On the convergent validity of risk sensitivity measures

Angela Radulescu^{1,2}, Katharine Holmes², and Yael Niv^{1,2,*}

¹Department of Psychology

²Princeton Neuroscience Institute

* Corresponding author: yael@princeton.edu

Abstract

There are a number of well-accepted ways to measure risk sensitivity, with researchers often making conclusions about individual differences based on a single task. Even though long-standing observations suggest that how risky outcomes are presented changes people's behavior, it is unclear whether risk sensitivity is a unitary trait that can be measured by any one of these instruments. To directly answer this question, we administered three tasks commonly used to elicit risk sensitivity within-subject to a large sample of participants on Amazon Mechanical Turk. Our findings revealed high individual variability in each measure, with little evidence of consistency among different tasks: many participants who were classified as "risk-averse" in one task were "risk-seeking" in another, and we observed no significant correlations between continuous measures of risk sensitivity as measured in each of the tasks. Our results cast doubt on the pervasive assumption that risk paradigms measure a single underlying trait, and suggest instead that behavior in risky situations is the result of heterogeneous, interacting and possibly task-dependent cognitive mechanisms.

On the convergent validity of risk sensitivity measures

Introduction

Life is fraught with decisions in which we need to evaluate risk, formally defined as the variance in the outcomes of one's actions. The way humans represent and process this variance may change depending on the situation. Intuitively, deciding among different insurance plans given detailed information about alternatives is quite different from deciding whether or not to go sky diving. In the first situation, we may perform explicit, meticulous calculations based on available data, while in the second we may rely more on affect (Loewenstein, Weber, Hsee, & Welch, 2001), or on bringing specific episodes to mind (Bornstein, Khaw, Shohamy, & Daw, 2015) in order to evaluate risk. This representational heterogeneity is reflected in the variety of measures that psychologists have used to elicit risk preferences (Schonberg, Fox, & Poldrack, 2011; Charness, Gneezy, & Imas, 2013), as well as in the different assumptions underlying models of risky decision-making (Glimcher & Fehr, 2013).

The psychological literature on risk has broadly distinguished between two types of risky decision-making: “decisions from description,” in which people are explicitly provided with information about risk in tasks framed as economic decisions, and “decisions from experience,” in which people need to learn action-outcome contingencies from trial and error, by sampling the payoff distribution associated with available actions.

Decisions from description are typically explained under expectation-based models of risk-taking (Schonberg et al., 2011). The most influential of these has been Prospect Theory (Kahneman & Tversky, 1979). According to Prospect Theory, individuals maximize the product of the subjective value of an outcome and the probability of its occurrence. Kahneman and Tversky famously extended expected utility theory to accommodate the finding that individuals tend to be risk averse in the domain of gains (they prefer a sure win to an uncertain win of the same expected value) and risk seeking in the domain of losses (they prefer an uncertain loss to a sure loss of the same expected

value), and have distortions in their subjective perception of probabilities. While illuminating in many cases, these types of are often silent as to the cognitive mechanisms that might have generated a particular representation of risk.

Experience-based paradigms stand in contrast to description-based studies in that they typically involve repeated choices among different alternatives (Weber, Shafir, & Blais, 2004; Hertwig, Barron, Weber, & Erev, 2004; Hertwig & Erev, 2009). Unlike in decisions from description, the participant must sample payoffs in order to learn from experience the distribution of possible outcomes and its associated uncertainty (i.e., risk). Decisions from experience have more recently been explained within reinforcement learning (RL) theory (Sutton & Barto, 1998), which models the dynamic, sequential nature of human decision-making. These models explicitly link psychological representations of risk to biologically plausible cognitive mechanisms of learning and decision-making (Niv, Edlund, Dayan, & O’Doherty, 2012; Niv, 2009).

In our own work investigating the cognitive and neural mechanisms of risky decision making in trial-and-error learning (Niv et al., 2012; Arkadir et al., 2016) we have often encountered a question regarding the choice to measure risk in a risk-from-experience task, rather than a an explicit choice task. Our reply has been that we do not believe these different tasks are measuring the same construct; however, we had no empirical evidence to this effect. Given that this question arises consistently from both readers and reviewers, here we report a systematic comparison, within subjects, of risk sensitivity elicited through several different experimental paradigms.

Our hypothesis that the different paradigms measure separable constructs is based on the previously-documented “description-experience gap” (Weber et al., 2004; Hertwig et al., 2004; Hertwig & Erev, 2009; Erev et al., 2010). Hertwig and Erev argue that humans overweight rare events in standard description paradigms but underweight rare events when sampling the underlying distribution over time. This suggests that experimental measures of risk sensitivity may not agree with each other or be consistent within

participant, as a trait view of risk might suggest.

To briefly illustrate the description-experience gap, take the example of a simple slot machine $S1$ that yields \$10 with 50% probability (and nothing otherwise). In a single-shot description paradigm, an experimenter might ask the participant to choose between playing this slot machine and a different slot machine $S2$ that yields \$5 with 100% probability. The expected value of playing the two slot machines is the same (\$5), so we would expect a participant who is risk averse to choose $S2$. Alternatively, the experimenter might repeatedly present the participant with the option to choose between $S1$ and $S2$ without making the probabilities known in advance. In this case, evaluation of risk occurs through sampling action-outcome contingencies and learning the value of each option (Niv et al., 2012; Arkadir et al., 2016). Importantly, RL theory predicts that option values would be continuously updated, with the consequence that participants may be both risk averse and risk seeking at different points in the task. Consider for instance a case in which the learner chooses $S1$ three times in a row and receives a reward all three times. According to RL theory, the expected value of $S1$ would be adjusted upward on each trial in proportion to the prediction error, and the learner would tend to select $S1$ again, as its value could well exceed \$5 (the value of $S2$). However, if choosing $S1$ leads to no reward for three trials in a row, the value of $S1$ relative to $S2$ would decrease, leading the learner to choose $S2$. Once the learner chooses $S2$, she gets no additional information about the value of $S1$, thereby she will continue to view $S1$ as inferior. This can render behavior seemingly risk averse, even though in this particular algorithm the learner is not learning about or representing risk at all. In other words, in experiential learning, momentary values of options depend on local sampling biases, which can lead to dynamic changes in apparent risk sensitivity that do not stem from evaluating risk at all.

There have also been attempts to assess attitudes towards risk in more ecologically relevant situations as opposed to economic gambles (Weber, Blais, & Betz, 2002; Blais, Weber, et al., 2006). The widely-used Domain-Specific Risk-Taking scale (DOSPERT)

proposed by Weber and colleagues, has been a tool of choice for characterizing domain and group differences in risk taking. DOSPERT has invited reflection on the domain-specificity of assessing risky behavior, in the service of nudging behavior outside the lab (Figner & Weber, 2011). However, it is unclear whether the parameters estimated using DOSPERT data reflect attitudes to risk defined mathematically, that is, to outcome variance, or to other factors that may drive risky behavior (Schonberg et al., 2011).

To assess the similarity and differences between these different measures of risk attitudes, and in particular to test whether explicit measures correspond to experience-based methods of eliciting risk preferences, we examine individual differences in risky decision-making in three different paradigms in a large sample of participants recruited from Amazon Mechanical Turk. A trait view of risk sensitivity would suggest that risk sensitivity within participant should be roughly the same regardless of the choice of paradigm. Alternatively, patterns of risk sensitivity could differ by paradigm. In particular, the “description-experience” gap suggests that we might expect an inverse relationship between risk sensitivity assessed from description or experience.

We estimate individual risk sensitivity from experience (a repeated choice RL task), description (the Holt and Laury method) and self-report (the DOSPERT scale). The results suggest that these three measures of risk sensitivity have little convergent validity: experience and description risk sensitivity are uncorrelated, with many participants labeled as “risk-averse” in one paradigm and “risk-seeking” in another, and only description-based risk-sensitivity measures account for a significant amount of variance in self-reported risk attitudes. We conclude from these results that risk sensitivity, at least as it is being measured in these laboratory tasks, does not reflect a unitary individual trait, but rather the task-dependent output of interacting decision-making mechanisms.

Methods

Our strategy was to pick several established risk paradigms, and from each extract a single continuous measure of risk sensitivity per participant. In line with the literature, we focused on the distinction between decisions from description and decisions from experience, and also included DOSPERT as a self-report measure of risk sensitivity.

Participants and power analysis

563 participants (259 female, 302 male, 2 did not select a gender; age distribution shown in Table 1) recruited through Amazon Mechanical Turk provided measures of risk sensitivity using three separate paradigms in one sitting that lasted 20 minutes on average. The large sample size enabled us to draw conclusions we would not be able to draw from fewer data points: a formal power analysis revealed that a small to medium correlation effect ($r = 0.2$) at a power level of 80% (Cohen, 1992) would require a sample size of at least 193 participants.

Risk sensitivity across three paradigms

We used two criteria to choose tasks for eliciting risk-sensitivity: (1) the task had to previously be validated as a reliable risk-sensitivity elicitation method; (2) the measure yielded a continuous dependent variable that assessed risk sensitivity relative to a task-specific risk neutrality baseline.

Task 1: Reinforcement Learning (decisions from experience). To evaluate risk sensitivity from experience, we examined behavior on a reinforcement learning task previously used to track the dynamics of risk sensitivity and identify its neural substrates (Niv et al., 2012; Arkadir et al., 2016). In a sequential decision-making task, participants made repeated choices between stimuli associated with different reward probabilities (Fig. 1A).

Of the four stimuli that appeared in the task, $S1$, $S3$ and $S4$ respectively yielded a

reward of 0, 2 and 4 cents every time they were chosen. $S2$ (hereby denoted the “risky” stimulus) yielded 4 cents or 0 cents each with 50% probability (Fig. 1B). There were 9 different randomly intermixed trial types distributed as shown in Table 2, adding up to a total of 150 trials. The task primarily consisted of 2-alternative choices between pairs of stimuli. Specifically, participants were presented with choices between $S1$ and $S2$ (trial type 1, $S2$ correct), between $S2$ and $S3$ (trial type 2, hereby denoted as “risk trials”, no predefined correct response, since both stimuli have the same expected value), between $S2$ and $S4$ (trial type 3, $S4$ correct), between $S1$ and $S3$ (trial type 4, $S3$ correct), and between $S3$ and $S4$ (trial type 5, $S4$ correct). To ensure that participants sampled all 4 stimuli, we also included “forced” trials in which only one of the four stimuli appeared on the screen and had to be chosen by the participant. The rewards associated with the risky stimulus were pseudo-randomized in blocks of 8. Participants’ performance was incentivized such that every cent earned in the task accrued towards a performance-based bonus of up to \$3.70 awarded at the end of the session. Finally, in addition to the 150 trials used for analysis, a training phase at the onset of the task included two of each of the choice trial types and four of each of the forced trial types, randomly intermixed, yielding a total of 176 trials.

Using a standard keyboard, participants indicated their preference for the left or right stimulus using the “B” and “M” keys. If a participant failed to respond within 1.5 seconds, the message “Too slow!” was displayed and reward for that trial was forfeited. If a participant pressed the wrong key on forced trials, the message “Wrong side!” was displayed, no reward was given and the task advanced to the next trial. Each outcome was presented for 0.5 seconds and the inter-trial interval varied between 1 and 2 seconds (uniform distribution). To ensure that participants did not take disruptive breaks, we imposed two monitoring criteria: the experiment was terminated if a participant missed more than 10 trials, or took longer than 20 seconds to respond on any of the trials.

In this setting, the measure of risk sensitivity is given by the proportion of times a

participant chose the risky stimulus $S2$ when $S3$ was presented as the alternative (trial type 2). Given that the two stimuli have the same expected value, a participant who is risk neutral should, on average, choose both equally often. In contrast, a risk-seeking participant should prefer $S2$ and a risk-averse participant should prefer $S3$.

Task 2: Holt and Laury’s method (decisions from description). To evaluate risk sensitivity in decisions from description, we used a variant of the paradigm popularized by Holt and Laury (Holt & Laury, 2002, 2005). The method consists of eliciting a series of choices between two gambles where the relative expected value of the options gradually changes (Fig. 2A). The pairs of gambles that we used are listed in Figure 2B (adapted from Holt & Laury 2002). Here Gamble A is considered the “safe” gamble, as its variance is lower than that of Gamble B. The number of times the participant chose Gamble A over Gamble B (i.e., the number of “safe”, low-risk choices) implies a particular shape for a participant’s utility function and can thus be used as a continuous measure of risk sensitivity ranging from highly risk loving (0-1 safe choices) to highly risk averse (9-10 safe choices, Holt & Laury 2002).

We presented each participant with the choices shown in Figure 2B one at a time, allowing 8 seconds for each choice, with time between trials varying uniformly between 1 and 2 seconds. Before starting the task, participants were instructed, by example, on how to read the options: they saw Gamble A from the first pair of choices displayed on the screen and were told that “this means that if you take this gamble, you have a 1 in 10 chance of winning \$2 and a 9 in 10 chance of winning \$1.60.” Participants were also informed that one of the chosen gambles will be randomly selected and played out at the end of the task, and that the outcome of this draw would count towards the overall bonus awarded at the end of the session. Participants thus had the opportunity to earn up to \$3.85 from this task.

Task 3: The Domain-Specific Risk-Taking Scale (DOSPERT). To obtain a valid measure of risk attitudes in ecologically-valid settings, we administered an online

version of the DOSPERT questionnaire (Weber et al., 2002; Blais et al., 2006).

Participants evaluated 30 activities that carry considerable levels of risk on the following three dimensions: expected benefits, perceived riskiness and preference for that activity. For instance, the item “Going down a ski run that is beyond your ability” appeared three times throughout the questionnaire, and the participants responded with a rating from 1 to 7 when asked: (1) to indicate how much benefit they thought they would obtain from the situation, (2) how risky they perceived the situation to be, and (3) the likelihood that they would engage in the described activity or behavior if they were to find themselves in that situation. Participants had 5 seconds to respond to each item. The ratings that participants reported for each of the three dimensions (expected benefit, perceived risk and preference) were then entered into a linear regression model that captures the trade-off between the benefit of engaging in a risky activity and the perceived risk associated with that activity in determining a participant’s preference for that activity. The measure of risk sensitivity is given by the estimated b coefficient, with $b < 0$ indicating risk aversion (that is, perceived risk decreasing preference for an activity) and $b > 0$ indicating risk seeking (as perceived risk increases the attractiveness of the activity). We also analyzed the intercorrelations with the other measures of risk sensitivity without decomposing preference into expected benefits and perceived risk, by entering the mean self-reported likelihood of engaging in risky activities directly into the correlation analysis.

Overall structure. The experiment began with a screen informing participants that they were to perform several separate “mini-experiments,” each preceded by its own set of instructions. Participants first performed the RL task, then they filled in the DOSPERT, and finally they completed the Holt and Laury task. Pauses between experiments were self-paced.

Mechanical Turk data-quality assurance

The tasks were adapted for online testing using the open source PsiTurk library (McDonnell et al., 2012). Mechanical Turk workers have successfully participated in tasks ranging from economic games to category learning, yielding data comparable to what is typically found in lab (Mason & Suri, 2012; Hauser, Rand, Peysakhovich, & Nowak, 2014; Crump, McDonnell, & Gureckis, 2013).

To further ensure data quality, for all analyses involving risk from experience, we excluded 136 participants who did not reach at least 65% overall accuracy in non-risky learning trials. The remaining participants showed similar performance to that of in-lab participants (Niv et al., 2012) (Fig. 3A).

For all analyses involving risk from description, we excluded 244 participants who showed choice inconsistency in the Holt and Laury paradigm (i.e. they either reversed after the switch point, or started choosing Gamble B and then reverted back to A). We note that the issue of choice inconsistency is not unique to this pool of participants. Charness and colleagues classify Holt and Laury's method as a "complex" risk elicitation paradigm and write that one of its main disadvantages is that it may be hard to convey to participants, and thus prone to noisier behavior (Charness et al., 2013; Dave, Eckel, Johnson, & Rojas, 2010). They further note that "this poses an obvious problem since the inference of risk preferences, and in turn, parameter estimation requires a unique switch point and such inconsistent behavior is difficult to rationalize under standard assumptions on preferences." (pp. 47). We nevertheless chose this paradigm as it is used more commonly than simpler methods to differentiate between different degrees of risk-seeking behavior (Charness et al., 2013), and relied on our ability to collect large amounts of data in order to obtain a large sample of participants who made consistent choices.

Finally, we excluded from DOSPRT analyses 112 subjects whose risk taking coefficients could not be estimated due to singularity in the regression model, or who had more than 10 extreme reaction times (defined as ± 3 standard deviations below or above

the mean).

These exclusion criteria ensured that for any pair of tasks, data from participants with noisier behavior in one or both tasks did not contaminate the convergent validity analyses.

Results

When making decisions from experience, participants quickly learned to select the best option in the majority of trials in which the risky stimulus appeared together with either the sure 0c or sure 4c stimuli (mean probability correct choice: 0.85 ± 0.15 *SD*, Fig. 3A left), as well as trials in which a pair of sure stimuli appeared together (mean probability correct choice: 0.88 ± 0.14 *SD*, Fig. 3A right). When the risky stimulus *S2* was presented together with the equal-expected-value sure stimulus *S3*, on average participants showed risk-averse behavior that stayed relatively constant over time, as is typically observed in such tasks (Niv et al., 2012) (Fig. 3B).

When making decisions from description, participants' risk sensitivity showed quite a bit of individual variability (Fig. 4, $M=6.31$, $SD=2.51$), but the number of safe choices differed significantly from the optimum of 4 (one-tailed t-test against 4, $t(318) = 16.44$, $p < 0.001$). Most participants switched from choosing Gamble A to Gamble B past the optimal switch point, revealing predominantly risk-averse tendencies at the population level.

Scoring the DOSPRT by regressing likelihood of engaging in a particular activity on perceived riskiness and expected benefit yielded an overall per-participant risk sensitivity index, as well as five domain-specific indices. Participants were risk averse overall ($M = -0.24$, $SD=0.31$, one-tailed t-test against 0, $t(450) = 16.17$, $p < 0.001$), and were most risk averse in the health domain ($M = -0.29$, $SD=0.72$) and least risk averse in the financial domain ($M = -0.15$, $SD=0.69$).

The main analysis of interest was the intercorrelation between pairs of risk sensitivity measures. Even though we were adequately powered for all three pairwise correlations ($N > 193$, see "Sample and power analysis" section), we observed no significant correlation

between risk sensitivity as measured in the RL task and risk sensitivity as measured in the Holt and Laury paradigm ($N = 261$, $r = 0.01$, $p = 0.82$, Fig. 5A), no significant correlation between risk sensitivity estimated from self-reported attitudes and risk sensitivity as measured in the Holt and Laury paradigm ($N = 257$, $r = 0.004$, $p = 0.95$, Fig. 5B), and no significant correlation between risk sensitivity estimated from self-reported attitudes and risk sensitivity as measured in the RL task ($N = 348$, $r = -0.06$, $p = 0.23$, Fig. 5C).

When entering the self-reported activity preference measure in the correlation analysis without decomposing it into expected benefits and perceived risk, we found a moderate correlation with risk sensitivity as measured in the Holt and Laury paradigm ($N = 257$, $r = -0.18$, $p = 0.003$), but no correlation with risk sensitivity as measured in the reinforcement learning task ($N = 348$, $r = -0.04$, $p = 0.40$).

Taken together, our findings suggest that the risk elicitation methods studied here have little convergent validity, and that individual risk sensitivity may not be consistent across elicitation methods. The same person may be classified as risk-taking in one setting and risk-averse in another, and only some graded measures predict each other, suggesting that risk sensitivity is not a trait, but rather the result of several cognitive processes that are perhaps differentially employed in different tasks.

Discussion

Previous research has shown that whether risk is experienced or described impacts people’s risk preferences during decision-making ([Hertwig & Erev, 2009](#)). In this study, we aimed to test whether this claim holds within individuals, and more specifically establish whether individual risk sensitivity is consistent across risk elicitation methods. A trait hypothesis of risk sensitivity predicts that different measures should agree with each other to some extent when labeling individuals as more or less risk taking. On the other hand, viewing risk sensitivity as resulting from interacting cognitive mechanisms differentially recruited in different decision-making contexts, it is plausible that the same individual be

risk taking in one task context and risk averse in another. Our results are consistent with the latter hypothesis, as we found no evidence for correlations between risk sensitivity during learning from experience and risk sensitivity in a description paradigm. The moderate correlation between self-reported preference for risky activities and risk sensitivity in decisions from description is in line with previous findings ([Weber et al., 2004](#); [Camilleri & Newell, 2009](#)) and suggests we are adequately powered; but it does not hold when accounting for perceived benefit by decomposing activity preference into expected benefit and perceived risk and estimating a risk-sensitivity coefficient. These findings suggest that risk sensitivity is not a stable trait, but rather the output of an interaction between several psychological mechanisms that may evaluate and respond to risk.

If risk sensitivity is not a trait, what factors might arbitrate between risk aversion and risk seeking within the same individual? Work on risk attitudes ([Weber et al., 2002](#)) has suggested that risky decisions are domain-specific: people treat risk differently in different content domains. One possible explanation is that individual differences in domain-specific risk attitudes may be explained by sampling biases over the course of one's lifetime: for instance, a young adult who has yet to get a speeding ticket may evaluate the risk associated with driving faster than the speed limit differently than her wiser future self. Indeed, one way reinforcement learning researchers ensure sampling of all options is by introducing so-called "forced-choice" trials ([Niv et al., 2012](#); [Arkadir et al., 2016](#)), as we have done in the RL task in this study. Indeed, several lines of research have emphasized the role of episodic sampling in evaluating choice alternatives, with recent evidence converging on the idea that the same mechanisms that support perceptual evidence accumulation are at play in value-based decisions ([Shadlen & Shohamy, 2016](#)). One can imagine a mechanism by which value construction and consequent risk evaluation for different options rely on consulting memories for individual past experiences related to those options ([Madan, Ludvig, & Spetch, 2017](#)). Another possibility is that in experience-based settings, where learning occurs over the course of a relatively short task,

idiosyncrasies in risky behavior within a single individual may arise simply due to experiential factors, such as the specific outcomes experienced, choice context (Louie, Glimcher, & Webb, 2015) or choice horizon (Wilson, Geana, White, Ludvig, & Cohen, 2013).

It is also possible that convergent validity of risk sensitivity in description and experience paradigms can be achieved, but under narrower experimental conditions than we have investigated in this work. A recent study by Madan and colleagues found a tight within-subject coupling between the two measures when presenting subjects with exactly the same contingencies. In both the description and experience conditions, participants always chose between a sure and a risky gamble (instead of two risky gambles). The described gambles were matched for outcome magnitude, variance, domain (gain vs. loss) and number of presentations to the stimuli in the experience condition (Madan et al., 2017). Moreover, no other options other than the sure vs. risky stimuli were presented in the experience condition, which raises the possibility that the participants reached learning asymptote much sooner than if other choice types had been interleaved throughout the task.

Other than the theoretical challenge our findings pose, one pragmatic concern is that psychological tasks are often tested as candidate biomarkers for psychopathology. In our own work, for instance, we used an experience paradigm very similar to the one we report here to examine risk taking in DYT1 dystonia, a rare motor disorder that manifests in an inability to suppress certain body movements. Animal models of DYT1 dystonia suggested that the pathology may arise from a specific alteration in neural learning (specifically, corticostriatal plasticity). This led us to predict that DYT1 patients would learn more from positive than from negative prediction errors, which in an experience-based paradigm, would lead to more risk taking compared to controls. Indeed, this is what we found, and moreover, we showed that this effect predicts dystonia symptom severity (Arkadir et al., 2016). But does this mean that we should expect DYT1 patients to be more reckless

investors, or go sky diving more often? The results of the current study suggest not. Importantly, we should not even necessarily expect DYT1 patients to be more risk taking than controls in a description setting. While we can make inferences from behavior on an RL task regarding underlying neurobiological function ([Arkadir et al., 2016](#)), we do not necessarily expect the pattern of results to transfer to other laboratory tasks.

There is a burgeoning interest in bridging the gap between neuroeconomic measures of decision making and the more applied domains of individual differences and clinical practice ([Hsu & Chiong, 2016](#)). Our work adds to recent efforts in the area of risky decision-making ([Camilleri & Newell, 2009](#); [Mamerow, Frey, & Mata, 2016](#); [Madan et al., 2017](#)), and highlights the potential of large subject samples on Amazon Mechanical Turk to better characterize the convergent validity of a wide set of measures tapping into a single psychological construct. The results we report are consistent with recent evidence that risk sensitivity is not a unitary trait ([Frey, Pedroni, Mata, Rieskamp, & Hertwig, 2017](#); [Pedroni et al., 2017](#)), and suggest that conclusions about individual differences are highly contingent on the choice of task.

References

- Arkadir, D., Radulescu, A., Raymond, D., Lubarr, N., Bressman, S. B., Mazzoni, P., & Niv, Y. (2016). Dyt1 dystonia increases risk taking in humans. *eLife*, *5*, e14155.
- Blais, A.-R., Weber, E. U., et al. (2006). A domain-specific risk-taking (dosPERT) scale for adult populations. *Judgment and Decision Making*, *1*(1), 33–47.
- Bornstein, A. M., Khaw, M. W., Shohamy, D., & Daw, N. D. (2015). What’s past is present: Reminders of past choices bias decisions for reward in humans. *bioRxiv*, 033910.
- Camilleri, A. R., & Newell, B. (2009). Within-subject preference reversals in description-and experience-based choice. In *Proceedings of the cognitive science society* (Vol. 31).
- Charness, G., Gneezy, U., & Imas, A. (2013). Experimental methods: Eliciting risk preferences. *Journal of Economic Behavior & Organization*, *87*, 43–51.
- Cohen, J. (1992). A power primer. *Psychological Bulletin*, *112*(1), 155.
- Crump, M. J., McDonnell, J. V., & Gureckis, T. M. (2013). Evaluating amazon’s mechanical turk as a tool for experimental behavioral research. *PloS one*, *8*(3), e57410.
- Dave, C., Eckel, C. C., Johnson, C. A., & Rojas, C. (2010). Eliciting risk preferences: When is simple better? *Journal of Risk and Uncertainty*, *41*(3), 219–243.
- Erev, I., Ert, E., Roth, A. E., Haruvy, E., Herzog, S. M., Hau, R., ... Lebiere, C. (2010). A choice prediction competition: Choices from experience and from description. *Journal of Behavioral Decision Making*, *23*(1), 15–47.
- Figner, B., & Weber, E. U. (2011). Who takes risks when and why? determinants of risk taking. *Current Directions in Psychological Science*, *20*(4), 211–216.
- Frey, R., Pedroni, A., Mata, R., Rieskamp, J., & Hertwig, R. (2017). Risk preference shares the psychometric structure of major psychological traits. *Science advances*, *3*(10), e1701381.

- Glimcher, P. W., & Fehr, E. (2013). *Neuroeconomics: Decision making and the brain*. Academic Press.
- Hauser, O. P., Rand, D. G., Peysakhovich, A., & Nowak, M. A. (2014). Cooperating with the future. *Nature*, 511(7508), 220–223.
- 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. *Trends in Cognitive Sciences*, 13(12), 517–523.
- Holt, C. A., & Laury, S. K. (2002). Risk aversion and incentive effects. *American economic review*, 92(5), 1644–1655.
- Holt, C. A., & Laury, S. K. (2005). Risk aversion and incentive effects: New data without order effects. *American Economic Review*, 902–904.
- Hsu, M., & Chiong, W. (2016). From laboratory to clinic and back: Connecting neuroeconomic and clinical measures of decision-making dysfunctions. In *Handbook of health decision science* (pp. 39–52). Springer.
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica: Journal of the Econometric Society*, 263–291.
- Loewenstein, G. F., Weber, E. U., Hsee, C. K., & Welch, N. (2001). Risk as feelings. *Psychological bulletin*, 127(2), 267.
- Louie, K., Glimcher, P. W., & Webb, R. (2015). Adaptive neural coding: from biological to behavioral decision-making. *Current opinion in behavioral sciences*, 5, 91–99.
- Madan, C. R., Ludvig, E. A., & Spetch, M. L. (2017). The role of memory in distinguishing risky decisions from experience and description. *The Quarterly Journal of Experimental Psychology*, 70(10), 2048–2059.
- Mamerow, L., Frey, R., & Mata, R. (2016). Risk taking across the life span: A comparison of self-report and behavioral measures of risk taking. *Psychology and Aging*.
- Mason, W., & Suri, S. (2012). Conducting behavioral research on amazon’s mechanical

- turk. *Behavior research methods*, 44(1), 1–23.
- McDonnell, J., Martin, J., Markant, D., Coenen, A., Rich, A., & Gureckis, T. (2012). *psiturk (version 1.02)[software]*. new york, ny: New york university.
- Niv, Y. (2009). Reinforcement learning in the brain. *Journal of Mathematical Psychology*, 53(3), 139–154.
- Niv, Y., Edlund, J. A., Dayan, P., & O’Doherty, J. P. (2012). Neural prediction errors reveal a risk-sensitive reinforcement-learning process in the human brain. *The Journal of Neuroscience*, 32(2), 551–562.
- Pedroni, A., Frey, R., Bruhin, A., Dutilh, G., Hertwig, R., & Rieskamp, J. (2017). The risk elicitation puzzle. *Nature Human Behaviour*, 1(11), 803.
- Schonberg, T., Fox, C. R., & Poldrack, R. A. (2011). Mind the gap: bridging economic and naturalistic risk-taking with cognitive neuroscience. *Trends in cognitive sciences*, 15(1), 11–19.
- Shadlen, M. N., & Shohamy, D. (2016). Decision making and sequential sampling from memory. *Neuron*, 90(5), 927–939.
- Sutton, R., & Barto, A. (1998). *Reinforcement learning: An introduction* (Vol. 116). Cambridge University Press.
- 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, 263–290.
- Weber, E. U., Shafir, S., & Blais, A.-R. (2004). Predicting risk sensitivity in humans and lower animals: risk as variance or coefficient of variation. *Psychological review*, 111(2), 430.
- Wilson, R. C., Geana, A., White, J. M., Ludvig, E. A., & Cohen, J. D. (2013). Exploration strategies in human decision making. *Reinforcement Learning and Decision Making*.

Age range	Count
18 – 24	80
25 – 34	257
35 – 44	129
45 – 54	67
55 – 64	24
65 – 74	4
No answer	2

Table 1
Age distribution of participants.

Trial type	Count
1 ($S1$ vs. $S2$)	20
2* ($S2$ vs. $S3$)	30
3 ($S2$ vs. $S4$)	20
4 ($S1$ vs. $S3$)	15
5 ($S3$ vs. $S4$)	15
6 ($S1$ forced)	12
7 ($S2$ forced)	14
8 ($S3$ forced)	12
9 ($S4$ forced)	12

Table 2

Distribution of trial types in the reinforcement learning task. The star denotes risk trials.

Figure 1. Reinforcement learning task (A) Task structure and example trials. ITI = inter trial interval. (B) Reward structure associated with each of the four stimuli. *S2* is the risky stimulus, and has the same expected value as *S3*.

A

Which gamble would you prefer?

Gamble A	Gamble B
1/10 of \$2 or 9/10 of \$1.60	1/10 of \$3.85 or 9/10 of \$0.10

B

Gamble A	EV(A)	Var(A)	Var(B)	EV(B)	Gamble B
1/10 of \$2 or 9/10 of \$1.60	1.64	0.014	1.27	0.48	1/10 of \$3.85 or 9/10 of \$0.10
2/10 of \$2 or 8/10 of \$1.60	1.68	0.026	2.25	0.85	2/10 of \$3.85 or 8/10 of \$0.10
3/10 of \$2 or 7/10 of \$1.60	1.72	0.034	2.95	1.23	3/10 of \$3.85 or 7/10 of \$0.10
4/10 of \$2 or 6/10 of \$1.60	1.76	0.039	3.38	1.60	4/10 of \$3.85 or 6/10 of \$0.10
5/10 of \$2 or 5/10 of \$1.60	1.80	0.040	3.52	1.98	5/10 of \$3.85 or 5/10 of \$0.10
6/10 of \$2 or 4/10 of \$1.60	1.84	0.038	3.38	2.35	6/10 of \$3.85 or 4/10 of \$0.10
7/10 of \$2 or 3/10 of \$1.60	1.88	0.034	2.95	2.73	7/10 of \$3.85 or 3/10 of \$0.10
8/10 of \$2 or 2/10 of \$1.60	1.92	0.026	2.25	3.10	8/10 of \$3.85 or 2/10 of \$0.10
9/10 of \$2 or 1/10 of \$1.60	1.96	0.014	1.27	3.48	9/10 of \$3.85 or 1/10 of \$0.10
10/10 of \$2 or 0/10 of \$1.60	2.00	0	0	3.85	10/10 of \$3.85 or 0/10 of \$0.10

Figure 2. Holt and Laury task (A) Sample trial in the description paradigm. Participants were asked to make repeated choices between Gamble A and Gamble B. In this example, Gamble A yields \$2 with 10% probability and \$1.60 with 90% probability, and Gamble B yields \$3.85 with 10% probability and \$0.10 with 90% probability. (B) Gambles used and their associated expected value and variance. Risk-neutral participants would choose Gamble A four times and then switch to Gamble B (grey arrow), simply maximizing the expected value (EV) of the gamble. Risk-sensitive participants would also take into account the risk or variance (Var) of each option.

Figure 3. Reinforcement learning task behavior. (A) Performance in non-risky choice trials, averaged over trials where the risky stimulus was present (left), and on trials where the risky stimulus was absent (right). Data were binned in blocks of 5 trials each. Dotted line denotes the per-participant performance criterion (chance is 50%). (B) Choice patterns in the risky choice trials. Dotted line denotes risk neutrality. Error bars denote one SEM.

Figure 4. Distribution of risk sensitivities in decisions from description. The dotted line denotes risk neutrality.

Figure 5. Correlations between risk sensitivity measures. (A) Risk from experience vs. risk from description. (B) Risk attitudes vs. risk from description. (C) Risk attitudes vs. risk from experience. Each datapoint represents a pair of measurements from the same subject. None of the correlations were significant at $p < .05$. Orange: line of best fit.