Jan, 2026
GRINS DISCUSSION PAPER SERIES DP N° 74/2026 ISSN 3035-5567

GRINS

Gendered decision making in explore-exploit tasks

DP N° 74/2026

Authors:
Francesca Barigozzi, Natalia Montinari, Elisa Orlandi

Jan, 2026
GRINS DISCUSSION PAPER SERIES DP N° 74/2026 **ISSN 3035-5567**

Gendered decision making in explore-exploit tasks

Francesca Barigozzi, Natalia Montinari, Elisa Orlandi

KEYWORDS

Explore-exploit tasks | Gender differences | Exploration behaviour | Competitive behaviour

JEL CODE

C91, D81, D03

ACKNOWLEDGEMENTS

This study was funded by the European Union - NextGenerationEU, in the framework of the GRINS - Growing Resilient, INclusive and Sustainable project (GRINS PE00000018). The views and opinions expressed are solely those of the authors and do not necessarily reflect those of the European Union, nor can the European Union be held responsible for them.

We express our sincere gratitude to the participants of the IDEE meeting at the University of Bologna, the 2022 European ESA Meeting in Bologna, and the seminar attendees at the University of Bath (UK) for their insightful comments and suggestions. We are grateful to Valeria Maggian and Michal Krawczyk for their valuable comments on an earlier version of the paper. We are particularly thankful to Stephanie Heger, Maria Bigoni, Alice Guerra, and Eva Ranheli for their invaluable feedback, and to Alycia Chin, Charley M. Wu, and Azzurra Ruggeri for their generosity in sharing their data and code. Our thanks also go to Riccardo Ghioni for his outstanding research assistance with the experiment code. This research received financial support from the PRIN project “Advancing Gender Equality through Experiments: New Developments and Applications (AGEENDA)” (CUP P2022R248L 002). Elisa Orlandi gratefully acknowledges additional f inancial support from the “Next (Second) Generation EU” programme, acronym Next2EU, within the GRINS project, SPOKE 3– “Households’ sustainability” (CUP UNISA D43C24002040006). All remaining errors are our own. The views and opinions expressed are solely those of the authors and do not necessarily reflect those of the European Union, and the European Union cannot be held responsible for them.

CITE THIS WORK

Author(s): Francesca Barigozzi, Natalia Montinari, Elisa Orlandi. Title: Gendered decision making in explore-exploit tasks. Publication Date: 2026.

Many real-world choices involve explore–exploit decisions: balancing the search for better opportunities against securing a possibly suboptimal outcome. Do gendered approaches to these decisions exist? We study this question in a pre-registered laboratory experiment with 419

participants (< 50% female). Behaviour is observed under a piece-rate scheme and a tournament setting. Participants complete three tasks: (i) the Grain Game (explore–exploit) in two otherwise identical environments—one with gains only and one with gains and losses; (ii) an incentivized risk-elicitation task (BRET); and (iii) a loss-aversion task, followed by a questionnaire eliciting individual characteristics and beliefs. We show that, when losses are not possible (gains only), women place a higher value on information and explore more than men (consistent with over-exploration); once losses are possible (gains and losses) and the cost of deviating from exploitation rises, gender differences in exploration disappear. Competition per se does not induce exploration, and the raw gender gap in competitive entry is accounted for by incentivized controls for beliefs/self-confidence. Move-level regressions show adjustment margins: men reduce exploration after realized negative payoffs (threshold response), whereas women reduce exploration at the treatment level when losses are possible. Overall, our evidence indicates an information-seeking mechanism that drives exploration when downside risk is not possible but is muted when losses become possible, whereas cross-gender differences in information processing and use persist, yielding distinct choice dynamics.

Finanziato dall'Unione europea NextGenerationEU
Italiadomani PIANO NAZIONALE DI RIPRESA E RESILIENZA
UNIVERSITÀ DEGLI STUDI DI SALERNO
Ministero dell'Università e della Ricerca
GRINS

Gendered decision making in explore-exploit tasks

Francesca Barigozzi* Natalia Montinari* Elisa Orlandi†

Abstract

Many real-world choices involve explore–exploit decisions: balancing the search for better opportunities against securing a possibly suboptimal outcome. Do gendered approaches to these decisions exist? We study this question in a pre-registered laboratory experiment with 419 participants ($\sim 50\%$ female). Behaviour is observed under a piece-rate scheme and a tournament setting. Participants complete three tasks: (i) the Grain Game (explore–exploit) in two otherwise identical environments—one with *gains only* and one with *gains and losses*; (ii) an incentivized risk-elicitation task (BRET); and (iii) a loss-aversion task, followed by a questionnaire eliciting individual characteristics and beliefs. We show that, when losses are not possible (*gains only*), women place a higher value on information and explore more than men (consistent with over-exploration); once losses are possible (*gains and losses*) and the cost of deviating from exploitation rises, gender differences in exploration disappear. Competition *per se* does not induce exploration, and the raw gender gap in competitive entry is accounted for by incentivized controls for beliefs/self-confidence. Move-level regressions show adjustment margins: men reduce exploration after realized negative payoffs (threshold response), whereas women reduce exploration at the treatment level when losses are possible. Overall, our evidence indicates an *information-seeking mechanism* that drives exploration when downside risk is not possible but is muted when losses become possible, whereas cross-gender differences in information processing and use persist, yielding distinct choice dynamics.

Keywords: Explore-exploit tasks, Gender differences, Exploration behaviour, Competitive behaviour

JEL Codes: C91, D81, D03, J16

*Department of Economics, University of Bologna
†Department of Economics and Statistics, University of Salerno
Email addresses: francesca.barigozzi@unibo.it; natalia.montinari2@unibo.it; elisa.orlandi15@unibo.it.
Acknowledgments: We express our sincere gratitude to the participants of the IDEE meeting at the University of Bologna, the 2022 European ESA Meeting in Bologna, and the seminar attendees at the University of Bath (UK) for their insightful comments and suggestions. We are grateful to Valeria Maggian and Michal Krawczyk for their valuable comments on an earlier version of the paper. We are particularly thankful to Stephanie Heger, Maria Bigoni, Alice Guerra, and Eva Ranheli for their invaluable feedback, and to Alycia Chin, Charley M. Wu, and Azzurra Ruggeri for their generosity in sharing their data and code. Our thanks also go to Riccardo Ghioni for his outstanding research assistance with the experiment code. This research received financial support from the PRIN project "Advancing Gender Equality through Experiments: New Developments and Applications (AGEENDA)" (CUP P2022R248L_002). Elisa Orlandi gratefully acknowledges additional financial support from the "Next (Second) Generation EU" programme, acronym Next2EU, within the GRINS project, SPOKE 3 – "Households' sustainability" (CUP UNISA D43C24002040006). All remaining errors are our own. The views and opinions expressed are solely those of the authors and do not necessarily reflect those of the European Union, and the European Union cannot be held responsible for them.

1 Introduction

Gender gaps in career trajectories often emerge where exploration and entry under uncertainty are pivotal—deciding how selectively to search and when to accept, taking on unfamiliar projects with uncertain payoffs, or entering tournaments for promotion.[1] Whether such gaps stem from differences in risk and loss preferences, beliefs about relative ability, information acquisition, or responses to competitive environments remains an open question with direct implications for how organizations design choices, frame stakes, and structure contests.

Risk attitudes capture tolerance for known variability in outcomes, but exploration adds an information-acquisition motive aimed at reducing uncertainty and improving subsequent decisions.[2] In addition, competition introduces an entry decision and changes how choices are framed. The way rewards are structured—how much more winners earn than others—and whether participants see their relative ranking can change how important the task feels and how people update their beliefs and information-seeking behaviour. At the same time, holding the information structure constant, competition *per se* does not have to increase exploration. It is therefore important to separate the decision to face competition (the entry margin) from the motive to explore. For these reasons, potential gender patterns in explore-exploit behaviour should not be reduced to risk preferences alone, and call for evidence on the roles of information acquisition and competitive framing under uncertainty.

We use a pre-registered lab experiment to examine if and how gendered patterns emerge in explore–exploit tasks and we provide causal evidence that the presence of losses—introducing the possibility of negative payoffs and thus increasing the opportunity cost of exploration—eliminates gender gaps in explore–exploit choices and in self-selection into competition. To do so, we adapt the Grain Game of Chin et al. [2023] where they study willingness to explore: in each move, participants *explore* new opportunities, *exploit* current gains, or *retreat*. Participants face two otherwise identical environments which vary only the payoff support, i.e., all realizations are nonnegative in *Gains-only* whereas negative realizations are also possible in *Gain&Loss*. We maintain this structure given the (mixed) evidence on gender differences in loss aversion [Gächter et al., 2022, Georgalos, 2024], and we introduce competition by nesting the Grain Game into the Niederle and Vesterlund [2007] framework. Participants therefore play under three payment schemes: (1) piece rate, (2) a within-group tournament, and (3) a choice between piece rate and tournament.

Niederle and Vesterlund [2007] study willingness to compete in a stereotyped context (a math-related task), whereas the Grain Game is gender-neutral: participants act as farmers maximizing harvests on a linear field; to reinforce neutrality, the on-screen character matches the player's gender. As the first study

[1] Young job seekers face an explore -exploit dilemma: how widely to search and when to accept. Early on they tailor a few applications and learn from interviews; later, one or two tentative offers with deadlines force a choice—accept now or keep searching at the risk of none.

[2] Models of exploratory choice require information-seeking motives over and above risk preferences to account for observed behaviour; see, e.g., Wilson et al. [2014], Knox et al. [2012].

on gendered decision-making in explore–exploit tasks, we use this abstract, gender-neutral environment to minimize confounds and focus on exploration attitudes.

We show that, in *Gains-only*, women explore more than men; once losses are possible (*Gain&Loss*), exploration converges and the raw gap in competitive entry vanishes. Because the design holds technology, information structure, and timing constant while varying only the payoff support, the pattern is consistent with women placing a higher marginal value on information when losses are not possible, and with the possibility of losses raising the cost of deviating from exploitation for women. Importantly, convergence in average exploration does not imply identical adjustment margins. Move-level regressions show distinct responses to information: men adjust exploration conditionally on realized outcomes—exploring less after negative payoffs and more after positive ones (a threshold response)—whereas women exhibit a treatment-level adjustment, reducing exploration once losses become possible irrespective of the most recent payoff. We do not find that incentivized measures of risk, loss aversion, or confidence account for the female exploration premium in *Gains-only* (though beliefs/confidence help explain the raw gap in competitive entry), supporting a mechanism centered on information seeking rather than on risk-based primitives.

As for players' aggregate payoffs, our findings reveal slight male-favoring gaps in payoffs across stages and treatments, with statistical significance only in the Choice stage of *Gains-only*. In that stage, women both progress farther and earn less, consistent with over-exploration relative to male players and a reinforcement-learning benchmark presented in Appendix I and discussed in Subsection 5.3.

Our results have potential implications for gender gaps in career trajectories. As mentioned above, explore–exploit maps to early-career search vs. acceptance and to entry into competitive tracks. If the possibility of losses suppresses information-seeking deviations from exploitation and equalizes competitive entry, then how institutions frame downside risk—through contract structures, feedback about failure probabilities, or the timing and the presence of losses—can materially affect gender gaps at precisely the margins where careers diverge. Our findings speak to the design of counseling, presentation of uncertainty in internal job markets, and how competitive opportunities are communicated (see Section 7).

To sum up, we offer four main contributions. First, we identify the role of the possibility of losses in gendered explore–exploit behaviour using a clean experimental design that holds technology, information structure, and timing constant while directly measuring rival channels (risk, loss aversion, and beliefs). Second, we provide mechanism-consistent evidence of *gender differences in information seeking*: when losses are not possible, women attach a higher marginal value to information and explore more; when losses become possible and the opportunity cost of exploration rises, exploration converges across genders, with distinct adjustment margins (men respond to realized losses; women adjust at the treatment level), as summarized in Table 8 of Section 6. Third, we quantify the magnitude and pattern of over-exploration in *Gains-only* and its disappearance under *Gain&Loss*, clarifying when gender differences in sequential choice

arise. Fourth, we help reconcile mixed evidence on gender and willingness to compete (Subsection 2): the tournament-entry gap is environment-dependent, appears only when losses are not possible, and is largely explained by confidence/beliefs.

The remainder of the paper is organized as follows. Section 2 reviews the related literature and motivates our hypotheses. Section 3 details the experimental design; Section 4 describes the data; Section 5 presents the main findings on exploration, tournament entry in stage 3 (Choice), and payoffs. Section 6 provides an overview of results, and Section 7 explores an application to job seekers' behaviour. Section 8 concludes.

2 Literature Review

Our paper contributes to two key strands of literature. The first examines the dynamics of the explore-exploit dilemma using experimental methods, focusing on how individuals navigate trade-offs between exploration and exploitation. The second investigates the determinants of gender differences in competitiveness, building on the seminal work of Niederle and Vesterlund [2007]. With respect to the first strand, we analyze a specific version of the explore-exploit dilemma, focusing on gender differences, which have not been studied in this context before. Regarding the second strand, we introduce a novel variation in the task used to study competitive behaviour, providing fresh insights into this well-researched area.

The exploration–exploitation dilemma, the trade-off between seeking new options and capitalizing on known ones, has been widely investigated across diverse areas, including animal foraging behaviour [Mehlhorn et al., 2015], computational models of reinforcement learning [Yogeswaran and Ponnambalam, 2012], procurement under incomplete information [Azoulay-Schwartz et al., 2004], organizational learning and performance [March, 1991, Gupta et al., 2006], and the neural mechanisms involved in decision-making [Daw et al., 2006, Blanchard and Gershman, 2018], just to name a few. For broader overviews, see also Mehlhorn et al. [2015], Berger-Tal et al. [2014].

In psychological, behavioural, and cognitive sciences, the exploration–exploitation dilemma has also been extensively analyzed through multi-armed bandit (MAB) frameworks, which model how individuals learn to allocate choices among multiple options with initially uncertain payoffs [e.g., Berry and Fristedt, 1985, Steyvers et al., 2009, Speekenbrink and Konstantinidis, 2015, Reverdy et al., 2014, Wu et al., 2018, Schulz et al., 2020, 2019, Meder et al., 2021]. Notably, a number of these papers focus on children, modelling how younger participants develop and adjust their exploration strategies [Meder et al., 2021, Wilson et al., 2021, Schulz et al., 2019].

The Grain Game developed by Chin et al. [2023], which we adapt in this paper, shares some similarities with Wu et al. [2018] in that participants explore a space where rewards are locally correlated, it is unique in that exploration is restricted to one step at a time. This constraint more closely mimics real-world

settings in which exploring vast spaces is neither free nor costless, reflecting the practical limitations people face when gathering information. We compare our results with those of Chin et al. [2023] in Section 6.

While several papers have utilized lab experiments to examine individual behaviours in explore-exploit tasks, to the best of our knowledge, this study is the first to specifically investigate gender differences within this framework.

Considering now more in details the second strand of the literature we contribute to. Niederle and Vesterlund [2007] established a foundational result in experimental economics, showing that men are more likely than women to self-select into competitive environments. This finding has been widely replicated—see Niederle and Vesterlund [2011], Azmat and Petrongolo [2014], Buser et al. [2014]—yet subsequent research has demonstrated that the magnitude and even the presence of this gender gap depend on various contextual factors. For example, the nature of the task influences competitiveness, with men often being more competitive in stereotypically male-dominated tasks, while gender-neutral or female-oriented tasks reduce or eliminate the gap [Dreber et al., 2014, Große and Riener, 2010, Shurchkov, 2012]. Cultural and societal differences also play a significant role, as cross-cultural studies have shown that gender gaps in competitiveness vary according to levels of gender equality and prevailing norms [Gneezy et al., 2009, Booth and Nolen, 2012]. Furthermore, differences in the conceptual frameworks used, such as variations in task difficulty or how competition is framed, highlight the contextual dependency of these differences [Buser et al., 2017]. This body of evidence highlight the importance of incorporating task-specific, cultural, and methodological dimensions into the study of gender and competitiveness. In this study, we contribute to this literature by employing the explore-exploit dilemma as a novel task to investigate gender differences in competitive behaviour, shifting the focus to decision-making under uncertainty in a gender-neutral context.

2.1 Hypotheses

Following the existing literature, we formulate the following hypotheses to address gendered differences in exploration and competition behaviours. These hypotheses were pre-registered prior to data collection to ensure transparency and credibility in our research design.[3] The first hypothesis addresses gender differences in behaviour within the context of the explore-exploit dilemma in absence of strategic uncertainty due to competition.

Hypothesis 1 ***Gender gap in exploration and treatments.*** *Women explore less than men, particularly in decision-making contexts that entail the possibility of facing losses.*

Mixed evidence exists on gender gaps in risk and loss aversion. While numerous influential studies

[3]The link for the pre-registration is available at the following OSF webpage: https://osf.io/24eah/?view_only=1a9950422ed54210a9da08ea2caa5654

document that women are, on average, more risk-averse than men [Croson and Gneezy, 2009, Eckel and Grossman, 2008, Dohmen et al., 2011], more recent meta-analytical and experimental work suggests that such differences are often context-dependent and sensitive to task design [Filippin and Crosetto, 2016]. Evidence on gender differences in loss aversion is likewise mixed: some studies find women to be more loss-averse [Georgalos, 2024], whereas others find no systematic gap [Gächter et al., 2022]. To the extent that exploration requires choosing under uncertainty or tolerating the possibility of negative payoffs, such differences—where they exist—could lead to a gender gap in the willingness to explore.[4]

Gender differences in exploration behaviours may also stem from early socialization processes and societal expectations regarding risk-taking. Research in developmental psychology suggests that gender norms, internalized from an early age, influence individuals' behaviours and decision-making processes in adulthood [Eagly and Karau, 2002, Fulcher and Coyle, 2011].

Additionally, women may adopt a more conservative approach due to early exposure to gendered play patterns that discourage risk-taking and adventurous decision-making [Kung, 2022, Harbin, 2023]. Women are also found to have higher levels of anxiety than men about spatial navigation [Lawton, 1994].

Our second and third hypotheses address gender differences in competitive environments, examining both behaviour within the explore-exploit dilemma (with a focus on the second stage of the experiment) and the decision to compete (with a focus on the third stage).

Hypothesis 2 ***More exploration under competition.*** *Controlling for individual characteristics, risk aversion, and loss aversion, individuals (both men and women) who choose to compete in the Choice stage engage in greater exploratory behaviour.*[5]

Hypothesis 3 ***Gender gaps in willingness to compete.*** *Controlling for individual characteristics, risk aversion, and loss aversion, women are less likely to choose competition compared to men.*

Competition has been shown to increase exploratory behaviour across genders, as individuals seek to outperform their peers and maximize their outcomes [Gill and Prowse, 2012, Mago et al., 2016]. Competitive settings often act as a motivational force, driving individuals to adopt strategies that involve greater exploration and risk-taking to gain a competitive edge [Gërxhani et al., 2023]. However, significant

[4]Higher risk or loss aversion could also motivate greater information seeking as a strategy to reduce uncertainty, potentially pushing in the opposite direction (i.e., towards more exploration among women). Our preregistered hypothesis nonetheless focused on the mechanism more consistently highlighted in the literature—namely, reduced willingness to engage in uncertain or potentially loss-generating choices.

[5]The pre-registered version of Hypothesis 2 stated: "*Controlling for individual characteristics, risk aversion and loss aversion, in competitive settings individuals (both men and women) explore more.*" While the two formulations are related, the version presented in the paper narrows the empirical focus to the Choice stage, where participants self-select into a competitive or non-competitive payment scheme. This allows us to test whether those who actively choose to compete also explore more—capturing the endogenous relationship between willingness to compete and exploratory behaviour. We acknowledge, however, that this design does not allow for a within-subject comparison of behaviour across incentive schemes (e.g., Tournament vs. Piece Rate), which would have been more directly aligned with the original formulation. This analytical choice reflects a practical constraint of the experimental design (i.e., fixed order and lack of random assignment in Choice, which is common in the experiments on competitiveness, see, e.g. Niederle and Vesterlund [2007]), and we have updated the hypothesis wording to match the operationalization in the data.

gender differences persist in competitive environments, with women being less likely to self-select into competition [Niederle and Vesterlund, 2011].

Women's reluctance to compete is often attributed to lower self-confidence, internalized social norms that discourage competitive behaviour in women, and perceived stereotypes about gender roles [Eagly and Karau, 2002, Francis, 2010]. Women may also experience a higher sensitivity to social penalties or a fear of underperformance in competitive contexts [Buser et al., 2017]. Moreover, studies highlight that men are more likely to engage in voluntary job-to-job mobility, reflecting a greater willingness to compete and explore new opportunities [Theodossiou and Zangelidis, 2009, Cortés et al., 2023].

Together, these findings highlight the complex interplay of individual traits, societal norms, and environmental factors in shaping exploratory and competitive behaviours across genders, making it crucial to account for these elements when analyzing decision-making in competitive settings.

Our fourth hypothesis is based on the attitude toward competitiveness of women attending STEM fields:

Hypothesis 4 ***STEM.*** *Women attending STEM fields of study are more competitive than men.*

The under-representation of women in STEM fields has led to increased scrutiny of competitiveness in male-dominated environments. Some studies suggest that women pursuing STEM careers tend to adopt more competitive behaviours as they navigate traditionally male-dominated spaces [Buser and Yuan, 2019, Cárdenas et al., 2012]. Moreover, the "gender-equality paradox" suggests that in countries with fewer gender-equal opportunities, women are more likely to pursue challenging STEM careers as a means of economic advancement [Stoet and Geary, 2018]. Women in these fields may develop resilience and adaptability, making them more competitive than their male counterparts in similar environments.

Finally, our last hypothesis refers to the impact of individual characteristics:

Hypothesis 5 ***Preferences affect exploration and competition.*** *Individual risk preferences, loss aversion, cognitive abilities, and self-confidence affect both the likelihood to explore and the likelihood to compete.*

A wide range of literature highlights the significant role of personality traits and economic preferences in shaping decision-making under uncertainty and competition [Camerer and Hogarth, 1999, Dohmen et al., 2010]. Higher risk aversion and loss aversion have been associated with a lower tendency to explore and compete, whereas cognitive abilities and self-confidence have a positive influence. Gender norms further intersect with these individual traits, influencing behaviours such as over-exploration and reluctance to engage in competition in specific environments [Chin et al., 2023]. The tendency to explore and compete is shaped not only by inherent traits but also by the external environment, which can encourage or inhibit certain behaviours.

3 Experimental Design

This study investigates gender differences in exploration and competition behaviours through a controlled incentivised laboratory experiment conducted at the BLESS (Bologna Laboratory for Experiments in Social Sciences)[6]. The experiment was pre-registered on OSF (https://osf.io/6kxsp) and received approval from the Bioethics Committee of the University of Bologna.

The participants took part in three computerised games presented in fixed order: the Grain Game [Chin et al., 2023], a Loss Aversion Task [Gächter et al., 2022] and the Bomb Risk Elicitation Task (BRET; Crosetto and Filippin 2013). These tasks were followed by a post-experimental questionnaire designed to collect additional data on socio-demographic characteristics, personality traits [Gosling et al., 2003], competitiveness and cognitive abilities [Frederick, 2005]. See Figure 1 for a blueprint of the experimental design.

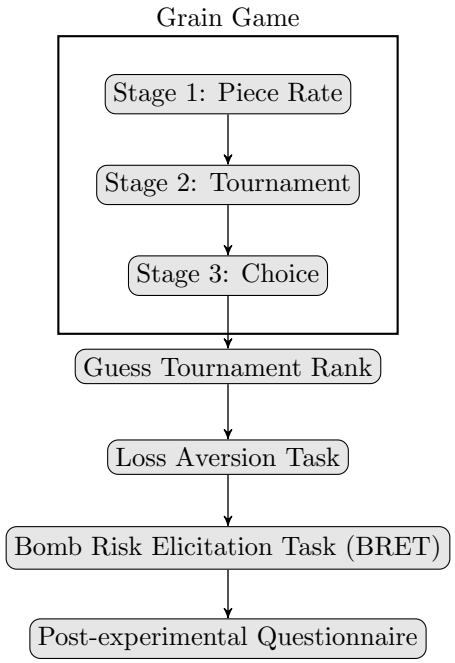


Figure 1: Experimental Design. *Note*: Participants first completed the Grain Game (composed of three stages: Piece rate, Tournament, and Choice), guessed their tournament rank, followed by the Loss Aversion Task, the BRET, and finally the post-experimental questionnaire.

3.1 The Grain Game

The Grain Game, adapted from Chin et al. [2023], placed participants in the role of farmers tasked with maximising their harvests by planting seeds in a linear field of 70 planting slots. During the course of the game, participants planted a total of 70 seeds, but made 69 decisions. This is because in the first round, participants automatically planted in position 1, the leftmost slot in the field, without making

[6]The experiment was conducted in Italian, an English version of the experimental instructions are reproduced in Appendix H.

a choice. From the second round onwards, they could decide whether to plant in the same slot as the previous round (exploit), in the slot immediately to the right (explore) or in the slot immediately to the left (retreat).[7] In each round, a seed had to be planted. De facto constraining the movement of participants to one step at a time at most. An illustration of the game interface, showing the planting field and participant options, is provided in Figure 2. More screenshots can be found in Appendix A.



Figure 2: Illustration of the Grain Game. *Note*: During the play, participants could see the total amount of accumulated points in the top left of the screen and the remaining number of seeds or rounds to play in the top right. They could also reread the instructions at any time by clicking on the question mark icon on the screen.

Payoffs in each slot were calculated as the sum of a fixed value derived from one of six predefined payoff sequences and random noise. The sequences were generated as random walks with errors drawn from a normal distribution with mean 0 and standard deviation 2.[8] The noise followed a symmetric distribution with probabilities assigned as follows: 0 points with a probability of 0.56, ± 1 point with a probability of 0.15 for each, and ± 2 points with a probability of 0.07 for each. Thus, the payoffs are locally correlated.[9]

As in Chin et al. [2023], the experiment included two treatments. In the Gain treatment, participants encountered only positive payoffs, with fixed values ranging from 1 to 25 points. In the Gain & Loss treatment, participants faced potential losses as payoffs ranged from -10 to 15 points. To account for the possibility of negative payoffs, participants in the Gain & Loss treatment received an initial endowment

[7] It follows that the furthest slot a participant could reach, given the 70 seeds and 69 decisions, is the planting slot in position 70.

[8] The sequences can be found in Appendix B.

[9] Framing the task as a farmer across a field gives the correlation structure an intuitive grounding: just as soil fertility tends to change only gradually from one plot to the next, the expected yield in neighbouring planting slots differs by only a few points. This narrative, therefore, creates ecological validity and helps participants form expectations about why nearby locations offer similar returns, even without explicitly teaching them about spatial correlation.

of 700 points.[10]

Differently than Chin et al. [2023], we used three payoff sequences for each treatment. The sequences in the Gain & Loss treatment were identical to those in the Gain treatment, with a -10 offset applied to all values. All sequences rewarded exploration in the sense that the global maximum of every payoff sequence was placed in the second half of the field. Participants played the Grain Game three times, with each stage using a different compensation scheme, following the experimental framework of Niederle and Vesterlund [2007]. Therefore, the first stage of the Grain Game, where participants are paid piece rate, represents the replication of Chin et al. [2023] with the only difference that we vary the sequences of payoffs encountered by participants within treatments, whereas in Chin et al. [2023], participants all face the same sequence. At the end of the game, only one of the three stages was randomly selected for payment and participants were informed of this rule before the experiment began.

In the first stage, participants earned 1 cent for each point they scored if this stage was selected for payment. In the second stage, participants were divided into gender-balanced groups of four, with two men and two women in each group. In this stage, only the participant with the highest score in the group earned 4 cents for each point if the stage was selected for payment, while the other three participants received nothing. If there was a tie, the winner was selected randomly among those with the highest scores.

In the third stage, participants could choose between the piece-rate scheme or the tournament scheme before starting the game. This choice is a proxy for competitiveness. If this stage was selected for payment, those who chose the piece-rate scheme earned 1 cent for each point as in the first stage. Those who chose the tournament scheme earned 4 cents for each point only if they outscored the top scorer from their group in the second stage (Tournament); otherwise, they received no payment. If their score tied with the prior top scorer, the winner was selected randomly.

Each participant played all three payoff sequences within their assigned treatment. Before starting each part, participants answered a set of control questions to ensure their understanding of the instructions (see Appendix H). The order of sequences was randomised across participants at the session level to mitigate potential order effects.

At the end of the Grain Game, participants were asked to guess their ranking in the Tournament stage. Correct guesses were incentivized as in Niederle and Vesterlund [2007].

3.2 Loss aversion task and BRET

After completing the Grain Game, in the second part of the experiment, participants were asked to perform two additional tasks: a loss aversion task and a task to measure their risk attitudes. Before

[10]The average payoff in the Gain treatment was 15 points and 5 points in the Gain & Loss treatment. The 700-point endowment ensured equal expected payoffs across the two treatments.

starting, participants were informed that their performance in only one of these two tasks would be randomly selected for payment of part 2.

To assess participants' loss aversion, we implemented the task developed by Gächter et al. [2022]. Participants were presented with six small-stakes lotteries. For each lottery, they decided whether to participate in a coin toss. In all six lotteries, participants could win €6 if the coin landed on tails, but they risked losing an increasing amount of money (from €2 in the first lottery to €7 in the sixth) if the coin landed on heads. If participants opted not to participate, they received €0. If the loss aversion task was selected for payment, one of the six lotteries was randomly chosen to determine their final earnings.

The Bomb Risk Elicitation Task (BRET), adapted from Crosetto and Filippin [2013], was used to measure participants' risk attitudes in a controlled and incentivised setting. Participants were presented with a grid of 100 cells (10 x 10) and asked to decide how many cells to colour red. Each red cell represented a potential gain of 5 cent. After their decision, the computer randomly selected one cell. If the chosen cell was grey, participants earned 5 cent for each cell they chose to colour red. However, if the chosen cell was red, they earned nothing. If the BRET task was selected for payment, their earnings were calculated based on the outcome of the task.

For an illustration of the Loss Aversion Task and BRET interfaces, see Appendix C.

3.3 Experimental Payments

At the end of the experiment, participants' final payment depended on the payoff from one randomly selected stage of the Grain Game (Piece Rate, Tournament, or Choice), on the payoff from one randomly selected additional task (Loss Aversion or BRET), plus a fixed 5 Euro show-up fee. Detailed payoff rules for each stage and task are summarized in Table 1. Formally:

$$\begin{aligned}\text{Final Payment} = &\underbrace{(\text{Payment from one of the three Grain Game stages})}_{\text{randomly chosen}} \\ &+ \underbrace{(\text{Payment from one of the two additional tasks})}_{\text{randomly chosen}} \\ &+ \underbrace{5\,\text{Euros}}_{\text{show-up fee}} .\end{aligned}$$

Table 1: Overview of the Experimental Compensation Scheme

Part	Description
Part 1: Grain Game	**Stage 1 (Piece Rate)**: 1 cent per point. **Stage 2 (Tournament)**: Groups of 4 (2 men, 2 women); highest score earns 4 cents/point (ties random). **Stage 3 (Choice)**: Each participant chooses either (i) to not compete and get Piece Rate payment (1 cent/point), or (ii) to compete and get Tournament payment (4 cents/point if beating the top scorer from Stage 2, else 0).
Part 2: Additional Tasks	**Loss Aversion Task**: One of six coin-toss lotteries randomly determines payoffs. **BRET**: Earn 5 cents per chosen red cell unless the random draw is red; then earnings are 0.
Show-up Fee	A fixed 5-Euro payment, independent of performance.

4 Data

Our pre-registered plan aimed to collect data from 432 participants (50% female). To achieve this, we invited 432 individuals to take part in the study, which was conducted at the Bologna Laboratory for Experiments in Social Sciences (BLESS) between June 2022 and March 2023. Recruitment was carried out through the ORSEE system [Greiner, 2015], ensuring random assignment to sessions and a balanced gender composition across experimental groups. Each session accommodated up to 32 participants, who were randomly assigned to gender-balanced groups of four (two men and two women). However, due to a software problem encountered in the Grain Game, data from 13 participants were lost, resulting in a final sample of 419 participants (206 women and 213 men) for the analysis presented in this paper. Importantly, the sample was well balanced not only in terms of gender but also with respect to participants' demographics (e.g., age, field of study, and nationality), ensuring that our findings are not driven by imbalances in observable characteristics. A table with descriptive statistics for all relevant variables is provided in Appendix D.

Table 2 summarises the distribution of participants across treatments and genders.

Table 2: Participants by Treatment and Gender

	Gain	**Gain & Loss**	**Overall**
Men	106	107	213
Women	104	102	206
Overall	210	209	419

While the main sample consists of 419 participants, data for the loss aversion task and the BRET are unavailable for the 62 participants from the first two sessions due to technical issues.[11] As a result, depending on the model specification used in the paper, the sample size might be smaller. In addition, in the post-experimental questionnaire, some individuals chose to skip certain questions, so there is missing

[11] In the first two sessions, if participants held the button down, some "exploit" actions were not recorded. This was a purely technical logging issue; we therefore imputed the missing exploit counts.

data from that source as well. However, the sample size used in each regression is always specified in the Results section, see also Appendix D.

Furthermore, when constructing the *Loss Aversion* (λ_{loss}) variable, our measure of loss aversion based on the loss aversion task, we exclude participants whose decisions are non-monotonic (following the procedure in Gächter et al. 2022). This exclusion generates additional missing values for those participants.

5 Results

In this section we present our results organized as follows: Section 5.1 focuses on the behaviour in the explore-exploit dilemma under piece rate. Section 5.2 analyses the willingness to compete and section 5.3 discusses the differences in earnings.

5.1 Exploration

5.1.1 Descriptive Evidence

We begin by examining descriptive evidence on exploration behaviour and whether it differs by gender and treatment. We use three different measures of *exploration*: (i) the average percentage of "explore" choices; (ii) the farthest position ("*Maximum position*") that participants reach at any point in the game, and (iii) the final position where they plant their last seed ("*Final position*").[12]

Figure 3 shows the average percentage of *Explore*, *Exploit*, and *Retreat* decisions across each stage of the Grain Game, disaggregated by gender (men/women) and by treatment condition (Gain vs. Gain & Loss).

Contrary to our pre-registered hypothesis 1, women do not explore less than men. Instead, we find evidence suggesting that women, in some specific settings, explore more than men, particularly in the Tournament and Choice stages of the Grain Game.

In the Tournament stage, women explore in 55.0% of decisions (SD = 24.9%), compared to 48.8% for men (SD = 25.7%). This difference is statistically significant ($t = -2.50$, $p = 0.0130$). Similarly, in the final choice stage, women explore in 56.9% of decisions (SD = 26.7%), compared to 50.9% for men (SD = 27.9%), and this difference is also statistically significant ($t = -2.28$, $p = 0.0234$).

When analysing by treatment, the gender difference is driven by the Gain treatment. In this setting, women explore 58.3% of decisions in the Tournament stage (compared to 49.4% for men; $t = -2.47$, $p = 0.0143$) and 60.6% of decisions in the final choice stage (compared to 49.2% for men; $t = -3.02$, $p = 0.0028$). In the Gain & Loss treatment, gender differences in exploration are not statistically significant.

[12] Because participants can retreat after planting in a new slot, *maximum position* and *final position* do not necessarily coincide.

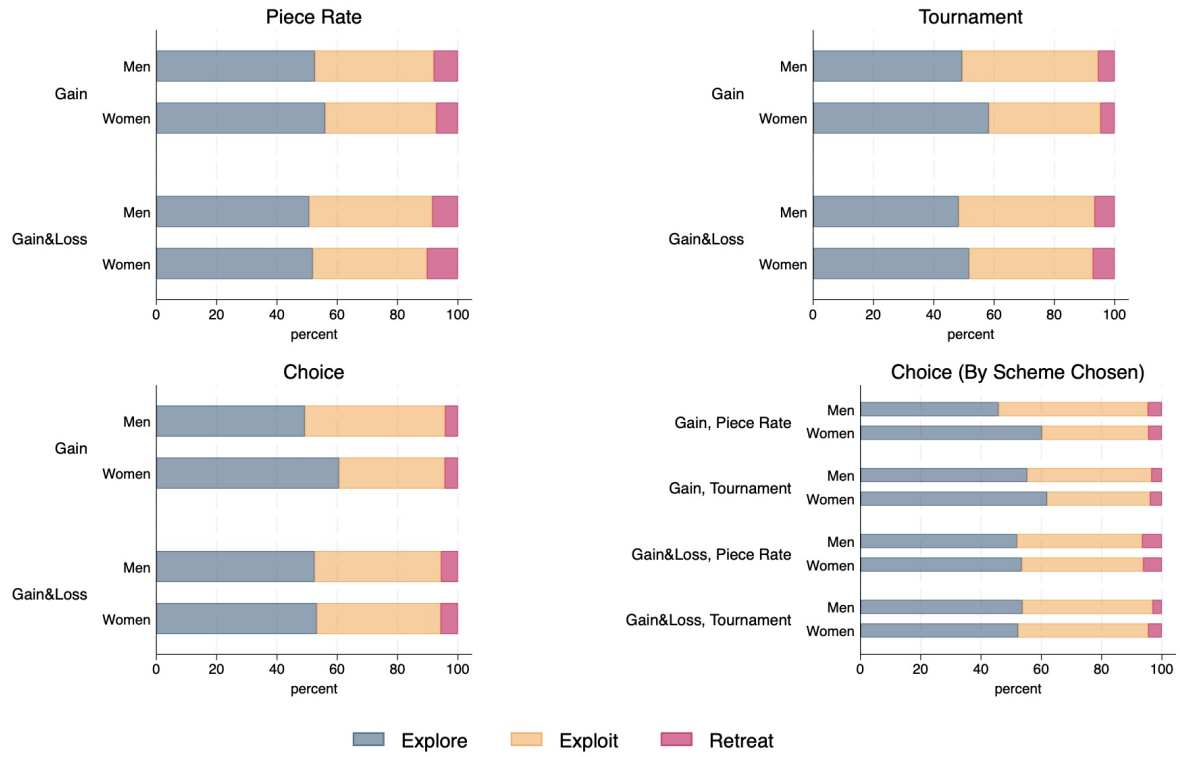


Figure 3: Explore, Exploit, Retreat % Across Stages, Gender and Treatment. *Note*: This figure displays the proportion of exploration, exploitation and retreat choices made by participants in each stage of the experiment, disaggregated by gender (men/women) and by treatment condition (Gain vs. Gain & Loss).

The other two measures, "*Maximum position*" and "*Final position*", confirm that in the Tournament and Choice stages under the Gain treatment, women tend to explore more than men; specifically, they reach a higher maximum position and conclude the game farther to the right. By contrast, in the Gain & Loss treatment, no significant gender difference emerges. Detailed results on those two additional measures and the percentage of explore are provided in Appendix E.

5.1.2 Regression Analysis

To assess whether women explore more or less than men controlling for individual characteristics, we focus on the *Choice* stage of the Grain Game.[13] This stage is particularly interesting because it allows us to verify if and how competition is associated with the players' choices. In this stage, as in every stage, each participant plants a total of 70 seeds and therefore makes 69 sequential decisions over time.

We define $Move_{it}$ as a categorical variable indicating the participant's i decision at round t, where $Move \in \{\text{Explore}, \text{Exploit}, \text{Retreat}\}$. To analyze these decisions, we estimate three random-effects multinomial logistic regressions with *Exploit* as the baseline outcome. Our general model can be written as:

[13] In Appendix F, we report the regression results for the *Piece Rate* and *Tournament* stages.

$$\ln\left(\frac{\Pr(Move_{it}=m)}{\Pr(Move_{it}=Explore)}\right) = \beta_{0,m} + \beta_{1,m}\, Female_i + \beta_{2,m}\, Gain\&Loss_i + \beta_{3,m}\,[Female_i \times Gain\&Loss_i] + \beta_{4,m}\, LagPayoff_{it-1} + \beta_{5,m}\, Threshold_{it-1} + \beta_{6,m}\, Compete_i + \boldsymbol{\gamma}'_m \mathbf{X}_i + u_{i,m}, \tag{1}$$

where $m \in \{Retreat, Explore\}$, $Female_i$ is a dummy taking value 1 for women and 0 otherwise, and $Gain\&Loss_i$ indicates whether participant i is assigned to the Gain & Loss treatment (= 1) or not (= 0). $LagPayoff_{it-1}$ is the payoff from the previous planting round, adjusted by adding 10 points in the Gain & Loss treatment to ensure comparability with the non-negative payoffs in the Gain setting. $Threshold_{it-1}$ is a dummy set to 1 if the adjusted payoff in the previous round was below 10 (i.e., a negative payoff in the Gain & Loss environment), and 0 otherwise. $Compete_i$ indicates whether participant i chose the competitive payment scheme in this final stage of the Grain game. Finally, $u_{i,m}$ denotes participant-level random intercepts, and $\mathbf{X}_i$ is a vector of additional participant-level controls (e.g., demographics, risk attitudes, personality traits) which differ across model specifications.[14]

We estimate three versions of Equation (1) that differ in the participant-level controls included in $\mathbf{X}_i$. Model A is the baseline and includes only the core regressors shown in Equation (1) ($\mathbf{X}_i$ is empty). Model B augments the baseline by adding socio-demographic characteristics (age, nationality, field of study and GPA), personality and cognitive measures (TIPI, CRT), piece-rate payoff, competitiveness, and self-reported risk attitudes. Finally, Model C replaces the self-reported risk attitude with the incentivized measure from the BRET task.[15] It also includes the loss-aversion parameter λ_{loss} from the loss-aversion task.[16] Table 3 presents the results from these three random-effects multinomial logit regressions.

Table 3: Random Effects Multinomial Logistic Results for *Move*

Dependent Variable	*Move* (Baseline = *Exploit*)					
	Model A		**Model B**		**Model C**	
	Retreat	Explore	Retreat	Explore	Retreat	Explore
Female	1.165	2.096*	1.199	2.330*	1.131	2.305
	(0.360)	(0.782)	(0.409)	(0.869)	(0.485)	(1.097)
Gain & Loss	0.996	1.171	1.147	1.190	0.730	1.323
	(0.340)	(0.430)	(0.424)	(0.443)	(0.297)	(0.564)
Female × (Gain & Loss)	1.250	0.432	1.070	0.353*	1.444	0.269*

(Continued on next page)

[14] As a robustness check, we also estimated these regressions while controlling for the payoff sequences that participants ultimately observed. For brevity, we do not include those additional results here, but they are available upon request. The interpretation of the hypothesis remains unchanged.

[15] Risk and loss-aversion were elicited *after* the Grain Game, so we checked whether they differ by treatment. They do not: BRET points ($t = 0.24$, $p = .81$), self-reported risk ($t = 0.03$, $p = .98$), and *Loss Aversion* (λ_{loss}) ($t = 0.87$, $p = .39$).

[16] Because some participants did not complete the BRET and the loss-aversion task and because non-monotonic (inconsistent) responses in the loss-aversion task were dropped following Gächter et al. [2022], Model C is estimated on a slightly smaller sample.

	Model A		Model B		Model C	
	Retreat	Explore	Retreat	Explore	Retreat	Explore
	(0.583)	(0.226)	(0.520)	(0.181)	(0.801)	(0.161)
Lag Payoff	0.722***	0.705***	0.723***	0.704***	0.729***	0.707***
	(0.022)	(0.021)	(0.022)	(0.021)	(0.024)	(0.024)
Threshold	0.265***	0.565***	0.276***	0.579***	0.277***	0.587***
	(0.055)	(0.074)	(0.058)	(0.078)	(0.064)	(0.094)
Piece Rate Payoff	–	–	0.9999	1.0039***	0.9999	1.0040***
			(0.0008)	(0.0009)	(0.0010)	(0.0012)
Risk (self-reported)	–	–	1.091	1.281**	–	–
			(0.084)	(0.097)		
Risk (Bret)	–	–	–	–	1.015	1.021*
					(0.010)	(0.011)
Loss Aversion (λ_{loss})	–	–	–	–	0.518**	0.545**
					(0.106)	(0.125)
Competitiveness (self-reported)	–	–	0.966	1.074	0.936	1.077
			(0.056)	(0.070)	(0.063)	(0.082)
CRT	–	–	1.227	0.977	1.295	1.023
			(0.152)	(0.128)	(0.181)	(0.160)
Extraversion	–	–	1.036	0.967	0.984	0.989
			(0.079)	(0.083)	(0.088)	(0.105)
Agreeableness	–	–	1.553***	1.331*	1.480**	1.444*
			(0.187)	(0.168)	(0.202)	(0.209)
Conscientiousness	–	–	1.028	0.928	0.989	1.000
			(0.095)	(0.094)	(0.105)	(0.119)
Emotional Stability	–	–	0.903	0.800*	1.042	0.883
			(0.081)	(0.082)	(0.099)	(0.099)
Openness	–	–	0.837	0.826	0.896	0.953
			(0.107)	(0.108)	(0.128)	(0.150)
GPA	–	–	0.950	0.921	0.875	0.957
			(0.148)	(0.144)	(0.152)	(0.170)
Born in Italy	–	–	0.263*	0.229*	0.192*	0.136**
			(0.152)	(0.134)	(0.135)	(0.106)
Age	–	–	0.987	1.016	0.988	1.041
			(0.033)	(0.042)	(0.047)	(0.055)
STEM	–	–	0.726	0.782	0.910	0.993
			(0.196)	(0.215)	(0.281)	(0.317)
Compete	0.627	1.383	0.600	0.999	0.449*	0.804
	(0.165)	(0.395)	(0.167)	(0.284)	(0.148)	(0.274)
Constant	4.513***	222.366***	5.823	16.359	21.436	5.098
	(2.134)	(114.842)	(10.166)	(31.442)	(42.361)	(11.517)
Model Statistics						
Observations	28,492		27,676		21,080	
Number of Groups	419		407		310	
Wald χ^2	193.85		242.75		216.16	
Log pseudolikelihood	−16139.104		−15671.713		−11755.302	

Notes: All entries are Relative Risk Ratios (RRRs) from random-effects multinomial logit regressions. Robust standard errors (clustered at participant level) are shown in parentheses. The dependent variable is *Move*, and the baseline (omitted) outcome is *Exploit*. Significance levels: $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$. While most variables are self-explanatory, note that *Loss Aversion* (λ_{loss}) is our measure of loss aversion where higher values indicate higher loss aversion, and *CRT* is the Cognitive Reflection Test score, where a higher value indicates more reflective (rather than intuitive) thinking.

The estimates from the random-effects multinomial logistic models indicate that in the *Gain-only* treatment, women are significantly more likely than men to explore (rather than exploit) in the final (Choice) stage of the Grain Game. In Models A and B of Table 3, the coefficient on *Female* in the "Explore vs. Exploit" comparison exceeds 2.0 and is statistically significant at the 5% level. In Model C, which relies on a smaller sample, the coefficient remains above 2.0 but is only marginally significant ($p \approx 0.08$). An RRR above 2.0 indicates that women are more than twice as likely as men to choose *Explore* over *Exploit*, and, notably, this differential persists even after we control for individual risk and loss aversion.

By contrast, the interaction term *Female*×*Gain & Loss* is below 1.0 in all three models and is significantly different from 1.0 in Models B and C. This pattern suggests that the higher propensity of women to explore, relative to men, is noticeably decreased when negative payoffs become possible. Furthermore, linear combination tests confirm that under *Gain & Loss* the net female effect fully vanishes in every specification, implying no statistically discernible difference in exploratory behaviour between women and men in the *Gain&Loss* tretament.[17] Hypothesis 1 is thus not supported.

Across all models, a larger payoff in the previous round strongly reduces participants' tendency to deviate from exploitation in the subsequent round, and if the adjusted payoff from the prior round falls below a certain threshold (i.e., is negative under the *Gain&Loss* treatment), individuals become even less inclined to explore.

Notably, choosing the competitive scheme (*Compete*) does not appear to increase exploration in any of the three model specifications. In Model C, *Compete* significantly reduces "Retreat vs. Exploit" (RRR ≈ 0.45, $p < 0.05$), indicating that those who select the competitive scheme are less likely to retreat (as opposed to exploit) than those in the noncompetitive setting. However, this effect does not translate into a higher relative probability of exploring. Thus, we find no evidence that individuals who self-select into competition systematically explore more, which does not support Hypothesis 2.

Turning to other variables, risk attitudes (both self-reported and measured via the BRET) consistently raise the relative probability of exploring rather than exploiting. In Model B, self-reported risk is strongly associated with "Explore vs. Exploit" (RRR ≈ 1.28, $p < 0.001$), and in Model C the incentivized risk measure (*RiskBret*) is smaller but still significant (RRR ≈ 1.02, $p < 0.05$). By contrast, *Loss Aversion*

[17] Specifically, we test whether the sum of the main *Female* coefficient and the *Female*×*Gain & Loss* interaction coefficient is zero in each model. For Model A, $\chi^2(1) = 0.07$ and $p = 0.7843$; for Model B, $\chi^2(1) = 0.23$ and $p = 0.6283$; and for Model C, $\chi^2(1) = 1.24$ and $p = 0.2659$. Thus, in all cases we fail to reject the null, indicating that the net female effect under *Gain & Loss* is not significantly different from zero.

(λ_{loss}) substantially reduces the RRR for both "Retreat vs. Exploit" and "Explore vs. Exploit" (RRRs of about 0.52–0.55), suggesting that more loss-averse participants are less likely to deviate from exploiting their current slot. These findings confirm that individual risk preferences are an important determinant of exploratory behaviour. Yet, our results also show that risk tolerance does not fully account for exploration in this setting. Even after controlling for risk and loss aversion, systematic gender differences remain (e.g. see Model B), suggesting that the task captures additional motives, such as information acquisition and uncertainty reduction, beyond classical risk attitudes [Wilson et al., 2014, Knox et al., 2012].

Among the remaining controls, *Agreeableness* increases the likelihood of leaving an exploited slot (whether via retreat or exploration), whereas being born in Italy lowers it (RRRs < 1). *Competitiveness*, *CRT*, *STEM*, *GPA*, and *Extraversion* show no consistent effects. Hence, Hypothesis 5 is supported, in that risk preferences, loss aversion, some personality traits, and certain demographic characteristics significantly influence exploration decisions.

5.2 Competition

5.2.1 Descriptive Evidence

We now turn to Hypothesis 3, which posits that, controlling for individual characteristics, women are less likely than men to enter competition, which in our setting means we expect women to be less likely to choose the competitive payment scheme in the final (Choice) stage of the Grain Game. In this stage, participants had to decide between receiving a piece-rate compensation of 1 cent per point or challenging the highest scorer from the previous (Tournament) stage, with a higher payoff of 4 cents per point awarded only if they outscored the tournament winner.

Looking at the descriptive statistics, in the *Gain-only* treatment, women are significantly less likely to compete compared to men (21% vs 36% for men, p=0.0183, $\chi^2 = 5.5548$), while no gender difference is found in the *Gain&Loss* treatment(27% for women vs. 28% for men p 1, $\chi^2 = 0.009$). See Figure 4.

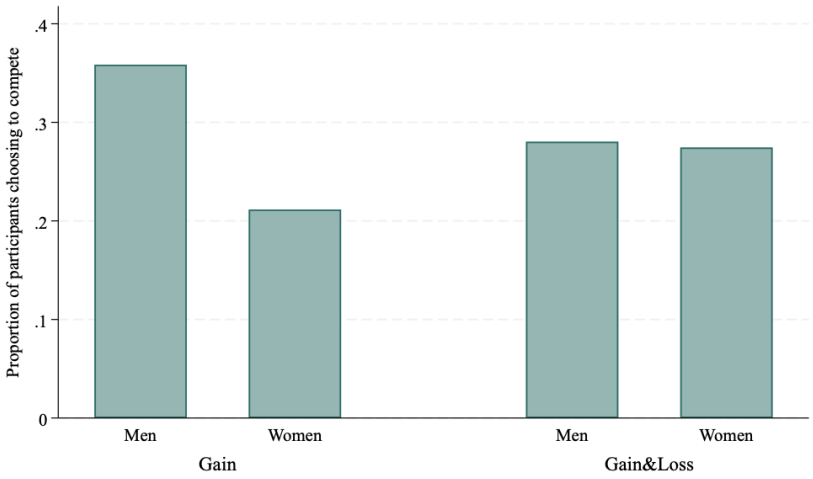


Figure 4: Decision to compete by Gender and Treatment. *Note*: This figure displays the proportion of participants who chose to compete in the Choice stage of the Grain Game disaggregated by gender (men/women) and by treatment condition (Gain vs. Gain & Loss).

5.2.2 Regression Analysis

To test Hypothesis 3, we estimate four logistic regressions of the binary decision to select the competitive payment scheme ($Compete = 1$) versus the piece-rate one ($Compete = 0$). Formally, let $\Pr(Compete_i = 1)$ denote the probability that participant i chooses to compete in the Choice stage. Our model is:

$$\begin{aligned} \ln\left(\frac{\Pr(Compete_i=1)}{\Pr(Compete_i=0)}\right) = \beta_0 \; &+ \; \beta_1 \, Female_i \; + \; \beta_2 \, Gain\&Loss_i \\ &+ \; \beta_3 \left(Female_i \times Gain\&Loss_i\right) \; + \; \beta_4 \, PieceRatePayoff_i \\ &+ \; \boldsymbol{\gamma}' \mathbf{X}_i \; + \; \varepsilon_i, \end{aligned} \tag{2}$$

where $Female_i$ is a dummy variable equal to 1 for women and 0 for men, $Gain\&Loss_i$ indicates whether participant i is assigned to the $Gain\&Loss$ treatment ($= 1$) or not ($= 0$). $PieceRatePayoff_i$ denotes participant i's payoff in the Piece Rate stage (a proxy for ability), $\mathbf{X}_i$ is a vector of additional controls, and ε_i is the error term. We estimate four versions of Equation (2) that differ in the participant-level controls included in $\mathbf{X}_i$. In Model 1, we use the parsimonious specification, with no additional participant-level controls. Model 2 adds self-assessed relative rank from the Tournament stage (a proxy for perceived ability). Model 3 introduces self-reported risk attitudes, a self-reported measure of competitiveness, cognitive reflection (CRT), personality traits (TIPI), and demographic controls (including a STEM-major indicator). Finally, Model 4 replaces self-reported risk with the incentivized Bomb Risk Elicitation Task (BRET) measure and incorporates the loss-aversion parameter λ_{loss}.[18] Table 4 presents the results from all four models.[19]

[18] Following Gächter et al. [2022], we exclude non-monotonic responses in the loss-aversion task. Because some participants did not complete these tasks or provided non-monotonic responses, the sample size is smaller in Model 4 than in the other specifications

[19] As a robustness check, we also ran regressions controlling for both piece-rate performance and tournament performance, and found no difference in the interpretation of the results.

Table 4: Logistic Regression Results on the Decision to Compete

Dependent Variable	*Compete* = 1 if Tournament, 0 if Piece Rate			
Model:	(1)	(2)	(3)	(4)
Panel A: AME, ×100				
Female	-7.28	-4.95	4.68	3.61
	(4.33)	(4.38)	(4.75)	(5.22)
Gain & Loss	-2.11	-1.03	-0.28	-0.55
	(4.33)	(4.27)	(4.23)	(4.78)
Piece rate payoff	0.03*	0.03	0.02	0.02
	(0.01)	(0.02)	(0.01)	(0.01)
Rank Belief	–	-10.13**	-9.80**	-11.78**
		(3.07)	(3.03)	(3.40)
Risk (self-reported)	–	–	2.88*	–
			(1.26)	
Competitiveness (self-reported)	–	–	3.49**	3.43**
			(1.01)	(1.11)
Risk (Bret)	–	–	–	0.31*
				(0.15)
Loss Aversion (λ_{loss})	–	–	–	-12.58**
				(4.44)
CRT	–	–	1.50	0.88
			(2.12)	(2.43)
Extraversion	–	–	0.87	2.44
			(1.49)	(1.63)
Agreeableness	–	–	2.42	4.27
			(2.24)	(2.48)
Conscientiousness	–	–	-0.52	0.10
			(1.66)	(1.79)
Emotional Stability	–	–	1.46	2.18
			(1.74)	(2.01)
Openness	–	–	-0.78	0.28
			(2.25)	(2.48)
GPA	–	–	-6.64**	-7.36*
			(2.52)	(2.89)
Born in Italy	–	–	6.98	6.03
			(7.47)	(7.65)
Age	–	–	0.75	0.94
			(0.70)	(0.91)
STEM	–	–	3.52	4.60
			(4.42)	(4.90)
Observations	419	419	407	310
Pseudo R^2	0.0210	0.0480	0.1169	0.1693
Panel B: Female AME (conditional on Gain vs. Gain&Loss), ×100				
AME(Female | Gain only)	-14.55*	-11.07	-1.68	0.35
	(6.23)	(6.23)	(6.26)	(7.73)
AME(Female | Gain & Loss)	-0.19	1.16	11.17	5.94
	(6.04)	(6.04)	(6.42)	(6.36)

Note: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Panel A reports average marginal effects (AME), multiplied by 100, with robust standard errors (also ×100) in parentheses. Panel B shows the AME of being *Female* conditional on (*Gain*) vs. (*Gain & Loss*), again multiplied by 100. While most variables are self-explanatory, note that *Loss Aversion* (λ_{loss}) is our measure of loss aversion where higher values indicate higher loss aversion, and *CRT* is the Cognitive Reflection Test score, where a higher value indicates more reflective (rather than intuitive) thinking.

As we can see in Panel A of Table 4 being female has no significant effect on the probability of competing in any specification. Panel B of Table 4 shows the average marginal effects (AMEs) of being female under the two treatments: *Gain-only* and *Gain&Loss*, separately. Focusing on Model (1), we see that in the *Gain-only* treatment, women are on average 14.6 percentage points less likely ($p < 0.05$) to choose to compete than men. By contrast, in *Gain&Loss* treatment, the AME of being female is statistically insignificant. Once we include participants' self-assessed rank (Model 2) and additional controls for risk preferences, competitiveness, cognitive reflection, personality traits, and demographics (Models 3 and 4), the gender gap observed in *Gain-only* disappears, dropping from −14.6 to −1.7 percentage points or smaller and being no longer statistically significant. In other words, the initial female disadvantage in choosing competition under *Gain-only* can be fully explained by differences in confidence, risk attitudes, personality traits, competitiveness, and demographic characteristics.

Turning to the other covariates, *Rank Belief* has a significant and negative effect on the decision

to compete in every specification where it is included.[20] Unsurprisingly, both *risk attitude* measures (self-reported risk (Model 3) and the BRET measure (Model 4)) also affect the decision to compete, with more risk-averse individuals being less likely to chose the competitive payment scheme. Similarly, *loss aversion* reduces the probability of competing, whereas *competitiveness* increases it. Finally, an increase in *GPA* is associated with a decrease in the probability of competing. This finding also supports the competition aspect of Hypothesis 5, as loss aversion, risk preferences, and self-confidence (measured by rank belief) all affect the likelihood of entering competition.

Finally, to test whether women attending STEM fields of study are more competitive than men, we run two additional specifications based on Equation (2). Specifically, we extend Model 3 and Model 4 by adding an interaction term between *STEM* (an indicator for studying a STEM field) and *Female* (let's rename these models Model 3a and Model 4a). These augmented models control for the same covariates as before (e.g, risk aversion, competitiveness, GPA, and personality traits).[21].

The results in Table 5 indicate that, once these covariates are accounted for, female STEM participants do not differ significantly from men (whether in STEM or not) in their probability of entering competition. All pairwise comparisons of marginal effects yield p-values above the conventional thresholds for statistical significance. Hence, we do not find evidence that women in STEM fields are more likely to choose the competitive payment scheme than men.

Table 5: Female STEM vs. Men: Predicted Probability (%) of Choosing Competition (Models 3a and 4a)

	Model 3a (N=407)	Model 4a (N=310)
Panel A: Predicted Probabilities (%)		
Male Non-STEM	23.3	23.5
	(3.8)	(4.6)
Male STEM	29.1	28.7
	(4.3)	(4.8)
Female Non-STEM	30.4	27.8
	(4.3)	(4.9)
Female STEM	30.4	31.4
	(5.3)	(5.1)
Panel B: Key Pairwise Differences (p.p.)		
Female STEM − Male STEM	+1.3	+2.7
	(6.9)	(7.2)
	[p=0.85]	[p=0.71]
Female STEM − Male Non-STEM	+7.1	+7.9
	(6.6)	(6.8)
	[p=0.28]	[p=0.25]

Note: Panel A shows the predicted probability (%) of choosing the competitive payment scheme for each subgroup. Panel B shows the corresponding pairwise differences in percentage points (p.p.). Standard errors are in parentheses; p-values in brackets. Both Model 3a and Model 4a include the same covariates as Models 3 and 4, respectively, with the addition of an interaction between *Female* and *STEM*.

[20] Recall that rank ranges from 1 (winner) to 4 (last place), so a one-point increase in believed rank (i.e. believing one performed worse) decreases the probability of competing by about 10–11 percentage points.

[21] The regressions results are reported in Appendix G

5.3 Participants' performance

5.3.1 Payoffs and maximum positions reached by treatment and gender

Table 6 reports average payoffs and maximum positions reached by gender and treatment for each of the three stages. While there are no significant differences between men and women in the Piece Rate and Tournament stages as for the payoffs, a marginally significant gap arises in the Choice stage, with men earning more points than women overall. At the treatment level, this marginal significance is present only in the *Gain-only* treatment and not in *Gain&Loss*. Notably, men's payoffs also exceed women's in earlier stages and both treatments, but these differences are never statistically significant.

Turning to the maximum position reached by participants, gender differences are statistically significant in the Gain treatment during both the Tournament and Choice stages. Women tend to reach positions farther from their initial ones compared to men. Consistently, across all stages and Treatments—except for the *Gain&Loss* treatment in the piece-rate stage—women systematically reach positions farther from their initial ones than men. However, these differences are never statistically significant in the *Gain&Loss* treatment.

To sum up, women tend to earn fewer points than men and reach positions farther from their initial ones in the *Gain-only* Treatment. These results suggest that women underperform relative to men due to a tendency to over-explore in the Gain Treatment, whereas no gender differences in overall performance are observed in the *Gain&Loss* Treatment.

Since a closed-form solution to the exploit-explore dilemma participants face in our game does not exist, we employed a reinforcement learning model to derive a benchmark outcome. This benchmark allows us to further evaluate participants' performance by comparing it to their output in the piece-rate

stage of the game, where no strategic interaction occurs among participants.

Table 6: Payoff and Maximum Position, by Gender, Treatment and Stage (Two-Sided p-values)

	Payoff			**Max Position**		
	Men (SE)	Women (SE)	p-value	Men (SE)	Women (SE)	p-value
Piece Rate						
Gain	836.64 (17.60)	828.57 (15.15)	0.7288	32.16 (1.59)	35.03 (1.71)	0.221
Gain & Loss	880.05 (13.03)	867.69 (13.72)	0.5141	30.58 (1.56)	30.23 (1.65)	0.876
Overall	858.45 (11.01)	847.94 (10.30)	0.4867	31.37 (1.11)	32.65 (1.20)	0.432
Tournament						
Gain	920.44 (13.78)	894.82 (13.16)	0.1803	31.73 (1.81)	38.09 (1.72)	0.012
Gain & Loss	912.95 (16.20)	896.09 (14.54)	0.4408	29.98 (1.73)	31.78 (1.73)	0.462
Overall	916.68 (10.62)	895.45 (9.77)	0.1426	30.85 (1.25)	34.97 (1.24)	0.020
Choice						
Gain	961.89 (11.10)	929.83 (14.73)	0.0829	32.37 (1.92)	39.67 (1.84)	0.007
Gain & Loss	947.35 (13.32)	929.43 (12.41)	0.3273	33.31 (1.86)	34.01 (1.90)	0.792
Overall	954.58 (8.67)	929.63 (9.62)	0.0544	32.84 (1.33)	36.87 (1.33)	0.033

Notes: Reported values are the means of payoff and max position (with standard errors in parentheses) for "Gain", "Gain & Loss" and "Overall" (pooled). Each row compares Men vs. Women. All p-values are two-sided from two-sample t-tests.

5.3.2 Reinforcement-learning benchmark in the piece-rate stage

We trained two reinforcement learning models, Q Learning and SARSA, to build a numerical benchmark for the first stage of the Grain Game. Each model plays the role of the farmer who learns through trial and error to maximise total points. At every slot, the farmer can *retreat* by moving one step left, *exploit* by planting in the same slot, or *explore* by moving one step right. The payoff in a slot equals the deterministic value of the underlying sequence plus a zero-mean random fluctuation, mirroring de facto the experimental conditions for the piece-rate stage (see Section 3).

The two models are trained on a grid of 1 728 parameter combinations. Each combination runs for 10 000 learning episodes followed by 5 000 test episodes. The full algorithmic details, grid search procedure, additional observations as well as the top performing configurations appear in Appendix I. We use the resulting outcomes as reference points to evaluate human performance, both in terms of final payoff and the maximum position achieved. We keep the configuration that attains the highest average payoff for each of the six payoff sequences. Recall that the first three sequences refer to the Gain Treatment, while the others refer to the Gain & Loss Treatment.[22]

Table 7 reports these six benchmarks both in terms of final payoff and maximum position achieved (second and sixth columns, respectively). Those outcomes are compared to the average experimental payoffs and the average maximum position achieved for male and female participants in each sequence of the piece-rate stage.[23] Participants fall short of the benchmark in all sequences except sequence 4, where

[22] See Appendix B for a description of the payoff sequences.

[23] Note that the information presented in Table 7 differs from that in Table 6, as the former reports performance by stage

they exceed it by about 350 points.[24] Since the algorithm's performance is unreliable in this sequence, we refrain from commenting on sequence 4.

When comparing human performance to the benchmark, we find that—excluding sequence 4—the benchmark payoff consistently exceeds that of the participants. Regarding the maximum position reached, we observe statistically significant deviations from the benchmark: over-exploration in sequences 3 and 6, with women tending to over-explore more than men, and under-exploration in sequence 5.

Across sequences, the difference in payoffs between male and female participants never exceeds approximately 40 points, and none of these differences is statistically significant.

Table 7: Piece-rate performance relative to reinforcement-learning (RL) benchmarks

	Payoff				**Maximum position**			
Seq	RL bench.	Gap: Overall	Gap: Men	Gap: Women	RL bench.	Gap: Overall	Gap: Men	Gap: Women
1	808.90	-62.43**	-61.92	-62.97*	35	-2.67	-3.38	-1.90
		(22.10) [87]	(34.07) [45]	(28.04) [42]		(1.82) [87]	(2.66) [45]	(2.49) [42]
2	933.68	-85.13***	-74.53**	-95.35***	29	2.93	2.07	3.75
		(14.58) [55]	(21.21) [27]	(20.23) [28]		(2.36) [55]	(3.19) [27]	(3.51) [28]
3	966.64	-36.61***	-29.20**	-44.02**	25	11.51***	8.74**	14.29***
		(8.37) [68]	(9.94) [34]	(13.49) [34]		(1.99) [68]	(2.51) [34]	(3.05) [34]
4	509.22	351.99***	362.63***	341.35***	43	-7.64***	-8.30**	-6.97*
		(24.23) [66]	(34.34) [33]	(34.61) [33]		(1.85) [66]	(2.42) [33]	(2.84) [33]
5	944.08	-100.67***	-83.81***	-120.08***	41	-14.08***	-12.71***	-15.67***
		(12.03) [71]	(16.79) [38]	(16.85) [33]		(1.94) [71]	(2.72) [38]	(2.78) [33]
6	977.57	-61.64***	-69.13***	-54.16***	15	14.31***	14.22***	14.39***
		(9.21) [72]	(13.75) [36]	(12.31) [36]		(1.96) [72]	(2.86) [36]	(2.72) [36]

Note. The RL benchmark payoff is the return generated by the best performing reinforcement learning algorithm described in Appendix I. The maximum-position benchmark is the furthest position reached by the same algorithm in each sequence. Each gap equals the experimental value minus the corresponding RL benchmark. Stars indicate two-sided t-test significance: * $p < 0.05$ ** $p < 0.01$ *** $p < 0.001$. Standard errors are in parentheses and observation counts in brackets.

6 Discussion

Our findings indicate that the gender gap in exploration is not driven by risk aversion or loss aversion. Rather, the pattern is *consistent with gender differences in information-seeking*: when losses are not possible (*Gain-only*), women explore more than men; once losses become possible (*Gain&Loss*) and the cost of deviating from exploitation rises, gender differences disappear. Competition *per se* does not induce exploration, and incentivized controls for risk, loss aversion, and confidence do not account for the exploration gap in *Gain-only* (although beliefs help to account for the raw gap in competitive entry).

To further qualify behaviour by gender, we focus on *Gain&Loss*, where the exploration gap vanishes,

and treatment, while the latter reports performance by sequence.

[24] Both algorithms learn poorly in sequence 4, so the positive gap there reflects a weakness of the benchmark rather than exceptional human performance.

and study responses to realized outcomes at the move level. Regression evidence (see Appendix J) shows that men reduce exploration after a realized negative payoff (a threshold response), whereas women reduce exploration at the treatment level when losses are possible in *Gain&Loss* (even absent a realized loss). These patterns explain why gender gaps disappear when losses are possible, without attributing the gap to gender-specific risk or loss preferences. Relative to *Gain-only*, men adjust exploration *conditionally* on realized outcomes—exploring more after positive payoffs and less after negative payoffs[25]—whereas women reduce exploration *unconditionally* once losses become possible at the treatment level.

For men, the behaviour observed in *Gain&Loss* fully aligns with Chin et al. [2023], which does not analyze gender heterogeneity. Conditioning on gender, however, women display a different pattern in our sample. Differences between our results and theirs may reflect sample composition and context (e.g., gender shares and demographics). Overall, our evidence is consistent with women placing a higher marginal value on information when losses are not possible and processing payoff information differently once losses are possible—i.e., adopting a different (and slightly less efficient according to our analysis in Subsection 5.3) decision rule in *Gain&Loss*.

Table 8 summarizes the main results, their interpretation, and the mechanism.

Our findings are partially consistent with Hypothesis 1: there is a gender gap in exploration and exploration declines when losses are possible. However, the sign of the gap is opposite to our preregistered prediction—women explore more than men in *Gain-only*. We view this as informative rather than anomalous. As discussed above, the pattern is consistent with women placing a higher marginal value on information when losses are not possible, which leads them to explore more. Hence, the results sharpen Hypothesis 1 by identifying when the gap emerges (in *Gain-only*) and why it vanishes: once losses become possible, the payoff to additional information-seeking falls.

[25] Men's overall share of exploration is similar across treatments; thus the change comes from outcome-contingent adjustments—more exploration after positive payoffs and less after negative payoffs.

Table 8: Summary of results and mechanism. Each row links to the formal hypotheses in Section 2.1.

Experimental design	Holds fixed	Activates (mechanism)	Key result	Hyp.
Environment switch: *Gain-only* → *Gain&Loss*	Task technology, information structure, timing	Possibility of losses	Women explore more than men in *Gain-only*; the exploration gap *closes* in *Gain&Loss*. This pattern is consistent with women's higher information-seeking when losses are not possibile.	1
Incentives: Piece rate; Compulsory Tournament; Self-Selection into Competition	Task, feedback, learning complexity	Competitive framing (information structure unchanged)	Competition does *not* raise exploration. The raw gender gap in competitive entry is present in *Gain-only* but *vanishes* in *Gain&Loss*.	2 and 3
Incentivized controls: risk tolerance, loss aversion, confidence	Environment, incentives	Rival channels measured (not activated)	Controlling for risk preferences, loss aversion, and confidence does *not* eliminate the gender gap in exploration in the *Gain-only* setting. These individual characteristics, and confidence in particular, instead help account for the gender gap in self-selection into competition.	3 and 5
Payoff dynamics: delayed reactions to outcomes and a loss-trigger threshold	All remaining features of the task	Higher subjective cost of deviating from exploitation when losses are possible	Men reduce exploration after a realized negative payoff (threshold response); women show no significant threshold response but reduce exploration at the *treatment* level in *Gain&Loss*. As losses become possible, behaviours converge and gender differences disappear.	1

Note: The table recaps the main empirical patterns and their mechanism-consistent interpretation. Pre-registered hypotheses are detailed in Section 2.1. Results summarized in the first row are displayed in Table 3. The second row is based on Table 3 and Figure 4. The third row is based on Tables 3 and 4. Appendix J reports the threshold regressions underlying the dynamics in the last row.

Regarding entry into competition, our descriptive statistics show that in the *Gain-only* treatment, women are less likely to choose the competitive scheme than men (21% vs. 36%). By contrast, there is no gap in the *Gain&Loss* treatment (27% vs. 28%). Once we control for individual characteristics, the gender gap in the *Gain-only* setting disappears. In particular, adding self-assessed rank beliefs (i.e. confidence in one's own performance) fully explains the raw gender difference in the propensity to compete. This finding does not support the idea that women are less likely to compete once individual-level heterogeneity is taken into account (Hypothesis 3). We also do not detect any evidence that women in STEM fields are more competitive than men (Hypothesis 4), since adding an interaction between being female and studying a STEM major yields no significant effect. The absence of a gender gap in competition contrasts with the results documented in Niederle and Vesterlund [2007] and subsequent studies, such as Niederle

and Vesterlund [2011] and Buser et al. [2014]. A possible explanation for this divergence lies in the specific task employed in our experiment, which differs substantially from the addition tasks used in Niederle and Vesterlund [2007] and the majority of studies replicating that paradigm. Our task was designed to be more abstract and gender-neutral, as suggested by Shurchkov [2012] and Gneezy et al. [2009], to minimize stereotype-related effects and ensure equal comfort across genders. This design choice may have mitigated the gendered behaviours typically observed in tasks perceived as stereotypically male-oriented.

Finally, we find substantial support for the influence of risk preferences, loss aversion, and cognitive or personality traits on both exploration and competition decisions (Hypothesis 5). Participants who are more risk-averse or more loss-averse generally explore less and show a lower inclination to enter competition.

Although we did not initially formulate a specific assumption regarding this aspect, our findings reveal slight differences in payoffs between women and men across the stages and treatments of the game, consistently favouring men. However, the gender gap in payoffs reaches statistical significance at the 10% level only in the Choice stage of the *Gain-only* treatment. Examining the maximum position reached by participants, we observe that women systematically reach positions farther from their initial ones than men, except in the *Gain&Loss* treatment during the piece-rate stage. Since women tend to earn fewer points than men and reach positions farther from their initial ones in the *Gain-only* treatment, this underperformance may be attributed to a tendency to over-explore in that treatment. Comparison with the benchmark performance derived from reinforcement learning models supports these observations.

7 An application: jobseekers' behaviours

In most countries, girls achieve higher academic performance than boys during their school years [Del Pero and Bytchkova, 2013]. Later in life, women attain higher levels of education than men [Bertrand, 2020] and, in many countries, surpass men in college achievement.[26] Despite these educational advantages, women are less likely to be employed than men, and when employed, they earn less [Bertrand, 2020].

To explain the early gender gaps emerging in the labor market, even among young workers with comparable academic skills and backgrounds, recent research has highlighted gendered preferences for job attributes. Women are more likely to seek family-friendly occupations with shorter commuting times (e.g., Le Barbanchon et al. 2021; Fluchtmann et al. 2024). Additionally, studies have emphasized gender differences in personality traits, with women generally being more risk-averse, less willing to compete, and less confident compared to men (e.g., Vesterlund 1997; Flinn et al. 2020). Among high-skilled young graduates, Cortés et al. [2023] find a clear gender difference in the timing of job offer acceptance, with women accepting jobs substantially earlier than men. The authors show that gendered job search

[26] See, for example, Conger and Long [2010] for the USA; Piazzalunga [2018] and Bovini et al. [2024] for Italy; Verbree et al. [2023] for the Netherlands; and Carroll [2023] for the UK.

strategies are driven by differences in risk preferences and overconfidence. Women, being more risk-averse, tend to have lower reservation wages, begin their job search earlier, and accept offers sooner. Conversely, men exhibit greater optimism regarding potential job offers, maintain higher reservation wages, and delay acceptance. The empirical evidence provided by Cortés et al. [2023] is further supported by findings from a laboratory experiment designed to test the previous intuition. In their setup, job seekers are uncertain about their skills and choose their reservation wages over multiple rounds.

Our intuition is that, beyond risk aversion and (under)confidence, inefficient job search behaviours could also stem from less effective strategies for navigating explore–exploit tasks. This dimension has received little attention so far. Although stylized and abstract, our lab experiment offers a novel perspective on such tasks, possibly complementing the evidence in Cortés et al. [2023]. In addition, our study may provide insights into whether gender differences in individual traits—such as willingness to compete—play a role in shaping behaviours in explore–exploit tasks.

In our setting, gender differences in exploration are *not* driven by risk or loss aversion. In incentivized elicitations we find no gender differences in these measures, and controlling for them does not remove the female exploration premium in *Gain-only*. Together, this evidence is consistent with women placing a higher marginal value on information when losses are not possible, rather than with differences in risk-based primitives.[27] This contrasts with Cortés et al. [2023], where gender gaps in job-seeking strategies depend on gender gaps in risk-related (and confidence) measures.

At the same time, this information-seeking motive does not survive the possibility of losses: in *Gain&Loss* the gender gap in exploration vanishes. Focusing on *Gain&Loss* allows us to study move-by-move responses to outcomes. We find that men adjust exploration conditionally on realized outcomes—exploring less after a negative payoff (threshold response)—whereas women exhibit a treatment-level adjustment, reducing exploration once losses are possible, regardless of the sign of the most recent payoff (see Appendix J). In *Gain-only* (in the Choice regime), women progress farther and earn less, which is consistent with over-exploration when the opportunity cost of exploration is low, i.e., when losses are not possible.

Overall, these results provide new evidence on gender differences in dynamic learning within explore–exploit environments and may supplement the findings of Cortés et al. [2023].

Our findings reveal that in the *Gains&Losses* treatment, all gender disparities disappear, with men and women exhibiting equal behaviour and performance in the explore-exploit task. Notably, the gap in willingness to compete is entirely eliminated for women from all academic backgrounds, including STEM and humanities. This result suggests that a riskier, high-stake framework involving both gains and losses curtails women's tendency to overexplore and mitigates their reluctance to compete. This pattern echoes the "gender-equality paradox" in STEM fields documented by Stoet and Geary [2018], who explore the

[27] We cannot rule out that differences in information seeking are partly mediated by lower confidence among women; however, incentivized confidence controls do not account for the exploration gap in *Gain-only*.

relationship between gender equality and women's representation in STEM disciplines. They show that countries with higher levels of gender equality have lower percentages of female STEM graduates. As a possible explanation, the authors propose that in wealthier, gender-equal countries, women are more likely to pursue careers aligned with personal preferences. Conversely, in poorer, less gender-equal countries, economic necessity drives women toward STEM fields due to their higher earning potential and career opportunities.

In the next section, we discuss policy implications inspired by the finding that environments with greater risks and high stakes reduce gender gaps. Proposed policies include career counselling programs emphasizing the high-stakes nature of early career decisions, mentorship initiatives to guide young women in navigating high-reward paths, and public awareness campaigns highlighting the long-term impact of initial job choices.

8 Conclusion

Our study examines gendered decision-making in explore–exploit tasks, a topic largely neglected in the literature on gender gaps in economic behaviour and outcomes. We show that gender differences in explore–exploit behaviour are state contingent and consistent with differences in both information seeking and information processing. In *Gains-only*, women explore more than men, suggesting that they place a higher marginal value on information; once losses are possible (*Gain&Loss*) and the opportunity cost of deviating from exploitation rises, exploration converges across genders. In *Gains-only*—especially under the Choice regime—women earn less despite greater exploratory effort, suggesting over-exploration relative to men and to a reinforcement-learning benchmark: women explore too much to fully exploit known alternatives yet not enough to discover and capitalize on unknown prospects.

Competition *per se* does not increase exploration, and the raw gender gap in competitive entry is largely accounted for by beliefs/confidence when losses are possible. In *Gain&Loss*, move-level regressions reveal distinct adjustment margins: men reduce exploration after realized negative payoffs (a threshold response), whereas women adjust at the treatment level once losses are possible. Thus, even when the amount of exploration is the same across genders, choice dynamics differ: women do not condition on realized payoffs, while men do, suggesting more payoff-contingent (and hence more sophisticated) adjustment among men. Overall, the possibility of incurring in losses attenuates gendered information-seeking deviations from exploitation, aligning both exploration and competitive entry across genders.

Future research should investigate whether alternative contexts or additional social and psychological factors amplify or mitigate these dynamics. Exploring such factors could provide a deeper understanding of the mechanisms behind gendered behaviours in high-stakes environments like *Gain&Loss*.

Our most striking result is the complete elimination of gender gaps, in both behaviours and outcomes,

in riskier, high-stakes environments with gains and losses. Such settings appear to reduce women's tendency to overexplore while fostering equal participation in competition across genders. This outcome offers valuable insights for policy interventions aimed at promoting gender equality in competitive domains.

These findings have practical implications for real-world explore–exploit problems. A particularly salient example is job search, where individuals must balance searching for better opportunities with exploiting existing ones.

A large body of empirical evidence documents gender gaps in labor market outcomes (among many others, Bertrand 2020, Le Barbanchon et al. 2021). Other papers emphasize that those gaps emerge from the very beginning: despite outperforming men academically, women face disadvantages in employment rates, salaries, and the quality of job positions from their initial entry into the workforce (Piazzalunga 2018, Bovini et al. 2024, Cortés et al. 2023). Inefficient search strategies among women may partially explain this phenomenon. In this respect, our findings suggest that gender gaps in search strategies could be mitigated by emphasizing the high-stakes nature of early labor market entry to job applicants.

This could be achieved through targeted interventions in schools and universities, such as career counselling programs that focus on the high-stakes implications of initial career choices, guiding young women in selecting roles and industries strategically. These programs could provide detailed insights into salary trajectories, job security, and career progression, emphasizing how these decisions shape long-term outcomes. Mentorship programs pairing women with experienced professionals could highlight the risks and rewards of different career paths, offering personalized advice to navigate high-stakes early decisions effectively. Public awareness campaigns could further stress that initial career choices carry significant long-term consequences, aligning with the high-stakes framework of gains and losses. By equipping women with the tools and knowledge to approach these critical decisions, such initiatives could improve early career choices.

Declaration of generative AI and AI-assisted technologies in the writing process

During the preparation of this work, the authors used ChatGTP 5 to improve language and readability, with caution. After using this tool, the authors reviewed and edited the content as needed and took full responsibility for the content of the working paper.

References

Ghazala Azmat and Barbara Petrongolo. Gender and the labor market: What have we learned from field and lab experiments? *Labour economics*, 30:32–40, 2014.

Rina Azoulay-Schwartz, Sarit Kraus, and Jonathan Wilkenfeld. Exploitation vs. exploration: choosing a supplier in an environment of incomplete information. *Decision support systems*, 38(1):1–18, 2004.

Oded Berger-Tal, Jonathan Nathan, Ehud Meron, and David Saltz. The exploration-exploitation dilemma: a multidisciplinary framework. *PloS one*, 9(4):e95693, 2014.

Donald A Berry and Bert Fristedt. Bandit problems: sequential allocation of experiments (monographs on statistics and applied probability). *London: Chapman and Hall*, 5(71-87):7–7, 1985.

Marianne Bertrand. Gender in the twenty-first century. In *AEA Papers and proceedings*, volume 110, pages 1–24. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, 2020.

Tommy C Blanchard and Samuel J Gershman. Pure correlates of exploration and exploitation in the human brain. *Cognitive, Affective, & Behavioral Neuroscience*, 18:117–126, 2018.

Alison Booth and Patrick Nolen. Choosing to compete: How different are girls and boys? *Journal of Economic Behavior & Organization*, 81(2):542–555, 2012.

Giulia Bovini, Marta De Philippis, and Lucia Rizzica. The origins of the gender pay gap: education and job characteristics. Technical report, SSRN Working Paper No. 5021036, 2024.

Thomas Buser and Huaiping Yuan. Do women give up competing more easily? evidence from the lab and the dutch math olympiad. *American Economic Journal: Applied Economics*, 11(3):225–252, 2019.

Thomas Buser, Muriel Niederle, and Hessel Oosterbeek. Gender, competitiveness, and career choices. *The quarterly journal of economics*, 129(3):1409–1447, 2014.

Thomas Buser, Noemi Peter, and Stefan C Wolter. Gender, competitiveness, and study choices in high school: Evidence from switzerland. *American economic review*, 107(5):125–130, 2017.

Colin F Camerer and Robin M Hogarth. The effects of financial incentives in experiments: A review and capital-labor-production framework. *Journal of risk and uncertainty*, 19:7–42, 1999.

Juan-Camilo Cárdenas, Anna Dreber, Emma Von Essen, and Eva Ranehill. Gender differences in competitiveness and risk taking: Comparing children in colombia and sweden. *Journal of Economic Behavior & Organization*, 83(1):11–23, 2012.

Matthew Carroll. Sex gaps in education in england: Research report. Technical report, Cambridge University Press & Assessment, 2023. Available via ERIC.

Alycia Chin, David Hagmann, and George Loewenstein. Fear and promise of the unknown: How losses discourage and promote exploration. *Journal of Behavioral Decision Making*, 36(3):e2309, 2023.

Dylan Conger and Mark C Long. Why are men falling behind? gender gaps in college performance and persistence. *The Annals of the American Academy of Political and Social Science*, 627(1):184–214, 2010.

Patricia Cortés, Jessica Pan, Laura Pilossoph, Ernesto Reuben, and Basit Zafar. Gender differences in job search and the earnings gap: Evidence from the field and lab. *The Quarterly Journal of Economics*, 138(4):2069–2126, 2023.

Paolo Crosetto and Antonio Filippin. The "bomb" risk elicitation task. *Journal of Risk and Uncertainty*, 47:31–65, 2013.

Rachel Croson and Uri Gneezy. Gender differences in preferences. *Journal of Economic literature*, 47(2): 448–474, 2009.

Nathaniel D Daw, John P O'doherty, Peter Dayan, Ben Seymour, and Raymond J Dolan. Cortical substrates for exploratory decisions in humans. *Nature*, 441(7095):876–879, 2006.

Angelica Salvi Del Pero and Alexandra Bytchkova. A bird's eye view of gender differences in education in oecd countries. Technical Report 149, OECD Publishing, 2013.

Thomas Dohmen, Armin Falk, David Huffman, and Uwe Sunde. Are risk aversion and impatience related to cognitive ability? *American Economic Review*, 100(3):1238–1260, 2010.

Thomas Dohmen, Armin Falk, David Huffman, Uwe Sunde, Jürgen Schupp, and Gert G Wagner. Individual risk attitudes: Measurement, determinants, and behavioral consequences. *Journal of the european economic association*, 9(3):522–550, 2011.

Anna Dreber, Emma Von Essen, and Eva Ranehill. Gender and competition in adolescence: task matters. *Experimental Economics*, 17:154–172, 2014.

Alice H Eagly and Steven J Karau. Role congruity theory of prejudice toward female leaders. *Psychological Review*, 109(3):573–598, 2002.

Catherine C Eckel and Philip J Grossman. Differences in the economic decisions of men and women: Experimental evidence. *Handbook of experimental economics results*, 1:509–519, 2008.

Antonio Filippin and Paolo Crosetto. A reconsideration of gender differences in risk attitudes. *Management Science*, 62(11):3138–3160, 2016.

Christopher Flinn, Petra Todd, and Weilong Zhang. Personality traits, job search and the gender wage gap. Technical report, Faculty of Economics, University of Cambridge, 2020. URL https://www.repository.cam.ac.uk/handle/1810/314752.

Jonas Fluchtmann, Anita M Glenny, Nikolaj A Harmon, and Jonas Maibom. The gender application gap: Do men and women apply for the same jobs? *American Economic Journal: Economic Policy*, 16(2): 182–219, 2024.

Becky Francis. Gender, toys and learning. *Oxford Review of Education*, 36(3):325–344, 2010.

Shane Frederick. Cognitive reflection and decision making. *Journal of Economic Perspectives*, 19(4): 25–42, 2005.

Megan Fulcher and Emily F Coyle. Breadwinner and caregiver: A cross-sectional analysis of children's and emerging adults' visions of their future family roles. *British Journal of Developmental Psychology*, 29(2):330–346, 2011.

Simon Gächter, Eric J Johnson, and Andreas Herrmann. Individual-level loss aversion in riskless and risky choices. *Theory and Decision*, 92(3):599–624, 2022.

Konstantinos Georgalos. Gender effects for loss aversion: A reconsideration. *Journal of Economic Psychology*, 105:102760, 2024.

Klarita Gërxhani, Jordi Brandts, and Arthur Schram. Competition and gender inequality: A comprehensive analysis of effects and mechanisms. *American Journal of Sociology*, 129(3):715–752, 2023.

David Gill and Victoria Prowse. A structural analysis of disappointment aversion in a real effort competition. *American Economic Review*, 102(1):469–503, 2012.

Uri Gneezy, Kenneth L Leonard, and John A List. Gender differences in competition: Evidence from a matrilineal and a patriarchal society. *Econometrica*, 77(5):1637–1664, 2009.

Samuel D Gosling, Peter J Rentfrow, and William B Swann Jr. A very brief measure of the big-five personality domains. *Journal of Research in Personality*, 37(6):504–528, 2003.

Ben Greiner. Subject pool recruitment procedures: organizing experiments with orsee. *Journal of the Economic Science Association*, 1(1):114–125, 2015.

Niels Daniel Große and Gerhard Riener. Explaining gender differences in competitiveness: Gender-task stereotypes. Technical report, Jena Economic Research Papers, 2010.

Anil K Gupta, Ken G Smith, and Christina E Shalley. The interplay between exploration and exploitation. *Academy of management journal*, 49(4):693–706, 2006.

S Julie Harbin. Gender differences in rough and tumble play behaviors. *International Journal of Undergraduate Research and Creative Activities*, 8(1):4, 2023.

W. Bradley Knox, A. Ross Otto, Peter Stone, and Bradley C. Love. The nature of belief-directed exploratory choice in human decision-making. *Frontiers in Psychology*, 2:398, 2012. doi: 10.3389/fpsyg.2011.00398.

Karson TF Kung. Recalled childhood gender-related play behaviour and current gender-related occupational interests in university students: Examining the mediating roles of gender compatibility, goal endorsement, and occupational stereotype flexibility. *Frontiers in Psychology*, 13:927–998, 2022.

Carol A Lawton. Gender differences in way-finding strategies: Relationship to spatial ability and spatial anxiety. *Sex roles*, 30:765–779, 1994.

Thomas Le Barbanchon, Roland Rathelot, and Alexandra Roulet. Gender differences in job search: Trading off commute against wage. *The Quarterly Journal of Economics*, 136(1):381–426, 2021.

Shakun D Mago, Anya C Samak, and Roman M Sheremeta. Facing your opponents: Social identification and information feedback in contests. *Journal of Conflict Resolution*, 60(3):459–481, 2016.

James G March. Exploration and exploitation in organizational learning. *Organization science*, 2(1): 71–87, 1991.

Björn Meder, Charley M Wu, Eric Schulz, and Azzurra Ruggeri. Development of directed and random exploration in children. *Developmental science*, 24(4):e13095, 2021.

Katja Mehlhorn, Ben R. Newell, Peter M. Todd, Michael D. Lee, Kate Morgan, Victoria A. Braithwaite, Daniel Hausmann, Klaus Fiedler, and Cleotilde Gonzalez. Unpacking the exploration–exploitation tradeoff: A synthesis of human and animal literatures. *Decision*, 2(3):191–215, 2015.

Muriel Niederle and Lise Vesterlund. Do women shy away from competition? do men compete too much? *The Quarterly Journal of Economics*, 122(3):1067–1101, 2007.

Muriel Niederle and Lise Vesterlund. Gender and competition. *Annu. Rev. Econ.*, 3(1):601–630, 2011.

Daniela Piazzalunga. The gender wage gap among college graduates in italy. *Italian Economic Journal*, 4 (1):33–90, 2018.

Paul B Reverdy, Vaibhav Srivastava, and Naomi Ehrich Leonard. Modeling human decision making in generalized gaussian multiarmed bandits. *Proceedings of the IEEE*, 102(4):544–571, 2014.

Eric Schulz, Charley M Wu, Azzurra Ruggeri, and Björn Meder. Searching for rewards like a child means less generalization and more directed exploration. *Psychological science*, 30(11):1561–1572, 2019.

Eric Schulz, Nicholas T Franklin, and Samuel J Gershman. Finding structure in multi-armed bandits. *Cognitive psychology*, 119:101261, 2020.

Olga Shurchkov. Under pressure: gender differences in output quality and quantity under competition and time constraints. *Journal of the European Economic Association*, 10(5):1189–1213, 2012.

Maarten Speekenbrink and Emmanouil Konstantinidis. Uncertainty and exploration in a restless bandit problem. *Topics in cognitive science*, 7(2):351–367, 2015.

Mark Steyvers, Michael D Lee, and Eric-Jan Wagenmakers. A bayesian analysis of human decision-making on bandit problems. *Journal of mathematical psychology*, 53(3):168–179, 2009.

Gijsbert Stoet and David C Geary. The gender-equality paradox in science, technology, engineering, and mathematics education. *Psychological science*, 29(4):581–593, 2018.

Ioannis Theodossiou and Alexandros Zangelidis. Should i stay or should i go? the effect of gender, education and unemployment on labour market transitions. *Labour Economics*, 16(5):566–577, 2009.

Anne-Roos Verbree, Lisette Hornstra, Lientje Maas, and Leoniek Wijngaards-de Meij. Conscientiousness as a predictor of the gender gap in academic achievement. *Research in Higher Education*, 64(3):451–472, 2023.

Lise Vesterlund. The effects of risk aversion on job matching: Can differences in risk aversion explain the wage gap? Unpublished manuscript, Iowa State University, 1997.

Robert C. Wilson, Andra Geana, John M. White, Elliot A. Ludvig, and Jonathan D. Cohen. Humans use directed and random exploration to solve the explore–exploit dilemma. *Journal of Experimental Psychology: General*, 143(6):2074–2081, 2014. doi: 10.1037/a0038199.

Robert C Wilson, Elizabeth Bonawitz, Vincent D Costa, and R Becket Ebitz. Balancing exploration and exploitation with information and randomization. *Current Opinion in Behavioral Sciences*, 38:49–56, 2021.

Charley M Wu, Eric Schulz, Maarten Speekenbrink, Jonathan D Nelson, and Björn Meder. Generalization guides human exploration in vast decision spaces. *Nature Human Behaviour*, 2(12):915–924, 2018.

Mohan Yogeswaran and SG Ponnambalam. Reinforcement learning: Exploration–exploitation dilemma in multi-agent foraging task. *Opsearch*, 49:223–236, 2012.

Appendix A

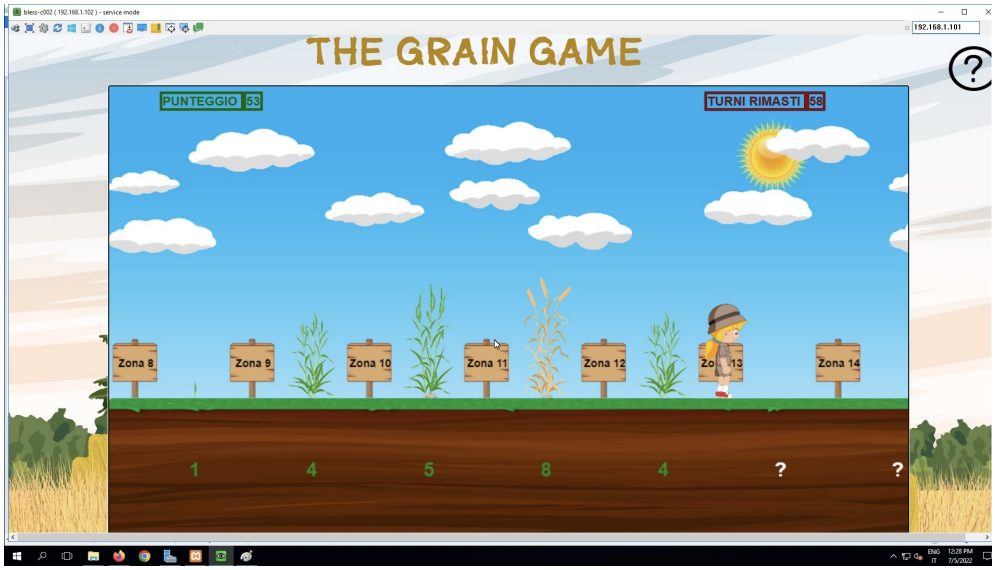


Figure A.1: Grain Game interface under the *Gain-only* treatment.

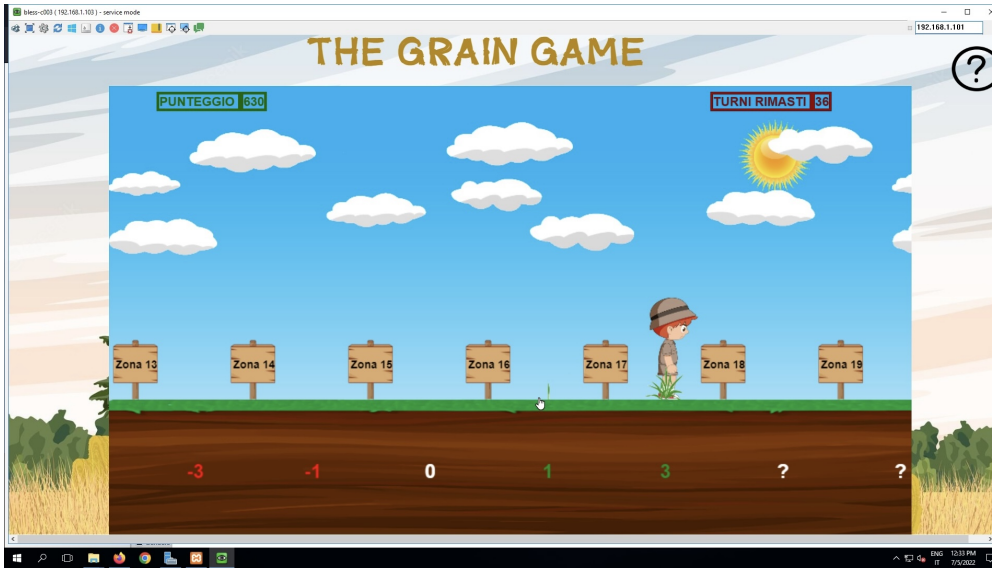


Figure A.2: Grain Game interface under the *Gain&Loss* treatment.

Appendix B

The payoff sequences used in the experiment were generated using a random walk process, with errors drawn from a normal distribution $\mathcal{N}(0, 2)$. These sequences were designed to reward exploration by ensuring that the global maximum of every sequence was placed in the second half of the field.

The sequences for the Gain treatment are shown below:

- **S1**: {7, 7, 7, 5, 0, 3, 2, 1, 4, 5, 8, 4, 7, 9, 10, 10, 12, 15, 18, 17, 17, 16, 12, 12, 11, 12, 13, 13, 11, 13, 15, 17, 15, 17, 17, 18, 16, 18, 17, 18, 18, 19, 22, 25, 23, 18, 20, 19, 20, 23, 24, 23, 23, 21, 24, 25, 25, 24, 23, 23, 23, 20, 19, 22, 21, 20, 17, 15, 17, 17}
- **S2**: {11, 8, 5, 8, 9, 11, 12, 13, 14, 11, 16, 15, 11, 10, 6, 7, 7, 10, 12, 13, 13, 11, 13, 13, 12, 11, 12, 10, 12, 12, 15, 14, 14, 13, 16, 13, 13, 15, 14, 16, 16, 13, 16, 15, 16, 20, 20, 21, 24, 23, 23, 23, 24, 25, 24, 22, 21, 18, 17, 17, 18, 18, 17, 19, 19, 18, 20, 17, 15, 19}
- **S3**: {7, 10, 14, 14, 15, 12, 13, 14, 15, 11, 12, 11, 12, 11, 12, 9, 10, 8, 9, 13, 13, 12, 15, 17, 13, 13, 16, 16, 13, 15, 17, 14, 16, 16, 17, 14, 14, 15, 15, 15, 18, 16, 11, 12, 13, 13, 11, 12, 10, 11, 12, 13, 14, 13, 15, 20, 22, 23, 25, 24, 21, 20, 21, 23, 20, 19, 19, 17, 14, 11}

The payoff sequences for the Gain & Loss treatment are derived by subtracting 10 points from the respective Gain sequences. For example:

- **S4 (Gain & Loss sequence derived from S1)**: {-3, -3, -3, -5, -10, -7, -8, -9, -6, -5, -2, -6, -3, -1, 0, 0, 2, 5, 8, 7, 7, 6, 2, 2, 1, 2, 3, 3, 1, 3, 5, 7, 5, 7, 7, 8, 6, 8, 7, 8, 8, 9, 12, 15, 13, 8, 10, 9, 10, 13, 14, 13, 13, 11, 14, 15, 15, 14, 13, 13, 13, 10, 9, 12, 11, 10, 7, 5, 7, 7}

Figure B.1 below provides a visual representation of the three payoff sequences for the Gain treatment, with a horizontal line at $Payoff = 10$. The threshold marks the transition point to negative payoffs in the Gain & Loss treatment, where a -10 offset was applied.

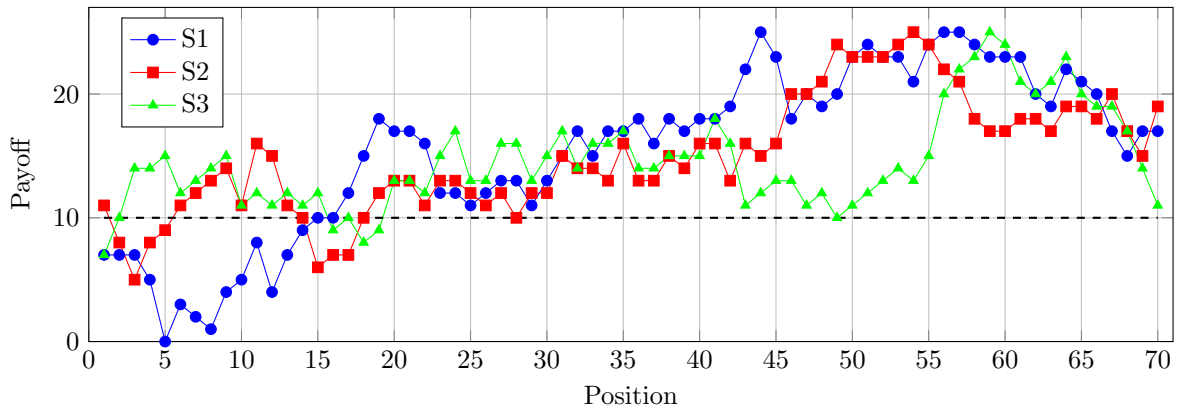


Figure B.1: Payoff sequences for the Gain treatment with a threshold line at $Payoff = 10$. *Note*: The threshold indicates the transition to negative payoffs in the Gain & Loss treatment, where a -10 offset was applied to all values.

Appendix C

This appendix provides illustrations of the interfaces used for the Loss Aversion Task (Figure C1) and the Bomb Risk Elicitation Task (BRET) (Figure C2).



Figure C1: Screen for the Loss Aversion Task. *Note*: Participants decided whether to accept or reject monetary gambles with potential gains and losses.



Figure C2: Screen for the Bomb Risk Elicitation Task (BRET). *Note*: Participants could adjust the number of red cells by selecting cells with their mouse and confirm their choice by clicking "Confirm".

Appendix D

In Table D1, we present the summary statistics for the main variables used in our analysis. The *Cognitive Reflection Test (CRT)* is measured following Frederick [2005], where a higher score indicates a greater tendency for reflective thinking rather than relying on intuitive responses. We define and construct the *TIPI* (Ten-Item Personality Inventory) as in Gosling et al. [2003], which captures concise measures of the Big Five personality traits. The measure of *Loss Aversion* (λ_{loss}) follows Gächter et al. [2022], where higher values indicate stronger aversion to losses; any non-monotonic decisions in this task are coded as missing observations, consistent with Gächter et al. [2022]. Finally, our measures of risk attitudes (*Risk (Bret)* and *Risk (Self-Reported)*) are such that higher values reflect lower levels of risk aversion.

Table D1: Descriptive statistics by Gender

	Overall (N=419)	**Men (N=213)**	**Women (N=206)**	**p-value**
Age				0.003
Mean (SD)	24.25 (3.33)	24.73 (3.68)	23.76 (2.85)	
Median [Min, Max]	24.00 [18.00, 52.00]	24.00 [19.00, 52.00]	24.00 [18.00, 41.00]	
Missing	3	3	0	
CRT				< 0.001
Mean (SD)	1.64 (1.10)	1.97 (1.02)	1.30 (1.09)	
Median [Min, Max]	2.00 [0.00, 3.00]	2.00 [0.00, 3.00]	1.00 [0.00, 3.00]	
Missing	3	1	2	
GPA				0.2
Mean (SD)	3.93 (0.87)	3.87 (0.91)	3.98 (0.82)	
Median [Min, Max]	4.00 [1.00, 5.00]	4.00 [1.00, 5.00]	4.00 [2.00, 5.00]	
Missing	3	2	1	
Risk (Self-Reported)				0.6
Mean (SD)	5.39 (1.83)	5.44 (1.94)	5.35 (1.73)	
Median [Min, Max]	5.00 [0.00, 10.00]	5.00 [0.00, 10.00]	5.00 [1.00, 10.00]	
Missing	1	1	0	
Risk (Bret)				0.2
Mean (SD)	46 (17)	45 (17)	48 (17)	
Median [Min, Max]	49 [1, 100]	46 [1, 100]	50 [1, 100]	
Missing	62	31	31	
Agreeableness				0.033
Mean (SD)	5.01 (1.11)	4.89 (1.12)	5.13 (1.10)	
Median [Min, Max]	5.00 [1.50, 7.00]	5.00 [2.00, 7.00]	5.00 [1.50, 7.00]	
Missing	1	1	0	
Extraversion				0.031
Mean (SD)	3.90 (1.54)	3.74 (1.47)	4.07 (1.60)	
Median [Min, Max]	4.00 [1.00, 7.00]	3.50 [1.00, 7.00]	4.00 [1.00, 7.00]	
Missing	1	1	0	
Conscientiousness				0.066
Mean (SD)	5.07 (1.33)	4.96 (1.29)	5.19 (1.35)	
Median [Min, Max]	5.50 [1.00, 7.00]	5.00 [1.50, 7.00]	5.50 [1.00, 7.00]	
Missing	1	1	0	
Emotional Stability				< 0.001
Mean (SD)	4.05 (1.38)	4.38 (1.37)	3.70 (1.31)	
Median [Min, Max]	4.00 [1.00, 7.00]	4.50 [1.00, 7.00]	3.50 [1.00, 7.00]	
Missing	1	1	0	
Openness				< 0.001
Mean (SD)	4.90 (1.15)	4.67 (1.06)	5.14 (1.19)	
Median [Min, Max]	5.00 [1.50, 7.00]	4.50 [2.00, 7.00]	5.00 [1.50, 7.00]	
Missing	1	1	0	
Competitiveness (self-reported)				< 0.001
Mean (SD)	5.89 (2.40)	6.48 (2.27)	5.29 (2.39)	
Median [Min, Max]	6.00 [0.00, 10.00]	7.00 [0.00, 10.00]	5.00 [0.00, 10.00]	
Missing	1	1	0	
Loss Aversion (λ_{loss})				0.2
Mean (SD)	1.77 (0.65)	1.81 (0.71)	1.72 (0.57)	
Median [Min, Max]	1.50 [0.86, 4.00]	2.00 [0.86, 4.00]	1.50 [0.86, 4.00]	
Missing	100	46	54	
Rank Belief				0.006
Mean (SD)	2.08 (0.81)	1.97 (0.86)	2.19 (0.74)	
Median [Min, Max]	2.00 [0.00, 4.00]	2.00 [0.00, 4.00]	2.00 [1.00, 4.00]	
Born in Italy	389 (93%)	198 (93%)	191 (93%)	0.8
Missing	1	1	0	
STEM	159 (39%)	95 (46%)	64 (31%)	0.003
Missing	7	5	2	

Notes: p-values in the last column reflect a Men vs. Women comparison. They come from Two Sample t-tests for continuous variables or Pearson's Chi-squared tests for categorical variables.

Appendix E

Table E1: Final and Max Position, by Gender, Treatment and Stage

	Final Position			Max Position		
	Men (SE)	Women (SE)	p-value (stars)	Men (SE)	Women (SE)	p-value (stars)
Piece Rate						
Gain	30.74 (1.69)	33.67 (1.82)	0.238	32.16 (1.59)	35.03 (1.71)	0.221
Gain & Loss	29.10 (1.63)	28.71 (1.74)	0.868	30.58 (1.56)	30.23 (1.65)	0.876
Overall	29.92 (1.17)	31.21 (1.27)	0.452	31.37 (1.11)	32.65 (1.20)	0.432
Tournament						
Gain	30.32 (1.90)	36.93 (1.81)	0.013	31.73 (1.81)	38.09 (1.72)	0.012
Gain & Loss	28.67 (1.79)	30.68 (1.80)	0.431	29.98 (1.73)	31.78 (1.73)	0.462
Overall	29.49 (1.30)	33.83 (1.29)	0.019	30.85 (1.25)	34.97 (1.24)	0.020
Choice						
Gain	31.03 (2.01)	38.75 (1.92)	0.006	32.37 (1.92)	39.67 (1.84)	0.007
Gain & Loss	32.34 (1.92)	32.75 (2.00)	0.883	33.31 (1.86)	34.01 (1.90)	0.792
Overall	31.69 (1.39)	35.78 (1.40)	0.039	32.84 (1.33)	36.87 (1.33)	0.033

Notes: Reported values are means of Final Position and Max Position, with standard errors in parentheses. All p-values come from two-sample t-tests comparing Men vs. Women within each row. "Overall" pools Gain and Gain & Loss within each stage.

Table E2: Explore Percent, by Gender, Treatment and Stage

	Explore Percent		
	Men (SE)	**Women** (SE)	**p-value**
Piece Rate			
Gain	0.526 (0.021)	0.560 (0.024)	0.284
Gain & Loss	0.507 (0.022)	0.519 (0.022)	0.697
Overall	0.516 (0.015)	0.540 (0.016)	0.293
Tournament			
Gain	0.494 (0.026)	0.583 (0.025)	0.014
Gain & Loss	0.483 (0.024)	0.517 (0.024)	0.311
Overall	0.488 (0.018)	0.550 (0.017)	0.013
Choice			
Gain	0.492 (0.028)	0.606 (0.025)	0.003
Gain & Loss	0.525 (0.026)	0.532 (0.027)	0.842
Overall	0.509 (0.019)	0.569 (0.019)	0.023

Notes: Reported values are the mean fraction of rounds in which participants explored, with standard errors in parentheses. All p-values come from two-sample t-tests comparing Men vs. Women in each row. "Overall" pools Gain and Gain & Loss within each stage.

Appendix F

Table F1: Random Effects Multinomial Logistic Results for *Move* (Piece Rate Stage)

Dependent Variable	*Move* (Baseline = *Exploit*)					
	Model A1		**Model B1**		**Model C1**	
	Retreat	Explore	Retreat	Explore	Retreat	Explore
Female	0.949	1.162	0.969	1.339	1.060	1.559
	(0.242)	(0.253)	(0.268)	(0.310)	(0.371)	(0.432)
Gain & Loss	1.099	0.936	1.079	0.970	1.135	1.155
	(0.291)	(0.200)	(0.303)	(0.212)	(0.371)	(0.288)
Female × (Gain & Loss)	1.511	0.968	1.547	0.977	1.463	0.847
	(0.543)	(0.298)	(0.567)	(0.304)	(0.652)	(0.311)
Lag Payoff	0.846***	0.853***	0.845***	0.850***	0.826***	0.819***
	(0.013)	(0.015)	(0.014)	(0.015)	(0.017)	(0.018)
Threshold	0.470***	0.604***	0.470***	0.589***	0.544***	0.622***
	(0.060)	(0.064)	(0.061)	(0.063)	(0.075)	(0.076)
Risk (self-reported)	–	–	0.966	1.030	–	–
			(0.057)	(0.049)		
Risk (Bret)	–	–	–	–	1.006	1.006
					(0.007)	(0.006)
Loss Aversion (λ_{loss})	–	–	–	–	0.645*	0.646**
					(0.112)	(0.089)
Competitiveness (self-reported)	–	–	0.946	1.037	0.936	1.031
			(0.042)	(0.039)	(0.047)	(0.042)
CRT	–	–	1.170	1.086	1.357**	1.261*
			(0.109)	(0.088)	(0.146)	(0.123)
Extraversion	–	–	1.085	0.974	1.025	0.922
			(0.070)	(0.051)	(0.078)	(0.055)
Agreeableness	–	–	1.114	1.173*	1.073	1.184*
			(0.107)	(0.088)	(0.120)	(0.100)
Conscientiousness	–	–	1.069	0.856*	1.051	0.827**
			(0.080)	(0.053)	(0.091)	(0.060)
Emotional Stability	–	–	0.972	0.939	1.068	0.985
			(0.070)	(0.059)	(0.087)	(0.069)
Openness	–	–	0.949	0.922	1.040	0.947
			(0.089)	(0.068)	(0.112)	(0.075)
GPA	–	–	0.836	0.916	0.810	0.950
			(0.094)	(0.086)	(0.103)	(0.094)
Born in Italy	–	–	0.804	0.491	0.511	0.482
			(0.348)	(0.199)	(0.225)	(0.213)
Age	–	–	1.035	1.052	1.080	1.105***
			(0.032)	(0.037)	(0.042)	(0.032)
STEM	–	–	0.926	1.080	0.935	1.144
			(0.185)	(0.180)	(0.222)	(0.214)
Constant	1.168	13.093***	0.739	12.540*	0.438	6.791
	(0.320)	(3.387)	(1.004)	(15.518)	(0.707)	(9.140)
Model Statistics						
Observations	28,492		27,676		21,080	
Number of Groups	419		407		310	
Wald χ^2	136.51		188.22		182.36	
Log pseudolikelihood	−21695.549		−21062.913		−15947.562	

Notes: All entries are Relative Risk Ratios (RRRs) from random-effects multinomial logit regressions. Robust standard errors (clustered at participant level) are shown in parentheses. The dependent variable is *Move*, and the baseline (omitted) outcome is *Exploit*. Significance levels: $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

Table F2: Random Effects Multinomial Logistic Results for *Move* (Tournament Stage)

Dependent Variable	*Move* (Baseline = *Exploit*)					
	Model A2		**Model B2**		**Model C2**	
	Retreat	Explore	Retreat	Explore	Retreat	Explore
Female	1.218	1.530	1.309	1.936*	1.083	2.228*
	(0.383)	(0.473)	(0.450)	(0.608)	(0.454)	(0.881)
Gain & Loss	1.338	0.933	1.454	0.874	1.093	1.030

(Continued on next page)

	Model A2		Model B2		Model C2	
	Retreat	Explore	Retreat	Explore	Retreat	Explore
	(0.419)	(0.283)	(0.487)	(0.258)	(0.410)	(0.357)
Female × (Gain & Loss)	1.018	0.732	0.954	0.666	1.044	0.491
	(0.451)	(0.313)	(0.433)	(0.273)	(0.555)	(0.241)
Lag Payoff	0.771***	0.763***	0.768***	0.759***	0.760***	0.758***
	(0.018)	(0.018)	(0.018)	(0.018)	(0.020)	(0.021)
Threshold	0.397***	0.670***	0.399***	0.681**	0.302***	0.683*
	(0.065)	(0.078)	(0.066)	(0.081)	(0.060)	(0.105)
Piece Rate Payoff	–	–	1.000	1.004***	1.001	1.004***
			(0.001)	(0.001)	(0.001)	(0.001)
Risk (Self-Reported)	–	–	0.911	1.120	–	–
			(0.061)	(0.066)		
Risk (Bret)	–	–	–	–	1.007	1.007
					(0.008)	(0.009)
Loss Aversion (λ_{loss})	–	–	–	–	0.551**	0.457***
					(0.101)	(0.079)
Competitiveness (self-reported)	–	–	0.946	1.024	0.909	0.987
			(0.049)	(0.053)	(0.054)	(0.059)
CRT	–	–	1.209	1.072	1.237	1.157
			(0.136)	(0.118)	(0.161)	(0.149)
Extraversion	–	–	0.982	0.977	0.948	0.953
			(0.076)	(0.067)	(0.085)	(0.076)
Agreeableness	–	–	1.320*	1.243*	1.320*	1.300*
			(0.149)	(0.132)	(0.166)	(0.155)
Conscientiousness	–	–	0.938	0.949	0.932	0.983
			(0.082)	(0.076)	(0.094)	(0.093)
Emotional Stability	–	–	0.913	0.925	0.914	0.961
			(0.075)	(0.079)	(0.085)	(0.092)
Openness	–	–	0.807	0.835	0.837	0.919
			(0.098)	(0.088)	(0.113)	(0.108)
GPA	–	–	0.989	0.922	1.020	0.983
			(0.137)	(0.123)	(0.161)	(0.152)
Born in Italy	–	–	0.225**	0.204**	0.183**	0.157**
			(0.121)	(0.106)	(0.113)	(0.105)
Age	–	–	1.023	1.043	1.024	1.069
			(0.036)	(0.038)	(0.046)	(0.050)
STEM	–	–	0.708	1.122	0.815	1.508
			(0.183)	(0.258)	(0.242)	(0.402)
Constant	1.661	58.353***	9.047	2.578	13.922	2.870
	(0.620)	(22.630)	(15.714)	(4.403)	(27.148)	(5.708)
Observations	28,492		27,676		21,080	
Number of Groups	419		407		310	
Wald χ^2	179.50		252.06		199.75	
Log pseudolikelihood	−18233.956		−17616.035		−13283.393	

Notes: All entries are Relative Risk Ratios (RRRs) from random-effects multinomial logit regressions. Robust standard errors (clustered at participant level) are shown in parentheses. The dependent variable is *Move*, and the baseline (omitted) outcome is *Exploit*. Significance levels: $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

Table F3: Random Effects Multinomial Logistic Results for *Move* (Tournament Stage with Compete control)

Dependent Variable	*Move* (Baseline = *Exploit*)					
	Model A3		Model B3		Model C3	
	Retreat	Explore	Retreat	Explore	Retreat	Explore
Female	1.140	1.593	1.267	1.932*	1.047	2.217*
	(0.367)	(0.495)	(0.442)	(0.607)	(0.450)	(0.877)
Gain & Loss	1.281	0.954	1.377	0.871	1.010	1.023
	(0.403)	(0.291)	(0.463)	(0.260)	(0.380)	(0.355)
Female × (Gain & Loss)	1.086	0.705	1.037	0.669	1.137	0.496
	(0.485)	(0.303)	(0.475)	(0.277)	(0.611)	(0.244)
Lag Payoff	0.771***	0.763***	0.768***	0.759***	0.760***	0.758***
	(0.018)	(0.018)	(0.018)	(0.018)	(0.020)	(0.021)
Threshold	0.397***	0.670**	0.399***	0.682**	0.302***	0.684*

(Continued on next page)

	Model A3		Model B3		Model C3	
	Retreat	Explore	Retreat	Explore	Retreat	Explore
	(0.065)	(0.078)	(0.066)	(0.081)	(0.060)	(0.105)
Compete (Tournament)	0.653	1.311	0.640	0.974	0.414**	0.893
	(0.166)	(0.296)	(0.172)	(0.220)	(0.138)	(0.250)
Piece Rate Payoff	–	–	1.000	1.004***	1.001	1.004***
			(0.001)	(0.001)	(0.001)	(0.001)
Risk (Self-Reported)	–	–	0.923	1.121	–	–
			(0.062)	(0.067)		
Risk (Bret)	–	–	–	–	1.010	1.007
					(0.008)	(0.009)
Loss Aversion (λ_{loss})	–	–	–	–	0.492***	0.451***
					(0.092)	(0.080)
Competitiveness (self-reported)	–	–	0.962	1.025	0.939	0.992
			(0.052)	(0.055)	(0.058)	(0.061)
CRT	–	–	1.217	1.072	1.248	1.159
			(0.137)	(0.118)	(0.164)	(0.150)
Extraversion	–	–	0.989	0.977	0.974	0.956
			(0.076)	(0.067)	(0.087)	(0.076)
Agreeableness	–	–	1.334*	1.243*	1.371*	1.306*
			(0.152)	(0.133)	(0.174)	(0.156)
Conscientiousness	–	–	0.937	0.949	0.923	0.983
			(0.081)	(0.076)	(0.092)	(0.092)
Emotional Stability	–	–	0.916	0.925	0.932	0.963
			(0.075)	(0.079)	(0.086)	(0.091)
Openness	–	–	0.800	0.835	0.828	0.920
			(0.097)	(0.088)	(0.112)	(0.108)
GPA	–	–	0.958	0.920	0.954	0.974
			(0.135)	(0.124)	(0.155)	(0.150)
Born in Italy	–	–	0.226**	0.205**	0.182**	0.158**
			(0.123)	(0.106)	(0.113)	(0.105)
Age	–	–	1.028	1.043	1.038	1.071
			(0.037)	(0.038)	(0.047)	(0.050)
STEM	–	–	0.722	1.122	0.857	1.509
			(0.187)	(0.258)	(0.256)	(0.402)
Constant	1.936	52.860***	7.494	2.535	9.764	2.690
	(0.751)	(20.931)	(13.025)	(4.353)	(18.887)	(5.374)
Observations	28,492		27,676		21,080	
Number of Groups	419		407		310	
Wald χ^2	185.44		256.03		214.48	
Log pseudolikelihood	−18232.316		−17615.063		−13280.770	

Notes: All entries are Relative Risk Ratios (RRRs) from random-effects multinomial logit regressions. Robust standard errors (clustered at participant level) are shown in parentheses. The dependent variable is *Move*, and the baseline (omitted) outcome is *Exploit*. Significance levels: $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

Appendix G

To test whether women attending STEM fields of study are more likely to choose the competitive payment scheme than men, we augment Equation (2) by adding the interaction term $\beta_5\big(Female_i \times STEM_i\big)$. Because $STEM_i$ is already included in the vector of controls $\mathbf{X}_i$, the only modification to the original model is the inclusion of this interaction. Formally, the augmented specification is:

$$\begin{aligned}\ln\!\left(\frac{\Pr(Compete_i=1)}{\Pr(Compete_i=0)}\right) = \beta_0 \;&+\; \beta_1\, Female_i \;+\; \beta_2\, Gain\&Loss_i \\ &+\; \beta_3\left(Female_i \times Gain\&Loss_i\right) \;+\; \beta_4\, PieceRatePayoff_i \\ &+\; \beta_5\left(Female_i \times STEM_i\right) \;+ \boldsymbol{\gamma}'\mathbf{X}_i \;+\; \varepsilon_i,\end{aligned} \tag{G.1}$$

where $\mathbf{X}_i$ includes the same additional controls as in Models 3 and 4.

In Table G1, we report the average marginal effects (AMEs) of the regressions estimated from Equation (G.1).

Table G1: Logistic Regression Results on the Decision to Compete

Dependent Variable	*Compete* = 1 if Tournament, 0 if Piece Rate	
Model:	(3a)	(4a)
Female	4.70	3.64
	(4.74)	(5.23)
Gain & Loss	-0.30	-0.58
	(4.23)	(4.80)
Piece rate payoff	0.02	0.01
	(0.01)	(0.02)
Rank Belief	-9.97**	-11.88**
	(3.06)	(3.51)
Risk (self-reported)	2.82*	–
	(1.27)	
Risk (Bret)	–	0.31*
		(0.14)
Loss Aversion (λ_{loss})	–	-12.58**
		(4.45)
Competitiveness (self-reported)	3.48**	3.43**
	(1.02)	(1.11)
CRT	1.65	0.91
	(2.12)	(2.45)
Extraversion	0.89	2.46
	(1.49)	(1.62)
Agreeableness	2.36	4.25
	(2.24)	(2.48)
Conscientiousness	-0.46	0.12
	(1.65)	(1.80)
Emotional Stability	1.55	2.19
	(1.73)	(2.00)
Openness	-0.86	0.21
	(2.26)	(2.53)
GPA	-6.66**	-7.40*
	(2.52)	(2.90)
Born in Italy	6.91	6.12
	(7.40)	(7.61)
Age	0.72	0.93
	(0.71)	(0.92)
STEM	3.25	4.53
	(4.39)	(4.86)
Observations	407	310
Pseudo R^2	0.1169	0.1693

Note: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. All entries are average marginal effects (AME), multiplied by 100, with robust standard errors (also ×100) in parentheses.

Appendix H

This appendix is organized as follows. In Section H.1 we give the full English translation of the instructions exactly as participants in the Gain&Loss treatment received them (the instructions were provided in written form and were also read aloud by the experimenter); in Section H.2 we show the concise help popup summary that appeared when they clicked the "?" icon; and in Section H.3 we report the control questions used to check participants' understanding of the task.[28]

H.1 Instructions shown to participants (English translation)

INTRODUCTION

You are participating in a study financed by the University of Bologna. During this study, you will earn a sum of money determined by the rules described in the following pages. Payment will be made via PayPal in a confidential manner in the days following the completion of the study.
Today's study will last approximately 1 hour. The study is composed of 2 parts. Your final earnings will consist of the earnings obtained in each of the two parts plus 5 Euros as a participation fee.
The rules we will follow to determine your earnings are different in each part, and you will receive the instructions for each part in sequence.
It is forbidden to communicate with other participants during the study. If you have any questions, raise your hand and an assistant will come to your workstation to answer in private.

INSTRUCTIONS PART 1

The first part is composed of 3 stages. At the end of Part 1, the computer will make a random draw and select one of the 3 stages that will be used to calculate your earnings in Part 1.

PART 1 – STAGE 1: INSTRUCTIONS

In the first stage, you will act as a farmer who has purchased a 70-hectare plot of land for growing wheat. Unfortunately, you do not know which parts of the land are more fertile (i.e., those capable of producing more wheat). You have 70 wheat seeds available, and you must decide where to plant them, one after another.

Each planted seed will generate points, and your goal is to obtain the highest number of points possible.

For each seed, you have two choices. You can plant a wheat seed in a hectare of land where you have never planted anything before (indicated by "?") or in a hectare where you have already planted. Note that if you decide to plant the seed in a hectare where you have already sown, you may receive a number of points slightly different from those you received in the previous turn.

At the end of each turn, you will receive points based on how much wheat has grown in the selected hectare. In each turn, you can obtain between -10 and 15 points per hectare. Since it is possible to obtain fewer than zero points, we will give you an initial endowment of 700 points.

On your screen, you will see an image like the one reproduced in Figure 1. At the top, you will see the total number of accumulated points and the number of seeds left to plant. You can also consult the instructions by clicking on the "?".

[28] To ensure the task was clearly understood, the experimenter provided the correct answers before the game began.

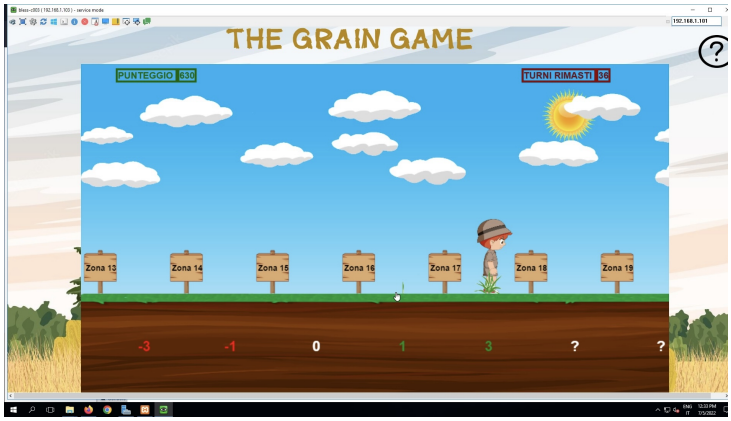


Figure 1

HOW TO PLANT THE SEEDS

1. To move right or left, use the arrow keys on your keyboard. Movements are possible one step at a time, both to the right and to the left.
2. To plant a seed, press the space bar.
3. You can plant each wheat seed in a hectare where you have already planted before, or you can move to a hectare to the right or to the left. The hectares where you can decide to plant are highlighted with a "?", while those where you have already planted are colored green (if you obtained a positive result) or red (if you obtained a negative result).
4. You can view the total number of points you have accumulated in the green box at the top left, and the number of seeds you still have to plant (i.e., remaining turns) in the red box at the top right.
5. You can use the "A" and "D" keys to scroll the field view

Remember: you have 70 wheat seeds available. The first seed has already been planted for you, and you can see it on your left. You must decide how to plant the remaining 69. Each planted seed will generate points. In each turn, you can obtain between -10 and 15 points per hectare. Since it is possible to obtain fewer than zero points, we will give you an initial endowment of 700 points.

EARNINGS FOR STAGE 1

At the end of Stage 1, we will sum all the points you have obtained. If Stage 1 is randomly drawn for payment in Part 1, you will receive 1 cent of a euro for every point obtained. We will call this payment system "piece rate."

Before starting, you will see a screen like the one reproduced in Figure 2, and you must select "NUOVA PARTITA" ("New Game").



Figure 2

After that, a screen like the one shown in Figure 3 will appear, and you will have to indicate your gender, enter the identification code you find at your workstation, and choose the level. For Stage 1, you

must choose the first level and then click “NUOVA PARTITA” (”New Game”).

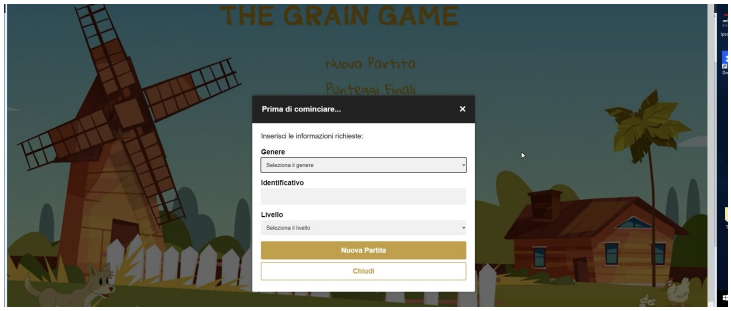


Figure 3

WHAT HAPPENS NOW?

We will give you a few minutes to answer the control questions at your workstation. If you have any questions, please raise your hand; an assistant will come to your workstation to answer. After answering any questions, Stage 1 will begin.

PART 1 – STAGE 2: INSTRUCTIONS

As in the first stage, you will act as a farmer who has purchased a 70-hectare plot of land to grow wheat, and you will have 70 wheat seeds available to plant one after another. Note that the fertility of the various hectares will be different from what you encountered in Stage 1, but you can still obtain between -10 and 15 points per hectare, and you will receive a new initial endowment of 700 points.

As in Stage 1, each planted seed will generate points, and your goal is to obtain the highest number of points possible.

For each seed, you have two choices. You can plant a wheat seed in a hectare of land where you have never planted anything before (indicated by “?”) or in a hectare of land where you have already planted. If you decide to plant the seed in a hectare where you have already sown, you may receive a number of points slightly different from those you received in the previous turn. Each planted seed will generate points, and in each turn you can obtain between -10 and 15 points per hectare.

EARNINGS FOR STAGE 2

Payment in this stage depends on your performance relative to the performance of a group of 4 participants. Groups are formed randomly at the beginning of this stage, and each participant remains in the same group for the entire duration of this stage and the next stage.

GROUP ALLOCATION

Each group is composed of 4 participants, 2 men and 2 women. The identity of the group members is never revealed, neither during nor after the study, so all decisions remain anonymous. In Stage 2, your earnings depend on the number of points you obtain compared to the other three people in your group. The person in your group who obtains the highest number of points is the winner of the tournament.
The winner receives 4 cents of a euro for every point obtained, while the other participants do not receive any payment. In the event of a tie, the final ranking is determined by a random draw. We will call this payment system a “tournament.”
At the end of Stage 2, you will know the total number of points you obtained, but you will not know whether you were among the tournament winners until the study is completed. If Stage 2 is randomly drawn for payment, your earnings depend on whether you were the tournament winner or not.
As in Stage 1, you will have to indicate your gender, enter the identification code you find at your workstation, and choose the level. For Stage 2, you must choose the second level and then click “NUOVA PARTITA” (”New game”).

WHAT HAPPENS NOW?
We will give you a few minutes to answer the control questions at your workstation. If you have any questions, please raise your hand; an assistant will come to your workstation to answer. After addressing any questions, Stage 1 will begin.

PART 1 – STAGE 3: INSTRUCTIONS
STAGE 3: ACTIVITY AND EARNINGS
As in the first and second stages, you will act as a farmer who has purchased a 70-hectare plot of land to grow wheat, and you will have 70 wheat seeds available to plant one after another. Note that the fertility of the land's various hectares will be different from what you experienced in Stage 1 and Stage 2, but each planted seed will generate between -10 and 15 points per hectare. As in the previous stages, you will have a new initial endowment of 700 points.
In Stage 3, you yourself will choose the payment method you prefer: you can choose piece rate payment (as in Stage 1) or tournament payment (as in Stage 2).
Summarizing:

• Earnings in Stage 3 if you choose the Piece Rate scheme: you receive 1 cent for each point accumulated.

• Earnings in Stage 3 if you choose the Tournament scheme: If in Stage 3 you obtain more points than the person who won the tournament in Stage 2, then you will receive 4 cents for each point (that is, 4 times the piece-rate payment); otherwise, you will receive no payment. If you obtain exactly the same number of points as the Stage 2 tournament winner, a random draw will select either you or the other winner. If in Stage 2 you were the winner, then you have to surpass your own previous score.

The group composition (2 men and 2 women) is the same as in Stage 2. If you choose the Tournament payment scheme, you will not be informed about the outcome of the tournament until the end of the study.
You will see a screen like the one in Figure 4 below, asking whether you want to choose piece rate payment or tournament payment for your performance in Stage 3.

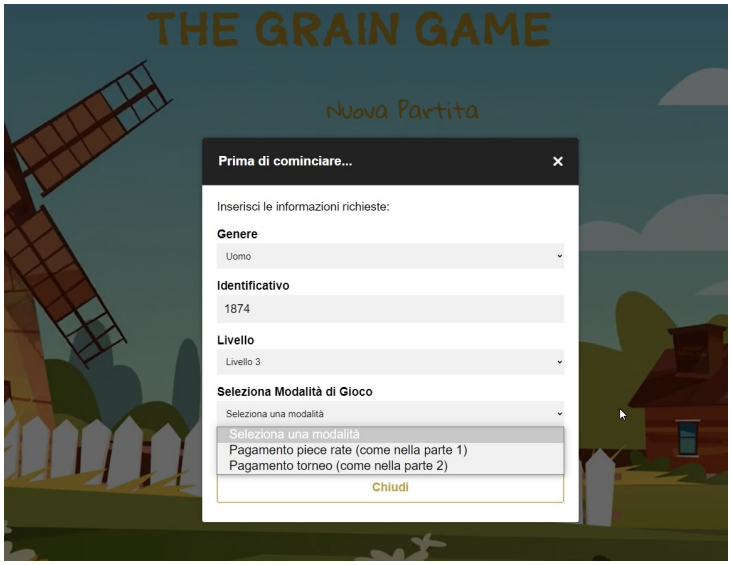


Figure 4

If Stage 3 is randomly drawn for final payment, your earnings depend on which payment choice you made and on your performance in this activity.
WHAT HAPPENS NOW?
We will give you a few minutes to answer the control questions at your workstation. If you have any questions, please raise your hand; an assistant will come to your workstation to answer. After addressing any questions, Stage 1 will begin.

INSTRUCTIONS FOR PART 2

Part 2 is composed of 2 stages. At the end of Part 2, the computer will make a random draw and select one of these 2 stages to calculate your earnings in Part 2.

PART 2 – STAGE 1: INSTRUCTIONS
In this stage, you will be asked to make some decisions. The decisions you make affect only your earnings and have no consequences for other participants. On your computer screen, you will see a table like the one reproduced in Figure 5 below. The 6 rows correspond to 6 decisions.



Figure 5

For each possible choice, there are different losses associated with a "heads" outcome of a coin toss. If you decide to toss the coin, you will win if the outcome is tails, according to the amounts shown in Figure 5. If you decide not to toss the coin, you will earn zero.
For example, in decision 4, if you decide not to toss the coin, you get 0. If you decide to toss the coin, you receive 6 if the result is tails; you get -5 if the result is heads.
If the result you obtain in this stage is negative, we will subtract the losses from your participation fee and from the other earnings obtained in Part 1.
To determine your earnings in this stage, one of the six rows will be randomly selected by the computer. If for the randomly selected row you indicated you want to toss the coin, the coin toss will be made by the computer, and you will learn your earnings associated with that outcome.

PART 2 – STAGE 2: INSTRUCTIONS
In this stage, you will be asked to make only one decision. The decision you make affects only your earnings and has no consequences for other participants. On your computer screen, you will see a table of 100 cells, as shown in Figure 6 below.



Figure 6

Your task is to decide how many cells you want to colour red.
You will receive 10 cents of a euro for each red cell. However, this earning is only potential because, after you have decided how many cells to colour red, the computer will randomly select one cell (out of the 100 in the table) for each participant.
If a gray cell is selected, you earn 10 cents of a euro for every cell you chose to colour red. If a red cell is selected, you earn 0 cents of a euro.
To choose how many cells to colour red, you must select them with your mouse. A message will appear showing the exact number of cells you have decided to colour red. If you are sure of your choice, press "Conferma" (Confirm); otherwise, press "Cancella" (Cancel) and make another choice.

H.2 Help Icon ("?") Instruction Recap (English Translation)

At any point during the grain-planting task, participants could click the "?" icon in the top bar to review a concise summary of the rules. Below is the English translation of that help popup.



Figure 7: Instructions recap for the Grain Game (Gain Treatment) as presented after clicking on Help Icon ("?")

> **Welcome!** In this game, you will take on the role of a farmer who has purchased a 70-hectare plot of land to grow wheat. Unfortunately, you do not know which parts of the land are the most fertile (i.e., those capable of producing more wheat). You have 69 wheat seeds at your disposal and must decide where to plant them (the first one has already been planted at the start of the field). Each seed planted will generate points, and your goal is to earn as many points as possible. Once you have planted the final seed, we will add up the total number of points you have scored and give you one cent of a euro for every point. For each seed, you can move one hectare to the left or right of where you last planted, or remain in the same area. At the end of each turn, you can earn between 0 and 25 points for each hectare.[29]
>
> **Controls:** - Use the arrow keys (← →) to move from one area of the field to another. - Press the space bar to plant a seed in the area you are currently in. - Use the A and D keys to scroll through the field view.

H.3 Control Questions (English translation)

Below are the control questions participants were asked to check their understanding of the task. The correct answers were provided by the experimenter before the game began.

Stage 1

1. For each seed planted, how many points can be obtained?
   (a) From -10 to 15
   (b) From 0 to 25
   (c) From -5 to 25
   (d) From -15 to 15
2. In choosing to plant a seed, how is it possible to move in the field?
   (a) For each new seed to be planted it is possible to move maximum one hectare at a time either forward or backward

[29] "Each hectare" is to be understood as a hectare in which a seed has been planted.

(b) For each new seed to be planted it is possible to move freely forward and backward, even several hectares at a time

(c) For each new seed to be planted it is possible to move one hectare at a time forward, while freely (even several hectares at a time) backward

(d) For each new seed to be planted it is possible to move one hectare at a time backward, while freely (even several hectares at a time) forward

3. What is it possible to do for each of the 70 decisions to be made?

   (a) Plant one seed at a time while staying in the same hectare or moving only to the right

   (b) Plant one seed at a time while staying in the same hectare or moving only to the left

   (c) Plant one seed at a time while staying in the same hectare and moving either to the right or to the left

   (d) Plant several seeds at once

4. How are the earnings of Stage 1 of the first part of the study determined?

   (a) The earnings are determined by the sum of the points obtained for each seed planted + 700 points of initial endowment

   (b) The earnings are determined by the highest number of points reached in a hectare + 700 point of initial endowment

   (c) The earnings are determined by randomly selecting a hectare and using the points obtained in that case + 700 points of initial endowment

Stage 2

1. Imagine you are the winner of the tournament, how is your payment determined?

   (a) I receive 1 cent for each point accumulated

   (b) I receive 4 cents for each point accumulated

   (c) I receive 2 cents for each point accumulated

2. Imagine you are **not** the winner of the tournament, how is your payment determined?

   (a) I receive 1 cent for each point accumulated

   (b) I receive 0 cents for each point accumulated

   (c) I receive 2 cents for each point accumulated

3. How many winners are there in the tournament in each group (consisting of 4 participants)?

   (a) Only 1, that is the participant with the highest score

   (b) 2, that is the two participants with the highest score

   (c) Only 1, that is the fastest participant to complete the task

   (d) Anyone who exceeds the threshold of 500 points

Stage 3

1. Imagine you have chosen the tournament payment system, how is your payment determined?

   (a) I receive 4 cents for each point accumulated **if** the points I obtain in Stage 3 are higher than those obtained by the winner of the tournament in my group in Stage 2, otherwise I receive 0 cents.

   (b) I receive 4 cents for each point accumulated **if** the points I obtain in Stage 3 are higher than those obtained by the other members of my group who, like me, chose the tournament in Stage 3, otherwise I receive 0 cents.

   (c) I receive 2 cents for each point accumulated **if** the points I obtain in Stage 3 are higher than those obtained by the winner of the tournament in my group in Stage 2, otherwise I receive 0 cents.

   (d) I receive 2 cents for each point accumulated **if** the points I obtain in Stage 3 are higher than those obtained by the other members of my group who, like me, chose the tournament in Stage 3, otherwise I receive 0 cents.

Appendix I

This appendix provides an overview of the reinforcement learning algorithm benchmark and the range of hyperparameters adopted.

Setup. We treat the Grain Game as a finite, episodic decision problem:

- The state variable $s \in \{0, \dots, 69\}$ is the index of the current planting slot.
- In each decision round t, the agent chooses an action $a_t \in \{\text{Retreat, Exploit, Explore}\}$. At $s = 0$, the available actions are limited to {Exploit, Explore}.
- An episode ends after $T = 69$ moves. The algorithm plays one of the six sequences, receiving the corresponding points in each planting slot plus the same zero-mean fluctuation as in the lab (see Section 3 and Appendix B).
- In the Gain&Loss setting, we add 700 points to the final cumulative reward to replicate the experimental endowment.

Action-selection rules. We implement two exploration mechanisms:

1. **ε-greedy.** Let $Q_t(s, a)$ denote the current estimate of the value of taking action a in state s at time t. With probability $1-\varepsilon$ the agent selects the greedy action $\arg\max_a Q_t(s_t, a)$; with probability ε it chooses one admissible action at random:

$$a_t = \begin{cases} \arg\max_a Q_t(s_t, a), & \text{prob. } 1-\varepsilon, \\ \text{uniform random action}, & \text{prob. } \varepsilon. \end{cases} \tag{I.1}$$

2. **ε-soft-max.** For every admissible action compute the *soft-max probability*

$$P_\tau(a) = \frac{\exp\big(Q_t(s_t, a)/\tau\big)}{\sum_{a'} \exp\big(Q_t(s_t, a')/\tau\big)}, \qquad \tau > 0 \text{ (temperature)}. \tag{I.2}$$

The same ε split is then applied, but in the exploration branch the agent takes the action with the *largest* $P_\tau(a)$ (no random draw):

$$a_t = \begin{cases} \arg\max_a Q_t(s_t, a), & \text{prob. } 1-\varepsilon, \\ \arg\max_a P_\tau(a), & \text{prob. } \varepsilon. \end{cases} \tag{I.3}$$

Small τ values make $P_\tau(a)$ concentrate on the best actions; large values flatten the distribution.

Value update. After executing a_t in state s_t the agent observes the reward r_t and the next state s_{t+1}. The *temporal–difference error* δ_t depends on the learning algorithm:

$$\delta_t^{\text{Q-Learning}} = r_t + \gamma \max_{a_{t+1}} Q_t(s_{t+1}, a_{t+1}) \; - \; Q_t(s_t, a_t), \qquad \text{(off-policy)} \tag{I.4}$$

$$\delta_t^{\text{SARSA}} = r_t + \gamma\, Q_t(s_{t+1}, a_{t+1}) \; - \; Q_t(s_t, a_t), \qquad \text{(on-policy)} \tag{I.5}$$

With SARSA, $\gamma \in [0, 1]$ is the discount factor and a_{t+1} is the action actually taken at the next step. With QL, $\gamma \in [0, 1]$ will be randomized by the algorithm.

The estimate for the *current state–action pair* (s_t, a_t) is then updated as

$$Q_{t+1}(s_t, a_t) \; = \; Q_t(s_t, a_t) + \alpha_t\, \delta_t,$$

where α_t is the learning rate.

The exploration rate ε and the learning rate α are either kept constant or set to decay as $\varepsilon_t = \alpha_t = 1/\sqrt{t+1}$.

Training and evaluation. For every combination of $(\varepsilon, \alpha, \gamma, \tau)$, exploration rule ($\varepsilon$-greedy or ε-soft-max) and update algorithm (Q-Learning or SARSA) the agent is trained for 10 000 episodes and tested for a further 5 000 episodes with learning disabled. The procedure is repeated independently for each of the six payoff sequences.

Hyper-parameter grid. The search ranges are

$$\varepsilon \in \{0.5,\ 1,\ \text{decay}\}, \quad \tau \in \{0.5,\ 1,\ 5\}, \quad \alpha \in \{0.2,\ 0.5,\ 0.8,\ \text{decay}\}, \quad \gamma \in \{0.1,\ 0.5,\ 0.9\}.$$

Crossing these values with the two exploration rules, the two updated algorithms and the six payoff sequences yield 1 728 distinct configurations. Table I1 lists the three best configurations for each payoff sequence. The full output and code are available upon request.

Table I1: Top 3 simulation results by payoff sequence

Payoff Sequence	Algorithm	Exploration Strategy	ε	γ	α	Avg. Payoff	Maximum Position
1	Q-Learning	ϵ-greedy	0.5	0.9	0.2	808.9	35
	Q-Learning	ϵ-greedy	0.5	0.9	0.5	808.3	32
	SARSA	ϵ-greedy	0.5	0.9	0.2	791.0	46
2	Q-Learning	ϵ-greedy	0.5	0.9	0.5	933.7	29
	Q-Learning	ϵ-greedy	0.5	0.9	0.2	932.9	21
	Q-Learning	ϵ-greedy	0.5	0.9	0.8	932.5	30
3	Q-Learning	ϵ-greedy	0.5	0.9	decay	966.6	25
	SARSA	ϵ-greedy	0.5	0.5	0.5	965.9	24
	Q-Learning	ϵ-greedy	0.5	0.5	0.2	965.7	16
4	Q-Learning	ϵ-greedy	0.5	0.9	decay	509.2	43
	Q-Learning	ϵ-greedy	0.5	0.9	0.5	501.1	24
	Q-Learning	ϵ-greedy	0.5	0.9	0.8	499.6	25
5	Q-Learning	ϵ-greedy	0.5	0.9	0.2	944.1	41
	Q-Learning	ϵ-greedy	0.5	0.9	0.8	943.8	20
	Q-Learning	ϵ-greedy	0.5	0.9	0.5	943.4	19
6	Q-Learning	ϵ-greedy	0.5	0.9	0.5	977.6	15
	Q-Learning	ϵ-greedy	0.5	0.9	decay	977.1	15
	Q-Learning	ϵ-greedy	0.5	0.5	0.5	976.2	17

Notes: Each model was trained for 10 000 episodes and tested for 5 000. For payoff sequences 4–6 (Gain & Loss treatment) the values include the 700-point endowment. "Maximum Position" is the furthest slot reached.

Appendix J

For the Choice stage, Table J1 and Table J2 report an additional specification in which we include the interaction between the *Gain & Loss* treatment and the threshold dummy (*Threshold*= 1 if the payoff in the previous round was negative). The regression is estimated separately for men and women.

To evaluate whether crossing the threshold changes the likelihood of *Explore* (versus *Exploit*) once losses are possible, we test, within each gender, the linear restriction

$$\beta_{\text{Threshold}} + \beta_{\text{Threshold}\times\text{Gain\&Loss}} = 0.$$

For men, the null hypothesis is rejected in every specification (Model A: $\chi^2(1) = 17.28$, $p < 0.001$; Model B: $\chi^2(1) = 16.87$, $p < 0.001$; Model C: $\chi^2(1) = 10.74$, $p < 0.01$), implying that a negative payoff in the preceding round reduces the relative probability of exploring (versus exploiting).

For women, the same joint test is never significant (Model A: $\chi^2(1) = 2.09$, $p = 0.148$; Model B: $\chi^2(1) = 2.10$, $p = 0.148$; Model C: $\chi^2(1) = 0.93$, $p = 0.335$).

Hence, under the *Gain & Loss* being below the threshold in previous turn implies a decrease in exploration for men (relative to exploitation), while women's exploration behaviour remains unchanged. For men the result is in line with what is found by Chin et al. [2023].

Table J1: Random Effects Multinomial Logistic Results for *Move* – Women Only (Choice stage)

Dependent Variable	*Move* (Baseline = *Exploit*)					
	Model A$_W$		**Model B$_W$**		**Model C$_W$**	
	Retreat	Explore	Retreat	Explore	Retreat	Explore
Gain & Loss	1.326	0.543	1.198	0.467*	0.896	0.356*
	(0.449)	(0.187)	(0.435)	(0.159)	(0.382)	(0.145)

Threshold	0.343**	0.598*	0.364*	0.668	0.186***	0.436**
	(0.131)	(0.137)	(0.148)	(0.157)	(0.085)	(0.132)
Gain & Loss × Threshold	0.805	1.139	0.761	1.017	1.920	1.740
	(0.380)	(0.415)	(0.377)	(0.373)	(0.969)	(0.743)
Lag Payoff	0.780***	0.780***	0.782***	0.780***	0.780***	0.768***
	(0.027)	(0.027)	(0.028)	(0.027)	(0.031)	(0.031)
Compete (Tournament)	0.515	1.131	0.560	0.848	0.788	0.840
	(0.194)	(0.410)	(0.228)	(0.309)	(0.393)	(0.374)
Piece Rate Payoff	–	–	1.000	1.003**	1.001	1.003*
			(0.001)	(0.001)	(0.002)	(0.001)
Risk (Bret)	–	–	–	–	1.007	1.001
					(0.011)	(0.012)
Risk (Self-Reported)	–	–	1.130	1.482***	–	–
			(0.114)	(0.153)		
Competitiveness	–	–	0.788**	0.994	0.742**	1.020
			(0.060)	(0.075)	(0.064)	(0.095)
CRT	–	–	1.487*	1.165	1.594*	1.228
			(0.233)	(0.194)	(0.290)	(0.237)
Extraversion	–	–	0.944	0.924	0.949	1.040
			(0.095)	(0.095)	(0.111)	(0.134)
Agreeableness	–	–	1.347*	1.241	1.326	1.532*
			(0.202)	(0.195)	(0.240)	(0.296)
Conscientiousness	–	–	0.994	0.886	1.095	0.923
			(0.096)	(0.106)	(0.123)	(0.133)
Emotional Stability	–	–	0.874	0.763*	0.944	0.830
			(0.105)	(0.089)	(0.130)	(0.121)
Openness	–	–	0.770	0.749	0.884	0.881
			(0.123)	(0.117)	(0.158)	(0.176)
GPA	–	–	1.229	1.289	1.085	1.185
			(0.282)	(0.257)	(0.292)	(0.297)
Born in Italy	–	–	0.874	0.639	0.954	0.316
			(0.444)	(0.451)	(0.615)	(0.267)
Age	–	–	1.040	1.049	1.029	1.075
			(0.058)	(0.070)	(0.072)	(0.084)
STEM	–	–	0.601	0.838	0.664	1.148
			(0.230)	(0.287)	(0.289)	(0.471)
Loss Aversion (λ_{loss})	–	–	–	–	0.763	0.424*
					(0.238)	(0.148)
Constant	1.998	99.670***	0.771	1.565	0.762	3.434
	(1.047)	(55.910)	(1.722)	(4.245)	(1.967)	(11.970)
Observations	14,008		13,668		10,132	
Groups	206		201		149	
Wald χ^2	85.67		139.76		115.93	
Log PL	−8513.96		−8289.73		−6103.07	

Table J2: Random Effects Multinomial Logistic Results for *Move* – Men Only (Choice stage)

Dependent Variable	*Move* (Baseline = *Exploit*)					
	Model $\mathbf{A}_M$		**Model $\mathbf{B}_M$**		**Model $\mathbf{C}_M$**	
	Retreat	Explore	Retreat	Explore	Retreat	Explore
Gain & Loss	0.961	1.292	1.199	1.332	0.619	1.547
	(0.350)	(0.589)	(0.480)	(0.627)	(0.281)	(0.807)
Threshold	0.292*	0.659	0.312*	0.670	0.327	0.939
	(0.142)	(0.234)	(0.152)	(0.241)	(0.187)	(0.407)
Gain & Loss × Threshold	0.605	0.520	0.584	0.497	0.734	0.407
	(0.361)	(0.241)	(0.353)	(0.236)	(0.505)	(0.225)
Lag Payoff	0.637***	0.610***	0.640***	0.610***	0.657***	0.633***
	(0.034)	(0.033)	(0.034)	(0.033)	(0.039)	(0.039)
Compete (Tournament)	0.794	1.820	0.525	1.042	0.259**	0.754
	(0.296)	(0.852)	(0.211)	(0.479)	(0.120)	(0.374)
Piece Rate Payoff	–	–	0.999	1.005**	0.999	1.005**
			(0.001)	(0.001)	(0.001)	(0.002)
Risk (Self-Reported)	–	–	1.046	1.109	–	–
			(0.119)	(0.125)		
Risk (Bret)	–	–	–	–	1.043*	1.055*
					(0.018)	(0.018)
Loss Aversion (λ_{loss})	–	–	–	–	0.493**	0.682
					(0.134)	(0.208)

Competitiveness	–	–	1.252*	1.141	1.227	1.122
			(0.115)	(0.124)	(0.129)	(0.133)
CRT	–	–	1.031	0.786	0.982	0.722
			(0.182)	(0.167)	(0.205)	(0.193)
Extraversion	–	–	1.110	0.976	1.012	0.946
			(0.128)	(0.137)	(0.137)	(0.157)
Agreeableness	–	–	1.751**	1.379	1.696*	1.363
			(0.325)	(0.296)	(0.365)	(0.320)
Conscientiousness	–	–	1.193	1.105	1.040	1.302
			(0.205)	(0.198)	(0.199)	(0.262)
Emotional Stability	–	–	0.960	0.846	1.135	0.867
			(0.127)	(0.148)	(0.160)	(0.147)
Openness	–	–	0.893	1.018	0.809	1.060
			(0.178)	(0.234)	(0.181)	(0.270)
GPA	–	–	0.750	0.595*	0.889	0.758
			(0.162)	(0.145)	(0.219)	(0.200)
Born in Italy	–	–	0.047**	0.063**	0.005***	0.008**
			(0.047)	(0.066)	(0.007)	(0.013)
Age	–	–	0.958	0.985	0.951	0.987
			(0.039)	(0.054)	(0.061)	(0.073)
STEM	–	–	0.933	0.740	1.257	0.872
			(0.348)	(0.338)	(0.541)	(0.445)
Constant	19.257***	1940.643***	68.240	581.495*	795.754*	89.917
	(14.832)	(1704.032)	(171.513)	(1712.102)	(2203.712)	(297.266)
Observations	14,484		14,008		10,948	
Groups	213		206		161	
Wald χ^2	137.69		203.03		208.08	
Log PL	−7418.08		−7167.44		−5502.25	

Notes: Entries are Relative Risk Ratios (RRRs); robust s.e. in parentheses, clustered at participant level. Significance: $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.