions of the Naïve sampler model as pre-registered are presented along with the corrected predictions (right-hand column).

Condition	Alternatives	v1	p1	v2	p2	EV	Simulation of choice rates based on the Naïve sampler model: original and (corrected) ⁶		Results		two-stage correct
							One-stage, k = 9	Two-stage, k = 9	Choice rate	Exp. Acc. rate	
Baseline	Safe	3	0.45	0	0.55	1.35	29% (32.5%)	37% (39%)	43%		39%
[Tax = 0]	Moderate-risk	2	0.97	-20	0.03	1.34	71% (67.3%)	63% (61%)	50%	1.90%	61%
	High-risk	1.5	0.94	-20	0.06	0.21	0% (0.2%)	0%	7%		0%
Moderate-taxation	Safe	3	0.45	0	0.55	1.35	47% (43%)	53% (49%)	58%		49%
[Tax = 0.4]	Moderate-risk	1.6	0.97	-20	0.03	0.95	53% (56.5%)	47% (51%)	36%	1.40%	51%
	High-risk	1.5	0.94	-20	0.06	0.21	0% (0.5%)	0%	6%		0%
High-taxation	Safe	3	0.45	0	0.55	1.35	51% (51%)	59% (60%)	63%		60%
[Tax = 0.8]	Moderate-risk	1.2	0.97	-20	0.03	0.56	4% (5%)	6% (4%)	14%	1.80%	4%
	High-risk	1.5	0.94	-20	0.06	0.21	45% (44%)	35% (36%)	23%		36%

Appendix B

The two-stage Inertia, SAMpling and Weighting model (I-SAW2)

I-SAW2 is a two parameter generalization of the 2-stage naïve sampler model. The first captures inertia: the probability that Agent *i* makes a new choice in trial $t > 1$ is n_i ; the agent simply repeats her last choice with probability $1 - n_i$. The term n_i is a property of Agent *i*, and is uniformly distributed in the population between 0 and ν . The second addition involves sensitivity to the average payoff. I-SAW2 allows for the possibility that the final choice in each trial is based on a weighting of the sample means, and the average payoff obtained from each of these options over all previous trials. Specifically, Agent *i*'s estimate of the value of option *j* in the second stage of trial *t* is $w_i(OP_{j,t}) + (1 - w_i)(SP_{j,t})$, where $OP_{j,t}$ is the average payoff over all observed trials, and $SP_{j,t}$ is the average payoff of the sampled trials (k_i trials for the safest alternative and $2k_i$ trials for the risky alternative). The weight w_i is a property of Agent *i*, and is uniformly distributed in the population from 0 to ω . Erev and Roth (2019) estimated the value of the three parameters to equal $\kappa = 8$, $\nu = 0.9$ and $\omega = 0.5$.

TABLE B1: The predictions of I-SAW2 for Experiment 1 (using the format of Table 1).

Tax	Safe	Low-Risk	High-Risk	Estimated accident rate (loss of 20)	MSD score in the prediction of the choice rates in 4 blocks
	Win 3 with $p = .45$; 0 otherwise EV = 1.35	Win 2–Tax if notE1; lose 20 if E1 P(E1) = 0.03 EV = 1.34–Tax(.97)	Win 1.5 if notE1 & notE2 lose 20 if E1 or E2 P(E1 or E2)=0.06 EV = 0.21		
0	38%	60%	2%	1.92%	
0.4	51%	46%	2%	1.54%	0.007
0.8	66%	8%	26%	1.80%	

Appendix C

The Accentuation of differences (AOD) model

The probability $Pr(i, t)$ of choosing alternative *i* in trial *t* is based on the soft-max rule:

$$Pr(i, t) = \frac{e^{\theta X_{i,t}}}{\sum_{j=1}^J e^{\theta X_{j,t}}} \tag{1}$$

Where: *J* is the number of alternatives (three in all the experiments analyzed here), here and thereafter. θ , captures the sensitivity to choose the alternative *X*, with the highest subjective expectations.

The subjective value of the alternative *i* in the current trial ($t+1$), is based on the weighted average of the current alternative *i* subjective value ($X_{i,t}$) at the previous trial – *t* and the “AOD” value of the current alternative $S_{AOD}(O_{i,t})$:

$$X_{i,t+1} = X_{i,t} + \alpha(S_{AOD}(O_{i,t}) - X_{i,t}) \tag{2}$$

Where: α is the learning rate (between 0 and 1). $O_{i,t}$ – is the observed payoff of alternative *i* in trial *t*. $X_{i,t}$ - is the subjective value of the alternative *i* in trial *t*. $S_{AOD}(O_{i,t})$ - is the Accentuation of Difference subjective value, which consist of the difference between the subjective utility (4) of the current alternative and its subjective similarity (5)

$$S_{AOD}(O_{i,t}) = f(O_{i,t}) - \eta * S_{i,t} * \overline{|f(O_{j,Z})|} \tag{3}$$

Where: $f(O_{i,t})$ – is the subjective utility of alternative *i* at trial *t*. $\overline{|f(O_{j,Z})|}$ – is the mean absolute value sum of all the subjective utilities at trial *t*. $S_{i,t}$ is the average similarity to other alternatives, and η is the weighting average between the similarity and the subjective utility.

The subjective utility is calculated as follows:

$$f(O_{i,t}) = \begin{cases} -|O_{i,t}|^\gamma, & O_{i,t} < 0 \\ O_{i,t}^\gamma, & O_{i,t} \geq 0 \end{cases} \tag{4}$$

Thus it is the obtained value of the respective alternative in the power of γ .

The subjective similarity to other alternatives is based on the absolute difference in subjective utility of the current alternative and the other (two in our settings) alternatives:

$$S_{i,t} = \frac{\sum_{j=1}^{J_i} e^{-\psi * |f(O_{i,t-1}) - f(O_{j,t-1})|}}{J - 1} \tag{5}$$

Where ψ is the scaling parameter of the subjective difference. Spektor et al. (2019) estimated a different distribution of the five free parameters to each of their five experiments. The current analysis use parameters that were randomly selected from uniform distribution around these distributions: $\alpha=.6, \psi=2.6, \theta=120, \gamma=3, \eta 1.6$

TABLE C1: The predictions of AOD for Experiment 1 (using the format of Table 1).

Tax	Safe	Low-Risk	High-Risk	Estimated accident rate (loss of 20)	MSD score in the prediction of the choice rates in 4 blocks
	Win 3 with p = .45; 0 otherwise EV = 1.35	Win 2–Tax if notE1; lose 20 if E1 P(E1) = 0.03 EV = 1.34–Tax(.97)	Win 1.5 if notE1 & notE2 lose 20 if E1 or E2 P(E1 or E2)=0.06 EV = 0.21		
0	34%	40%	26%	2.78%	
0.4	37%	36%	27%	2.70%	0.034
0.8	38%	32%	29%	2.73%	