The role of feedback in risk-sensitivity

Ohad Dan^{1,2} and Yonatan Loewenstein^{1,2,3,4}

¹ Department of Cognitive Sciences, The Hebrew University, Jerusalem, Israel; ² The Federmann Center for the Study of Rationality, The Hebrew University, Jerusalem, Israel; ³ The Alexander Silberman Institute of Life Sciences, The Hebrew University, Jerusalem, Israel; ⁴ The Edmond and Lily Safra Center for Brain Sciences, The Hebrew University, Jerusalem, Israel

Abstract

Over the last century, myriad versions of the bandit task were used to study operant conditioning in humans and other animals. However, the overwhelming majority of these variations utilized one of two types of feedbacks, partial and full feedback, revealing to participants the single outcome of the chosen alternative or the outcomes of all alternatives respectively. While ecologically relevant, when restricting the feedback method to these two methods alone, observed behavioral phenomena could potentially be confounded with specific effects that the feedback method itself might induce, for example attitude towards risk. Here we introduce a new form of feedback. In a 2-armed bandit task, the reverse feedback reveals to participants only the outcome of the unchosen alternative. In a behavioral experiment, human participants were incentivized to maximize their per-trial reward while exploring the reward-distribution associated with two alternatives. Randomly assigning participants to a specific type of feedback, we find that participants in the partial and reverse feedback condition demonstrated behavior consistent with risk aversion and risk seeking. This result is intriguing for two reasons. First, in gains-domain, humans are considered risk-averse and it is hence surprising to observe a robust demonstration of risk seeking. Second, We present risk-sensitivity as a casual outcome of the utilized feedback. Since in most ecological and lab environments humans utilize the partial feedback and demonstrate risk-aversion, our finding sheds new light on the common perception of risk-sensitivity as an inherent, rather than induced, characteristic. Utilizing a simple reinforcement learning model, we explain the emergent risk preference as an outcome of learning in the specific environment we use. We present the relation of our paradigm to prospect theory, relate our finding to existing literature, and discuss the new light our novel feedback-mechanisms shed on conclusions drawn from previous paradigms.

Keywords: Decision-making, learning from experience, risk-attitude

Introduction

Consider a restaurant manager whose task is to choose one of two chefs to cook at her restaurant daily. It is her first day on the job and she has no prior knowledge of the quality of the chefs' performance. The success or failure of her chosen chef's daily delights is directly reflected in the sum of a day's tips but, if her job wasn't difficult enough, the owner's policy is to aggregate these tips and count them after long periods of time. So, while she aims to optimize the daily tips, she cannot immediately utilize the information encapsulated in them to improve her decision making. Nevertheless, her daily decisions do yield a reliable signal she may exploit. The other chef, the one she did not choose for the day, does not linger idly but rather heads to the single other restaurant in town and cooks there for that day. The owner of the other restaurant does, fortunately, count the tips daily and does not mind sharing the other chef's daily outcomes. So, each day, while she cannot learn how well did the chef she did choose do, the manager does learn about the performance of the chef she did not choose.

Should the manager complain to the restaurant's owner that she cannot efficiently optimize the restaurant's earnings given these conditions? Learning on the consequences of her actions in this way is arguably unconventional, but is it any worse compared to more conventional forms of feedback? Certainly, learning each day how well did both chefs do is advantageous as it provides more information. But, if the only options available to the manager are learning on a single chef, either the one she chose or the one she did not choose, should she have any preference for either of the feedback-methods? In terms of the amount of information, the two methods are identical, namely both procedures provide a single sample from the two tip-distributions associated with the two chefs per decision. Setting aside the preferences of the manager, should the owner of the restaurant expect any difference in the overall behavior of the manager given the symmetry between these two conditions? Naively, not. For example, the owner may expect the manager to demonstrate the well-documented tendency to be risk-averse and, *ceteris paribus*, prefer the chef whose tip-distribution has smaller variance (Kahneman & Tversky, 1979; Tversky & Kahneman, 1992). However, surprisingly, this is not the case.

For more than a century, the theoretical framework of operant-conditioning has been used to study the decision-making and learning processes of humans and other animals (Mongillo et al., 2014; Shteingart & Loewenstein, 2014; Thorndike, 1927). Crucial to learning is the

concept of feedback, namely informing the decision-maker on some aspect of the outcomes of its previous actions. One of the most prevalent experimental paradigms utilized to study decision making and learning is the k -armed bandit (Bergemann & Valimaki, 2006). In this task, k alternatives are presented to a decision-maker that in each trial is asked to choose one. In response to its choice, some feedback on the outcomes of its choice is provided. Traditionally, feedback is provided in one of two ways. Partial feedback entails providing information solely on the outcome associated with the chosen alternative (see Fig 1.a). Full feedback entails providing information both on the outcome granted by the chosen alternative and all other outcomes, the forgone payoffs associated with the unchosen alternatives (see Fig 1.a). The methodological choice in these methods of feedback is not arbitrary and could be justified by the observation that the full and partial feedback represent the vast majority of ecological-feedback settings. This may be best exemplified by how unnatural is the scenario presented above, describing the restaurant-manager that should learn on the quality of her available actions (choices between chefs) by observing the outcome of her unchosen action. Yet, the conclusions of studies that utilize “standard” feedback methods are not formalized as being feedback specific. For example, behavioral models aiming to characterize the algorithms human employ in learning from feedback are often formulated in a generalized form that does not highlight the feedback method. Due to the generalized form of behavioral models, applying any form of feedback to a model is a trivial modification. Here we test whether trivial changes in the feedback method also yield trivial changes in behavior. To this end, we test both in simulations and in behavior the outcomes of utilizing a novel type of feedback.

Here we introduce the reverse-feedback, the chiral counterpart of the partial feedback. As illustrated by the chefs example, the reverse-feedback entails learning only the forgone payoff (i.e. and not on the outcome of the chosen action, Fig 1.a). Given the unintuitive nature of this feedback we first ask whether human participants can effectively utilize this feedback. Second, we test how does the reverse feedback shape participants’ attitude towards risk. Risk may be defined as a measure of reward uncertainty (Jaeger et al., 2013; Tversky & Fox, 1995). Hence, attitude towards risk is the subjective perception and valuation of uncertainty (Arrow, 1965; Hillson & Murray-Webster, 2007). Broadly categorized, risk preference may be classified into 3 complementary attitudes which determine how subjective utility is affected

by uncertainty: subjective utility that is elevated by uncertainty is termed risk-seeking, decreased subjective utility induced by uncertainty is termed risk-aversion and, uncertainty-indifference is termed risk-neutrality (Hillson & Murray-Webster, 2007; Kahneman & Tversky, 1979). Importantly, attitudes towards risk depend on the sign of the uncertain prospects, namely whether rewards are manifested as gains or losses (Kahneman & Tversky, 1979; Tversky & Kahneman, 1992). For conciseness, we focus here on gains (but see Discussion).

Standard economic theory and empirical findings assert that humans are in general risk-averse (O'donoghue & Somerville, 2018). This assertion explains why humans are willing to pay a risk-premium and forgo substantial value to diminish uncertainty (Arnott & Bernstein, 2002; Mehra & Prescott, 1985; see however caveats of this approach Rabin & Thaler, 2001). While some previous studies attributed risk-aversion to specific features of learning processes (e.g. reliance on small samples and diminishing sensitivity to absolute payoffs, Erev et al., 2008) the majority of the literature describes risk-aversion, either implicitly (Fairlie & Holleran, 2012; Rosen et al., 2003; Weber et al., 2002) or explicitly (Borghans et al., 2008; Yechiam & Ert, 2011), as a measurable characteristic (Blais & Weber, 2006; Harrison et al., 2008; Pennings & Smidts, 2000) that to a large extent is stable within individual over time, namely a trait. Further supporting this stance of trait-like characteristic for attitude towards risk are studies which identify genetic (Zhong et al., 2009) and neurobiological (De Martino et al., 2010; Hsu et al., 2005; Levy et al., 2010; Smith et al., 2002; Trepel et al., 2005) basis for individual risk-attitude. Utilizing the reverse-feedback in simulations and behavioral experiments, here we aim to test whether attitude towards risk must be considered as an inherent characteristic or whether it may also result from a process of learning.

Results

To test what attitude towards-risk is expected under a reverse-feedback regime, we simulated behavior using a standard reinforcement learning model (see Methods). Briefly, the model makes stochastic choices based on a (Q-) value associated with each alternative, values which represent the subjective estimated valuation of an alternative and are updated with every new observation. Importantly, both in the decisions and learning processes, the model does not consider any measure of risk nor does it track the variance associated with the alternatives. Explicitly having no attitude towards risk a model may nevertheless demonstrate behavior that is consistent with risk-seeking or risk-aversion (see Biele et al., 2009; Erev & Barron, 2005).

A common way to formulate a choice problem between alternatives associated with different risk is to consider a repeated 2-alternative forced-choice task between a “risky” and “safe” alternatives (see Erev et al., 2010), where the safe alternative provides a constant payoff r_{safe} , and the risky alternative provides a high payoff r_{high} in probability p ($0 \leq p \leq 1$) and in probability $1 - p$ a low payoff r_{low} ($r_{high} \geq r_{low}$). In this framework, the relative number of choices in the two alternatives is considered as the empirical risk-preference, where more choices in the risky (safe) alternative are considered as a more risk-seeking (aversion) attitude. In this paradigm, if one of the alternatives is associated with higher expected value, choice patterns are less easily interpretable and hence, for simplicity, we consider a case where the expected value of the alternatives is equal. Namely, we choose parameter values which satisfy the equation:

$$p * r_{high} + (1 - p) * r_{low} = r_{safe}$$

Testing risk preference in simulations using these conditions subjected to different feedback methods reveals two main trends (Fig, 1.b). First, the proportion of choices in the risky alternative (greater risk-seeking) increases with p , i.e. increases as the higher reward becomes more common (see Discussion). A second trend is that for any given p the proportion of choices in the risky alternative is lowest in the partial-feedback, greater in full-feedback, and greatest with reverse-feedback (see Discussion).

Fig 1. Simulated risk attitude in the reverse, partial and full feedbacks. **a. Different feedback types.** In a two-armed bandit task, participants choose between two alternatives. A choice in alternative A yields a reward r_A according to a mechanism unknown to the participant (here, a stationary reward distribution associated with each alternative). In the reverse feedback, choice in A yields a reward drawn from A 's reward-distribution but this reward remains unknown to the participant. In contrast, the forgone reward r_B , the one that could have been gained had alternative B been chosen, is revealed but its reward not granted to the participant. Reverse feedback (white squares) is contrasted with (darken squares) the partial feedback in which a choice in A reveals the reward associated with A (r_A), and full feedback in which a choice in A reveals the rewards associated with A and B (r_A and r_B). **b. Simulated risk attitude in the different feedbacks.** In the simulations, one alternative yielded a constant reward (r_{safe}) and the other high and low rewards in probability p and $1-p$ respectively, such that the two alternatives had equal expected values. Each line

represents the proportion of choices in the risky alternative out of the 100 simulated trials for a specific feedback (error bars are standard error of the mean over 100 repetitions). The slope of the lines crucially depends on the simulated learning rate (here $\eta=0.4$, other values not shown) such that greater learning rates lead to steeper slopes.

To test how humans actually behave under the reverse-feedback regime, we conducted a series of experiments comparing participants' preferences in partial and reverse feedbacks. In the experiments, participants repeatedly chose between two alternatives which provided rewards according to specific, stationary reward distributions (Methods). Participants were guided to maximize their reward in each trial but were not instructed on any feature of the reward distributions. In the first set of experiments, we verified that participants were able to effectively learn and utilize the reverse feedback. To this end, one of the alternatives was configured with a higher expected value compared to the other. Importantly, to avoid trivial identification of the alternatives as being better or worse (e.g. if the "better" alternative provided values greater than 100 and the "worse" alternative lower than 10), the range of reward-values was identical for the better and worse alternatives (see Methods). This paradigm was tested in two independent experiments which utilized different reward distributions: in the first, each reward distribution yielded just one of two values, high or low (Fig 2, left), and in the second, rewards were drawn from a bi-modal distribution with high and low modes (Fig 2, right; see Methods). In both experiments, participants robustly demonstrated an ability to utilize the reverse feedback as was evident by their successful identification and maintained exploitation of the better alternative. In both the first and second experiments, the 37 participants of each experiment (different groups) have shown a steady increase in the proportion of choices in the better alternative (Fig 2). When considering the latter half of the experiment, post an initial learning phase, participants chose the better alternative in 67% and 72% of the trials in the first and second experiments, significantly more than they chose the worse alternative ($z=4.41$, $p<.001$ and $z=3.47$, $p<.001$ respectively). We repeated these experiments using the more standard partial feedback and found that the ratio of choices in the better alternative was comparable to those demonstrated in the reverse feedback paradigm, 71% and 78% (38 and 37 participants, $z=7.31$, $p<.001$ and $z=5.15$, $p<.001$ respectively, a different group of participants, Fig 2, see Methods). Confirming that

participants can effectively utilize the reverse feedback, we next tested the emerging risk-preference induced by the different feedback types.

Fig. 2. Participants can efficiently utilize the reverse feedback. *Top panels* present the reward distributions associated with the two alternatives in each experiment. These distributions were of two types: binary, on the left, where each alternative yielded just one of two values, and on the right, distributions which yielded a range of values (see Methods). *Middle panels*, proportion of choices in the better alternative per each trial. Shaded area is standard error of the mean, over participants per each trial. Black dashed line is chance level (0.5). *Bottom panels*, overall proportion of choices in the better alternative in the last half of the experiment (after an initial learning period). Error bars are standard error of the mean over the choice proportions of individual participants.

The following set of experiments was identical to the previous, apart from the reward distributions associated with the two alternatives. If previously one of the alternatives was better, now the alternatives had identical expected values, but they differed in the amount of their reward uncertainty. Specifically, the reward distribution of one of the alternatives, the “risky” one, was associated with a greater variance compared to the other alternative, the “safer” one (see Methods). We found that unlike the previous “better-worse” experiments where participants demonstrated similar preferences in the two feedback methods, here participants in the partial-feedback condition demonstrated behavior consistent with risk-aversion while participants in the reverse-feedback condition demonstrated behavior consistent with risk-seeking. We reproduced this result in two different experiments that utilized different reward distributions (different participants in each experiment). In the first experiment, the reward distribution was defined by two values: “low” and “high” for the safe alternative, and “lower” and “higher” for the risky alternative (Fig 3 left, Methods). In the second experiment, the safe alternative was associated with a “narrower” reward distribution and the risky alternative with a “wider” one (Fig 3 right, Methods). In the first experiment, 38 participants in the partial- and 40 in the reverse-feedback conditions chose the risky alternative 40.0% and 55.0% of the trials; In the second experiment, 40 participants in the partial- and 37 in the reverse-feedback conditions chose the risky alternative 43.7% and 53.4% of the trials (Fig 3, middle panels). To compare participants’ risk-preference after an initial learning phase, we calculated the proportion of choices in the risky alternative in the second half of the experiments. We found that participants in the reverse feedback chose the risky alternative significantly more than those of the partial feedback both in the first, $t(78)= 3.56$, $p<.001$, and second experiments, $t(77)=2.04$, $p=0.045$ (Fig. 3, bottom panels).

Fig. 3. In comparison to the partial feedback, participants in the reverse feedback are risk seekers. *Top panels* present the reward distributions associated with the two alternatives in each experiment. These distributions were of two types: binary, on the left, where each alternative yielded just one of two values, and on the right, distributions which yielded a range of values (see Methods). *Middle panels*, proportion of choices in the riskier alternative per each trial. Shaded area is standard error of the mean, over participants per each trial. Black dashed line is chance level (0.5). *Bottom panels*, overall proportion of choices in the riskier alternative in the last half of the trials (after an initial learning period). Error bars are standard error of the mean over the choice proportions of individual participants.

Discussion

Risk-aversion is so prevalent in human behavior that it is considered a general human trait, and deviations from it are considered anomalous or mistakes (Bowman, 1982; Ida, 2010; Zuckerman, 2007). Here, in a setting of learning from experience, we utilized a novel mechanism, the reverse feedback, for reporting to participants some of the outcomes of their choices, the reward they could have gotten had they chosen the other alternative. Naively surprising, we show that using this feedback humans demonstrate behaviors consistent with risk-seeking. Using simple computational models, we show how this pattern is, in fact, an expected outcome of learning.

In simulations, we show that under the specific experimental conditions we utilized, risk-aversion is an emergent property of the reverse-feedback. First, we showed that given two alternatives with equal expected value, a safe one that yields constants payoff and a risky one which yields a higher reward in probability p , preference for the risky alternative increases with p . To elucidate this, consider a problem where p is very small. While the two alternatives are indeed equally valuable in their expectation, in the vast majority of the trials, when sampled, the risky alternative will yield r_{low} which is inferior to the guaranteed payoff r_{safe} of the safe alternative. Hence, in most trials the safe alternative will be perceived as superior. Conversely, if p is very high (approaches 1) then in the majority of choices the risky alternative would yield r_{high} , which by construction is greater than r_{safe} , and hence the risky alternative would seem advantageous. A second pattern observed in simulations is that for a given p the proportion of choices in the risky alternative is greater in the reverse-feedback condition, compared to that of the partial-feedback. Given that the problem parameters are constant across feedback types, why does this pattern emerge? The difference in overall behavior may be interpreted as stemming from which “beliefs” are reinforced by either of the feedbacks. In both cases, a learning model prescribes a higher choice-probability to the alternative which presently seems better (has a greater Q-value). In the partial feedback, if the risky alternative presently seems better it will be chosen, until by chance r_{low} will be drawn. This will alter the valuation of the risky alternative to be considered as less valuable, which will decrease the probability of choice in it. Then, if current evidence is such that the safe alternative is better, future choices in the safe alternative will consistently yield r_{safe} , providing evidence consistent with this belief and further reinforcing it. Contrarywise, in the reverse feedback, if

the risky alternative presently seems better it will more likely be chosen thus revealing the forgone payoff of the safe alternative which will consistently still seem inferior, hence reinforcing the continued choices in the risky alternative. This preference as an outcome of learning may explain the risk-preference we identified when comparing behavior in the partial and reverse feedback.

These results are surprising in several ways. First, we show that minor changes in the decision settings yield substantial changes in human behavior. The changes we introduce in the feedback mechanism are on the one hand transparent from an algorithmic perspective, but they lead to atypical behavioral patterns on the other. Future studies should address this susceptibility, both in the way experimental methods are designed and in the ways conclusions formulated, stressing the dependence of outcome on the learning process in general and the feedback method in particular. Second, in a repeated choice context, our results shed new light on current definitions and measurement methodologies of attitude towards risk. As we demonstrate, given the unique but not inconceivable appropriate conditions, risk-seeking must not be an anomaly as it trivially results from standard theory of learning.

Does this mean we believe people will start seeking risk in the ecological scenarios? Probably not. In addition to the vast literature on humans' risk-aversion, evolutionary theories suggest mechanisms that should have reinforced risk-aversion in adult animals. Rather, our results offer a new viewpoint on some basic perceptions of human cognitive constructs. According to the common view humans are risk averse. When we observed risk-seeking behavior in our experiments, we did not conclude that participants developed an inherent preference for risk. Instead, we interpreted this as a mere outcome of the unconventional feedback utilized. However, the same argument may be made for claiming that humans are not in fact inherently risk-averse but rather, similarly, that risk-aversion is a mere outcome of a specific form of feedback. In many experimental settings and real-world scenarios, decision-makers utilize partial feedback, having access solely to the outcomes of the alternatives they end up choosing. This is often the case, either because the information of the chosen alternative is more readily available or, more profoundly, because the forgone payoffs are unknowable (e.g. how happy one would have been had one chosen a different life-partner, career-path, etc.). In such scenarios, as we show in simulations (see also Hertwig et al., 2004; Hertwig & Erev,

2009), the resulting behavior of a learning algorithm is an effective preference for alternatives with lower variance, namely risk-aversion. And so, even an agnostic prediction of the prevailing risk preference of humans would be that they are risk-averse, simply as an outcome of feedback and learning. Importantly, this kind of argument is limited to scenarios of learning from experience. Decisions from description, where rewards and probabilities are known prior to decisions (Kahneman & Tversky, 1979), do not incorporate learning and hence are not susceptible to such influences. However, such scenarios are, arguably, less frequent in the real world (Hertwig & Erev, 2009).

A different conceptual framework for interpreting our results is that of reference-dependent risk attitude, or “gain-loss utility” (Kőszegi & Rabin, 2006, 2007). One of the classical findings of prospect theory (1979) is that, while in gain domains decision makers are risk averse, in the loss domain they are risk seeking. However, gains and losses are relative terms and it is not clear what to relate them to. The more recent theory of reference-dependent utility states that gains and losses of an outcome are evaluated relative to a reference point, and that this reference point is determined by “recent beliefs” about that outcome (Kőszegi & Rabin, 2007). In this perspective, an alternative interpretation of our results is that in the reverse feedback participants act as risk seekers since their perception of the environment is framed in a loss domain. For participants, the unearned rewards reported in the reverse feedback may be anchored as a reference point relative, and lacking information on the earned reward, the unearned ones could be perceived as losses. Once in the loss domain, participants risk aversion is expected to yield the observed risk seeking behavior (Rabin & Thaler, 2001). In this perspective, our experimental framework provides a novel environment for testing unintuitive predictions of prospect theory and reference-dependent utility, and the empirical evidence we observe in our experiments is in line with the predictions of these theories.

As evident from the unnatural example we utilized to illustrate the reverse feedback, encountering it in ecological settings is not likely. Hence, there is a good reason to continue using the partial and full feedback in future experimental paradigms. However, an experimental arsenal which in a setting of learning only utilizes these conventional forms of feedback might be limited in its ability to disentangle innate characteristics from the influences induced by the environment itself. As an experimental tool, the novel method for feedback-delivery we utilized provide a useful mechanism to dissect these components. More

broadly, we make the trivial claim that behavior is critically shaped by feedback but demonstrate the nontrivial extent to which this claim may be utilized. Namely, that risk-attitude, which is often perceived and discussed as an innate trait, may in fact be induced in decision maker by their environment, and specifically as a function of their learning feedback mechanism.

Methods

In all behavioral experiments, participants sequentially chose between two alternatives, choices that were followed by condition-specific information. We utilized a between-subject design whereby each participant was randomly assigned to receive a single feedback, either partial or reverse. The nature of the feedback utilized in participants' assigned condition was explained to the participants without contrasting it with the other type of feedback. Participants were instructed to maximize their per-trial reward and were incentivized to do so by choosing the outcome of one random trial at the end of the experiment as an additional bonus payment. We included in the analysis only and all of the participants who completed their assigned task, choices in 80 consecutive trials.

The experiments were run on the online MTurk platform. In all experiments, following the instructions, a screen with several alternatives represented by buttons was displayed. A choice, clicking one of the buttons, disabled all buttons for 1.5 seconds and revealed a condition-specific reward. To equalize the motoric cost of all choices, progressing to the next trial required pressing a reset button located equidistantly relative to all other alternatives. In all conditions, each trial provided the single reward associated with the chosen alternative. As entailed by the specific allocated feedback type, information was revealed to participants following each choice. In the partial feedback condition, information was provided on the reward of the chosen alternative and in the reverse feedback condition information was provided on the reward of the unchosen alternative.

The rewards revealed to participants were drawn from condition-specific reward distribution. To verify that participants can successfully utilize the reverse feedback a better/worse paradigm was utilized, according to which the expected value of one alternative was greater than the other. This paradigm was employed in two version in two different experiments. In the first, choice yielded rewards (r) with just one of two values (2 or 5):

$$\text{worse} := \begin{cases} P(r = 2) = 0.8 \\ P(r = 5) = 0.2 \end{cases}, \text{better} := \begin{cases} P(r = 2) = 0.2 \\ P(r = 5) = 0.8 \end{cases}$$

In the second version the reward distributions were the integer-rounded (square brackets below) draws from each alternative bi-modal normal distribution:

$$\text{worse} := \begin{cases} P(r \sim [N(10, 1)]) = 0.8 \\ P(r \sim [N(20, 1)]) = 0.2 \end{cases}, \text{better} := \begin{cases} P(r \sim [N(10, 1)]) = 0.2 \\ P(r \sim [N(20, 1)]) = 0.8 \end{cases}$$

To test for risk-preference a safer/riskier paradigm was utilized according to which both alternatives had equal expected values, but one was associated with greater variability. This paradigm was utilized in two different experiments. In the first, choices yielded “high” and “low” rewards (r) in the safe alternative and “higher” and “lower” rewards in the risky alternative:

$$\text{safe} := \begin{cases} P(r = 9) = 0.5 \\ P(r = 11) = 0.5 \end{cases}, \text{risky} := \begin{cases} P(r = 5) = 0.5 \\ P(r = 15) = 0.5 \end{cases}$$

In the second experiment, rewards were sampled from a discrete uniform distribution that was “narrow” for the safe alternative and “wide” for the risky:

$$\text{safe} := r \sim [U\{28, 32\}], \text{risky} := r \sim [U\{10, 50\}]$$

To simulate behavior we utilized a Q-learning model (Sutton & Barto, 2018), a reinforcement learning algorithm that has been widely used to model sequential decision-making behavior in humans and animals (Barto et al., 1991; Mongillo et al., 2014). Applied to our tasks (formally a one-state MDP), Q-learning describes how a Q-value associated with each action, its expected average reward, is updated given new observations. The initial value of an action is set by the first observation it yields and, inspired by experimental findings (Shteingart et al., 2013), the agent starts its sequence of choices by sampling each alternative once in random order (a form of optimism, Szita & Lőrincz, 2008). Following initialization, at each subsequent trial t , upon its choice the agent is informed on the trial’s outcome. Namely, some rewards $r_t(a)$ associated with the action a are revealed and the agent utilizes this information to update the Q-values of specific actions. The updated value $Q_{t+1}(a)$ of a particular action a , is a weighted average of the previous value $Q_t(a)$ and the reward r_t associated with that action in that trial:

$$Q_{t+1}(a) = Q_t(a) + \eta(r_t(a) - Q_t(a))$$

where $0 < \eta \leq 1$ is the learning rate. The values of actions that did not receive updated information in a trial remain unchanged. Given the alternatives' present values, the agent chooses one of the actions using an action-selection rule. Motivated by experimental findings (Shteingart et al., 2013), we utilize an ϵ -softmax action-selection rule, which assigns a choice probability ($\Pr[a_i]$) for an action i given all the actions' (a) values:

$$\Pr[a_i] = (1 - 2\epsilon) \frac{e^{\beta Q_t(a_i)}}{\sum_a e^{\beta Q_t(a)}} + \epsilon$$

Where $\beta > 0$ and $0 \leq \epsilon \leq 1$ are exploration parameters. Where not otherwise specified, we utilized in simulations the parameter values $\eta=0.4$, $\beta=52$, and $\epsilon=0.2$, as extracted from behavior in similar experimental conditions by Shteingart et al., 2013.

References

- Arnott, R. D., & Bernstein, P. L. (2002). What Risk Premium Is “Normal”? *Financial Analysts Journal*, 58(2), 64–85.
- Arrow, K. J. (1965). *Aspects of the theory of risk-bearing*. Yrjö Jahnssonin Säätiö.
- Barto, A. G., Sutton, R. S., & Watkins, C. J. C. H. (1991). Learning and Sequential Decision Making. In *Learning and Computational Neuroscience* (pp. 539–602).
- Bergemann, D., & Valimaki, J. (2006). *Bandit problems*.
- Biele, G., Erev, I., & Ert, E. (2009). Learning, risk attitude and hot stoves in restless bandit problems. *Journal of Mathematical Psychology*, 53(3), 155–167.
- Blais, A.-R., & Weber, E. U. (2006). A domain-specific risk-taking (DOSPERT) scale for adult populations. *Judgment and Decision Making*, 1(1).
- Borghans, L., Duckworth, A. L., Heckman, J. J., & ter Weel, B. (2008). The economics and psychology of personality traits. *Journal of Human Resources*, 43(4), 972–1059.
- Bowman, E. H. (1982). Risk seeking by troubled firms. *Sloan Management Review (Pre-1986)*, 23(4), 33.
- De Martino, B., Camerer, C. F., & Adolphs, R. (2010). Amygdala damage eliminates monetary loss aversion. *Proceedings of the National Academy of Sciences*, 107(8), 3788–3792.
- Erev, I., & Barron, G. (2005). On Adaptation, Maximization, and Reinforcement Learning Among Cognitive Strategies. *Psychological Review*, 112(4), 912–931.
- 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., Ert, E., & Yechiam, E. (2008). Loss aversion, diminishing sensitivity, and the effect of experience on repeated decisions. *Journal of Behavioral Decision Making*, 21(5), 575–597.
- Fairlie, R. W., & Holleran, W. (2012). Entrepreneurship training, risk aversion and other personality traits: Evidence from a random experiment. *Journal of Economic Psychology*,

33(2), 366–378.

Harrison, G. W., Rutström, E. E., & others. (2008). Risk aversion in the laboratory. *Research in Experimental Economics*, 12(8), 41–196.

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.

Hertwig, R., & Erev, I. (2009). The description-experience gap in risky choice. In *Trends in Cognitive Sciences* (Vol. 13, Issue 12, pp. 517–523). Elsevier Current Trends.

Hillson, D., & Murray-Webster, R. (2007). *Understanding and managing risk attitude*. Gower Publishing, Ltd.

Hsu, M., Bhatt, M., Adolphs, R., Tranel, D., & Camerer, C. F. (2005). Neural systems responding to degrees of uncertainty in human decision-making. *Science*, 310(5754), 1680–1683.

Ida, T. (2010). Anomaly, impulsivity, and addiction. *Journal of Socio-Economics*, 39(2), 194–203.

Jaeger, C. C., Weblar, T., Rosa, E. A., & Renn, O. (2013). *Risk, uncertainty and rational action*. Routledge.

Kahneman, D., & Tversky, A. (1979). Prospect Theory: An Analysis of Decision under Risk. *Econometrica*, 47(2), 263–292.

Kőszegi, B., & Rabin, M. (2006). A model of reference-dependent preferences. *The Quarterly Journal of Economics*, 121(4), 1133–1165.

Kőszegi, B., & Rabin, M. (2007). Reference-dependent risk attitudes. *American Economic Review*, 97(4), 1047–1073.

Levy, I., Snell, J., Nelson, A. J., Rustichini, A., & Glimcher, P. W. (2010). Neural representation of subjective value under risk and ambiguity. *Journal of Neurophysiology*, 103(2), 1036–1047.

Mehra, R., & Prescott, E. C. (1985). The equity premium: A puzzle. *Journal of Monetary Economics*, 15(2), 145–161.

Mongillo, G., Shteingart, H., & Loewenstein, Y. (2014). The Misbehavior of Reinforcement

- Learning. *Proceedings of the IEEE*, 102(4), 528–541.
- O’donoghue, T., & Somerville, J. (2018). Modeling risk aversion in economics. *Journal of Economic Perspectives*, 32(2), 91–114.
- Pennings, J. M. E., & Smidts, A. (2000). Assessing the construct validity of risk attitude. *Management Science*, 46(10), 1337–1348.
- Rabin, M., & Thaler, R. H. (2001). Anomalies: risk aversion. *Journal of Economic Perspectives*, 15(1), 219–232.
- Rosen, A. B., Tsai, J. S., & Downs, S. M. (2003). Variations in Risk Attitude across Race, Gender, and Education. In *Medical Decision Making* (Vol. 23, Issue 6, pp. 511–517). SAGE Publications.
- Shteingart, H., & Loewenstein, Y. (2014). Reinforcement learning and human behavior. In *Current Opinion in Neurobiology* (Vol. 25, pp. 93–98).
- Shteingart, H., Neiman, T., & Loewenstein, Y. (2013). The role of first impression in operant learning. *Journal of Experimental Psychology: General*, 142(2), 476–488.
- Smith, K., Dickhaut, J., McCabe, K., & Pardo, J. V. (2002). Neuronal substrates for choice under ambiguity, risk, gains, and losses. *Management Science*, 48(6), 711–718.
- Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction*. MIT press.
- Szita, I., & Lőrincz, A. (2008). The many faces of optimism. *Proceedings of the 25th International Conference on Machine Learning - ICML ’08*, 1048–1055.
- Thorndike, E. L. (1927). The Law of Effect. *The American Journal of Psychology*, 39(1/4), 212.
- Trepel, C., Fox, C. R., & Poldrack, R. A. (2005). Prospect theory on the brain? Toward a cognitive neuroscience of decision under risk. *Cognitive Brain Research*, 23(1), 34–50.
- Tversky, A., & Fox, C. R. (1995). Weighing risk and uncertainty. *Psychological Review*, 102(2), 269.
- Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, 5(4), 297–323.

- Weber, E. U., Blais, A.-R., & Betz, N. E. (2002). A domain-specific risk-attitude scale: measuring risk perceptions and risk behaviors. *Journal of Behavioral Decision Making*, 15(4), 263–290.
- Yechiam, E., & Ert, E. (2011). Risk Attitude in Decision Making: In Search of Trait-Like Constructs. *Topics in Cognitive Science*, 3(1), 166–186.
- Zhong, S., Chew, S. H., Set, E., Zhang, J., Xue, H., Sham, P. C., Ebstein, R. P., & Israel, S. (2009). The heritability of attitude toward economic risk. *Twin Research and Human Genetics*, 12(1), 103–107.
- Zuckerman, M. (2007). Sensation seeking and risky behavior. In *Sensation seeking and risky behavior*. American Psychological Association.