Myopia drives reckless behavior in the face of over-taxation

Mikhail S. Spektor[1,2] and Dirk U. Wulff[3,4]

[1]Universitat Pompeu Fabra

[2]Barcelona Graduate School of Economics

[3]University of Basel

[4]Max Planck Institute for Human Development

2,700 words

Author Note

Mikhail S. Spektor, Department of Economics and Business, Universitat Pompeu Fabra and Barcelona Graduate School of Economics, Barcelona, Spain. Dirk U. Wulff, Faculty of Psychology, University of Basel, Basel, Switzerland, and Max Planck Institute for Human Development, Berlin, Germany.

We thank Laura Wiles for editing the manuscript.

Correspondence: Mikhail S. Spektor, Department of Economics and Business, Universitat Pompeu Fabra and Barcelona Graduate School of Economics, Ramon Trias Fargas, 25–27, 08005 Barcelona, Spain.

Email: mikhail@spektor.ch

Abstract

Governments often use taxes to discourage undesired behaviors or encourage desired ones. One target of such interventions are reckless behaviors such as texting while driving, which in most cases are harmless but sometimes lead to catastrophic outcomes. Past research has demonstrated how interventions can backfire when taxes for specific options are set too high and other reckless behaviors remain untaxed. In the case of experience-based decisions, this undesirable outcome arises from people choosing as if they underweighted rare events. A popular explanation for this as-if underweighting lies in individuals basing their decisions on a small sample of past experiences. Here, we reevaluate the adverse effect of overtaxation using a different theoretical account. We show that a reinforcement-learning model that weights more recently observed outcomes stronger than past ones can provide an equally good account of individuals' behavior. Furthermore, using this model, we show that individuals fall into two classes with qualitatively distinct patterns of behavior. We conclude that targeted interventions tailored at the one third of individuals who act myopically by disregarding catastrophic outcomes quickly after they have been experienced can be more effective than an omnibus intervention based on taxation.

Keywords: safety-enhancement, reliance on small samples, cognitive modeling

Behavior in everyday life comprises "risks", implying that we cannot be certain about the outcomes that ensue our actions. These risks in the world are often structured such that a particular course of action results in small gains almost all of the time but in catastrophic losses in rare occasions that can easily offset any previously accumulated gains. Choosing such courses of action is dangerous, yet people recklessly engage in such behaviors, for instance, by texting while driving or riding a bicycle without a helmet, to name but two. A recent paper (Yakobi et al., 2020, henceforth: YCNE) investigated the effectiveness of monetary incentives (i.e., taxation) as a means to regulate reckless behavior. Specifically, they were interested in situations where moderate taxation of a moderately risky option would lead to the desired effect of swaying people toward a safer option, but excessive taxation could backfire by driving people toward an even more risky, non-taxed option. In such cases, the effect of taxation would result in a U-shaped pattern of reckless behavior.

YCNE studied this potential U-shaped pattern in two experiments using a decisions-from-experience task. In this task, participants made repeated decisions between three initially unknown options, comprising one relatively safe option (e.g., \$3 with $Pr = .45$, otherwise nothing), one moderately risky option that was subject to a tax (\$2-tax with $Pr = .97$, otherwise -\$20), and one inferior, highly risky but non-taxed option (\$1.5 with $Pr = .94$, otherwise -\$20). On each choice, participants would see the outcomes of all three options, allowing them to learn about the underlying properties of the options, but only the outcome of the chosen option impacted the participant's bonus. Varying the level of taxation between three amounts (representing no, moderate, and excessive taxation), a U-shaped pattern emerged in study 1 and in one of the two conditions of study 2. To explain their results, YCNE put forth the "reliance on small samples" hypothesis (Erev & Roth, 2014), according to which individuals draw a small sample of observed outcomes from memory. Reliance on small samples is known to effectively produce (as-if) underweighting of rare events, which can explain people's preference for the risky options as long as they provide the better outcome most of the time and the catastrophic outcome is rare enough.

The research reported by YCNE produces important insights about how policies based on economic incentives can backfire when individuals base their decisions on experience, contributing to a growing body of literature on experienced-based decision making and the description–experience gap (see Wulff et al., 2018, for a recent meta analysis). Moreover, YCNE suggest a cognitive mechanism to explain these results, namely that people base their decisions on small samples of past experiences. However, this need not be the only explanation. In this article, we put forth an alternative explanation of the presented effects that arguably rests on weaker assumptions and allows us to conveniently study the effects of taxation on the level of the individual, the level at which cognitive processes are at work and interact with potential policy interventions.

Models of reckless behavior

To identify which psychological processes best describe people's behavior, YCNE contrasted several different models based on the reliance-on-small-samples hypothesis (Erev & Roth, 2014) with the natural-mean heuristic (called full-data model in YCNE), a model that takes all previous experiences into account and predicts deterministically the choice of the option that so far has yielded the highest average outcome.[1] Quantitatively, the sampling models clearly outperformed the natural-mean heuristic across all conditions (see YCNE Tables 1 and 2). However, qualitatively, the natural-mean heuristic captured the patterns of results rather well (see YCNE Figures 3 and 5), suggesting that models that take into account many or even all samples could be able to accurately describe people's behavior. Consistent with this assessment, the estimated sample-size parameter κ in the sampling-based models obtained by YCNE in their post-hoc comparisons varied between 24 and 47, implying that people drew large samples of prior experiences from memory. Moreover, the overall best-performing model, I-SAW2, is an ensemble model that averages the

[1] For some comparisons, they also included the *accentuation-of-differences model* (Spektor et al., 2019). However, for the focus of the present investigation, this model is not of relevance.

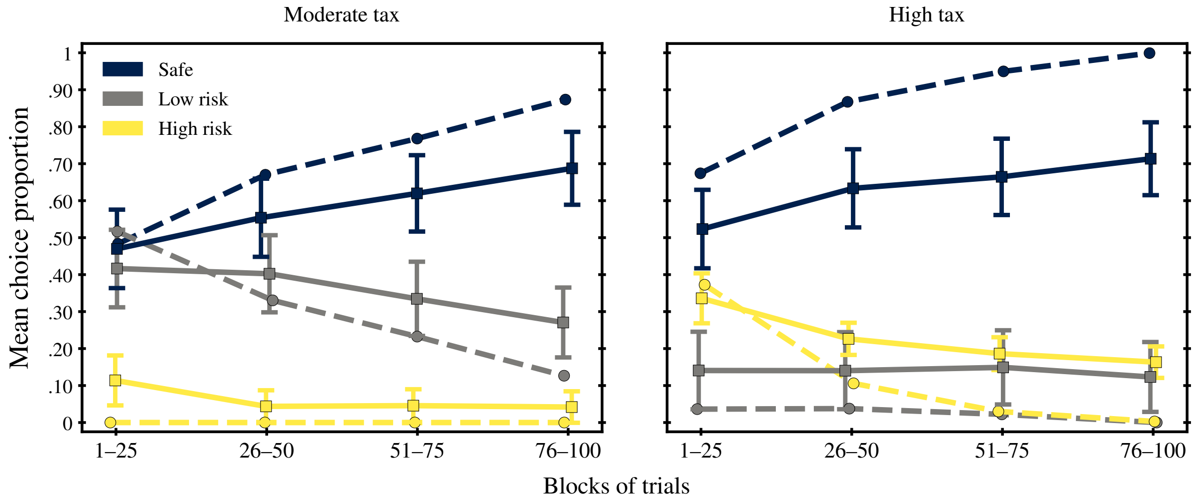

Figure 1

Aggregated choice proportions and predicted choice proportions of the natural-mean heuristic in Experiment 1. Solid lines indicate participants' choices and dashed lines indicate the choice probabilities predicted by the model. Error bars indicate the 95% CI.

predictions of a two-stage sampling model with those of the natural-mean heuristic.

Two behavioral tendencies in the data are likely responsible for the, nonetheless, poor quantitative fit of the natural-mean heuristic: stochasticity of choices and (as-if) underweighting of rare events. Lacking both properties, the natural-mean heuristic predicts that people quickly develop a strong preference for the safe option as the cumulative likelihood of experiencing catastrophic events grows large. However, people's preferences develop more moderately and favor either of the two risky options more strongly than predicted by the natural-mean heuristic (see Figure 1). Sampling models elegantly account for these patterns using a single mechanism: The sampling process introduces stochasticity, which renders choice proportions less extreme, as well as (as-if) underweighting, which accounts for higher-than-expected preference for the risky options under taxation. However, sampling models are neither the only account of the data nor, as we will illustrate next, necessarily the most plausible.

An alternative account of YCNE's data exists in recency-based models as formalized in the framework of reinforcement learning (Sutton & Barto, 1998). Such models also account for probabilistic choice and (as-if) underweighting of rare events. However, they produce these phenomena using a somewhat different psychological mechanism than sampling models. Specifically, reinforcement-learning models assume that people keep track of a long-run reward expectation Q_i of option i that is updated at each time t with incoming reward (or punishment) $R_{t,i}$. If people observe a better-than-expected reward, they adjust Q upward and vice versa. A popular and simple implementation of this mechanism is given by the temporal-difference model (Gershman, 2015):

$$Q_{t+1,i} = (1 - \alpha) \times Q_{t,i} + \alpha \times R_{t-,i}$$

The learning rate α controls the degree to which the expectations are updated. In its most basic form, α is constant over time, which inevitably results in recency; In other words, that recent experiences receive more weight than earlier ones. The extent of recency varies with the value of α. For instance, $\alpha = .10$ implies that an experience ten epochs ago retains about 38% of its original weight, whereas the same experience's weight essentially drops down to zero under $\alpha = .90$. Therefore, α controls the number of experiences that effectively influence choice, and with that, quite analogously to the sampling models, the degree of (as-if) underweighting of rare events. Higher levels of α also introduce more stochasiticity into the choice process, although reinforcement-learning models typically include extra parameters to allow for additional sources of stochasiticity in choice behavior.

We assessed whether the recency-based account can accurately describe the data of YCNE, including people's responses to varying levels of taxation. We fitted the temporal-difference model to the aggregate choice proportions of YCNE's three between-subject conditions. The model's predictions were derived by determining, across all participants, the proportion of trials for which Q_i was highest. A single parameter (α) was used to fit, overall, 14 independent choice proportions (6 from Experiment 1, 4 from each condition from Experiment 2). A learning rate of $\alpha = .15$ yielded the best fit with a resulting mean squared error of .006 (see Figure 2). Most

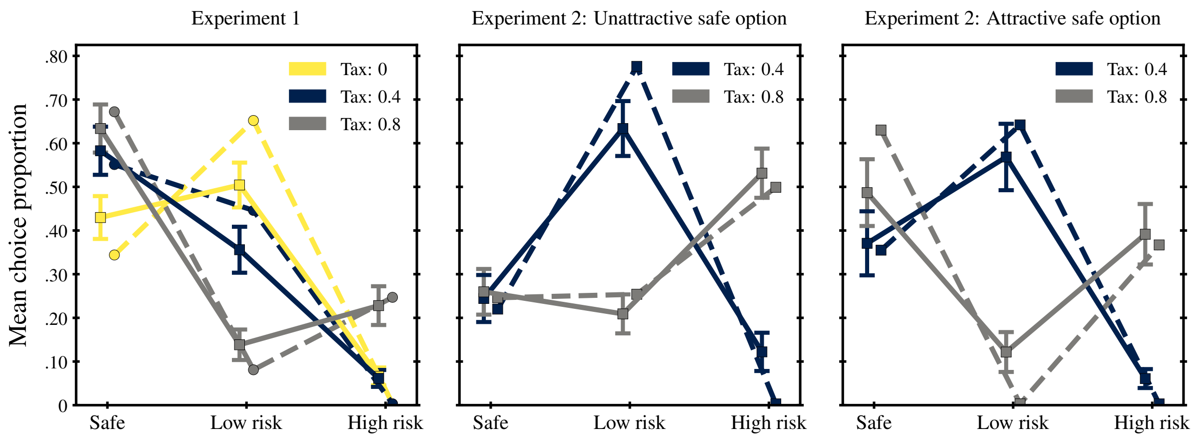

Figure 2

Aggregated choice proportions and predicted choice proportions of the reinforcement-learning model across the different conditions. Solid lines indicate participants' choices and dashed lines indicate the choice probabilities predicted by the model. Error bars indicate the 95% CI.

predictions fell within the 95% confidence interval of the observed choice proportions and the model accurately accounted for the qualitative patterns of taxation. Moreover, when we used the model to predict the data of one experiment on the basis of the respectively other experiment, we observed mean squared errors of .006 (Experiment 1) and .010 (Experiment 2), outperforming all sampling-based models evaluated by YCNE except for the I-SAW2 model, which achieved a slightly better performance in Experiment 2.

The recency-based account formalized in the temporal-difference model predicts the data at least as well as sampling-based accounts. This good performance primarily stems from the model's ability to adjust the number of outcomes that is effectively used, much like the sampling-based accounts. Nonetheless, there exist crucial differences between the two accounts that have implications for both theory and practice. First, the recency-based account can be considered more (cognitively) parsimonious. In contrast to the sampling-based models, it does not require an explicit representation of all past experiences or a process of sampling from memory. Instead, people have only to memorize a single value Q and carry out only a minimal set of operations for each new choice. Second, in contrast to the sampling-based accounts, choices in the recency-based account always reflect all experienced information, even if their influence becomes negligible the further away they are. Third, whereas in sampling-based accounts each experience has equal sway in the long run, the recency-based account predicts that recent outcomes will influence choices more than earlier ones.

While the first two points concern theoretical aspects of recency-based accounts, the third concerns people's behavior and can therefore be evaluated. In the next section, we apply the recency-based model at the individual level to evaluate whether it can account for people's trial-by-trial behavior and to investigate individual differences in the reliance on recent information.

Individual differences in recency and reckless behavior

Aggregate analyses of human behavior always bear the risk of misrepresenting the mechanisms that actually are at work at lower levels of analysis, sometimes leading to drastically wrong conclusions (e.g., Regenwetter & Robinson, 2017; Wulff & van den Bos, 2018). This can be particularly problematic in cases in which the identified mechanism serves as the basis for behavioral interventions. To circumvent such problems and to assess the role of recency, we fitted the temporal-difference model to the individual-level data. To account for both variability in recency and in the stochasticity of people's behavior (Hey & Orme, 1994), the model was additionally equipped with an ϵ-greedy (Sutton & Barto, 1998) choice rule which maps subjective expectations Q to choice probabilities. It predicts a choice of the option with the highest subjective expectation with probability $1 - \epsilon$ and a randomly selected option with the error probability ϵ. Thus, two parameters were fit to 300 trinary choices in Experiment 1 and 200 trinary choices in Experiment 2, maximizing the trial-by-trial choice likelihoods separately for each individual. Across the three between-subject conditions, the temporal-difference model described 89.8% (221 out of 246) of individuals better than chance, determined on the basis of the Bayesian information

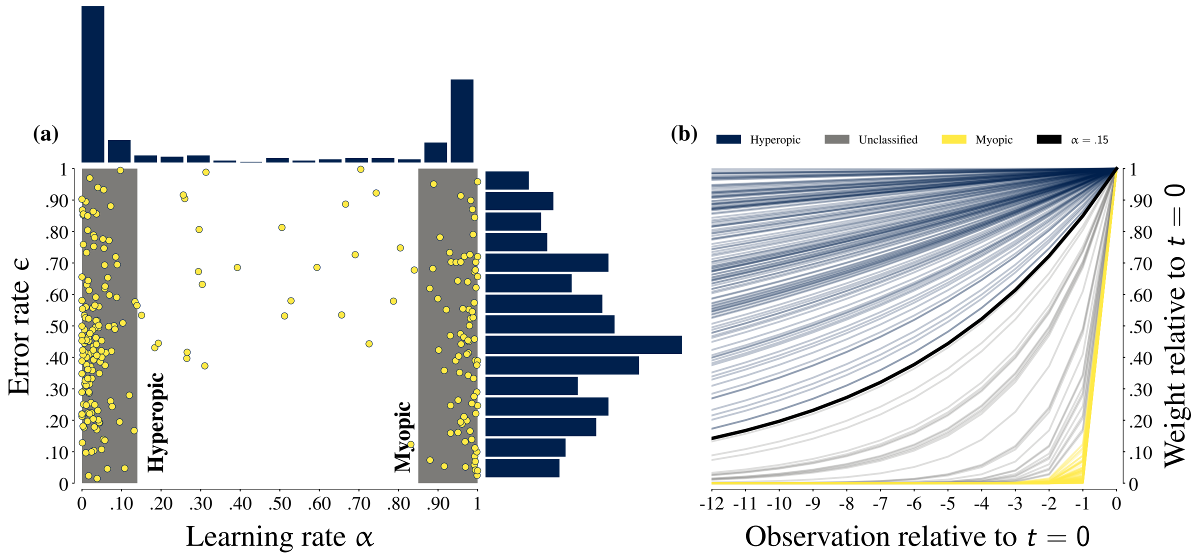

Figure 3

(a) *Distribution of individual-level learning rates α and error rates ϵ across the three between-subject conditions. Parameters were estimated using maximum-likelihood estimation. Shaded areas represent classification according to* $0 \leq \alpha \leq .15$ *(hyperopic) or* $.85 \leq \alpha \leq 1$ *(myopic).* ***(b)*** *Relative weights of past experiences implied by the estimated learning rates. Each line represents one individual and the weight attached to each observation (up to 12 observations into the past).*

criterion (Schwarz, 1978). This suggests that the model provides an adequate account of the individual-level data.

The estimated parameters revealed substantial individual differences: Learning rates followed a bimodal distribution (see Figure 3a), such that a vast majority of people fell into two clearly distinct classes: myopic and hyperopic learners. Myopic learners (29%) are characterized by a high learning rate of $\alpha = [.85, 1]$, implying that only the last one or two observations build the basis of their choices. Hyperopic learners (59%), on the other hand, are characterized by a low learning rate of $\alpha = (0, .15]$, implying that even the most distant experiences are still factored into their choices. The distribution of error rates, by contrast, was clearly unimodal and reflected a maximization rate of 69% ($1 - \epsilon + \epsilon \times \frac{1}{3}$), which is in line with previous research (Harless & Camerer, 1994). Furthermore, error rates did not covary with learning

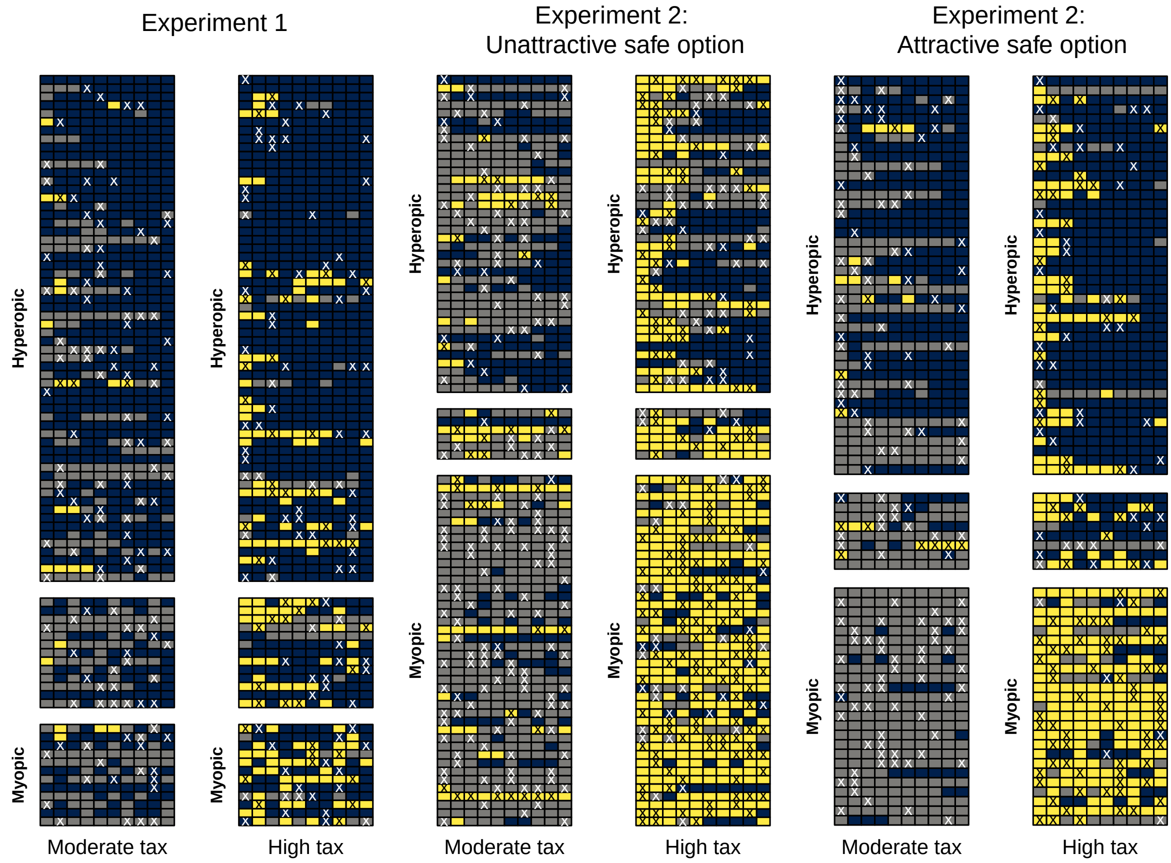

Figure 4

Modal choices in bins of 10 trials for each participant in each experiment, ordered by the learning rate from low (top) to high (bottom). Dark blue represents a modal choice of the safe option, gray represents of the low-risk option, and yellow of the high-risk option. Individuals are grouped according to their classification, hyperopic ($0 \leq \alpha \leq .15$), myopic ($.85 \leq \alpha \leq 1$) or unclassified. Crosses indicate that individuals suffered an accident in the corresponding bin.

rates ($r = .01$), suggesting that the estimated learning rates reflect systematic differences in people's tendency to focus on recent experiences and are not merely the result of identifiability problems known from reinforcement-learning models (Spektor & Kellen, 2018).

To confirm whether the individual differences in learning rates reflect systematic differences in behavior, we plotted the modal choices of all individuals ordered by learning rate separately for all conditions (Figure 4). This revealed that, whereas most

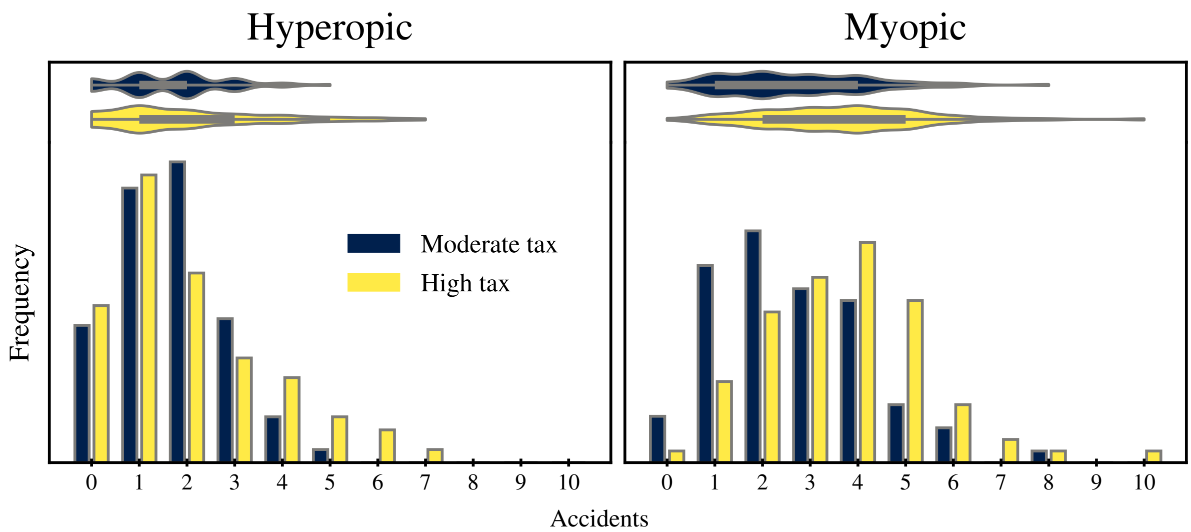

Figure 5

Distribution of experienced accidents, split by taxation amount. Individuals are grouped according to their classification, hyperopic ($0 \leq \alpha \leq .15$) or myopic ($.85 \leq \alpha \leq 1$). Shaded areas in the violin plots indicate the central 50% interval.

hyperopic individuals quickly learned to choose the safe option (dark blue), especially under high taxation, most myopic learners exhibited continued preference for whichever risky options offered the better outcome most of the time (gray = moderate risk, yellow = high risk). These patterns were most pronounced in the condition with an attractive safe option in Experiment 2.

Does the continued preference for risky options mean that myopic individuals might not have learned at all? Using a mixed-effects regression accounting for participant random effects nested within condition, we evaluated whether myopic individuals reacted to the observation of an accident. We found that immediately after the observation of an accident, preferences for the safe option were substantially elevated ($OR = 2.80, p < .001$), but not one ($OR = 0.64, p = .12$) or two ($OR = 1.20, p = .40$) trials later, relative to all other trials. Thus, consistent with the high learning-rates estimates, myopic individuals learned about and reacted to accidents, but then discounted them very quickly as they continued.

Considering only moderate and high taxation situations, the data, furthermore, showed that myopic individuals experienced, on average, 3.15 accidents, whereas hyperopic individuals experienced only 1.85 accidents (see Figure 5). More importantly, compared to moderate taxation, myopic individuals suffered 1.01 accidents more under high taxation, whereas hyperopic individuals suffered only 0.30 more accidents. These analyses suggest that myopic individuals not only suffered considerably more accidents in general, but also that they were much more susceptible to the negative effects of over-taxation.

Discussion

When people make decisions based on experience they tend to choose as if they underweight small-probability events. This phenomenon forms the basis of what has become known as the description–experience gap and it is the key to understanding people's responses to taxation. As-if underweighting of rare events implies that people tend to prefer the option that yields the best outcome most of the time (Wulff et al., 2015). Under excessive taxation of moderately risky behaviors, an even more reckless option can suddenly become the better option (most of the time) and, thus, the preferred choice alternative. The present investigation shows that different mechanisms can provide a good qualitative and quantitative account of how taxation affects behavior in such settings. Moreover, our analyses uncovered substantial differences between individuals that quite likely are of greater import than the question of which mechanism best accounts for people's behavior. Specifically, we uncovered two distinct groups of people, myopic and hyperopic learners, who responded to the experience of accidents in qualitatively different ways. Accounting for these individual differences is crucial for understanding the behavior and for deriving effective policies to prevent accidents from reckless behavior due to over-taxation.

It is a well-known fact that aggregate behavior is at risk of misrepresenting people's actual behavior (Regenwetter & Robinson, 2017; Wulff & van den Bos, 2018) and, in the present investigation, it actually does. The learning rate obtained by fitting the recency model to the aggregate choice proportions of both studies suggests a steady decay of the weight of past experiences, where an experience ten epochs ago receives about 20% of its original weight (see Figure 3b). However, there are almost no individuals who were accurately described by such a weighting scheme. Instead, individuals seem to assign to past outcomes a weight that is either well above that of the aggregate estimate, or one that is essentially zero. These differences not only imply that people effectively base their decision on different experiences, they also suggest that people may have relied on different mechanisms.

Hyperopic individuals, indeed, might have recruited a recency-based mechanism with gradually diminishing weights as formalized in the temporal-difference model, they could, alternatively, have also recruited a stochastic variant of the natural-mean heuristic or a sampling-based model with a large sample size. All three mechanisms should be able to account well for the behavior of hyperopic individuals because in a stable environment, a large sample of both *recent* and of *random* samples will be representative of all observations. The behavior of myopic individuals, on the other hand, can be explained by neither the natural-mean heuristic nor a sampling-based model; only the recency-based account can capture the high weight given to the single most recent outcome.

Even though the recency-based account, as the only model, seems to fit the behavior of both types of individuals, we do not think that it provides a complete account of the psychology of hyperopic and myopic individuals. For instance, recency is often attributed to either memory limitations or adaptations to assumed changes in the environment (see Bornstein et al., 2017; Wulff & Hertwig, 2019; Wulff et al., 2018; Wulff & Pachur, 2016). However, neither of these two represents a compelling account of the two extreme forms of recency observed here, namely practically no recency (hpyeropic) and maximum recency (myopic). Rather, it is likely that other factors not included in the recency-based account, such as risk preferences or goals (see, e.g., Hertwig et al., 2019), also play a role. For example, the differences between hyperopic and myopic individuals can also be construed as the pursuit of short- versus long-run goals (see Lopes, 1981; Wulff et al., 2015) or as maximization versus probability matching strategies (Gaissmaier & Schooler, 2008; van den Bos, 2009). In any case, based on our findings and similar results of previous studies (Spektor et al., 2019; Spektor & Kellen, 2018), strong individual differences seem to be a robust phenomenon that should be studied in future research using larger, more diagnostic data sets.

Finally, the presence of two classes of people relying on potentially different mechanisms has crucial implications for policy development. We have shown that myopic individuals are already at a much greater risk of suffering accidents than hyperopic individuals and that this gap widens under higher levels of taxation. A targeted policy addressing myopic individuals, for instance, by using targeted boosts (Hertwig & Grüne-Yanoff, 2017), might be more effective than an omnibus policy addressing everyone equally. The data suggest that had myopic individuals acted like hyperopic individuals, a total of 205 accidents would have been prevented. In contrast, placing people under a moderate (rather than a high) level of taxation prevented only 123 accidents. Even more accidents (272) would have been prevented by the combination of both; In other words, when everyone was placed under a moderate level of taxation and everyone had acted in a hyperopic fashion. This suggests that the overall best policy to prevent reckless behavior and accidents likely recruits both targeted and omnibus strategies.

References

Bornstein, A. M., Khaw, M. W., Shohamy, D., & Daw, N. D. (2017). Reminders of past choices bias decisions for reward in humans. *Nature Communications*, *8*(1), 15958. https://doi.org/10.1038/ncomms15958

Erev, I., & Roth, A. E. (2014). Maximization, learning, and economic behavior. *Proceedings of the National Academy of Sciences of the United States of America*, *111*, 10818–10825. https://doi.org/10.1073/pnas.1402846111

Gaissmaier, W., & Schooler, L. J. (2008). The smart potential behind probability matching. *Cognition*, *109*(3), 416–422. https://doi.org/10.1016/j.cognition.2008.09.007

Gershman, S. J. (2015). Do learning rates adapt to the distribution of rewards? *Psychonomic Bulletin & Review*, *22*(5), 1320–1327. https://doi.org/10.3758/s13423-014-0790-3

Harless, D. W., & Camerer, C. F. (1994). The predictive utility of generalized expected utility theories. *Econometrica*, *62*(6), 1251. https://doi.org/10.2307/2951749

Hertwig, R., & Grüne-Yanoff, T. (2017). Nudging and boosting: Steering or empowering good decisions. *Perspectives on Psychological Science*, *12*(6), 973–986. https://doi.org/10.1177/1745691617702496

Hertwig, R., Wulff, D. U., & Mata, R. (2019). Three gaps and what they may mean for risk preference. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *374*(1766), 20180140. https://doi.org/10.1098/rstb.2018.0140

Hey, J. D., & Orme, C. (1994). Investigating generalizations of expected utility theory using experimental data. *Econometrica*, *62*(6), 1291. https://doi.org/10.2307/2951750

Lopes, L. L. (1981). Decision making in the short run. *Journal of Experimental Psychology: Human Learning & Memory*, *7*(5), 377–385. https://doi.org/10.1037/0278-7393.7.5.377

Regenwetter, M., & Robinson, M. M. (2017). The construct–behavior gap in behavioral decision research: A challenge beyond replicability. *Psychological Review*, *124*(5), 533–550. https://doi.org/10.1037/rev0000067

Schwarz, G. (1978). Estimating the dimension of a model. *The Annals of Statistics*, *6*(2), 461–464. https://doi.org/10.1214/aos/1176344136

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. https://doi.org/10.1037/rev0000122

Spektor, M. S., & Kellen, D. (2018). The relative merit of empirical priors in non-identifiable and sloppy models: Applications to models of learning and decision-making. *Psychonomic Bulletin & Review*, *25*(6), 2047–2068. https://doi.org/10.3758/s13423-018-1446-5

Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction*. Cambridge, MA, MIT Press.

van den Bos, W. (2009). Better than expected or as bad as you thought? The neurocognitive development of probabilistic feedback processing. *Frontiers in Human Neuroscience*, *3*. https://doi.org/10.3389/neuro.09.052.2009

Wulff, D. U., & Hertwig, R. (2019). Uncovering the anatomy of search without technology. In A. Schulte-Mecklenbeck, A. Kuehberger, & J. G. Johnson (Eds.), *A handbook of process tracing methods (2. ed.)* Psychology Press.

Wulff, D. U., Hills, T. T., & Hertwig, R. (2015). How short- and long-run aspirations impact search and choice in decisions from experience. *Cognition*, *144*, 29–37. https://doi.org/10.1016/j.cognition.2015.07.006

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. https://doi.org/10.1037/bul0000115

Wulff, D. U., & Pachur, T. (2016). Modeling valuations from experience: A comment on Ashby and Rakow (2014). *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *42*(1), 158–166. https://doi.org/10.1037/xlm0000165

Wulff, D. U., & van den Bos, W. (2018). Modeling choices in delay discounting. *Psychological Science*. https://doi.org/10.1177/0956797616664342

Yakobi, O., Cohen, D., Naveh, E., & Erev, I. (2020). Reliance on small samples and the value of taxing reckless behaviors. *Judgment and Decision Making*, *15*(2), 266–281.