



Life
Course
Centre

WORKING
PAPER
SERIES

No. 2024-17

May 2024

The gender reference point gap: Evidence from a representative sample

Jonathan Levy

Olivia Ru

Agnieszka Tymula

The Australian Research Council Centre of Excellence
for Children and Families over the Life Course
Phone +61 7 3346 7477 **Email** lcc@uq.edu.au
lifecoursecentre.org.au



Australian Government
Australian Research Council



THE UNIVERSITY
OF QUEENSLAND
AUSTRALIA



THE UNIVERSITY OF
SYDNEY



THE UNIVERSITY OF
WESTERN
AUSTRALIA



THE UNIVERSITY OF
MELBOURNE

Research Summary

Why was the research done?

This study explores the differences between men and women when it comes to their willingness to take risks. Previous research has shown that women tend to be more risk averse than men, and this gender gap can affect women's economic opportunities. Typically, attitudes toward risk have been seen as unchangeable traits, which has made it challenging to grasp why these gender gaps exist and how they might be addressed.

What were the key findings?

We estimate the “reference point” for men and women by conducting an incentivized experiment with 579 individuals, who are representative of the US population along several dimensions. The reference point is the dollar amount that a person would consider as neither a loss, nor a gain in a situation involving risk—in other words, a break-even value. We measured women’s and men’s reference point using several measures, including self-reported salary expectations, task-specific earnings expectations, and estimated the reference point from decisions in a lottery choice task. Across this multitude of reference point measurements, we consistently demonstrate that men have a higher reference point than women. We find that the reference point is positively correlated with financial comfort. We also find that risk tolerance is positively correlated with the reference point.

What does this mean for policy and practice?

We find evidence to suggest that common experience closes the gender reference point gap. This finding is particularly relevant to reducing the gender pay gap. By acknowledging that individuals may have disparate reference points in assessing their self-worth, organizations and policymakers can develop interventions that challenge existing gender-based biases. For instance, implementing transparent salary structures, conducting regular pay equity audits, and establishing standardized criteria for performance evaluations can contribute to creating a shared experience, which in turn may reduce the gender gap in risk attitudes by equalizing reference points and expectations.

Citation

Levy, J., Ru, O., & Tymula, A. (2024). 'The gender reference point gap: Evidence from a representative sample', Life Course Centre Working Paper Series, 2024-17. Institute for Social Science Research, The University of Queensland.

The authors

Jonathan Levy

University of Sydney

Email: jonathan.levy@sydney.edu.au

Olivia Ru

University of Sydney

Email: olru9317@uni.sydney.edu.au

Agnieszka Tymula

University of Sydney

Email: agnieszka.tymula@sydney.edu.au

Acknowledgements/Funding Sources

This work was supported by the Australian Research Council's Centre of Excellence for Children and Families over the Life Course (Project ID CE200100025) and ARC Discovery Grant (Project ID DP240101927).

DISCLAIMER: The content of this Working Paper does not necessarily reflect the views and opinions of the Life Course Centre. Responsibility for any information and views expressed in this Working Paper lies entirely with the author(s).

This work is licensed under a [Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License](https://creativecommons.org/licenses/by-nc-nd/4.0/).



The gender reference point gap: Evidence from a representative sample *

Jonathan Levy[†] Olivia Ru[‡] Agnieszka Tymula[‡]

Abstract

While women have been found to be more risk averse than men, there is no consensus regarding the underlying mechanism that drives these differences. Drawing inspiration from the efficient coding literature, we hypothesize that women's lower tolerance to financial risk is driven by lower reference points. We measured financial reference points in a US-representative sample of 579 adults using a range of unincentivized and incentivized survey questions, and additionally estimated it from behavior. In line with our predictions, we found that women have lower reference points, regardless of how they are measured, and that this translates to more risk-averse behavior. Our results suggest that, rather than being endowed with different risk attitudes, men and women may have different reference points. We discuss possible reasons for this and its implications for policy.

JEL Classification: C90, D87, D91, J16

Keywords: Reference point, risk attitude, neuroeconomics, gender, inequality, experiment

* This work was supported by the Australian Research Council's Centre of Excellence for Children and Families over the Life Course (Project ID CE200100025) and ARC Discovery Grant (Project ID DP240101927). We are grateful for valuable feedback from Alexander Svenson and research assistance from Sura Majeed.

[†]Corresponding author. University of Sydney, School of Economics, Camperdown NSW 2006, Australia. jonathan.levy@sydney.edu.au

[‡]University of Sydney, School of Economics, Camperdown NSW 2006, Australia.

1. Introduction

Many economists have found that women take fewer financial risks than men. There are fewer female entrepreneurs (Guzman and Kacperczyk 2019, Howell and Nanda 2019, Ewens 2022), fewer women in managerial positions (Bertrand and Hallock 2001, Blau and Kahn 2017, Eckel *et al.* 2021), fewer female investors (Ewens and Townsend 2020), and women shy away from initiating negotiations (Hernandez-Arenaz and Iriberry 2019, Babcock and Laschever 2021, Biasi and Sarsons 2021, Recalde and Vesterlund 2022) and risk in economic experiments (Eckel and Grossman 2008, Charness and Gneezy 2012). Women receive 34% less than their male counterparts in the financial industry due to their increased avoidance to risk (Goldin 2014). While the existing explanations of the gender difference in risk attitude generally rely on women being different than men, for example having a different emotional response to risk (Eckel and Grossman 2008), we propose that the reason why women take fewer financial risks than men is driven by a lower reference point that may be an efficient response to the reward distributions experienced by women.

In recent years, research in neuroeconomics has argued that people's decisions, including seemingly irrational ones that result in money left on the table, can be understood through the lens of efficient coding (e.g., Louie and Glimcher 2012, Louie *et al.* 2015, Khaw *et al.* 2017, Rustichini *et al.* 2017, Polanía *et al.* 2019, Landry and Webb 2021, Glimcher and Tymula 2023, Payzan-LeNestour *et al.* 2023, Robson *et al.* 2023). In essence, we could think of many previously considered choice anomalies as an efficient adjustment to the environment (Glimcher 2022, Page 2022). The key biological facts that motivate the efficient coding models are that people have limited neural resources to encode value and that they encode value with noise (Webb 2019). As a result, to minimize decision errors, our brains' value-

encoding is hardwired to constantly adjust to the distributions of experienced and expected payoffs.

In the context of risk-taking, previous investigations that build on the efficient coding hypothesis showed that changing payoff distributions affects risk attitudes measured in economic experiments in line with the models' predictions (Frydman and Jin 2021, Guo and Tymula 2021). The key idea is that the efficient utility function adapts to the reward statistics of the environment. To visualize this, consider Ann and Nick whose expected payoff distributions are depicted in the top panel of Figure 1. For simplicity we assume that these distributions have the same variance but differ in their mean with the mean payoff for Ann being lower than that for Nick. Ann and Nick's efficient subjective value (SV) functions, measured by the number of action potentials (y -axis) their brains produce in response to payoffs of different size (x -axis), are drawn in the bottom panel of Figure 1. First notice, that in line with biological constraints, the range of brain activity (y -axis) is bounded between SV_{min} and SV_{max} . Moreover, Ann's brain activity is centered at a lower value (RP_{Ann}) and her subjective value function (red, dashed) is optimized to better distinguish between smaller payoffs (the SV difference between x_2 and x_1 is large). For Ann distinguishing between higher payoffs ($x_4 - x_3$) is harder even though the difference between the payoffs is objectively the same (on the x -axis). Glimcher and Tymula (2023) call the payoff that corresponds to the midpoint of subjective value, a reference point (here RP_{Ann} is Ann's reference point and RP_{Nick} is Nick's). Crucially, like in the random utility framework, brain activity is stochastic. And this stochasticity is more likely to dominate over Ann's true preferences, leading to a decision mistake for high payoffs (x_3 or x_4) than for low payoffs (x_1 or x_2). Nick is the opposite. This simple example shows that Ann and Nick's subjective value functions reduce overall mistakes by improving discriminability between the most frequently encountered payoffs. But steepening the bounded SV in one area means it must be flatter, decreasing discriminability

and increasing the mistake rate elsewhere. Here, for the reward distributions we selected, the efficient coding results in an S-shaped SV function – steepest in the middle around the midpoint of the distribution, like a standard value function in Prospect Theory. Clearly, such efficient coding also affects risk attitudes creating another difference between Ann’s and Nick’s behavior. Consider a risky alternative that pays either x_2 or x_3 . These payoffs are on the concave part of Ann’s subjective value function, so she will avoid risk. Nick, on the other hand will seek risk because these payoffs are on the convex part of his subjective value function. In general, as the expected payoff distribution shifts to the right (i.e., as payoffs increase on average), the subjective value function shifts to the right, and this neuroeconomic reference point increases as well (compare RP_{Ann} and RP_{Nick}). In this example, due to the lower reference point, Ann is more likely to be more risk averse than Nick.




Figure 1: Probability density and subjective value functions for Ann and Nick

The idea that the amount of risk that an individual tolerates depends on the reference point is not new (Kahneman and Tversky 1979). The novelty is that the neuroeconomic framework provides a new lens to understand individual heterogeneity. It has been well documented that women experience different payoff distributions to men. Women receive lower wages (Roussille 2024), less recognition for the same work (Bohren *et al.* 2019, Coffman *et al.* 2021), and tend to engage less in financial decision-making on behalf of the household (Wagner and Walstad 2023). At the same time, they are more likely to take responsibility for household tasks that require decisions between low value items such as groceries which has been shown to

explain gender differences in inflation expectations (D'Acunto *et al.* 2020). In line with the efficient coding literature, one could then hypothesize that the experience of lower payoff distributions will result in women having a lower reference point. Given the established link between risk attitudes and the reference point, one could further hypothesize that the reason why women shy away from risk in the financial domain is because they have a lower financial reference point. This is precisely the hypothesis that we test in our study.

To test the hypothesis that women have a lower financial reference point, we conducted an incentivized experiment with 579 individuals, who are representative of the US population along several dimensions. We measured women's and men's financial reference point using several measures, including self-reported salary expectations, task-specific earnings expectations, and estimated the reference point from decisions in a lottery choice task using a canonical structural model from the efficient coding family (Glimcher and Tymula 2023). Across this multitude of reference point measurements, we consistently demonstrate that men have a higher financial reference point than women. In line with the intuition that reward distributions influence the reference point, we find that the reference point is positively correlated with financial comfort. Consistent with a variety of reference dependent models (e.g., Kahneman and Tversky 1979, Glimcher and Tymula 2023) we find that risk tolerance is positively correlated with the reference point.

We contribute to the literature in several ways. We provide a novel explanation for the gender difference in risk attitudes that is inspired by the emerging research on efficient coding (Louie and Glimcher 2012, Louie *et al.* 2015, Khaw *et al.* 2017, Rustichini *et al.* 2017, Polanía *et al.* 2019, Landry and Webb 2021, Glimcher and Tymula 2023, Payzan-LeNestour *et al.* 2023, Robson *et al.* 2023). As such, our explanation is rooted in how the brain actually processes value and makes choices, rendering it biologically realistic. Our explanation aids in

understanding why the gender gap in risk attitudes persists and how to translate this finding into practical and implementable policy interventions. Deciding whether and, if so, how to rectify gender differences in risk attitudes has been challenging using traditional economic approaches that treat risk attitudes as static primitives. If risk attitudes are fixed, forcing women to take more risk is not desirable because it would decrease their utility. However, if risk attitudes are not fixed, as has been demonstrated in many experiments and observational studies, we can think of ways to improve women's financial outcomes by encouraging them to take more risk. Before we start thinking about designing policy, we need a theory that explains shifts in risk attitudes. In this paper, we follow the logic of the efficient coding literature, and find evidence to suggest that gender differences in risk attitudes are driven by the reference point. This implies that by closing the gender reference point gap, for example by equalizing the reward distributions that men and women experience, we may also close the gender gap in risk attitudes. Indeed, we find evidence that after completing the same task and learning the distribution of possible rewards which is the same for both genders, the gender gap in task earnings expectations disappears. While our focus is on gender, the same logic would apply to any group that has historically experienced lower payoff distributions and leaves money on the table by making more risk averse decisions.

Our study is the first to compare a variety of reference point measures in one, large representative sample of the US population. A small number of studies (e.g., Terzi et al. 2016, Baillon et al. 2020, Rees-Jones and Wang 2022) estimated the usage of different reference point rules from behavior in lottery choice tasks and concluded that none of the so-far suggested reference point rules (like max-min, min-max, status quo etc.) are used by people all the time. Here, we use behavior to estimate a new type of reference point that is neurobiologically motivated and well-known to be calculated in the brain when people make decisions. Moreover, our sample is much bigger allowing us to study heterogeneity in this reference point.

Other papers elicited what could be considered another reference point measure by asking people to report their salary expectations (Reuben *et al.* 2017, Fernandes *et al.* 2021, Briel *et al.* 2022, Cortés *et al.* 2022). Unlike those before us we estimate reference points in multiple ways in one study. This includes both stated measures and measures garnered from choices made in a lottery choice task. This enables a novel examination of the correlation between different reference point measures. By doing so we can determine the validity of different reference point measures, something which should be useful for future research in this area. We find that reference points estimated from behavior are positively correlated with incentivized task earning expectations but not with unincentivized salary expectations.

The rest of this paper is structured as follows. Section 2 outlines our experimental design and empirical approach. Section 3 reports the results and section 4 concludes.

2. Methods

2.1. Ethics statement

Our protocols and procedures were approved by the University of Sydney Human Research Ethics Committee (HREC) (application number 2023/503). Two pilot studies ($n = 10$) were conducted prior to pre-registering the study to confirm our survey did not contain any errors. They were also used to establish the time limit for responses in the lottery choice task. Before commencing the main stage of data collection, the experiment was preregistered at the Open Science Framework (“[Closing the reference point gap](#)”: 2023-10-17 12:38 PM | Last Updated: Last Updated: 2023-10-18 04:08 PM).

2.2. Participants

Using the online recruitment platform Prolific, we recruited a representative sample of the adult American population. To generate a representative sample, Prolific divides the sample into three demographics of age, sex, and race and recruits from each subgroup in the same proportion as the national population according to 2022 US census data. Table A.1 in the Appendix displays the breakdown of the sample by age, sex, and race, and compares our sample to the American population according to 2022 US Census data. Our sample is representative of the American population.

Using G*power we conducted a power analysis and calculated a minimum sample size of $N = 578$ to detect a small effect size Cohens $d = 0.25$, $\alpha = 0.05$, $1 - \beta = 0.85$ for a gender difference in reference points for an independent sample two-sided t-test (Faul *et al.* 2007).

2.3. Tasks

For robustness we measured participants' reference points using three different tasks. These tasks are summarized below.

2.3.1. Salary expectations

At the end of the survey, we asked participants for their expected salary (in USD) in one year and in five years' time.

2.3.2. Lottery choice task

The lottery choice task consisted of 107 trials which were presented in a random order (see Table A.2 in the Appendix). In each trial participants were presented with a choice between a positive safe payoff and a risky lottery which offers a positive dollar amount or \$0. The position

of the two options – whether a lottery or a sure amount was seen on the left or right of the screen, was randomly determined on each trial. Participants were required to submit their choice for each trial within a 10 second time limit.¹ If a participant failed to submit their choice within the time limit or submitted an empty response, that trial would automatically be assigned a payoff of \$0. Therefore, it was in the best interest of the participant to submit a choice within the time limit. We enforced a time limit to ensure that participants did not take extended breaks between questions and to limit external interferences which may impact their preferences. Participants were not permitted to move backwards through the survey or change their response after submission. Participants were provided five practice trials to familiarize themselves with the 10 second time limit.

The payoff amounts and probabilities of the lotteries and sure amounts varied between trials and were all in the domain of gains. Participants were presented safe amount options of \$2, \$3, \$5, \$7, \$10, \$15, \$20, \$25, \$30, \$40, \$50, \$60, and \$70. Risky lotteries offered participants a chance (10%, 25%, 50%, 75%, and 90%) between a positive payoff (\$5, \$10, \$20, \$50, and \$90) or \$0. We included five trials where one option stochastically dominated the other option. These trials provided participants with a choice between a \$5 sure amount and a lottery that paid \$5 or \$0 with varying probabilities. A violation of first-order stochastic dominance occurred in these trials when an individual chose the lottery over the sure amount.

To ensure incentive compatibility and induce truth-telling, participants were informed that one of their own decisions in the lottery choice task would be randomly selected to be played, and one in ten participants would receive the resulting payoff amount as a bonus payment. This has

¹ From a pilot study without response time restrictions imposed on the lottery choice task, the average response time was approximately 5 seconds, and 95% of choices were made within 8 seconds. Therefore, we decided that a ten second response limit was not restrictive and provides ample time for participants to record a choice.

been argued to be an effective method of incentivizing behavior (Starmer and Sugden 1991, Cubitt *et al.* 1998, Bardsley *et al.* 2010, Baillon *et al.* 2020).

2.3.3. *Reservation wage and expectation of earnings in the lottery choice task*

All participants were asked to report the amount of money that would make them indifferent between participating and not participating in a typical one-hour-long Prolific study (i.e., their reservation wage). This question was unincentivized and was asked at the start of the experiment before participants received the instructions for the lottery choice task (see Figure 2).

After participants received the instructions for the lottery choice task, we asked them to predict the average amount of money they expect to earn in this task. To induce truth-telling, these questions were incentivized using a binarized scoring rule – an incentive-compatible belief elicitation method (Hossain and Okui 2013, Erkal *et al.* 2020). Under the binarized scoring rule participants' chance of receiving a bonus payment of \$1 increases with the accuracy of their predicted average earnings. A computer was used to randomly choose a number between 1 and 100 and if that number was less than their probability of receiving the payment, as determined by the binarized scoring rule the participant would receive a chance to earn the \$1 bonus payment. Participants were notified that one of their prediction trials would be randomly chosen to be paid. All participants were provided with a link located in the question text that would direct them to a separate web page explaining how the binarized scoring rule was used to calculate their probability of receiving the \$1 bonus payment. This question was asked twice, before and after the participant completed the lottery choice task. When this question was asked before the lottery choice task, participants were asked for a forward-looking prediction of their average earnings from all future lottery choice tasks. When asked after the lottery choice task,

participants were asked for a backward-looking prediction of their average earnings from all previous lottery choice tasks (see Figure 2).

2.4. Other procedural details

The study was programmed as an online questionnaire in Qualtrics. Figure 2 presents the sequence of tasks.

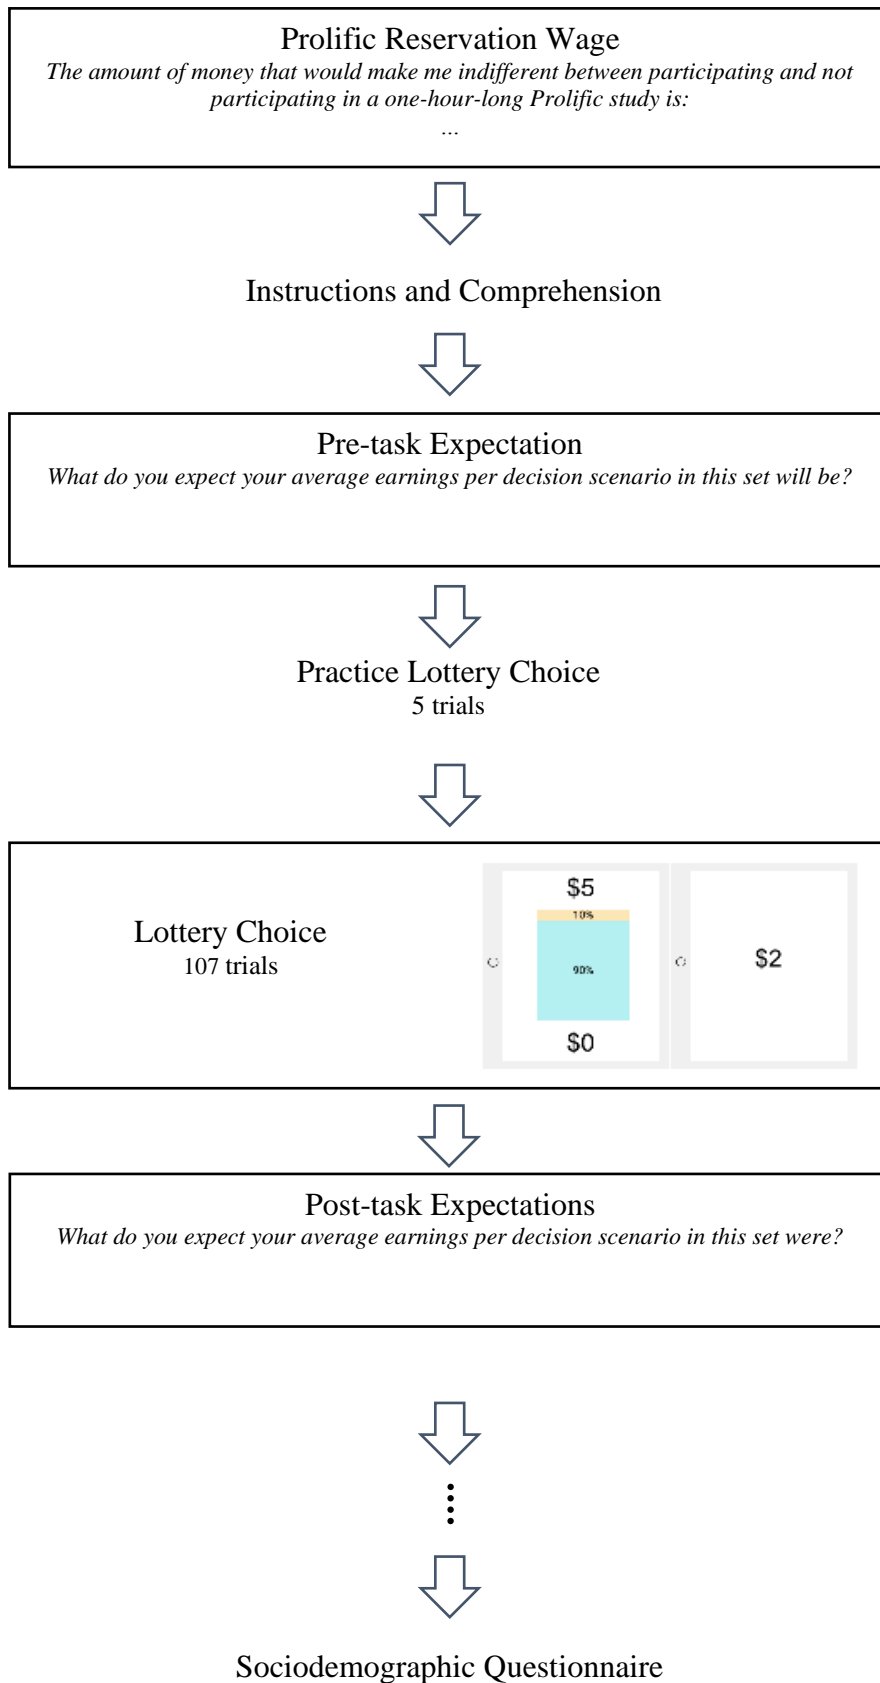


Figure 2: Experimental procedure outline

First, participants were provided with instructions outlining the nature of the task and presented with an example trial (see Appendix C). To avoid influencing participants' reported earnings expectations the example lotteries did not specify numerical payoffs. Instead, participants saw letters in the place of probabilities and payoff amounts (see Figure 3).




Figure 3: Example of a trial in the instructions

Following the main instructions, participants were required to answer three comprehension questions. After each incorrect answer participants received feedback on the correct answer. If a participant failed to answer one of the three questions correctly after two attempts, they had to withdraw from the study. If they passed the comprehension questions, they continued to the lottery choice task.²

Participants finished the session by completing a series of questionnaires about their demographics and socioeconomic factors (see Appendix D). Related to socioeconomic status, we measured financial comfort by asking participants to rate their financial situation on a 6-

² After completing the lottery choice task, participants completed additional tasks which are beyond the scope of this paper and the results will be reported in a separate paper. Participants did not know the details of those tasks and as such they should not affect their earlier decisions studied in this paper.

point Likert scale (0, very poor; 1, poor; 2, just getting along; 3, comfortable; 4, very comfortable; and 5, prosperous).

All participants received a participation fee of \$8. Additionally, one in ten participants were randomly selected to receive the bonus payment. Bonus payments were calculated by adding together the outcome of one randomly chosen trial from one of their lottery choices and one randomly chosen prediction from the expected earnings task. Possible bonus payments ranged from \$0 to \$91. Payments were sent within 7 days of completing the survey and paid through Prolific. All payments and dollar amounts listed in the survey were in USD.

2.5. Empirical approach to measuring the reference point

2.5.1. Stated measures of the reference point

We estimate the reference point using several stated measures which include expectations about salary in one years' time, expectations about salary in five years' time, their prolific reservation wage, pre- and post-task earning expectations. We use one-sided t-tests and OLS regressions to estimate gender differences in these reference point measures.

2.5.2. Reference point estimated from observed choice

We estimate the reference point using the recently proposed modeling approach in Expected Subjective Value Theory (ESVT) (Glimcher and Tymula 2023). ESVT is based on the neuroscientific understanding about how value signals are efficiently encoded in the brain. The main intuition behind ESVT is that the utility function adapts to the payoff expectation to efficiently encode value (Steverson *et al.* 2019, Bucher and Brandenburger 2022). Since the brain does not have unlimited resources (action potentials) to encode the utility of payoffs, it

adjusts dynamically so that the subjective value function³ is most sensitive to the payoff ranges that the brain is expecting to encounter. In this vein, the model is very similar to range normalization models (Padoa-Schioppa and Rustichini 2014, Kontek and Lewandowski 2018) and other models that originated from the efficient coding hypothesis (e.g., Polanía et al. 2019). ESVT has been shown to implement behaviors captured by Prospect Theory (PT), offering new interpretations for risk taking, reflection in risk attitudes, probability weighting, the endowment effect, and the Allais paradox (Glimcher and Tymula 2023). For our purposes, the biggest benefit of ESVT is its biological validity and the ease of estimating the reference point using the standard maximum likelihood procedure.

We assume that the utility of a payoff $x \in \mathbb{R}_+$ is given by:

$$u(x) = \frac{x}{x + M} \quad (1)$$

where M is the reference point. The utility function takes values between 0 and 1 ($u \in [0, 1]$) consistent with the idea that decision makers are bounded in the range of subjective values that they can biophysically assign to payoffs. Note that when $M = x$, the utility is equal to 0.5. This means that this function assigns the midpoint of its biologically restricted utility values to the reference point. (Rayo and Becker 2007, Woodford 2012, Robson *et al.* 2023).

We fit decisions of our participants with a logistic choice function, where the probability of choosing lottery A is:

$$P(A) = \frac{1}{1 + e^{-Z}} \quad (2)$$

where $Z = \frac{u(A) - u(B)}{\mu}$, and μ captures noise.

³ Neuroeconomists use the term “subjective value” to distinguish it from utility to capture that the former is usually thought of as cardinal and the latter ordinal. Throughout the paper we use the term utility as is the norm in economics.

The log-likelihood function is then given by:

$$LL(\boldsymbol{\theta}) = \sum_{n=1}^N \sum_{i=1}^I y_{ni} \ln(P_{ni}(A_i)) + (1 - y_{ni}) \ln(1 - P_{ni}(A_i)) \quad (3)$$

where N is the number of participants, I is the number of trials, $y_{ni} = 1(0)$ is an indicator function denoting the choice of lottery $A(B)$ for participant n in trial i , and $\boldsymbol{\theta}$ is the vector of behavioral parameters to be estimated.

To examine gender effects, we use a dummy variable which is equal to 1 if the participant is female and equal to 0 otherwise. For each parameter θ_n in our model, we specify:

$$\theta_n = \theta_0 + \theta_{Female} \times Female_n + X'_n \beta \quad (4)$$

where X is a vector of controls and β are the associated coefficients. For each parameter θ , we report the point estimate for the maximum likelihood estimation. The standard errors are clustered at the individual level.

2.6. Choice data quality

Inspecting the choice data, we find that participants were attentive to the decisions they made. Participants responded rationally to incentives and chose the risky lottery over the sure amount more often as the probability of receiving the payoff and the magnitude of the lottery payoff increased and less often as the magnitude of the sure amount increased (see Table A.3 in the Appendix).

Participants violated first order stochastic dominance on only 4.21% of the trials (122 out of 2895 trials on which one option dominated the other one). The average number of violations per individual was 0.211 (SD = 0.637) with a maximum possible value of 5. Participants also

consistently made decisions within the time limit⁴, missing on average less than one trial out of 107 (0.739, SD = 4.112), which is equivalent to 0.7% of total trials completed by participants. In our sample, 77% (443 out of 579) completed all trials. Among the participants who missed trials, the median number of missed trials was 1, and the mean number of missed trials was 3.147. The maximum number of missed trials made by a single participant was 78, with over 95% of participants missing 3 or fewer trials. Participants who violated FOSD more than 50% of the time or missed more than 20% of trials were excluded from the analysis that follows.

3. Results

3.1. Preliminary results

3.1.1. Sample demographic characteristics

We analyze data from 535 participants⁵ whose ages range from 19 to 79 years old (mean = 46.181, SD = 15.807). Our sample is well-balanced in terms of gender – 260 participants are male and 275 are female. Our sample is also well-balanced across gender in terms of their age and race (see Table 1) but compared to male participants female participants are significantly less likely to be employed and reported a significantly lower level of financial comfort (see Table 1). Additionally, female participants are significantly more likely to have completed tertiary education than male participants. We control for these differences in our analysis. Our sample is representative of the US population on gender, age, and race (see Table A.1 in the Appendix).

⁴ Missing responses occurred when a participant did not submit their choice within 10 seconds, or accidentally submitted a response before choosing an option.

⁵ We recruited 579 participants, of which 44 participants were excluded from the analysis. This includes: 7 participants who did not report a binary gender, 23 participants who after completing the lottery choice task reported incentivized earning expectations above the maximum possible amount of \$90, 2 participants who missed more than 20% of trials, and 12 participants who violated first-order stochastic dominance by choosing a lottery that pays at most \$5 over a sure payoff of \$5 more than 50% of the time.

Table 1: Descriptive statistics

	Mean Male	Mean Female	Difference	<i>p</i> -value
Age	46.231	46.135	0.096	0.944
White	0.735	0.716	0.018	0.637
Black	0.127	0.131	-0.004	0.891
Asian	0.054	0.065	-0.012	0.572
Hispanic	0.058	0.058	-0.000	0.981
Native American	0.000	0.007	-0.007	0.169
Employed	0.735	0.658	0.076	0.055
Financial	2.685	2.382	0.303	0.001
University	0.654	0.731	-0.077	0.054
Obs.	260	275	--	--

Note: Calculated from non-missing values from a sample of 260 males and 275 females. See Table A.4 in the Appendix for variable definitions. Two-sided *p*-values are presented.

3.1.2. Gender differences in risk attitudes

On average male participants chose the risky lottery more frequently than female participants – (30.31% versus 26.67% of the time, one-sided t-test: *p*-value = 0.013). Females also had a more concave, structurally estimated utility function (see Table 2, and estimation details in Appendix B). For men, the exponent in the power utility function was equal to 0.42. For women, it was significantly lower, by approximately 0.07-0.09 (see the coefficients of r_{Female} in Table 2). This result held both in the structural models without the probability weighting function (models (1) and (2)) and with the probability weighting function (models (3) and (4)).⁶ Gender differences persisted after we controlled for differences in socioeconomic variables (see Table 2 models (2) and (4)). Furthermore, model (2) in Table 2 indicated that participants who reported higher financial comfort, were younger, not employed, and not white, exhibited greater risk tolerance, denoted by a higher value of r .

⁶ We did not find gender differences in probability weighting (γ_{Female} is not significantly different from zero).

Table 2: Utility curvature estimates

	(1)	(2)	(3)	(4)
r_{Female}	-0.086*** (0.026)	-0.080*** (0.026)	-0.079*** (0.024)	-0.067*** (0.025)
r	0.419*** (0.018)	0.418*** (0.040)	0.417*** (0.017)	0.364*** (0.040)
Controls	No	Yes	No	Yes
γ_{Female}			-0.034 (0.066)	-0.065 (0.085)
γ			1.011*** (0.045)	1.322*** (0.178)
Controls			No	Yes
μ_{Female}	-0.131* (0.072)	-0.124* (0.070)	-0.127* (0.071)	-0.115* (0.069)
μ	0.737*** (0.054)	0.727*** (0.052)	0.736*** (0.054)	0.720*** (0.051)
Obs.	56976	56976	56976	56976
AIC	53653.470	53389.938	53655.130	53243.345
BIC	53689.271	53470.491	53708.832	53386.551

Notes: θ_{Female} is the mean difference between females and males in parameter θ . Controls include age, level of financial comfort, and indicator variables for white, employed, and completed tertiary education. Robust standard errors clustered on individual in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

3.2. Gender differences in the reference point

Consistent with the theoretical link between the reference point and risk tolerance, empirical literature finds that women tend to be more risk-averse than men (Eckel and Grossman 2008, Charness and Gneezy 2012) – with differences marked in less egalitarian societies (Olofsson and Rashid 2011, Liu and Zuo 2019, Friedl *et al.* 2020) – and expect lower earnings than men (Filippin and Ichino 2005, Briel *et al.* 2022). Therefore, we hypothesize that women will have a lower reference point than men.

Across all reference point measures, we found that women had a lower reference point than men (see Figure 4). Both 1-year and 5-year salary expectations were higher for men than for women (one-sided t-test: p -values are 0.0013 and 0.0302, respectively, see Figure 4 panel A).

One year from when the study was conducted women expected to make \$44,426, which is approximately 21% less than what men expected to make. The gender gap grew to approximately 31% in expected salary in five years' time with women expecting \$92,114 and men expecting \$132,799. Compared to male participants, female participants stated on average 16% lower Prolific reservation wages (one-sided t-test: p -value = 0.0028, see Figure 4 panel B). Moreover, before completing the lottery choice task and while not yet knowing the possible distribution of earnings, female participants reported 19% lower expected earnings from the lottery choice task than male participants (one-sided t-test: p -value = 0.0894).

In addition to the stated measures of the reference point, we estimated the reference point from the decisions that people made in the lottery choice task using the ESVT model. Overall, for the whole sample we obtain an estimate of \$6.02 for the reference point. Moving on to gender differences, we again find that men have a significantly higher reference point than women. The reference point for men was \$17.48, and for women was \$3.10 (p -value < 0.001, Figure 4 panel B). Controlling for gender differences in the reference point, improves the structural model fit – the loglikelihood changes from -28475.64 to -28283.32 and the BIC score changes from 56995.09 to 56665.19.




Figure 4: Gender gap in different reference point measures

3.3. Other determinants of the reference point

It is possible that other factors, in addition to gender, shape the reference point, or that other socioeconomic variables that differ between genders are the true reason why we observe men and women having different reference points. If the latter true, then controlling for such variables should reduce the gender gap in the reference point. As shown in Table 1, in our sample, women and men differed in the attained education level, employment status, and financial comfort. Therefore, it is entirely possible that some combination of these variables rather than gender per se is the reason for the gender gap in the reference point. To check this, we first use OLS estimations with the different stated reference point measures as dependent variables (see Table 3) and these key socioeconomic variables as independent variables. Then, we repeat this analysis using a structurally estimated reference point (see Table 4). Additionally, since we found a relationship between risk taking and age and race (unreported coefficients in the analysis in Table 2), we include these in our regressions even though these variables do not differ across gender.

Table 3: Determinants of the reference point

DV:	1-year salary expectation in \$K		5-year salary expectation in \$K		Reservation wage		Pre-task exp.	
Female	-11.553*** (3.829)	-7.217** (3.313)	-40.685* (21.621)	-28.538 (21.680)	-1.576*** (0.566)	-1.277** (0.577)	-1.504 (1.118)	-0.777 (1.128)
Age		0.195* (0.107)		-0.353 (0.698)		0.015 (0.019)		0.057 (0.036)
White		3.676 (3.710)		-23.440 (24.281)		0.261 (0.646)		-3.853*** (1.264)
University		20.235*** (3.867)		29.486 (25.306)		-0.509 (0.674)		-3.020** (1.317)
Employed		14.615*** (3.682)		34.230 (24.100)		1.353** (0.641)		1.897 (1.254)
Financial		15.499*** (1.696)		40.506*** (11.098)		0.497* (0.295)		1.369** (0.578)
Constant	55.980*** (2.745)	-21.321*** (7.296)	132.799*** (15.501)	13.167 (47.751)	10.086*** (0.406)	7.192*** (1.271)	8.094*** (0.801)	5.209** (2.485)
Obs.	535	535	535	535	535	535	535	535

Notes: Robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 4: Structural estimate of the reference point

DV: Reference point	(1) ESVT without controls	(2) ESVT with controls
Female	-14.385*** (3.750)	-11.473*** (4.413)
Age		-0.090** (0.045)
White		-3.758 (2.847)
Tertiary		-0.503 (1.349)
Employed		-0.387 (1.315)
Financial		0.674 (0.604)
Constant	17.482*** (3.687)	22.352*** (5.514)
μ_{Female}	0.128*** (0.020)	0.093*** (0.027)
μ	0.116*** (0.014)	0.119*** (0.016)
Obs.	56976	56976
AIC	56959.284	56584.637
BIC	56995.086	56665.190

Notes: Robust standard errors clustered on individual in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

The gender reference point gap persisted across most measures even when we included control variables. No matter which measure of the reference point we use, adding control variables reduced the magnitude of the estimated gender difference in the reference point as indicated by the size of the coefficients for the *Female* dummy variable. The gender gap in expected income one year from now reduced by \$4,336 to \$7,217 and in expected income five years from now by \$12,147 to \$28,538. The gender gap in the Prolific study reservation wage dropped by \$0.30 to \$1.28 and the gap in the expectations about earnings in our lottery choice task that participants are about to complete approximately halved. The difference in the reference point gap estimated from behavior dropped by \$2.91 to \$11.47 (see Table 4).

This general reduction in the gender reference point gap across all measures, suggests that part of our initially estimated gender-gap was due to socioeconomic differences between men and women.

Recent literature on efficient coding argues that individual utility functions adjust to experienced and expected payoff distributions (Woodford 2012, Polanía *et al.* 2019, Frydman and Jin 2021, Glimcher 2022, Page 2022, Robson *et al.* 2023). From an efficient coding perspective, individuals with histories of lower payoffs would as a result have a lower reference point. Consistent with this theoretical link, previous literature has shown that past experiences in the financial domain inform future expectations, and socioeconomically disadvantaged individuals are more pessimistic over expected outcomes (Kuhnen and Miu 2017) and typically more risk averse (Haushofer and Fehr 2014). Much research has demonstrated that women earn less for the same quality of work and work experience. We therefore hypothesize that the gender differences in reference points and risk attitudes are driven to some extent by differences in financial comfort.

In line with our hypothesis, for all stated reference point measures the coefficient estimates for financial comfort (*Financial*) were positive and significant. This suggests that those who are more prosperous tend to have a higher reference point. The fact that we see financial comfort having a significantly positive impact on the reference point, and that the gender difference in the reference point diminished when we controlled for it is consistent with men in our representative sample exhibiting higher financial comfort than women (see Table 1).

3.4. Pre- and post-task earnings expectations

We asked participants about their lottery choice task earning expectations twice – once before they completed the task and before they knew the distribution of possible earnings, and the second time immediately after they completed the task. As discussed earlier and illustrated in Figure 4, when considering a lottery choice task in which they did not know the payoff distributions, men expected to earn more than women. This is particularly striking because the accuracy of these earnings expectations was incentivized.

In general, these pre-task expectations did not accurately represent participants' real earnings. The mean absolute value of the difference between participant's true and expected earnings was \$17.73 (SD = 7.163). Participants expected to make less (\$7.32 on average) than what they would be really paid in expectation (\$21.51), consistent with our study having higher expected earnings than a typical Prolific study. We did not find a gender difference in the accuracy of pre-task expectations (one-sided t-test: p -value = 0.342).

A policy-relevant question is whether after completing the task, the gender difference in expected task earnings persisted or whether this common experience eliminated it. Prior research has provided evidence to suggest that risk attitudes adjust to changes in reward distributions (Guo and Tymula 2021, Payzan-LeNestour *et al.* 2023). To this end, we believe that exposing participants to a common lottery choice task will result in reference point convergence.

We found that after completing the task, participants were significantly more accurate in their earnings expectations with a mean absolute value of the difference between participant's true and expected earnings equal to \$12.73 (SD = 9.022), indicating that they correctly updated

their earnings expectations. On the aggregate level, earnings expectations increased by \$9.53, from \$7.32 to \$16.85 (one-sided paired t-test: p -value < 0.001). This upward shift in expectations from pre-task to post-task completion likely reflects participants learning and adjusting their expectations to the higher-than-expected payoff amounts. The deviation between post-task expectations and true earnings was significantly smaller for men than women (\$11.92 versus \$13.49, one-sided t-test: p -value = 0.0219).

Returning to the main question about the gender difference, we found that after participants completed the lottery choice task, the gender difference in men’s and women’s earning expectations reduced from \$1.50 to -\$0.20 and its significance disappeared (\$16.95 versus \$16.74, one-sided t-test: p -value = 0.562, see Figure 5). When we include the controls, the difference increases (with women reporting higher earnings expectations) but remains insignificant. This suggests that the common experience of completing the same task removed the gender gap in expected earnings.




Figure 5: Pre- and post-task earnings expectations

3.5. The relationship between risk attitudes and the reference point

Recall that we found evidence of gender differences in both risk attitudes and the reference point, which are robust across various estimation methods. A range of models predicts that observed tolerance to risk is partly driven by the reference point (Kahneman and Tversky 1979, Glimcher and Tymula 2023). In what follows we investigate whether risk tolerance is indeed correlated with the reference point.

To determine whether our stated reference point measures are correlated with risk tolerance we include them as the power utility curvature covariates in our structural model estimation (Table 5). The coefficient estimates indicate that 1-year salary, pre-task and post-task expectations were positively correlated with the utility curvature parameter, while the 5-year salary expectation and reservation wage were not. This means that both task-specific stated reference points correlate with the level of risk tolerance. This association between task earnings expectations and risk taking gets more significant after people completed the task but even when people do not know the task-specific distribution of possible earnings, those who expect to earn more, will subsequently take more risk. Interestingly, salary expectations one year from the study, but not the Prolific reservation wage, correlate with utility curvature.

Table 5: Relationship between utility curvature and stated reference points

DV: r	(1)	(2)	(3)	(4)	(5)
1-year salary	0.000* (0.000)				
5-year salary		0.000 (0.000)			
Reservation			-0.001 (0.001)		
Pre-task exp.				0.001* (0.001)	
Post-task exp.					0.001*** (0.000)
Constant	0.362*** (0.016)	0.374*** (0.014)	0.388*** (0.017)	0.369*** (0.014)	0.356*** (0.015)
μ	0.669*** (0.036)	0.669*** (0.036)	0.668*** (0.036)	0.670*** (0.036)	0.670*** (0.036)
Obs.	56976	56976	56976	56976	56976
AIC	53805.294	53828.239	53817.962	53804.713	53771.424
BIC	53832.145	53855.091	53844.813	53831.564	53798.275

Notes: Robust standard errors clustered on individual in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

3.6. Correlation across reference points

Next, we investigate the correlation between different reference point measures. Previous literature on risk attitudes found that self-reported risk attitudes differ across domains (Weber 2010) but also that behaviorally estimated risk attitudes are highly correlated across domains (Levy and Glimcher 2011, Cheung *et al.* 2022). In our experiment, all reference point measures were in the financial domain, however, they measure different financial aspects and were elicited using various methods, meaning that it is not obvious whether we should expect them to be correlated or not. We measured both the long-term reference point (salary in five-years' time) as well as the short-term reference point (earning expectations from a current task). We measured the reference point using stated measures and from behavior estimating a structural reference-dependent utility model.

First, we focus on the correlations between our stated measures of the reference point. Table 6 presents a correlation matrix. We found significant positive correlations between 1-year salary expectations and 5-year salary expectations. We also found significant positive correlations between the Prolific reservation wage and 1-year salary expectations. In terms of task earnings expectations, the expectations reported before completing the lottery choice task correlated with the Prolific reservation wage, but the post-task earnings expectations did not. Additionally, we found a significant positive correlation between pre- and post-task earnings expectations, meaning that even though people update their earnings expectations, those who initially expected to earn more in the lottery choice task also expected to earn more after completing the task.

Table 6: Correlation matrix of stated reference point measures

	1-year salary	5-year salary	Reservation	Pre-task exp.	Post-task exp.
1-year salary	1.000				
5-year salary	0.350***	1.000			
Reservation	0.095**	0.024	1.000		
Pre-task exp.	0.012	0.011	0.103**	1.000	
Post-task exp.	0.026	0.038	0.035	0.424***	1.000

Notes: Pearson's correlation coefficients. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Second, we check whether the behaviorally estimated reference point is correlated with each of the stated reference point measures. If our stated measures are capturing participants' behavioral reference point, we expect to find a positive correlation. To examine this, we estimate the ESVT model with stated reference point measures included as controls (see Table 7).

Table 7: Relationship between stated and behaviorally estimated reference points

DV: <i>M</i>	(1)	(2)	(3)	(4)	(5)
1-year salary	0.019 (0.019)				
5-year salary		0.013 (0.010)			
Reservation			-0.015 (0.076)		
Pre-task exp.				0.249** (0.118)	
Post-task exp.					0.173** (0.075)
Constant	5.141** (2.326)	4.662** (1.982)	6.164** (2.630)	5.190** (2.141)	4.111** (1.803)
Obs.	56976	56976	56976	56976	56976
AIC	57320.416	57229.385	57353.040	57126.677	57084.196
BIC	57347.267	57256.237	57379.891	57153.529	57111.047

Notes: Robust standard errors clustered on individual in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Pre- and post-task expectations are positively correlated with the behavioral reference point estimate. We did not find a correlation between the reference point estimated from behavior and the reservation wage, 1-year salary expectation, and 5-year salary expectation. Overall, we conclude that not all financial reference points are correlated. However, the closer they are in their economic meaning, the more likely the correlation is to emerge. For example, all measures of the reference point related to the lottery choice task are correlated, no matter whether we measure it by eliciting task earnings expectations or by estimating it from behavior on the same task. Stated reference points that are not task specific (like income expectations or the reservation wage) correlate with one another but are unlikely to be good proxies for a task-specific behavioral reference point.

4. Conclusion

A multitude of reasons have been suggested for why women are more risk averse than men. Evolutionary economic models suggest that risk seeking behavior leads to greater reproductive success among men (Robson 1996, Dekel and Scotchmer 1999, Niederle and Vesterlund 2007). More contemporary theories direct attention to environmental and social factors. For example, stereotype threat effects (Carr and Steele 2010), differences in socialization between genders and patriarchal cultural norms (Liu and Zuo 2019, Andersen *et al.* 2022). What all these explanations have in common is the idea that our behavior is the best response to the environment we are exposed to. Our explanation for why there may be gender differences in risk attitudes complements this line of reasoning. Reference points may depend on societal norms and environmental factors. It is entirely possible that women have lower expectations as they have been historically given less opportunities to succeed, and this gender difference feeds into the gender difference in risk attitudes.

Following this line of thought, the general reduction in the gender reference point gap across all measures after controlling for socioeconomic factors, suggests that part of our initially estimated gender-gap was due to socioeconomic differences between men and women. We found those with a higher degree of financial comfort had a significantly higher reference point. This finding is consistent with previous literature which has shown that past experiences in the financial domain inform future expectations, and socioeconomically disadvantaged individuals are more pessimistic over expected outcomes (Kuhnen and Miu 2017) and typically more risk averse (Haushofer and Fehr 2014).

This paper focused on gender disparities in risk attitudes and reference points. However, the implications of our findings extend beyond gender to encompass any demographic facing economic disparities. While the notion that risk attitudes are malleable rather than fixed is not

entirely novel (Post *et al.* 2008, Malmendier and Nagel 2011, Imas 2016), its application to enrich our comprehension of economic disadvantage remains underexplored. Haushofer and Fehr (2014) argue that economic inequality sets in motion a self-reinforcing pattern – poverty breeds heightened stress levels, fostering impatience and risk aversion, ultimately leading to decisions yielding diminished expected returns. By illustrating the gender example, we propose a novel mechanism through which poverty perpetuates itself – by establishing a lower reference point.

The efficient coding literature provides a promising approach for contextualizing and understanding our research findings. It has been well established both theoretically and empirically that the recently experienced distributions of rewards have a measurable impact on choice by influencing the encoding of value in the brain (Kobayashi *et al.* 2010, Louie *et al.* 2013, Rustichini *et al.* 2017, Zimmermann *et al.* 2018, Guo and Tymula 2021, Glimcher and Tymula 2023). Efficient coding, while generally beneficial for decision-making processes, can sometimes lead to suboptimal outcomes. For instance, shifting reward distributions can alter the perception of risk and reward, leading to suboptimal financial decisions (Payzan-LeNestour *et al.* 2016, Frydman and Jin 2021). By optimizing for efficient value coding in the historically expected reward distributions, efficient coding also makes people prone to mistakes and suboptimal decisions when choices are unexpected and surprising. From a policy perspective, the implication seems to be that financially suboptimal decisions may not be a manifestation of a preference or lack of ability and that to make the best decisions possible choosers need to be well adapted to the distributional properties of the decision environment.

Consistently, we find evidence to suggest that common experience closes the gender reference point gap. This finding is particularly relevant to reducing the gender pay gap. By acknowledging that individuals may have disparate reference points in assessing their self-

worth, organizations and policymakers can develop interventions that challenge existing gender-based biases. For instance, implementing transparent salary structures, conducting regular pay equity audits, and establishing standardized criteria for performance evaluations can contribute to creating a shared experience, which in turn may reduce the gender gap in risk attitudes by equalizing reference points and expectations. In this way, leveraging insights from the convergence of reference points can play a crucial role in designing targeted interventions to address and mitigate gender-based disparities in earnings.

We included several covariates relating to sociodemographic factors, but we were not able to eliminate the gender gap in reference points. This leaves the question of what really drives the gender reference point gap partially unanswered. To fully answer this question, it may be fruitful for future research to identify covariates that may be relevant which were not included in our survey. Recent research has shown how gender norms may play a pivotal role in generating gender inequality (Exley and Kessler 2022, Coffman *et al.* 2023, Gangadharan *et al.* 2024). Therefore, it may be interesting for future research to explicitly control for the extent to which participants internalize or conform to gendered social norms.

References

- Andersen, S., Badarinza, C., Liu, L., Marx, J., and Ramadorai, T., 2022. Reference Dependence in the Housing Market. *American Economic Review*, 112 (10), 3398–3440.
- Babcock, L. and Laschever, S., 2021. *Women Don't Ask: Negotiation and the Gender Divide*. Princeton University Press.
- Baillon, A., Bleichrodt, H., and Spinu, V., 2020. Searching for the Reference Point. *Management Science*, 66 (1), 93–112.
- Bardsley, N., Cubitt, R., Loomes, G., Moffatt, P., Starmer, C., and Sugden, R., 2010. *Experimental Economics*. Princeton University Press.
- Bertrand, M. and Hallock, K.F., 2001. The gender gap in top corporate jobs. *ILR Review*, 55 (1), 3–21.
- Biasi, B. and Sarsons, H., 2021. Flexible Wages, Bargaining, and the Gender Gap. *The Quarterly Journal of Economics*, 137 (1), 215–266.
- Blau, F.D. and Kahn, L.M., 2017. The Gender Wage Gap: Extent, Trends, and Explanations. *Journal of Economic Literature*, 55 (3), 789–865.
- Bohren, J.A., Imas, A., and Rosenberg, M., 2019. The Dynamics of Discrimination: Theory and Evidence. *American Economic Review*, 109 (10), 3395–3436.
- Briel, S., Osikominu, A., Pfeifer, G., Reutter, M., and Satlukal, S., 2022. Gender differences in wage expectations: the role of biased beliefs. *Empirical Economics*, 62 (1), 187–212.
- Bucher, S.F. and Brandenburger, A.M., 2022. Divisive normalization is an efficient code for multivariate Pareto-distributed environments. *Proceedings of the National Academy of Sciences*, 119 (40).
- Carr, P.B. and Steele, C.M., 2010. Stereotype Threat Affects Financial Decision Making. *Psychological Science*, 21 (10), 1411–1416.
- Charness, G. and Gneezy, U., 2012. Strong Evidence for Gender Differences in Risk Taking. *Journal of Economic Behavior & Organization*, 83 (1), 50–58.
- Cheung, S.L., Tymula, A., and Wang, X., 2022. Present bias for monetary and dietary rewards. *Experimental Economics*, 25 (4), 1202–1233.
- Coffman, K.B., Collis, M.R., and Kulkarni, L., 2023. Whether to Apply. *Management Science*.
- Coffman, K.B., Exley, C.L., and Niederle, M., 2021. The Role of Beliefs in Driving Gender Discrimination. *Management Science*, 67 (6), 3551–3569.
- Cortés, P., Pan, J., Reuben, E., Pilosoph, L., and Zafar, B., 2022. Gender Differences in Job Search and the Earnings Gap: Evidence from the Field and Lab. *The Quarterly Journal of Economics*.
- Cubitt, R.P., Starmer, C., and Sugden, R., 1998. On the Validity of the Random Lottery Incentive System. *Experimental Economics*, 1 (2), 115–131.
- D'Acunto, F., Malmendier, U., and Weber, M., 2020. *Gender Roles and the Gender Expectations Gap*. Cambridge, MA.
- Dekel, E. and Scotchmer, S., 1999. On the Evolution of Attitudes towards Risk in Winner-Take-All Games. *Journal of Economic Theory*, 87 (1), 125–143.
- Eckel, C., Gangadharan, L., J. Grossman, P., and Xue, N., 2021. The gender leadership gap: insights from experiments. In: *A Research Agenda for Experimental Economics*. Edward Elgar Publishing.
- Eckel, C.C. and Grossman, P.J., 2008. Chapter 113 Men, Women and Risk Aversion: Experimental Evidence. 1061–1073.
- Erkal, N., Gangadharan, L., and Koh, B.H., 2020. Replication: Belief elicitation with quadratic and binarized scoring rules. *Journal of Economic Psychology*, 81, 102315.

- Ewens, M., 2022. *Race and Gender in Entrepreneurial Finance*. Cambridge, MA.
- Ewens, M. and Townsend, R.R., 2020. Are early stage investors biased against women? *Journal of Financial Economics*, 135 (3), 653–677.
- Exley, C.L. and Kessler, J.B., 2022. The Gender Gap in Self-Promotion. *The Quarterly Journal of Economics*, 137 (3), 1345–1381.
- Faul, F., Erdfelder, E., Lang, A.-G., and Buchner, A., 2007. G*Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior Research Methods*, 39 (2), 175–191.
- Fernandes, A., Huber, M., and Vaccaro, G., 2021. Gender differences in wage expectations. *PLOS ONE*, 16 (6), e0250892.
- Filippin, A. and Ichino, A., 2005. Gender wage gap in expectations and realizations. *Labour Economics*, 12 (1), 125–145.
- Friedl, A., Pondorfer, A., and Schmidt, U., 2020. Gender differences in social risk taking. *Journal of Economic Psychology*, 77, 102182.
- Frydman, C. and Jin, L.J., 2021. Efficient Coding and Risky Choice. *The Quarterly Journal of Economics*, 137 (1), 161–213.
- Gangadharan, L., Gertsberg, M., Levy, J., and Xiao, E., 2024. *Who wants to move first?*
- Glimcher, P.W., 2022. Efficiently irrational: deciphering the riddle of human choice. *Trends in Cognitive Sciences*, 26 (8), 669–687.
- Glimcher, P.W. and Tymula, A.A., 2023. Expected subjective value theory (ESVT): A representation of decision under risk and certainty. *Journal of Economic Behavior & Organization*, 207, 110–128.
- Goldin, C., 2014. A Grand Gender Convergence: Its Last Chapter. *American Economic Review*, 104 (4), 1091–1119.
- Guo, J. and Tymula, A., 2021. Waterfall illusion in risky choice – exposure to outcome-irrelevant gambles affects subsequent valuation of risky gambles. *European Economic Review*, 139, 103889.
- Guzman, J. and Kacperczyk, A. (Olenka), 2019. Gender gap in entrepreneurship. *Research Policy*, 48 (7), 1666–1680.
- Haushofer, J. and Fehr, E., 2014. On the psychology of poverty. *Science*, 344 (6186), 862–867.
- Hernandez-Arenaz, I. and Iriberry, N., 2019. A Review of Gender Differences in Negotiation. *In: Oxford Research Encyclopedia of Economics and Finance*. Oxford University Press.
- Hossain, T. and Okui, R., 2013. The binarized scoring rule. *Review of Economic Studies*, 80 (3).
- Howell, S. and Nanda, R., 2019. *Networking Frictions in Venture Capital, and the Gender Gap in Entrepreneurship*. Cambridge, MA.
- Imas, A., 2016. The Realization Effect: Risk-Taking After Realized Versus Paper Losses. *American Economic Review*, 106 (8), 2086–2109.
- Kahneman, D. and Tversky, A., 1979. Prospect Theory: An Analysis of Decision under Risk. *Econometrica*, 47 (2), 263.
- Khaw, M.W., Glimcher, P.W., and Louie, K., 2017. Normalized value coding explains dynamic adaptation in the human valuation process. *Proceedings of the National Academy of Sciences*, 114 (48), 12696–12701.
- Kobayashi, S., Pinto de Carvalho, O., and Schultz, W., 2010. Adaptation of Reward Sensitivity in Orbitofrontal Neurons. *The Journal of Neuroscience*, 30 (2), 534–544.
- Kontek, K. and Lewandowski, M., 2018. Range-Dependent Utility. *Management Science*, 64 (6), 2812–2832.
- Koszegi, B. and Rabin, M., 2006. A Model of Reference-Dependent Preferences. *The Quarterly Journal of Economics*, 121 (4), 1133–1165.

- Kuhnen, C.M. and Miu, A.C., 2017. Socioeconomic status and learning from financial information. *Journal of Financial Economics*, 124 (2), 349–372.
- Landry, P. and Webb, R., 2021. Pairwise normalization: A neuroeconomic theory of multi-attribute choice. *Journal of Economic Theory*, 193, 105221.
- Levy, D.J. and Glimcher, P.W., 2011. Comparing Apples and Oranges: Using Reward-Specific and Reward-General Subjective Value Representation in the Brain. *The Journal of Neuroscience*, 31 (41), 14693–14707.
- Liu, E.M. and Zuo, S.X., 2019. Measuring the impact of interaction between children of a matrilineal and a patriarchal culture on gender differences in risk aversion. *Proceedings of the National Academy of Sciences*, 116 (14), 6713–6719.
- Louie, K. and Glimcher, P.W., 2012. Efficient coding and the neural representation of value. *Annals of the New York Academy of Sciences*, 1251 (1), 13–32.
- Louie, K., Glimcher, P.W., and Webb, R., 2015. Adaptive neural coding: from biological to behavioral decision-making. *Current Opinion in Behavioral Sciences*, 5, 91–99.
- Louie, K., Khaw, M.W., and Glimcher, P.W., 2013. Normalization is a general neural mechanism for context-dependent decision making. *Proceedings of the National Academy of Sciences*, 110 (15), 6139–6144.
- Malmendier, U. and Nagel, S., 2011. Depression Babies: Do Macroeconomic Experiences Affect Risk Taking?*. *The Quarterly Journal of Economics*, 126 (1), 373–416.
- Niederle, M. and Vesterlund, L., 2007. Do Women Shy Away From Competition? Do Men Compete Too Much? *The Quarterly Journal of Economics*, 122 (3), 1067–1101.
- Olofsson, A. and Rashid, S., 2011. The White (Male) Effect and Risk Perception: Can Equality Make a Difference? *Risk Analysis*, 31 (6), 1016–1032.
- Padoa-Schioppa, C. and Rustichini, A., 2014. Rational Attention and Adaptive Coding: A Puzzle and a Solution. *American Economic Review*, 104 (5), 507–513.
- Page, L., 2022. *Optimally irrational: The good reasons we behave the way we do*. . Cambridge University Press.
- Payzan-LeNestour, E., Balleine, B.W., Berrada, T., and Pearson, J., 2016. Variance After-Effects Distort Risk Perception in Humans. *Current Biology*, 26 (11), 1500–1504.
- Payzan-LeNestour, E., Pradier, L., and Putniņš, T.J., 2023. Biased risk perceptions: Evidence from the laboratory and financial markets. *Journal of Banking & Finance*, 154, 106685.
- Polanía, R., Woodford, M., and Ruff, C.C., 2019. Efficient coding of subjective value. *Nature Neuroscience*, 22 (1), 134–142.
- Post, T., van den Assem, M.J., Baltussen, G., and Thaler, R.H., 2008. Deal or No Deal? Decision Making under Risk in a Large-Payoff Game Show. *American Economic Review*, 98 (1), 38–71.
- Prelec, D., 1998. The Probability Weighting Function. *Econometrica*, 66 (3), 497.
- Rayo, L., and Becker, G.S., 2007. Evolutionary Efficiency and Happiness. *Journal of Political Economy*, 115 (2), 302–337.
- Recalde, M.P. and Vesterlund, L., 2022. Gender Differences in Negotiation and Policy for Equalizing Outcomes. In: *Bargaining*. Cham: Springer International Publishing, 455–475.
- Rees-Jones, A. and Wang, A., 2022. *An Approach to Testing Reference Points*. Cambridge, MA.
- Reuben, E., Wiswall, M., and Zafar, B., 2017. Preferences and Biases in Educational Choices and Labour Market Expectations: Shrinking the Black Box of Gender. *The Economic Journal*, 127 (604), 2153–2186.
- Robson, A.J., 1996. The Evolution of Attitudes to Risk: Lottery Tickets and Relative Wealth. *Games and Economic Behavior*, 14 (2), 190–207.

- Robson, A.J., Whitehead, L.A., and Robalino, N., 2023. Adaptive utility. *Journal of Economic Behavior & Organization*, 211, 60–81.
- Roussille, N., 2024. The Role of the Ask Gap in Gender Pay Inequality. *The Quarterly Journal of Economics*.
- Rustichini, A., Conen, K.E., Cai, X., and Padoa-Schioppa, C., 2017. Optimal coding and neuronal adaptation in economic decisions. *Nature Communications*, 8 (1), 1208.
- Starmer, C. and Sugden, R., 1991. Does the Random-Lottery Incentive System Elicit True Preferences? An Experimental Investigation. *American Economic Review*, 81 (4), 971–978.
- Steverson, K., Brandenburger, A., and Glimcher, P., 2019. Choice-theoretic foundations of the divisive normalization model. *Journal of Economic Behavior & Organization*, 164, 148–165.
- Terzi, A., Koedijk, K., Noussair, C.N., and Pownall, R., 2016. Reference Point Heterogeneity. *Frontiers in Psychology*, 7.
- Wagner, J. and Walstad, W.B., 2023. Gender Differences in Financial Decision-Making and Behaviors in Single and Joint Households. *The American Economist*, 68 (1), 5–23.
- Webb, R., 2019. The (Neural) Dynamics of Stochastic Choice. *Management Science*, 65 (1), 230–255.
- Weber, E.U., 2010. On the coefficient of variation as a predictor of risk sensitivity: Behavioral and neural evidence for the relative encoding of outcome variability. *Journal of Mathematical Psychology*, 54 (4), 395–399.
- Woodford, M., 2012. Prospect Theory as Efficient Perceptual Distortion. *American Economic Review*, 102 (3), 41–46.
- Zimmermann, J., Glimcher, P.W., and Louie, K., 2018. Multiple timescales of normalized value coding underlie adaptive choice behavior. *Nature Communications*, 9 (1), 3206.

Appendix

A. Tables

Table A.1: Sample characteristics compared to US Census data

Variable	Proportion of total sample	Proportion of final sample	US Census (2022)
Female	0.501	0.514	0.504
Male	0.487	0.486	0.496
Age			
19 – 24	0.098	0.101	0.089
25 – 34	0.187	0.181	0.179
35 – 44	0.193	0.194	0.174
45 – 54	0.171	0.178	0.160
55 – 64	0.207	0.198	0.167
> 64	0.143	0.148	0.229
White	0.715	0.725	0.755
Black	0.142	0.129	0.136
Asian	0.057	0.060	0.063
Hispanic	0.057	0.058	0.191
Native American	0.003	0.004	0.130

Table A.2: List of all decision scenarios. Each row is one decision scenario in which the participant is choosing between a lottery that pays $\$x_{risky}$ with probability p_{risky} and a sure payout of $\$x_{sure}$.

	p_{risky}	x_{risky}	x_{sure}
1	0.1	5	2
2	0.1	5	3
3	0.1	5	5
4	0.1	10	5
5	0.1	10	7
6	0.1	20	7
7	0.1	20	10
8	0.1	20	15
9	0.1	50	20
10	0.1	50	25
11	0.1	50	30
12	0.1	50	40
13	0.1	90	30
14	0.1	90	40
15	0.1	90	50
16	0.1	90	60
17	0.1	90	70
18	0.25	5	2
19	0.25	5	3
20	0.25	5	5
21	0.25	10	3
22	0.25	10	5
23	0.25	10	7
24	0.25	20	7
25	0.25	20	10
26	0.25	20	15
27	0.25	50	15
28	0.25	50	20
29	0.25	50	25
30	0.25	50	30
31	0.25	50	40
32	0.25	90	25
33	0.25	90	30
34	0.25	90	40
35	0.25	90	50
36	0.25	90	60
37	0.25	90	70
38	0.5	5	2

39	0.5	5	3
40	0.5	5	5
41	0.5	10	2
42	0.5	10	3
43	0.5	10	5
44	0.5	10	7
45	0.5	20	5
46	0.5	20	7
47	0.5	20	10
48	0.5	20	15
49	0.5	50	10
50	0.5	50	15
51	0.5	50	20
52	0.5	50	25
53	0.5	50	30
54	0.5	50	40
55	0.5	90	20
56	0.5	90	25
57	0.5	90	30
58	0.5	90	40
59	0.5	90	50
60	0.5	90	60
61	0.5	90	70
62	0.75	5	2
63	0.75	5	3
64	0.75	5	5
65	0.75	10	2
66	0.75	10	3
67	0.75	10	5
68	0.75	10	7
69	0.75	20	2
70	0.75	20	3
71	0.75	20	5
72	0.75	20	7
73	0.75	20	10
74	0.75	20	15
75	0.75	50	5
76	0.75	50	7
77	0.75	50	10
78	0.75	50	15
79	0.75	50	20
80	0.75	50	25
81	0.75	50	30

82	0.75	90	10
83	0.75	90	15
84	0.75	90	20
85	0.75	90	25
86	0.75	90	30
87	0.75	90	40
88	0.75	90	50
89	0.75	90	60
90	0.9	5	5
91	0.9	10	2
92	0.9	10	3
93	0.9	20	2
94	0.9	20	3
95	0.9	20	5
96	0.9	50	2
97	0.9	50	3
98	0.9	50	5
99	0.9	50	7
100	0.9	50	10
101	0.9	50	15
102	0.9	90	5
103	0.9	90	7
104	0.9	90	10
105	0.9	90	15
106	0.9	90	20
107	0.9	90	25

Table A.3: Regression of lottery characteristics on probability of choosing lottery

DV: Choose lottery	(1)	(2)
	OLS	Logit
x_{risky}	0.004*** (0.000)	0.024*** (0.001)
p_{risky}	0.835*** (0.022)	4.767*** (0.149)
x_{sure}	-0.008*** (0.000)	-0.049*** (0.002)
Constant	-0.138*** (0.010)	-3.668*** (0.117)
Obs.	56976	56976

Robust standard errors clustered at the individual level in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A.4: Variable definitions

Variable	Definition
Age	Age at last birthday
Female	= 1 if female, = 0 if male
White	= 1 if White, = 0 otherwise
Black	= 1 if Black, = 0 otherwise
Asian	= 1 if Asian, = 0 otherwise
Hispanic	= 1 if Hispanic, = 0 otherwise
Native American	= 1 if Native American, = 0 otherwise
Employed	= 1 if employed, = 0 otherwise
Financial	= 0 if Very poor, = 1 if Poor, = 2 if Just getting along, = 3 if Comfortable, = 4 if Very comfortable, = 5 if Prosperous
University	= 1 if they are university educated, = 0 otherwise

B. Estimating risk attitudes

We measure risk attitudes using multiple methods. Our first measure of risk attitudes is non-parametric and involves calculating simple counts in which participants choose the lottery over the sure amount. A participant who chooses the lottery more frequently is more risk tolerant. To examine the gender effect, we first identify whether there are gender differences in the proportion of risky choices made. To control for socioeconomic factors and other confounding factors, we run a simple OLS model, using the following specification:

$$y_i = \beta_0 + \beta_1 Female_i + \beta_2 X_i + e_i \quad (\text{B. 1})$$

where y_i is a dummy variable equal to 1 if participant i chooses the lottery, and equal to 0 if participant i chooses the sure amount. $Female_i$ is a dummy variable equal to 1 if participant i is female, and equal to 0 if participant i is male. β_1 represents the difference in the frequency with which men and women choose the lottery. X_i is a vector of controls including age, race, employment status, self-reported financial comfort, the gender discrimination scale score, the perceived discrimination scale score, and whether someone completed tertiary education. e_i is an error term. We use robust standard errors, which are clustered at the individual level.

Our second method of estimating risk attitudes involves structural estimation. We structurally estimate a power utility function defined as:

$$u(p, x) = px^r \quad (\text{B. 2})$$

where $r < 1$ indicates a risk-averse individual, $r = 1$ indicates a risk-neutral individual, and $r > 1$ indicates a risk-loving individual.

We estimate the model separately without probability weighting (i.e., $w(p) = p$) and with probability weighting using the one-parameter specification (i.e., $w(p) = e^{-(-\ln(p))^\gamma}$) as proposed by (Prelec 1998), where γ determines the shape of the probability weighting function and is bounded between 0 and 1. As γ approaches 0, this causes a more inverse S-shaped function (greater over (under)-weighting of low (high) probabilities).

C. Instructions for lottery choice task

Task Instructions

You will complete tasks which provide you with a chance to earn additional money. At the end of the study, you will also complete several questionnaires.

How much will I get paid?

If you complete the study, you will receive a participation fee of **\$8 for sure**.

In addition, we will randomly select **one out of every ten** participants to receive a bonus payment. If you are one of the randomly selected participants, you will be paid one of your choices in the lottery choice task and one of your predictions in the prediction task.

Payment will be made through Prolific.


Understanding the lottery choice task

In the lottery choice task, you will make 107 decisions between lotteries shown on the screen.

There are no wrong decisions because different people have different preferences.

Understanding the sure amount display

We will show the sure amount as in the example below:




If you pick this option, you will receive \$X for sure (with 100% chance).

Understanding the lottery choice task

Understanding the lottery display

We will show lotteries as in the example below:




If you pick this lottery, you will earn one of the two dollar amounts which are written above and below the yellow-and-blue rectangle. In this case you would earn either \$Y or \$W. This means that if you pick the lottery, your outcome is not certain. You know you would get \$Y or \$W but you are not sure which one. In the experiment, the letters will be replaced with numbers of dollars.

The probability of receiving each amount of money is written inside the rectangle and is proportional to the size of the corresponding colored area.

- The chance with which you would receive the amount on the top (\$Y) is written in the top and yellow part of the rectangle (p%)
- The chance with which you would receive the amount on the bottom (\$W) is written in the bottom and blue part of the rectangle (100%-p%)

In the task, the probabilities will be real numbers that will change from one decision scenario to another.

Click next to see some examples to get a better understanding.



First notice, that both lotteries above pay either \$Y or \$W.

Second, have a look at the lottery on the left only. The yellow box that indicates the probability of getting \$Y is smaller than the blue box that indicates the probability of getting \$W. This means that if you pick the lottery on the left, the chance of receiving \$Y is smaller than the chance of \$W. The precise probabilities are written inside the colored boxes – 10% chance of \$Y and 90% chance of \$W.


Third, compare the lottery on the left with the lottery on the right. In the lottery on the right, the yellow box is larger than the blue box, indicating that the chance of receiving \$Y (90%) is greater than the chance of receiving \$W (10%).

In the lottery choice task, you will choose between lotteries and sure amounts. You will indicate your choice by clicking on your preferred option and pressing continue. For each decision

scenario you will have 10 seconds to indicate and submit your choice by pressing the next button. After you press the next button, you will see the next decision scenario.

If a decision scenario in which you did not make a choice and press the next button within 10 seconds gets picked for payment, your payoff for that decision scenario will be \$0. Therefore, it is in your best interest to make your choice within 10 seconds.

Example of a decision scenario



In the above hypothetical example of a decision scenario, you can choose between:

- a lottery which pays \$Y $p\%$ of the time and \$W $100\%-p\%$ of the time (on the left), and
- receiving \$X for sure (on the right).

In the task, the amounts and probabilities will be real numbers that will change from one decision scenario to another.

You will complete sets of the lottery choice task. Then, for **one out of ten** participants, **one decision scenario** will be **randomly selected**. If this happens for you, your choice in this

decision scenario will be realized for payment. Only one of your decisions made throughout the study can be selected for payment. This means that it is in your best interest to treat every decision as independent from the other choices you make. In other words, whenever you face a decision scenario, you should decide as if this is the only decision that you are making.

Payment example



Suppose that this example decision scenario was randomly selected for payment. How much would you earn?

There are three possible cases:

1. If you picked the option on the left, the computer would play the lottery. It would randomly pick a number between 0 and 100. If the number is smaller or equal to p , you would get $\$Y$, if it is larger than p , you would get $\$W$.
2. If you picked the option on the right, you would get $\$X$ for sure.

3. If you did not make a choice and press the next button within 10 seconds, you would get \$0 for sure.

Once you feel you understand each component of the visual lottery display, please click next to answer comprehension questions.

D. Sociodemographic questionnaire

1. Which state do you currently live in?
2. What is your ZIP Code?
3. How many people reside in your household (including yourself)?
4. Is English your first language?
 - a. Yes
 - b. No
5. Which cultural heritage best describes you?
 - a. Hispanic or Latinx
 - b. White alone
 - c. Black or African American alone
 - d. American Indian and Alaska Native alone
 - e. Asian
 - f. Native Hawaiian and other Pacific Islander
 - g. Some Other Race alone
 - h. Multiracial
6. What is your gender?
 - a. Male
 - b. Female
 - c. Non-binary
 - d. Other
 - e. Prefer not to say
7. What year were you born?
8. Please indicate which best describes your current relationship status?

- a. Never married
- b. Separated
- c. Divorced
- d. Widowed
- e. Married
- f. De facto

9. Do you have any children?

- a. Yes
- b. No

10. What is the highest educational degree you have received?

- a. None
- b. GED
- c. High school Diploma
- d. Associates/Junior College (AA)
- e. Bachelor's Degree (BA, BS)
- f. Masters degree (MA, MS, MBA)
- g. PHD
- h. Professional degree (DDS, JD, MD)

11. From the options below what is the highest level of education your **Father** has achieved?

- a. None
- b. GED
- c. High school Diploma
- d. Associates/Junior College (AA)
- e. Bachelor's Degree (BA, BS)

- f. Masters degree (MA, MS, MBA)
- g. PHD
- h. Professional degree (DDS, JD, MD)

12. From the options below what is the highest level of education your **Mother** has achieved?

- a. None
- b. GED
- c. High school Diploma
- d. Associates/Junior College (AA)
- e. Bachelor's Degree (BA, BS)
- f. Masters degree (MA, MS, MBA)
- g. PHD
- h. Professional degree (DDS, JD, MD)

13. What is your current employment status?

- a. Full time
- b. Part-time
- c. Casual employee
- d. Self-Employed
- e. Unemployed and looking for work
- f. Unemployed and not looking for work
- g. Retired
- h. Student

14. Thinking about last month, on average how much was your usual **weekly** income from all sources before tax and other deductions?

15. How much annual income do you expect to earn **next year**?

16. How much annual income do you expect to earn in **5 years time**?
17. How much annual income do you expect to earn in **10 years time**?
18. Given your current needs and financial responsibilities, would you say you and your families are
- a. Prosperous
 - b. Very comfortable
 - c. Comfortable
 - d. Just getting along
 - e. Poor
 - f. Very poor
19. How variable has your monthly income been over the past year? (scale from Very stable (1) to Very unstable (100))