

Economics Working Paper Series

2025/004

The Grant Proposal Game: An Experimental Study on Herding and Divergent Behaviors in Competition

Yevgeny Mugerman, Eyal Winter, and Tomer Yafeh

The Department of Economics Lancaster University Management School Lancaster LA1 4YX UK

© Authors All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission, provided that full acknowledgement is given.

LUMS home page: http://www.lancaster.ac.uk/lums/

The Grant Proposal Game:

An Experimental Study on Herding and Divergent Behaviors in Competition

Yevgeny Mugerman,* Eyal Winter,** and Tomer Yafeh***

Conflicts of interest: The authors declare that they have no competing interests.

Ethics approval and consent to participate: This research was conducted in compliance with ethical standards. All experimental procedures involving human participants were approved by the Hebrew University's Ethics Review Board. Informed consent was obtained from all individual participants involved in the study.

Availability of data and material: The replication material for the study is available at <u>https://doi.org/10.17632/ffk9nphmxv.1</u>.

Authors' contributions: All authors contributed to the study conception and design.

We acknowledge the helpful comments and suggestions from the anonymous referees, the Guest Editor Ido Erev, and the Advisory Editor, Roberto A. Weber. We are also grateful for the feedback on previous versions from Sergiu Hart, Ilan Kremer, Daniel Seidmann, and Jörgen Weibull. Eyal Winter wishes to acknowledge financial support from the ESRC through research grant no. ECA7674.

^{*}Graduate School of Business Administration, Bar-Ilan University, and the Federmann Center for the Study of Rationality, The Hebrew University of Jerusalem (<u>mugi@huji.ac.il</u>).

^{**(}*Corresponding author*) Department of Economics, Lancaster University and Department of Economics, and the Federmann Center for the Study of Rationality, The Hebrew University of Jerusalem (eval.winter@mail.huji.ac.il).

^{***}Business School, and the Federmann Center for the Study of Rationality, The Hebrew University of Jerusalem (tomer.yafeh@mail.huji.ac.il).

Herding and Divergent Behaviors in Competition: An Experimental Study

Abstract

We explore strategic betting in competitive environments with multiple participants and potential winners. We examine two scenarios: an "inclusive" low-competition scenario with many winners and an "exclusive" high-competition scenario with few winners. Using a simple model, we illustrate the strategic insights in these scenarios and present experimental results that align with our predictions. In the experiment, participants made repeated bets with feedback on past results and their payoffs. In the inclusive scenario, all but the worst guessers were rewarded, while in the exclusive scenario, only the top guessers received rewards. Our findings show that in the inclusive scenario, participants exhibit herding behavior by coordinating their bets, while in the exclusive scenario, they diversify their bets across multiple options. The main general insight of our findings is that in moderate competitions, one tends to join the majority to avoid standing out in case of failure, whereas in intense competitions, one tends to differentiate oneself from one's peers to ensure that success stands out. This insight is relevant for a broad domain of strategic interactions.

Keywords: Experimental Economics, Competitive Decision-Making, Herding Behavior, Divergence Strategy, Inclusive vs. Exclusive Scenarios, Multi-Winner Competition

JEL Classification: C9, D7, D8, L1, M3, G1

1 Introduction

Consider the dynamics of a competition or contest with multiple winners, where contesters are ranked by their performance, and a select fraction of top performers are rewarded with prizes. A critical factor in these competitions is the strategic bets each contester makes, heavily influenced by uncertain factors. Take, for example, a research grant competition, where applicants submit proposals, and their success hinges on aligning with the funding body's undisclosed preferences for topics or research methodologies. Each applicant, in essence, gambles on these preferences. This logic extends to other scenarios like government service procurements or job interviews, where strategic positioning is key.

This paper conducts a study of the strategic choices behind these bets in different competitive environments. We categorize these environments into two types: inclusive environments, characterized by low-intensity competition with a higher proportion of winners, and exclusive environments, marked by high-intensity competition where only a few emerge victorious. Our primary finding, supported by both theoretical analysis and experimental data, suggests a distinct pattern: in inclusive environments, players are inclined to coordinate their bets, leading to herding behavior. Conversely, in exclusive environments, participants are more likely to diverge, spreading their bets across various options.

To illustrate the intuition behind this finding, consider a competitive scenario involving three contestants and two potential bets. For example, the bets could involve submitting a proposal for either an empirical (E) or a theoretical (T) research project. The grant-awarding body, which is interested exclusively in one type of proposal, selects between the two options with equal probability. Winners receive a payoff of 1, while those not selected receive nothing.

In an inclusive environment, where the body intends to fund two out of three proposals, priority is given to submissions aligned with the body's interest. Among the remaining submissions, the body selects randomly to ensure two winners. If all three contestants submit identical proposals (all choose E or all choose T), two winners are chosen at random from the three submissions.

In any equilibrium of this game, all three contestants would choose the same bet, each achieving an expected payoff of 2/3. For example, if all contestants choose T, any

unilateral deviation—where a contestant chooses E instead—would reduce their expected payoff to 1/2, as their success would now depend entirely on the grant body's preference aligning with their distinct choice.

A mixed strategy equilibrium also exists, in which contestants randomize between E and T with equal probabilities. However, if the probability of one state of the world—say T—were slightly greater than 1/2, the game would have a unique Nash equilibrium where all players choose T.

Suppose the scenario evolves such that the granting body decides to award a grant to only one of the three submissions. In this configuration, a uniform strategy where all players opt for the same bet ceases to be an equilibrium. By diverging from the common choice, a player could increase her payoff from 1/3 to 1/2. The only pure Nash equilibria¹ in this setup involve one player choosing E while the other two opt for T, or vice versa. This phenomenon highlights a critical insight: In moderate competitions we should protect ourselves from potential failure, whereas in intense competition we should leverage the potential success.

To explain this insight, consider a researcher that is already informed about the bets made by all her competitors, and that the group is split between the two options very unequally, say, 95% choose T and 5% choose E. If the competition is moderate, e.g., 75% of the proposals will be awarded a grant,² our researcher should protect herself from potential failure (in her guessing). If she herds with the majority, no matter whether her guess is correct or not, she will be awarded the grant with a probability of close to 75% (though less than 75% if her guess is incorrect, and more than 75% if it is correct). If she goes with the minority, she is granted the prize with a probability of 100% if her guess is successful, and with zero probability if it is not. Since there is an equal probability of her guess being correct and incorrect, herding with the majority to protect herself against failure is the superior option. Consider now the case of intense competition, in which only 5% of the submissions will be awarded. In this case, going with the majority offers no protection at all. If she fails, she is awarded with probability

¹ There is also one mixed equilibrium in which all three players assign equal probability for the two bets. ² For clarity, consider the following example: Suppose there are 100 applications—95 for T and 5 for E—and 75 grants to distribute. If the agency prefers T, 70 grants will be randomly allocated among those who submitted T, and 5 grants will remain unallocated. Conversely, if the agency prefers E, all 5 applicants who submitted E will receive funding, and the remaining 70 grants will be randomly distributed among those who submitted T.

zero. If she succeeds, she is awarded with a probability slightly greater than 5%. However, if she joins the minority and succeeds, she will be awarded with probability 1. Hence if competition is intense, she is better off leveraging her potential success rather than protecting herself from failure.

While we provide in the appendix a more general result that builds on the above insight, our main interest in this paper is testing this insight experimentally. In game theoretic terms, our insight above implies that in the inclusive environment, there are two Nash equilibria in the guessing game, in both of which all players choose the same bet. In contrast, in exclusive environments, they split equally between the two bets (almost equally if the number of players is odd). Clearly, playing these equilibria in both environments requires coordination. Ex ante, before players have made their decisions, they do not know what other players have chosen (or will choose). Even if they understand the benefits of herding in the inclusive environment, and believe that everyone else understands it, they still cannot know whether herding means choosing bet A or bet B. Hence, subjects need to play the game repeatedly within the same group, receive feedback about what other players did, form beliefs, and act based on these beliefs. Support for this insight would then be convergence to a unanimous choice of bet in the inclusive environment and a major split in the exclusive one. The experimental design is straightforward yet revealing. Participants are asked to guess the outcome of a bet between two named individuals, "Jane" and "Jill," repeated across thirty rounds. After each round, we display the results of the bet, all previous rounds' guess distributions both in absolute numbers and percentages, and the participant's performance and earnings overall and in every round played so far. Participants with a correct prediction are ranked above the rest. In the inclusive variant, the top 80% of participants are declared winners, while in the exclusive variant, only the top 20% are rewarded. Results show a pronounced tendency for guess convergence in the inclusive environment, and a tendency for differentiation in the exclusive environment.

This paper contributes to the literature by exploring how the inclusivity of environments influences herding behavior. Previous research has focused on uncertainty and information access as key drivers of herding, with less attention given to the effects of the degree of competition on the prospects of herding. Our study introduces the concept of environmental inclusivity as a determinant of herding or divergence, providing empirical evidence to support this novel perspective.

The paper is structured as follows: Section 2 reviews the relevant empirical and theoretical literature on herding in competitive environments. Section 3 describes the experimental design. Section 4 presents the experimental results. Section 5 outlines a simple reinforcement learning model, while the formal Nash equilibrium analysis of the static game is detailed in Appendix C. Finally, Section 6 discusses our findings and concludes the paper.

2 Related Literature

2.1 Market Entry Games and Coordination in Competitive Environments

This paper contributes to the literature on coordination in large groups, an area significantly influenced by the work of Amnon Rapoport. His series of papers on market entry games (Rapoport, 1995; Rapoport, Seale, Erev, & Sundali, 1998; Rapoport, Seale, & Winter, 2000; Rapoport, Seale, & Winter, 2002; Rapoport & Seale, 2008) serves as a key inspiration for our study. These studies emphasize the critical role of coordination in competitive environments, applicable to scenarios such as traffic congestion, retail competition, and school applications. Notably, they reveal that aggregate player behavior in large-scale games quickly aligns with equilibrium predictions, despite individual deviations.

While our strategic environment differs from that of market entry games, both contexts underscore the unexpected role of coordination in competitive, non-cooperative settings and the tendency for experimental behavior to converge towards Nash equilibria. Rapoport's work also highlights how individual decision-making processes adapt to changing conditions, providing valuable insights into cognitive mechanisms in strategic environments.

2.2 Theories of Herding and Contrarian Behavior

Various theoretical models have been proposed to explain herding and contrarian behavior, primarily focusing on the influence of uncertainty. Banerjee (1992) introduced a seminal model explaining herding, where an agent, observing the decisions of others, chooses to make the same decision, assuming that others may possess relevant information. This can lead to a scenario where agents ignore their own information, resulting in less informative decisions and potentially leading to inefficient equilibrium states. Banerjee termed this phenomenon the "herd externality."

Herding behavior is also prevalent in financial markets. Investors may change planned investment decisions after observing others, thinking that the herd behavior could reflect relevant information. Scharfstein and Stein (1990) describe how investment managers, categorized into "smart" (with informative signals) and "dumb" (with noisy signals), mimic each other to appear knowledgeable. This imitation is driven by the desire to align with others, as shared failures have less impact on their reputations. Compensation schemes further incentivize imitation, promoting financial herding (Bikhchandani & Sharma, 2000).

Avery and Zemsky (1998) applied this model to stock markets, showing that multidimensional uncertainty can cause herding. When three or more dimensions of uncertainty exist regarding an asset, herd behavior becomes likely, leading to indistinguishable markets composed of well-informed and poorly informed traders, potentially resulting in asset price bubbles.

Contrarian behavior³ arises when investors deliberately choose to go against the herd, especially in situations of extreme uncertainty. Contrarianism has typically been studied in cases of information asymmetry, particularly when private information is available (e.g., Medrano & Vives, 2001), or when agents have doubts regarding the rationality of others, and contrarian behavior then emerges in an attempt to compensate for the potential mistakes of others (Drehmann et al., 2005) Meanwhile, herding tends to occur with signals indicating extreme states or limited private information. Herding trades are self-reinforcing, leading to a persistent majority, whereas contrarian trades are self-defeating, as widespread adoption of contrarian behavior eventually nullifies its uniqueness. Both herding and contrarianism contribute to price volatility and reduced liquidity (Avery & Chevalier, 1999; Park & Sabourian, 2011).

Relative position concerns can also drive herding. Cole et al. (2001) demonstrated that concerns about relative wealth can influence investment decisions, leading to either herding or contrarian behavior. DeMarzo et al. (2008) proposed an equilibrium model where agents' future investment opportunities depend on relative wealth, causing them to herd to minimize the risk of ending up significantly poorer than their peers. This

³ In the context of this paper, contrarian or divergent behavior is interpreted as the inverse of herding, rather than as a separate and independent phenomenon. This approach reflects our focus on studying varying degrees of herding rather than a dual framework contrasting herding and contrarianism. While some references to the literature on contrarian behavior are included, they serve to provide context rather than to imply the existence of a distinct force separate from herding.

herding behavior distorts asset prices and can create bubbles, particularly if agents have strong relative wealth concerns and heterogeneous preferences

2.3 Experimental and Real-World Evidence

Experimental literature supports these theoretical findings on herding behavior. Cipriani and Guarino (2005, 2009) and Drehmann et al. (2005) confirmed that flexible pricing prevents herding, and that multidimensional uncertainty increases herding. Drehmann et al.'s (2005) study, which included financial consultants, showed no significant differences from the general population. Cipriani and Guarino (2005) initially used university students, but their subsequent study with financial market professionals confirmed the results (Cipriani & Guarino, 2009). Bracht et al. (2010) provided experimental evidence that poor information generates herding behavior, with highly informed individuals breaking this behavior only in about one-third of cases.

Payoff externalities can also encourage or discourage herding, and are often prevalent in real-world environments in which herding occurs. Drehmann et al.'s (2007) experimental findings demonstrate that network effects offer a prominent example of such externalities. Where they constitute a positive externality, herding is more likely to occur, and to persist for longer, whereas with negative externalities, herding is less likely to persist. Andersson et al. (2014) provide direct experimental evidence that a payoff structure that rewards predictions that are in line with the herd increases the influence of the majority group, while an individual reward for making accurate predictions reduces the majority group's influence.

Herding behavior has been extensively studied in real-world contexts. Foucault and Frésard (2019) explored the impact of informational costs on product differentiation, identifying herding tendencies in pre-IPO companies and demonstrating that IPOs mitigate these costs, thereby promoting increased differentiation. Research has shown that herding mutual funds often fail to outperform non-herding funds and may underperform relative to anti-herding funds, particularly among managers with significant career concerns (Koch, 2017; Jiang & Verardo, 2018). Herding behavior has been observed in various domains, such as the influence of ethnic background on long-term savings decisions (Mugerman et al., 2014), the patterns in online rating dynamics (Sunder et al., 2019), and the strategic adaptations of business groups during the COVID-19 pandemic (Espinosa-Mendez & Maguieira, 2023).

Contrarian behavior has also been observed in a range of real-world settings. For instance, competitive pressure can drive firms to pursue innovation, with early innovation and patenting enabling firms to expand market share and profit margins (Leyva-de la Hiz et al., 2021). Firms experiencing earnings pressure from shareholders or analysts have also exhibited contrarian behavior, particularly in markets where they possess significant power yet face intense competition (Zhang & Gimeno, 2017).

In the context of pension funds, Hamdani et al. (2017) found that, in the absence of performance-based fees, competitive pressures among fund managers led to herding behavior, often characterized by insufficient risk-taking. Conversely, Mugerman et al. (2022) demonstrated that heightened salience of risk could increase mutual fund risk exposure, potentially encouraging contrarian behavior. Additionally, Yousaf and Yarovaya (2022) documented evidence of herding in cryptocurrency markets, highlighting its prevalence in emerging financial ecosystems.

Our paper contributes to the existing literature on herding behavior by introducing an experimental perspective that examines herding under varying levels of competitive intensity. By exploring strong and weak competition conditions, our study provides new insights into how reward structures and strategic interactions shape the degree of herding behavior. This approach extends prior theoretical and empirical findings, offering a deeper understanding of the dynamics of herding in competitive environments.

3 Experiment

3.1 Illustrative example – grant proposal game

Before delving into the experiment, we present a simple three-person example to highlight the main strategic considerations in the two environments. We use the conventional three-player game representation, where Player 1 selects a row, Player 2 selects a column, and Player 3 selects a matrix. The three-coordinate payoff vectors represent the (expected) payoffs awarded to Players 1, 2, and 3, respectively.

This example is similar to the one introduced in the Introduction section. The scenario involves three contestants choosing between two potential bets: submitting a proposal for either an empirical (E) or a theoretical (T) research project. The grant-awarding body, which is exclusively interested in one type of proposal, selects between the two

options with equal probability. Winners receive a payoff of 1, while those not selected receive nothing.

In the inclusive environment, two grants are awarded out of three. In the exclusive environment, only one grant is awarded out of three. To set up the game (as shown in Figure 1), note the following payoff dynamics:

- Inclusive environment: Players in the majority group (2 out of 3) earn an expected payoff of 3/4, while the minority player expects to earn 1/2. If all players herd on the same choice, each has an expected payoff of 2/3.
- Exclusive environment: The minority player (1 out of 3) earns an expected payoff of 1/2, while each majority player has an expected payoff of 1/4. Herding in this case results in an expected payoff of 1/3 for each player.



Figure 1: The Grant Proposal Game

Notes: The figure illustrates the Nash equilibrium partitions for the grant proposal game in two environments: inclusive (k = 2), where two grants out of three are awarded, and exclusive (k = 1), where only one grant out of three is awarded. In this representation: Player 1 selects a row, Player 2 selects a column, and Player 3 selects a matrix. Each strategy profile corresponds to a specific allocation of payoffs among the three players. The three-coordinate payoff vectors denote the (expected) payoffs received by Players 1, 2, and 3, respectively. Strategy profiles marked with ** indicate Nash equilibria for the environment in question.

Figure 1 illustrates how the Nash equilibrium partitions the set of strategy profiles in the two environments. In the inclusive environment, the Nash equilibria (denoted by **) are the pure strategy profiles (T,T,T) and (E,E,E), where all players herd on the same choice. Conversely, in the exclusive environment, the Nash equilibrium includes

the remaining six strategy profiles (denoted by **), where players strategically split their choices to maximize individual payoffs.

The extension of this 3-person game for an arbitrary number of players is presented in Appendix C. We note that these n-person games are all symmetric constant sum games.

3.2 Experimental setting

To empirically investigate the strategic environments, we designed an experimental setup where participants, grouped together, engaged in a guessing game.⁴ Each participant guessed the outcome of a bet between two named individuals, "Jane" and "Jill," repeated across thirty rounds.⁵ Note that the equilibria in both the inclusive and exclusive games depend on coordination among players. As such, a relevant experimental setup requires dynamic interaction, allowing players to gradually coordinate their strategies through a process of learning. To facilitate this learning process, we provided participants with the following information after each round: (1) The outcome of the bet for that round; (2) The distributions of guesses from all previous rounds, displayed in both absolute numbers and percentages; (3) The participant's performance and earnings, both overall and for each round played so far. This feedback enabled participants to adjust their strategies over time, fostering the coordination necessary to approach equilibrium outcomes

Each participant was assigned to either the inclusive or exclusive environment. In the inclusive environment, in every round, if the participant's guess was correct, and he or she was among the top 80% of guessers, the participant received ten bonus points. The exclusive environment was similar, but the top 20% of guessers received 50 bonus points. At the end of the experiment, a lottery was conducted, in which participants had

⁴ This project included an additional experiment where participants were asked to guess the outcome of a coin toss across multiple rounds, with the distribution of guesses from prior rounds displayed. Participants who made correct predictions were ranked above the rest. In the inclusive variant, the top 90% of participants were declared winners, while in the exclusive variant, only the top 10% were rewarded. The conclusions of this design align with those of our main experiment, reinforcing our findings on herding and divergent behaviors. For further details on this experiment, please see Appendix B ("Coin Experiment").

⁵ To avoid a focal point affecting the results of our experiment, we alternated the order of the named individuals on which participants were betting. Half of our participants were asked to choose between "Jane and Jill," and the other half between "Jill and Jane," with the same order preserved throughout the session. The names were selected to be similar in gender, language, length, and starting letter to prevent any bias or preference for one over the other.

a 20% probability of receiving a bonus payment based on the points they had earned (at a ratio of 30 points = 1 GBP).⁶

The experiment was conducted using oTree, with all participants of a group present in the same session. Participants were recruited through the online platform Prolific, which also handled payment at the end of the experiment.⁷ Upon enrolling in the experiment, each participant received a link to join the session, and all players in a session played simultaneously. Participation was fully anonymous, and a participant could see the distribution of choices in a given round, but not the choices of any other individual player. This was done to reduce the likelihood of interference from social aspects that could encourage or discourage herding, such as the desire to punish another participant out of spite for making the wrong choice, or the altruistic desire to help others. For screenshots of the experimental interface that was presented to the participants, see Appendix A. To further ensure that our findings were the result of the competitive environment rather than other social and emotional factors, we include in Appendix B the results of another specification of the experiment, in which participants were only told the outcomes of the coin tosses at the conclusion of the entire game. This specification was also conducted in an in-person laboratory environment, reinforcing the validity of our online experimental setting.

⁶ The lottery approach is a well-established and valid method in incentivized experimental settings, supported by a substantial body of empirical research. Charness et al. (2016) demonstrated that employing a lottery system, where a subset of individuals is paid, can be as effective, if not more so, than the traditional 'pay all' method. Our lottery payments are more nuanced in that we ensure the independence of other participants' performances. A participant becomes eligible for payment only if they satisfy two sequential criteria: (1) meeting the requirements stipulated in the experiment, and (2) being selected as a winner in the lottery. Moreover, the key advantages of the lottery payment method apply to our structure as well. A similar approach was recently employed by Mugerman et al. (2024).

⁷ Prolific is typically used for individual surveys rather than simultaneous experiments involving large numbers of participants. Our experiment employs a simultaneous, interactive setup with participants drawn directly from Prolific's subject pool. Participants were directed to a virtual waiting room, a built-in feature of oTree. Once the waiting room had filled with twenty-five participants or two minutes had elapsed (whichever was sooner), we launched the session. We contacted Prolific to disable the naivety distribution feature, which typically acts as a throttling mechanism for recruitment. By disabling this feature, we allowed a significantly larger number of participants to sign up for the experiment immediately. As a result, we consistently managed to recruit enough participants within two minutes of activating our study on Prolific offers a large (hundreds of thousands of eligible individuals), international, and fully anonymous recruitment pool. Participants signed up and started the experiment within moments of its publication, and every stage of the game was timed to no more than thirty seconds. It is therefore extremely unlikely that any kind of chats or other forms of external coordination occurred during the game, as participants could not arrange beforehand to participate together and had neither the means nor the time to find out who they were playing against once they had started playing.

3.3 Experimental procedures

- Briefing and consent: Participants were briefed about their rights, emphasizing voluntary participation, anonymity, and data usage limited to academic research.
- Experiment participation: Participants were randomly assigned, approximately half to the exclusive environment, and the remainder to the inclusive environment. Within each environment, participants were randomly assigned, half to the "Jill and Jane" treatment, with the remainder assigned to the "Jane and Jill" treatment.
- Results and payment: The final experiment results and corresponding payments were disclosed to participants.

3.4 Reward structure

- Participants were rewarded with points for every correct guess they made, if the overall accuracy of their guesses exceeded a certain threshold. Participants were clearly informed of their earnings immediately following every round.
- These points were converted to British pounds at a predefined rate which was communicated to participants beforehand.⁸
- Following the conclusion of the experiment, a lottery was held, in which 20% of participants were awarded a bonus payment according to the points they had earned throughout the experiment.
- The only difference in reward structure between the two environments is as follows:
 - In the inclusive environment, 80% of the participants with the most accurate guesses earned ten points per correct guess.
 - In the exclusive environment, the top 20% earned fifty points per correct guess.⁹
 - In every other way, the two environments were identical.

3.5 Participants

A total of 191 participants were recruited through Prolific, of which 20 failed to participate in a satisfactory manner (defined as participating in less than twenty rounds),¹⁰ giving a sample of 171 participants who provided us with data. Participants

⁸ For their participation in the experiment, participants were assured a minimum compensation of six British pounds per hour.

⁹ As the probability of winning is considerably lower in the exclusive environment, higher rewards were offered as compensation in this environment.

¹⁰ If a participant did not actively play in a given round, the data recorded a missing value for that specific player-round combination. This made it straightforward to determine whether a participant was indeed participating throughout the experiment.

were divided into eight groups of 21-25 participants each (see Table 1 for more details). Four groups played in the inclusive environment, and four in the exclusive environment. In each environment, two groups were assigned to the "Jane and Jill" treatment, and the other two to the "Jill and Jane" treatment. Session scheduling ensured consistency in terms of weekdays and hours to minimize potential external influences.

Participants were directed to a virtual waiting room, which is a built-in feature offered by oTree. As soon as the waiting room had filled up with the participants, we launched the session. Upon the completion, participants were given a completion code and redirected back to Prolific, where they automatically received the base participation fee. Participants who had signed up but failed to complete the experiment were not paid, with the exception of those who had encountered technical difficulties, who were awarded full or partial compensation depending on the time and effort spent on the experiment. To ensure proper participation and guarantee that all information provided was fully visible to the participants, participation was restricted to desktop and laptop computers only, and all participants' input was monitored throughout the experiment using the oTree platform.¹¹ Bonus payments were calculated and paid out manually, also through Prolific.

Environment	Treatment	Groups	Participants
Inclusive	Jane and Jill	2	42
	Jill and Jane	2	41
	Total	4	83
Exclusive	Jane and Jill	2	45
	Jill and Jane	2	43
	Total	4	88

Table 1: Participant Distribution by Group

¹¹ To verify this, several trial runs of the experiment were conducted, both by the researchers and with live participants recruited in the same way as those who participated in the experiment itself.

Total

Notes: The table presents the distribution of participants among the environments and treatments. "Groups" indicates the number of groups assigned to the treatment, while "Participants" indicates the total number of participants assigned to each environment-and-treatment combination.

Table 2 shows a summary of the descriptive statistics of the participants in the experiment. These statistics were collected by Prolific, and not all participants provided information for every category. Table 2 shows that there are no statistical differences in the control variables between the treatment groups.

3.6 Prediction

We anticipate pronounced herding behavior in the inclusive environment. Specifically, we predict that when a majority choice emerges, the size of this majority will be considerably larger in the inclusive environment, approaching 100% of participants. Conversely, in the exclusive environment, majorities are expected to be smaller, closer to an evenly split 50%.

	Inclusive	Exclusive	Total
Participants	83	88	171
Female	0.27 (0.45)	0.44 (0.49)	0.35 (0.48)
Age	30.97 (9.94)	33.38 (11.73)	32.25 (10.96)
Ethnicity			
Asian	0.8 (0.27)	0.08 (0.28)	0.08 (0.27)
Black	0.16 (0.37)	0.24 (0.43)	0.21 (0.41)
Mixed	0.08 (0.27)	0.05 (0.21)	0.06 (0.24)
White	0.65 (0.48)	0.59 (0.50)	0.62 (0.49)
Native English speaker	0.60 (0.49)	0.55 (0.50)	0.58 (0.50)

Table 2: Descriptive statistics

Student	0.42 (0.50)	0.27 (0.45)	0.34 (0.48)
Employment status			
Unemployed	0.18 (0.37)	0.23 (0.43)	0.21 (0.41)
Not in paid work	0.01 (0.12)	0.05 (0.22)	0.03 (0.18)
Due to start new job	0.01 (0.12)	0 (0)	0.01 (0.08)
Employed part-time	0.13 (0.34)	0.16 (0.37)	0.15 (0.35)
Employed full-time	0.55 (0.50)	0.51 (0.50)	0.53 (0.50)
Other	0.10 (0.31)	0.04 (0.19)	0.07 (0.26)

Notes: The table shows the observations in the inclusive and exclusive environments. The table shows the mean and standard deviation (in parentheses) for the demographic characteristics: female indicator (Female); age in years (Age); Ethnicity, with categories including Asian, Black, Mixed, and White; indicator for whether the participant is a native English speaker; indicator for whether the participant is a student; and employment status, including unemployed, not in paid work (and not looking for work), due to start a new job, employed part-time, employed full-time, and other.

Furthermore, we foresee greater consistency in choices within the inclusive environment. Participants are likely to cluster around either "Jill" or "Jane" and maintain this choice across rounds. In contrast, the exclusive environment should exhibit more dynamic guess patterns, with participants frequently altering their choices in an attempt to stand apart from a dominant majority.

To quantify these behaviors, we propose the following statistical models:

1) clustering_{*i*} = $\alpha_1 + \beta_1 \times \text{inclusivity}_i + \beta_2 \times x_i + \epsilon_i$

2) fluctuation_{*i*} = $\alpha_2 + \gamma_1 \times \text{inclusivity}_i + \gamma_2 \times x_i + \zeta_i$

In these models:

- "Clustering" measures the tendency of individuals to align with either "Jill" or "Jane," as observed in the majority. Practically, this is defined as the proportion of participants in a given round whose choice aligns with the majority of the previous rounds.
- "Fluctuation" assesses the likelihood of a participant switching their choice between "Jill" and "Jane" from one round to the next. To measure this, we use

the proportion of participants in a given round who changed their decision compared to the previous round.

- "Inclusivity" is a binary variable denoting the environment type, which equals
 1 in the inclusive environment and 0 in the exclusive environment.
- *x* represents a range of control variables.
- ε and ζ are error terms for Models (1) and (2), respectively.

We hypothesize that in Model 1 (clustering), the coefficient β_1 will be positive, indicating higher clustering in the inclusive environment. Conversely, in Model 2 (fluctuation), γ_1 is expected to be negative, reflecting increased fluctuation in the exclusive environment. Furthermore, we anticipate the experiment to involve a learning process, so the magnitude of both coefficients is expected to grow in the later rounds of the game compared to its early rounds.

4 **Results**

4.1 Clustering

Figure 2 simply presents the distribution of bets made by participants that conformed to the majority and the minority in the preceding round. We identified the decision that received the majority of bets in the round before a given round; then, if the participant's choice in a given round was the same option as the one that received the majority of bets, that decision was assigned to the "majority" group; otherwise, it was assigned to the "minority" group. Naturally, the data only contain rounds 2-30, omitting round 1, as well as any instances immediately following an even 50% split between "Jill" and "Jane."

The figure reveals that participants in the inclusive environment stuck to the majority choice 1,385 times (62.87%), whereas their counterparts in the exclusive environment only did so 1,250 times (54.28%). This disparity grows stronger as the game progresses. In the early game (rounds 2-10), 58.20% and 56.53% of participants conform to the majority in the inclusive and exclusive environments respectively. By the late game, this grows to 66.16% for the inclusive environment, and shrinks to 53.03% in the exclusive environment.

Note that in both environments there is a certain bias in favor of the majority, which increases in the inclusive environment and decreases in the exclusive one. This bias is potentially driven by the standard force of herding, i.e., the perception that the majority knows better than oneself (see Bracht et al., 2010).

Moreover, we calculated the expected value of the majority size using a binomial distribution. Specifically, for 25 participants, the expected majority size is approximately 0.5804, based on the following formula:

$$\sum_{k=0}^{25} \left(\frac{1}{2}\right)^{25} \binom{25}{k} \max(k, 25-k)$$

Figure 2 illustrates this expected value of the majority size with a dotted horizontal line. This value falls between the observed majority sizes in the exclusive and inclusive environments, providing a useful baseline for comparison. It allows us to evaluate which environment aligns more closely with the equilibrium prediction under the assumption that "Jill" and "Jane" employ random strategies with equal probabilities.

The results indicate that, overall, the expected value lies roughly at the midpoint of the range representing the majority sizes observed in the two environments. This suggests that the equilibrium predictions perform comparably in both contexts. However, in the later stages of the experiment, the Nash equilibrium prediction appears to align more closely with the data in the exclusive environment.

Figure 3 presents the average majority size for participants in the two environments. Majority size is defined as the portion of participants in a given group and round who selected the majority choice. For each round, we identified the guess that received the majority of votes. A majority size of 1 indicates a unanimous selection (of either "Jane" or "Jill") by all participants in a given round. A majority size of 0.5 signifies an equal distribution of participants between both guesses. This figure shows the average majority size for each round of the game for the inclusive and exclusive environments. This reveals a growing divergence in majority size between the two environments, from less than five percentage points in the first round of the game to over fifteen percentage points in its final round.

Figure 2: Majority and Minority Decisions by Game Stage



Note: This figure shows the distribution (in percentages) of instances in which a player decided to make a bet conforming with the majority and the minority in the preceding round (absolute number of observations appears underneath each column). Capped ranges indicate 95% confidence intervals. From left to right, each two pairs of columns show percentages for all rounds (2-30), for the early game only (rounds 2-10), for the mid-game (rounds 11-20), and for the late game (rounds 21-30) respectively. The share of majority decisions on the left of each column pair (in blue), and the share of minority decisions is on the right (in red). The dashed vertical line indicates the expected majority size (58.04%), derived from a binomial distribution.

In line with our predictions, the inclusive environment displayed a tendency for choices to cluster towards larger majority sizes. In other words, participants exhibited a stronger preference for either "Jane" or "Jill." Meanwhile, the exclusive environment exhibited a more evenly distributed pattern of guesses. Consistent with our prediction that participants would undergo a learning process, we see that majority sizes in both environments are relatively similar in rounds 1-10. In rounds 11-30, however, there is a noticeable divergence, with majority sizes in the exclusive environment tending to be closer to 0.5 than those in the inclusive environment, and there is not a single round in the second half of the game in which the average majority size in the exclusive environment was larger than in the inclusive environment.

Figure 3: Majority Size by Round and Treatment



Note: This figure shows the average majority size for the inclusive environment (in blue) and the exclusive environment (in red). Each point represents the simple average of four observations (representing the four groups in each environment). The x axis represents the round number, starting from 1 and ending at 30. The y axis represents the majority size (0.5 - half of all participants selected "Jane," half selected "Jill;" 1.0 - all participants selected "Jane," or all participants selected "Jill").

To more formally test this, we ran T test¹² comparing the majority size in each round between the two treatment groups. The inclusive environment shows larger majority sizes on average, of 66.11%, compared to 60.88% in the exclusive environment. This difference is significant at the 1% level. Furthermore, while in rounds 1-10, the difference between the environments is just 3.35 percentage points, and not statistically significant, in rounds 21-30 the difference more than doubles to 6.88 percentage points, statistically significant at the 1% level.

Table 3: Clustering (probit)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Rounds	2-30	2-30	2-10	11-20	21-30	2-20	11-30

 $^{^{12}}$ A skewness and kurtosis test conducted on the distribution of majority sizes indicates a deviation from normal distribution at a 95% confidence level. Consequently, to accommodate the non-normality of the data, we opted for a T test instead of a Z test for our analysis. The T test is more appropriate in this context as it is less sensitive to deviations from normal distribution, especially in smaller sample sizes.

Inclusive	0.2419***	0.2505***	0.0477	0.2719*	0.4115***	0.1660*	0.3378***
dummy	(0.0760)	(0.0806)	(0.1431)	(0.1403)	(0.1377)	(0.1000)	(0.0980)
Demographic controls	No	Yes	Yes	Yes	Yes	Yes	Yes
Session fixed	Yes	Yes	Yes	Yes	Yes	Yes	Yes
effects							
Number of observations	4,506	4,187	1,311	1,435	1,441	2,746	2,876
Pseudo R ²	0.0108	0.0107	0.0096	0.0246	0.0232	0.0078	0.0204

Notes: The table presents probit marginal effect coefficients for the treatment effect on whether participants' choice conformed to the majority choice in the preceding round. The dependent variable is an indicator variable for whether a participant's choice belongs to the majority in the previous round. Column (1) shows the basic model with session fixed effects. Column (2) shows the basic model with session fixed effects. Columns (3), (4), and (5) show the coefficient for only the early (rounds 2-10), mid (rounds 11-20), and late game (rounds 21-30) respectively. Column (6) shows the coefficient excluding the late game (rounds 2-20), and column (7) shows the coefficient excluding the early game (rounds 11-30). The inclusive dummy equals 1 if the environment played is the inclusive environment, and 0 if it is the exclusive environment. Standard errors appear in parentheses. * and *** denote significance at the 10% and 1% levels respectively.

Table 3 presents the marginal effect coefficients for clustering (that is, whether a participant's guess in a given round matches the majority choice in the preceding round). The observations are individual decision points; that is, each observation is a decision made by a specific participant in a specific round of the game. The dependent variable is an indicator which equals 1 if the participant's choice in a given round is identical to the majority choice in the previous round, and 0 if the participant's choice is identical to the minority choice in the previous round. Columns (1) and (2) show the full results (rounds 2-30, as there is no preceding for round 1) of all games played, without and with additional demographic control variables. Column (3) shows the early-game results (rounds 2-10), Column (4) shows the mid-game results (rounds 11-20), and Column (5) shows the late-game results (rounds 2-30). Column (6) shows the results for the early and mid-game (rounds 2-20), and Column (7) shows the results for the mid and late game (rounds 11-30). All instances where the preceding round was evenly split between "Jill" and "Jane" are omitted.

Consistent with our predictions, the results show that overall, participants in the inclusive environment were 24-25% more likely to stay with the majority, statistically significant at the 1% level. Columns (3-7) illustrate the learning process observed in

the game. While in the early game, the difference between the environments is negligible, it rises to 27% in the mid-game phase, and grows to 41% in the late game, with these differences being significant at the 10% and 1% levels respectively. In the mid and late game combined, participants in the inclusive environment are over 33% more likely to stay with the majority, significant at the 1% level.

4.2 Fluctuation

A Z test comparing the share of participants in each round who changed their decision compared to their decision in the preceding round (either from "Jill" to "Jane" or from "Jane" to "Jill") found no statistically significant differences between the two environments. Table 4 displays the marginal effect coefficients for fluctuation (participants' decision in a given round to change their choice compared to their choice in the preceding round, either from "Jill" to "Jane" or vice versa). The dependent variable is an indicator variable which equals 1 if the participant's decision changed in a given round compared to the previous round, and 0 if the decision did not change. Since participants were told the outcome of their bets after each round, and people are more likely to "double down" on winning bets (see e.g., Gilovich et al., 1985; Abe et al., 2020), we sought to examine the differential effect of a correct bet on the individual's propensity to change their bet in the following round in the two environments. To this end, we included an independent variable showing the interaction between the environment and the accuracy of the participant's guess in the previous round. This variable is equal to 1 if the participant's previous bet was correct and if the participant is playing in the inclusive environment, and 0 otherwise.

Columns (1) and (2) show the full results (rounds 2-30, as there is no preceding for round 1) of all games played, without and with additional demographic control variables. Column (3) shows the early-game results (rounds 2-10), Column (4) shows the mid-game results (rounds 11-20), and Column (5) shows the late-game results (rounds 21-30). Column (6) shows the results for the early and mid-game (rounds 2-20), and Column (7) shows the results for the mid and late game (rounds 11-30).

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Rounds	2-30	2-30	2-10	11-20	21-30	2-20	11-30

Table 4: Fluctuation (probit)

Inclusive*correct	-0.3914***	-0.4016***	-0.2164	-0.7656***	-0.2065	-0.5088***	-0.4773***
	(0.0743)	(0.0773)	(1386)	(0.1334)	(0.1333)	(0.0957)	(0.0937)
Inclusive	0.3129***	0.2484***	0.0841	0.3800**	0.2691*	0.2498**	0.3167***
dummy	(0.0832)	(0.0880)	(0.1590)	(0.1506)	(0.1503)	(0.1091)	(0.1059)
Correct dummy	0.0383	0.0492	-0.0164	0.1417	0.0044	0.0741	0.0717
	(0.0512)	(0.0526)	(0.0944)	(0.0903)	(0.0910)	(0.0650)	(0.0637)
Demographic controls	No	Yes	Yes	Yes	Yes	Yes	Yes
Session fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Number of observations	4,778	4,453	1,386	1,524	1,543	2,910	3,067
Pseudo R ²	0.0107	0.0150	0.0138	0.0291	0.0244	0.0174	0.0185

Notes: The table presents probit marginal effect coefficients for the treatment effect on participants' choice to change their decision compared to the preceding round. The dependent variable is an indicator variable for choosing to change one's decision relative to the previous round. Column (1) shows the basic model with session fixed effects. Column (2) shows the basic model with session fixed effects and additional control variables. Columns (3), (4), and (5) show the coefficient for only the early (rounds 2-10), mid (rounds 11-20), and late game (rounds 21-30) respectively. Column (6) shows the coefficient excluding the late game (rounds 2-20), and column (7) shows the coefficient excluding the early game (rounds 11-30). The inclusive*correct variable equals 1 if the environment played is the inclusive environment and if the participant's guess in the previous round was correct, and 0 otherwise. The inclusive environment. The correct dummy equals 1 if the participant's bet in the previous round was correct, and 0 if it was incorrect. Standard errors appear in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels respectively.

Interestingly, while the results show an unexpectedly high overall propensity to fluctuate in the inclusive environment, they also show that, having placed a correct bet in a given round, participants in the inclusive environment were considerably more likely to stick to that decision in the following round. The coefficient for the latter effect appears to be considerably greater than that of the former, and varies widely, ranging from 20-76%, and averaging 39% for the entirety of the game. Also notably, we did not observe a straightforward learning process for fluctuation, which was higher in the midgame stage (rounds 11-20) than in either the early (rounds 2-10) or late game (rounds 21-30). This likely arises due to the fact that participants' decisions and learning process are influenced by the realization of their bet in the previous round, and not just the distribution of bets. Suppose that while herding starts to form the majority choice of bet

is wrong. Reinforcement will now push some people to join the minority (which guessed correctly in the last period). Indeed, this would be more prevalent among those who have not yet gained intuition about the advantage of being with the majority in this environment. This is in contrast to market entry games (e.g., Rapoport et al., 2000), where the environment is deterministic, facilitating faster coordination. Further supporting this explanation, in our alternative experiment specification, where participants were not told the realization of their bets until the completion of the game, the inclusive dummy consistently receives a negative coefficient (for details, see Appendix B).

5 Reinforcement Learning Model

We present here a simple reinforcement learning model of the strategic environment.¹³ We compare the behavior observed in the experimental game to the dynamics of the reinforcement learning model.

It is important to note that Nash equilibrium is a static concept, referring to a stable situation where no player has an incentive to unilaterally deviate. The process by which such equilibria emerge remains an open question in game theory. Some dynamic models, such as best response dynamics, fictitious play, or evolutionary dynamics, are known to converge to Nash equilibria in specific types of games.

We adopt a more suitable approach for an experimental study by designing a reinforcement learning model, following the framework of Erev and Roth (1998). This model incorporates assumptions about sequential dependencies in players behavior.

Let p_t^A , p_t^B be the propensities of playing A and B in period t, respectively.

Let q_t^A , q_t^B be the probabilities of playing A and B in period t respectively.

 w_t is the reinforcement weight at period t. This weight declines with t.

The probabilities of playing the two options are proportional to their respective propensities as follows:

$$q_t^A = \frac{p_t^A}{p_t^A + p_t^B}, \ q_t^B = (1 - q_t^A)$$

¹³ The formal Nash equilibrium of the static game is detailed in Appendix C.

We now have to define how the propensities change over time.

Let x(t) be the option chosen by a player in period t, and let y(t) be the option not chosen in period t:

$$x(t) \in \{A, B\}, y(t) \in \{A, B\}.$$

Let us assume for simplicity that $p_{t+1}^{y(t)} = p_t^{y(t)}$; that is, only the propensity of the choice played changes from one round to the next, while the propensity of playing the option not chosen remains unchanged in the following round.

We then define:

$$p_{t+1}^{x(t)} = p_t^{x(t)} + \delta^+ w_t [\pi(t) - \pi(t/y(t))]$$

$$w_{t+1} = (1 - d) w_t$$

$$\delta^+ = \delta \quad \text{if } \pi(t) - \pi\left(\frac{t}{y(t)}\right) \ge 0, \text{ and } 1 \text{ if } \pi(t) - \pi(t/y(t)) < 0.$$

where $\pi(t)$ is the payoff the player received in period t, $\pi(t/y(t))$ is the payoff the player would have received had he played the other option y(t), and δ is the weight of the impact of loss. d is the discount effect of the learning process.

Initial conditions are as follows:

$$p_1^B = 1$$

 p_1^A can be easily derived from the true distribution of bets participants made in the first round of the experiment.

The parameters of the model are, w_1 , d, and δ . Their values are determined using the experimental data.

To evaluate this model, we conducted a simulation of our experimental setting. In this simulation, a group of twenty simulated participants played thirty rounds of the game, betting between "1" and "0" in each round. The result of the bet was determined randomly by the computer. This was repeated 1,000 times each for the inclusive and exclusive environments.

We used mean squared error function to optimize for the fitness of the model with the experimental results, yielding the following optimal values: $q_1^1 = 0.675$, $w_1 = 0.05$, d = 0.1, $\delta = 0$.

Initial probabilities $(q_1^1 \text{ and } q_1^0)$ for round 1 were set to 67.5% for "1" and 32.5% for "0," to reflect the tendency of participants to prefer the first option presented to them, as observed in the real experiment. The reinforcement weight (w_1) in round 1 was set to 0.05, to prevent the magnitude of the rewards (10 and 50 points in the inclusive and exclusive environments respectively) from dominating the outcome. Learning (d) was set to a slow-intermediate rate of 0.1. The weight impact of gain (δ) was set to 0, as the behaviour of the simulated participants was closest to that of the real participants when they adjusted their propensities exclusively according to loss. Propensities and probabilities in rounds 2-30 were calculated according to our model. We then calculated the average majority size obtained in each round for the inclusive and exclusive environments.

The outcome of the simulation is presented in Figure 4. In both environments, the simulation starts with a majority size of approximately 67.5%. However, they diverge, with the majority size in the inclusive environment growing slightly, reaching approximately 70% by round 11, subsequently staying within the 69.2-70.6% range until round 30. Meanwhile, in the exclusive environment, following an initial jump to 77% in round 2, the majority size gradually shrinks, stabilizing at around 58.8% by round 14.¹⁴ This initial jump in the exclusive environment occurs as participants overcompensate, with too many seeking to escape the majority in the first round, resulting in the creation of an even bigger majority in subsequent rounds. As part of their learning process, participants gradually overcome this tendency to overcompensate, eventually settling on the smaller majority size as expected.

The experimental results show a somewhat smaller difference between the two environments than the simulation results, as well as a greater variance between rounds. The variance can be easily explained as occurring due to the considerably smaller sample size in the experimental results (191 participants in eight groups) than in the simulation (50,000 participants in 2,000 groups). The difference in means is more challenging to explain, but the apparent upward trend in the experimental results for the inclusive environment indicates that the learning process may be longer than the thirty rounds played in the game. However, this would likely be impractical to test in our

¹⁴ Our model predicts that the majority size in the exclusive environment will still be larger than a purely random 50%, as the realization of coin flips, particularly in the early rounds, may create a slight preference for one bet over the other.

experimental setting, as the repetitive tasks required, in conjunction with the fairly small amounts of money offered, would likely exceed many participants' attention spans, and could also result in information overload (as participants are shown the outcomes of every single preceding round played).





Note: This figure shows the average majority size for the inclusive environment (in blue) and the exclusive environment (in red). A continuous line represents the true results of the experiment, and a dashed line represents the results of the simulation. Each point represents the simple average of four observations (representing the four groups in each environment). The x axis represents the round number, starting from 1 and ending at 30. The y axis represents the majority size (0.5 – half of all participants selected "Jane," half selected "Jill;" 1.0 - all participants selected "Jane," or all participants selected "Jill").

6 Discussion and Conclusion

Amnon Rapoport was perhaps one of the most "cognitive" experimental economists. His primary interest lay in decision-making situations, whether individual or interactive, that pose a cognitive or analytical challenge to the parties involved. This contrasts with many other experimental economists who are more inspired by strategic situations that introduce moral dilemmas or generate emotional reactions in participants. In private conversations with the second author of this paper, Rapoport often emphasized that moral sentiments and emotions are challenging to study due to their susceptibility to multiple interpretations, both from experiment participants and researchers. We speculate that the game we studied here aligns with the nature of those Rapoport was particularly drawn to. Like many of Rapoport's games, it is far from trivial, yet a relatively short learning process in the lab yields strategic behavior consistent with equilibrium predictions.

Our experimental results underscore the significant influence of the strategic environment on herding and divergence in betting behavior. Consistent with our theoretical predictions, we observed a pronounced inclination for players to align with the majority in an inclusive environment, whereas in an exclusive setting, participants actively sought to distinguish their choices from the group norm. While a complete convergence to the Nash equilibria was not observed, the prevalence of a larger majority group in the inclusive environment was notable. This majority choice demonstrated stability, often persisting throughout the game once established in the early rounds. Conversely, in the exclusive environment, a marked tendency for diversification was evident, with participants not only balancing their guesses more evenly but also more likely to switch between "Jane" and "Jill" after making a winning bet—a behavior less observed in the inclusive setting.

Interestingly, the increasing preference for majority alignment in the inclusive environment as the game progresses is evidence of a learning process with dual aspects: First, it allows players to gradually understand the strategic considerations of the game. Second, it enables them to coordinate a choice around which a majority will gradually grow.

The classic game theoretic hawk-dove game offers an interesting analogy, particularly in the exclusive scenario where players compete intensely for a single funding opportunity. This dynamic creates a situation in which players must balance aggressive (hawk) and conservative (dove) strategies. In this context, divergence is incentivized, and strategic imbalance arises because one player benefits from diverging while the others align. The exclusive environment fosters strategic differentiation by introducing externalities, which may be positive or negative depending on the decision in question. Beyond the obvious context of contests, the potential applications of the model seem to be much broader, offering insights into social, economic, and political realms. These include the following:

1. Policy design for fostering competition in innovation: Consider a scenario where drug companies are competing to develop a drug using one of two technologies, T1 and T2. One of these technologies allows for much faster development, but firms initially do not know which is superior, assigning equal prior probabilities to each. Each firm must select a technology and begin development. Firms that choose the advanced technology develop faster than those opting for the inferior one, with development order determined randomly within each group. The government can influence the competition by deciding on the number of licenses (or production capacities) to issue, allocated on a first-come, first-served basis. If the government sets a high capacity constraint (allowing more than half of the firms to obtain licenses), the environment mirrors our inclusive scenario, leading all firms to choose the same technology. Conversely, if the government imposes a low capacity constraint (limiting licenses to fewer than half of the competing firms), it ensures that both technologies are tested, as firms split their decisions evenly to hedge their bets and maximize their chances of success.

2. Relative wealth and investment decisions: Whether in the context of economic success, intellectual reputation, or social prominence, status is often relative within a group. The nature of this status – inclusive or exclusive – profoundly influences agents' strategies. In an inclusive status environment, agents herd their bets seeking to be conventional. If their choice of behavior turns out to be wrong, they are protected against being dragged to the bottom of the social scale as would have happened had they been part of a small minority that acted wrongly. However, when status involves intense competition, people will be motivated to act unconventionally, as the only option to gain status is by being part of a small minority that acted correctly. Hence, in exclusive settings, people tend to diverge in an attempt to stand out. For instance, in a real estate market in the wake of a stock market boom, those who participated in the boom find their purchasing power significantly enhanced, whereas non-participants face reduced affordability. The set of winners is relatively large (inclusive environment), so investors will choose to herd with the market. Conversely, in high-end markets, only the most successful investors can afford luxury properties, while

others are left out, indicating a stark difference in outcomes based on investment strategies (exclusive environment). As a result, more investors will go against the market.

3. Financial markets and auction systems: Our findings can inform the design of financial markets and auction systems by providing insights into how competitive intensity affects herding and diversification. This understanding can help distinguish herding behavior that emerges for strategic reasons from that driven by psychological or informational factors. Such insights can aid in creating mechanisms to mitigate irrational market bubbles and promote more stable and efficient markets.

4. Political lobbies: Consider a group of lobbies that seek to make campaign contributions to one of two political leaders, in exchange for the leader adopting a policy favoring their interests. Only one of these leaders can be elected and each has the same chance of being elected. Each lobby has to decide which of the two candidates to support with campaign contributions. To avoid public criticism, the elected leader can only cave in to lobbies' pressure on a certain number of policies. Again, this example matches well with the games. Here, winning means that a lobby manages to gets the elected leader to approve the policy it is interested in. The implications of our model depend on the level of corruption prevailing in the underlying political environment. If the leader is willing to approve a large number of corrupt policies so that most of the lobbies will be served, then the equilibrium will involve herding and all lobbies will support the same leader. With less corruption, only a minority of the lobbies can be served, and campaign contributions will be split between the two candidates more equally.

In conclusion, by understanding how competitive intensity influences strategic decision-making, policymakers can design more effective regulations and incentives that promote a balanced and innovative environment.

Our study provides a comprehensive analysis of how competitive intensity influences strategic decision-making, with broad implications for both theoretical advancements and practical applications. By extending Rapoport's legacy of cognitive and analytical exploration in experimental economics, this paper contributes to the understanding of strategic behavior in competitive environments.

Bibliography

Abe, N., Nakai, R., Yanagisawa, K., Murai, T., & Yoshikawa, S. (2020). "Effects of sequential winning vs. losing on subsequent gambling behavior: analysis of empirical data from casino baccarat players." *International Gambling Studies*, *21*(1), 103–118.

Andersson, M., Hedesström, M., & Gärling, T. (2014). "A Social-Psychological Perspective on Herding in Stock Markets." *Journal of Behavioral Finance*, *15*(3), 226–234.

Avery, C.N., & Chevalier, J.A. (1999). "Herding over the career." *Economics Letters*, 63, 327–333.

Avery, C., & Zemsky, P. (1998). "Multidimensional Uncertainty and Herd Behavior in Financial Markets." *The American Economic Review*, 88(4), 724–748.

Banerjee, A.V. (1992). "A Simple Model of Herd Behavior." *The Quarterly Journal of Economics*, 107(3), 797–817.

Bikhchandani, S., & Sharma, S. (2000). "Herd Behavior in Financial Markets." *IMF Staff Papers*, 47(3), 279–310.

Bracht, J., Koessler, F., Winter, E., & Ziegelmeier, A. (2010). "Fragility of Information Cascades: An Experimental Study using Elicited Beliefs." *Experimental Economics*, *13*(2), 121–145.

Charness, G., Gneezy, U., & Halladay, B. (2016). "Experimental methods: Pay one or pay all." *Journal of Economic Behavior & Organization*, *131*, 141–150.

Cipriani, M., & Guarino, A. (2005). "Herd Behavior in a Laboratory Financial Market." *American Economic Review*, 95, 1427–1443.

Cipriani, M., & Guarino, A. (2009). "Herd Behavior in Financial Markets: An Experiment with Financial Market Professionals." *Journal of the European Economic Association*, 7(1), 206–233.

Cole, H.L., Mailath, G.J., & Postlewaite, A. (2001). "Investment and concern for relative position." *Review of Economic Design*, *6*, 241–261.

DeMarzo, P., Kremer, I., & Kaniel, R. (2008). "Relative Wealth Concerns, and Financial Bubbles." *Review of Financial Studies (RFS), 21*(1), 19–50.

Drehmann, M., Oechssler, J., & Roider, A. (2005). "Herding and Contrarian Behavior in Financial Markets – An Internet Experiment." *American Economic Review*, 95, 1403–1425.

Drehmann, M., Oechssler, J., & Roider, A. (2007). "Herding with and without payoff externalities — an internet experiment." *International Journal of Industrial Organization*, 25(2), 391–415.Erev, I., & Roth, A.E. (1998). "Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria." *The American Economic Review*, 88(4), 848–881.

Espinosa-Mendez, C., & Maguieira, C.P. (2023). "Herding Behavior of Business Groups During COVID-19: Evidence in Chile." *Investigación Economica*, 82(324), 98–125.

Foucault, T., & Frésard, L. (2019). "Corporate Strategy, Conformism, and the Stock Market." *The Review of Financial Studies*, *32*(3), 905–950.

Gilovich, T., Vallone, R., & Tversky, A. (1985). "The hot hand in basketball: On the misperception of random sequences." *Cognitive Psychology*, *17*(3), 295–314.

Hamdani, A., Kandel, E., Mugerman, Y., & Yafeh, Y. (2017). "Incentive Fees and Competition in Mutual Funds: Evidence from a Regulatory Experiment." *Journal of Law, Finance, and Accounting*, 2(1), 49–86.

Jiang, H.J., & Verardo, M. (2018). "Does Herding Behavior Reveal Skill? An Analysis of Mutual Fund Performance." *The Journal of Finance*, *73*(5), 2229–2269.

Koch, A. (2017). "Herd Behavior and Mutual Fund Performance." *Management Science*, 63(11), 3849–3873.

Leyva-de la Hiz, D.I., Aragon-Correa, A., & Earle, A.G. (2021). "Innovating for Good in Opportunistic Contexts: The Case for Firms' Environmental Divergence." *Journal of Business Ethics*, *176*, 705–721.

Medrano, L.A., & Vives, X. (2001). "Strategic Behavior and Price Discovery." *The RAND Journal of Economics*, *32*(2), 221–248.

Mugerman, Y., Sade, O., & Shayo, M. (2014). "Long term savings decisions: Financial reform, peer effects and ethnicity." *Journal of Economic Behavior and Organization*, *106*, 235–253.

Mugerman, Y., Steinberg, N., & Wiener, Z. (2022). "The exclamation mark of Cain: Risk salience and mutual fund flows." *Journal of Banking & Finance, 134*.

Mugerman, Y., Winter, E., & Yafeh, T. (2024). "Performance Monitoring and the Incentives for Exertion of Effort." *The Quarterly Journal of Finance*, *14*(02).

Park, A., & Sabourian, H. (2011). "Herding and Contrarian Behavior in Financial Markets." *Econometrica*, 79(4), 973–1026.

Rapoport, A. (1995). "Individual strategies in a market entry game." *Group Decision and Negotiation*, *4*, 117–133.

Rapoport, A., & Seale, D. A. (2008). "Coordination success in non-cooperative large group market entry games." In: Plott, C. R., & Smith, V. L. (Eds.), *Handbook of Experimental Economic Results*, 273–295. New York: North-Holland.

Rapoport, A., Seale, D.A., Erev, I., & Sundali, J.A. (1998). "Equilibrium play in large market entry games." *Management Science*, *44*, 119–141.

Rapoport, A., Seale, D., & Winter, E. (2000). "An Experimental Study of Coordination and Learning in Two-Market Entry Games." *Economic Theory*, *16*, 661–687.

Rapoport, A., Seale, D.A., & Winter, E. (2002). "Coordination and Learning Behavior in Large Groups with Asymmetric Players." *Games and Economic Behavior*, *39*, 111–136.

Scharfstein, D.S., & Stein, J.C. (1990). "Herd Behavior and Investment." *The American Economic Review*, 80(3), 465–479.

Sunder, S., Kim, K.H., & Yorkston, E.A. (2019). "What Drives Herding Behavior in Online Ratings? The Role of Rater Experience, Product Portfolio, and Diverging Opinions." *Journal of Marketing*, 83(6), 93–112.

Yousaf, Y., & Yarovaya, L. (2022). "Herding Behavior in Conventional Cryptocurrency Market, Non-Fungible Tokens, and Defi Assets." *Finance Research Letters*, 50.

Zhang, Y., & Gimeno, J. (2017). "Earnings Pressure and Competitive Behavior: Evidence from the U.S. Electricity Industry." *Academy of Management Journal*, 53(4).

Appendix A – Experiment Screenshots

The experimental screens and data shown in this Appendix are derived from a trial run conducted with dummy participants for demonstrative purposes. These screenshots are illustrative and do not represent actual game sessions with real participants. In the actual experiments, participants received complete data and feedback for all rounds as described in the main text.

Initial Screen

Thank you for agreeing to participate in this study.

This study will take approximately 30 minutes. Through your participation, you will be able to earn points based on your responses as part of an experiment combining elements of luck, interpersonal interaction, and answering questions. At the end of the experiment, you will participate in a lottery with the other participants in this game, in which you will have a 1/5 chance of receiving a bonus payment based on the number of points earned, at a ratio of: 30 points = 1 GBP.

Please press the space bar to proceed to the next screen

Prolific ID Screen

Prolific ID

Please enter your prolific ID:

Continue

Instruction Screen (Version 1 – Jill first, exclusive environment)

In this experiment a bet will be made 30 times consecutively between two individuals, Jill and Jane. For each bet, you will have 30 seconds to guess whether the winner will be Jill or Jane. In each round, if you guessed correctly, and if you were ranked among the **20% of guessers** (1) whose guesses are the closest to the true outcomes of all the bets up to this point, you will earn 10 bonus points for this round. Otherwise, you will not win any points for this round. At the side of the screen, you will see your ranking in the game, as well as the distribution of all participant guesses between Jill and Jane for every preceding bet, in both numbers and percentages.

The outcome of each bet will be displayed immediately following the end of the countdown.

The game will begin once all participants have announced that they are ready. There may be a delay of up to 120 seconds (two minutes) until this happens.

Please press the space bar when you are ready to begin.

Instruction Screen (Version 2 – Jane first, exclusive environment)

In this experiment a bet will be made 30 times consecutively between two individuals, Jane and Jill. For each bet, you will have 30 seconds to guess whether the winner will be Jane or Jill. In each round, if you guessed correctly, and if you were ranked among the **20% of guessers** (1) whose guesses are the closest to the true outcomes of all the bets up to this point, you will earn 10 bonus points for this round. Otherwise, you will not win any points for this round. At the side of the screen, you will see your ranking in the game, as well as the distribution of all participant guesses between Jane and Jill for every preceding bet, in both numbers and percentages.

The outcome of each bet will be displayed immediately following the end of the countdown

The game will begin once all participants have announced that they are ready. There may be a delay of up to 120 seconds (two minutes) until this happens.

Please press the space bar when you are ready to begin.

Instruction Screen (Version 3 – Jill first, inclusive environment)

In this experiment a bet will be made 30 times consecutively between two individuals, Jill and Jane. For each bet, you will have 30 seconds to guess whether the winner will be Jill or Jane. In each round, if you guessed correctly, and if you were ranked among the **80% of guessers (4)** whose guesses are the closest to the true outcomes of all the bets up to this point, you will earn 10 bonus points for this round. Otherwise, you will not win any points for this round. At the side of the screen, you will see your ranking in the game, as well as the distribution of all participant guesses between Jill and Jane for every preceding bet, in both numbers and percentages.

The outcome of each bet will be displayed immediately following the end of the countdown.

The game will begin once all participants have announced that they are ready. There may be a delay of up to 120 seconds (two minutes) until this happens.

Please press the space bar when you are ready to begin.

Instruction Screen (Version 4 – Jane first, inclusive environment)

In this experiment a bet will be made 30 times consecutively between two individuals, Jane and Jill. For each bet, you will have 30 seconds to guess whether the winner will be Jane or Jill. In each round, if you guessed correctly, and if you were ranked among the **80% of guessers (4)** whose guesses are the closest to the true outcomes of all the bets up to this point, you will earn 10 bonus points for this round. Otherwise, you will not win any points for this round. At the side of the screen, you will see your ranking in the game, as well as the distribution of all participant guesses between Jane and Jill for every preceding bet, in both numbers and percentages.

The outcome of each bet will be displayed immediately following the end of the countdown.

The game will begin once all participants have announced that they are ready. There may be a delay of up to 120 seconds (two minutes) until this happens.

Please press the space bar when you are ready to begin.

Round 115

Time left to complete this page: 0:29

In this experiment a bet will be made 30 times consecutively between two individuals, Jane and Jill. For each bet, you will have 30 seconds to guess whether the winner will be Jane or Jill. In each round, if you guessed correctly, and if you were ranked among the **80% of guessers (4)** whose guesses are the closest to the true outcomes of all the bets up to this point, you will earn 10 bonus points for this round. Otherwise, you will not win any points for this round. At the side of the screen, you will see your ranking in the game, as well as the distribution of all participant guesses between Jane and Jill for every preceding bet, in both numbers and percentages.

The outcome of each bet will be displayed immediately following the end of the countdown.

bet 1 out of 30:

⊖ Jane ⊖ Jill

Choose one of the options, then press the spacebar to proceed.

Round 1 – Results

Jill

2 (40%)

Time left to complete this page: 0:14

In this experiment a bet will be made 30 times consecutively	Bet #	Jane
between two individuals, Jane and Jill. For each bet, you will have 30 seconds to quess whether the winner will be Jane or Jill.	1	3 (60%)
In each round, if you guessed correctly, and if you were ranked among the 80% of guessers (4) whose guesses are the closest to the true outcomes of all the bets up to this point, you will earn 10 bonus points for this round. Otherwise, you will not win any points for this round. At the side of the screen, you will see your ranking in the game, as well as the distribution of all participant guesses between Jane and Jill for every preceding bet, in both numbers and percentages. The outcome of each bet will be displayed immediately following the end of the countdown.		
Bet 1 out of 30		
The outcome of the bet is Jane		
You guessed Jane		
You are ranked 1 out of 5		
Therefore, you did earn 10 points in this round.		

Press the spacebar to continue.

¹⁵ The following screenshots use Version 4 as an example. The differences in instructions between the versions remain consistent throughout the entire game. While not all rounds of the game are shown here, each round follows the same structure.

Round 2

n this	experiment	a bet will be	e made 30 times consecutively	Bet #	Jane	llit	
oetwee nave 30	en two indivi) seconds to	duals, Jane quess whe	and Jill. For each bet, you will ther the winner will be Jane or Jill.	1	3 (60%)	2 (40%)	
mong o the f arn 10 ny po our ra particip pet, in The ou ollowi	the 80% of true outcom) bonus poir ints for this nking in the boant guesses both number tcome of ea ng the end of	guessers (es of all the nts for this r round. At the game, as v between J ers and pero ch bet will l of the count	(4) whose guesses are the closest e bets up to this point, you will round. Otherwise, you will not win he side of the screen, you will see well as the distribution of all ane and Jill for every preceding centages. be displayed immediately tdown.				
Det 2 o Jane Jill Choose	ut of 30: e e one of the	e options, t	hen press the spacebar to				
Det 2 o Jane Jill Choose Drocee Bet #	ut of 30: e one of the d . Jane	e options, t	then press the spacebar to Outcome	Ranking	Points won in round	Total points	

Round 1716

Time left to complete this page: 0:01

bet 17 out

⊖ Jane ⊖Jill

n this e	xperiment	a bet will b	e made 30 times consecutively		Bet #	Jane	л	Bet #	Jane	л
have 30	seconds to	iduals, Jane o guess wh	e and Jill. For each bet, you will ether the winner will be Jane or		1	3 (60%)	2 (40%)	16	0 (0%)	0 (0%)
iil. In ea	ach round,	if you gues	sed correctly, and if you were		2	4 (80%)	1 (20%)			
anked a	among the	80% of gu	Jessers (4) whose guesses are the		3	0 (0%)	0 (096)			
will earr	10 bonus	points for	this round. Otherwise, you will not			0.000	0.000			
vin any	points for	this round.	At the side of the screen, you will		4	o (use)	0 (0%)			
ee you	r ranking ir	n the game, r. between	, as well as the distribution of all		5	0 (0%)	0 (0%)			
participa pet, in b	oth numbe	ers and per	centages.		6	0 (0%)	0 (096)			
The out	come of ea	ch bet will	be displayed immediately		7	0 (0%)	0 (0%)			
ollowin	g the end (of the coun	itdown.		8	0 (0%)	0 (0%)			
oet 17 d	out of 30:				9	0 (0%)	0 (0%)			
Jane					10	0.000	0.000			
) lill					10	u (USE)	n (nar)			
hoose	one of the	e options	then press the spacebar to		11	0 (0%)	0 (096)			
proceed	d.	e options,	the press of space of the		12	0 (0%)	0 (096)			
					13	0 (0%)	0 (096)			
					14	0 (0%)	D (096)			
					15	0 (0%)	0 (0%)			
Bet #	Jane	Jill	Outcome	Ranking		Poir	nts won in	round	Total	points
16	0 (0%)	0 (0%)	Jill (your guess was incorrect)	20% (1 a	out of 5) 0 p	pints		20 p	oints
15	0 (0%)	0 (0%)	Jill (your guess was incorrect)	20% (1 a	out of 5) 0 pr	oints		20 p	oints
14	0 (0%)	0 (0%)	Jane (your guess was incorrect)	20% (1 a	out of 5) 0 pr	oints		20 p	oints
13	0 (0%)	0 (0%)	Jane (your guess was incorrect)	20% (1 a	out of 5) 0 pr	pints		20 p	oints
12	0 (0%)	0 (0%)	Jane (your guess was incorrect)	20% (1 a	out of 5) 0 po	pints		20 p	oints
11	0 (0%)	0 (0%)	Jane (your guess was incorrect)	20% (1 a	out of 5) 0 po	pints		20 p	oints
10	0 (0%)	0 (0%)	Jill (your guess was incorrect)	20% (1 d	out of 5) 0 po	pints		20 p	oints
9	0 (0%)	0 (0%)	Jane (your guess was incorrect)	20% (1 a	out of 5) 0 p	pints		20 p	oints
8	0 (0%)	0 (0%)	Jill (your guess was incorrect)	20% (1 a	out of 5) 0 p	pints		20 p	oints
7	0 (0%)	0 (0%)	Jane (your guess was incorrect)	20% (1 a	out of 5) 0 p	pints		20 p	oints
6	0 (0%)	0 (0%)	Jane (your guess was incorrect)	20% (1 a	out of 5) 0 p	pints		20 p	oints
5	0 (0%)	0 (0%)	Jill (your guess was incorrect)	20% (1 a	out of 5) 0 p	pints		20 p	oints
4	0 (0%)	0 (0%)	Jill (your guess was incorrect)	20% (1 a	out of 5) 0 pr	oints		20 p	oints
3	0 (0%)	0 (0%)	Jill (your guess was incorrect)	20% (1 a	out of 5) 0 pr	oints		20 p	oints
2	4 (80%)	1 (20%)	Jane (your guess was correct)	20% (1 a	out of 5) 10	points		20 p	oints
1	3 (60%)	2 (40%)	Jane (your guess was correct)	20% (1 a	out of 5) 10	points		10 p	oints

¹⁶ Note that this is a screenshot for demonstrative purposes from a trial run of the experiment with five dummy participants (and no real humans playing other than the authors). Hence, the tables for this and subsequent rounds appear to contain data for only rounds 1 and 2, and only five participants. In the real experiment, participants would have been shown complete data for every single round, with a much higher number of participants. However, these screens were displayed only to participants on their personal devices, and thus screenshots were not taken as the experiment was in progress.

<u>Round 30</u>

Time left to complete this page: 0:22

n this e	experiment	a bet will b	e made 30 times consecutively		Bet #	Jane	л	Bet #	Jane	JII.	
etweei ave 30	n two indiv seconds to	iduals, Jane o guess wh	e and Jill. For each bet, you will ether the winner will be Jane or		1	3 (60%)	2 (40%)	16	0 (0%)	0 (0%)	
il. In ea anked	ach round, among the	if you gues 80% of cu	used correctly, and if you were		2	4 (80%)	1 (20%)	17	0 (0%)	0 (0%)	
losest	to the true	outcomes	of all the bets up to this point, you		в	0 (0%)	0 (0%)	18	0 (0%)	D (096)	
rill earr rin any	points for	points for this round.	this round. Otherwise, you will not . At the side of the screen, you will		4	0 (0%)	0 (0%)	19	0 (0%)	0 (0%)	
ee you articin	r ranking ir ant ouesse	n the game s between	as well as the distribution of all lane and till for every preceding		5	0 (0%)	0 (0%)	20	0 (0%)	0 (0%)	
et, in t	oth numb	ers and per	centages.		6	0 (0%)	0 (0%)	21	0 (0%)	0 (0%)	
he out ollowin	come of ea og the end	och bet will of the cour	be displayed immediately itdown.		7	0 (096)	0 (0%)	22	0 (096)	0 (0%)	
et 30 d	out of 30:				8	0 (096)	0 (0%)	23	0 (0%)	0 (096)	
lana					9	0 (096)	0 (0%)	24	0 (096)	0 (0%)	
) Jane) Jill					10	0 (0%)	0 (0%)	25	0 (096)	0 (0%)	
hoose	one of th	e options,	then press the spacebar to		12	0 (0%)	0 (0%)	25	0 (090)	0 (0%)	
rocee	d.				12	0 (0%)	0 (0%)	28	0 (096)	0 (0%)	
					14	0 (096)	0 (0%)	29	0 (096)	0 (096)	
					15	0 (096)	0 (0%)		- from	- Jacob	
Bet #	Jane	Jill	Outcome	Ranking	9	Poir	nts won in	round	Tota	points	
29	0 (0%)	0 (0%)	Jill (your guess was incorrect)	20% (1	out of S) 0 pc	pints		20 p	pints	
28	0 (0%)	0 (0%)	Jane (your guess was incorrect)	20% (1	out of S) 0 po	pints		20 p	pints	
27	0 (0%)	0 (0%)	Jill (your guess was incorrect)	20% (1	out of S) 0 po	pints		20 p	oints	
26	0 (0%)	0 (0%)	Jane (your guess was incorrect)	20% (1	out of S) 0 po	pints		20 p	pints	
25	0 (0%)	0 (0%)	Jane (your guess was incorrect)	20% (1 out of 5) 0 points					20 points		
24	0 (0%)	0 (0%)	Jill (your guess was incorrect)	20% (1	out of S) 0 pc	pints		20 p	pints	
23	0 (0%)	0 (0%)	Jane (your guess was incorrect)	20% (1	out of S) 0 pc	pints		20 p	pints	
22	0 (0%)	0 (0%)	Jill (your guess was incorrect)	20% (1	out of S) 0 pc	pints		20 p	pints	
21	0 (0%)	0 (0%)	Jane (your guess was incorrect)	20% (1	out of S) 0 pc	pints		20 p	pints	
20	0 (0%)	0 (0%)	Jane (your guess was incorrect)	20% (1	out of S) 0 pc	pints		20 p	pints	
19	0 (0%)	0 (0%)	Jane (your quess was incorrect)	20% (1	out of S) 0 pc	pints		20 p	pints	
18	0 (0%)	0.(0%)	Ell (upur quess was incorrect)	20% (1	out of S	0.0	nints		20 n	nints	
17	0 (0%)	0.(0%)	III (your guess was incorrect)	2095 (1	out of 5) 0m	ninte		20 p	ninte	
16	0 (0%)	0 (0%)	SII (your guess was incorrect)	2010 (1	out of 5	, op	ainta		20 p	uints	
10	U (U%)	U (U%)	Jill (your guess was incorrect)	20% (1	out of a) Up	oints		20 p	DINTS	
15	0 (0%)	0 (0%)	Jill (your guess was incorrect)	20% (1	out of S) 0 pc	oints		20 p	oints	
14	0 (0%)	0 (0%)	Jane (your guess was incorrect)	20% (1	out of S) 0 po	pints		20 p	oints	
13	0 (0%)	0 (0%)	Jane (your guess was incorrect)	20% (1	out of S) 0 pc	pints		20 p	pints	
12	0 (0%)	0 (0%)	Jane (your guess was incorrect)	20% (1	out of S) 0 po	pints		20 p	oints	
11	0 (0%)	0 (0%)	Jane (your guess was incorrect)	20% (1	out of S) 0 pc	pints		20 p	pints	
10	0 (0%)	0 (0%)	Jill (your guess was incorrect)	20% (1	out of S) 0 po	pints		20 p	pints	
9	0 (0%)	0 (0%)	Jane (your guess was incorrect)	20% (1	out of S) 0 po	pints		20 p	pints	
8	0 (0%)	0 (0%)	Jill (your guess was incorrect)	20% (1	out of S) 0 pc	pints		20 p	pints	
7	0 (0%)	0 (0%)	Jane (your guess was incorrect)	20% (1	out of S) 0 pc	pints		20 p	pints	
6	0 (0%)	0 (0%)	Jane (your guess was incorrect)	20% (1	out of S) 0 pc	pints		20 p	pints	
5	0 (0%)	0 (0%)	Jill (your guess was incorrect)	20% (1	out of 9	а) (р	oints		20 n	oints	
4	0 (0%)	0 (0%)	Jill (your puess was incorrect)	20% (1	out of	0.0	pints		20 m	oints	
-	0 (090)	0.090	El (unor numero una incontect)	2006 (4	out of t	a o-	ninte		20 -	oint-	
•	0 (0%)	0 (0%)	an your guess was incorrect)	2078 (1	out of t	, up	units		20 p	oms	
2	4 (80%)	1 (20%)	Jane (your guess was correct)	20% (1	out of !	oj 10	points		20 p	oints	
1	3 (60%)	2 (40%)	Jane (your guess was correct)	20% (1	out of 9	5) 10	points		10 p	oints	

Round 30 - Results

		er und buge								
In this e	xperiment a	a bet will be	e made 30 times consecutively		Bet #	Jane	л	Bet #	Jane	л
betweer	n two indivi	duals, Jane	and Jill. For each bet, you will the lane or		1	3 (60%)	2 (40%)	16	0 (0%)	00
Jill. In ea	seconas ta sch round, i	f you guess	sed correctly, and if you were		2	4 (80%)	1 (20%)	17	0 (0%)	0.0
ranked a	among the true of	80% of gu	essers (4) whose guesses are the		3	0 (0%)	0 (0%)	18	0 (0%)	00
will earr	10 bonus	points for t	his round. Otherwise, you will not		4	0 (0%)	0 (0%)	19	0 (0%)	0(0
win any see you	points for t r ranking in	this round. the game.	At the side of the screen, you will as well as the distribution of all		5	0 (096)	0 (0%)	20	0 (0%)	0.0
particip	ant guesses	between J	ane and Jill for every preceding		6	0 (096)	0 (0%)	21	0 (0%)	0.0
bet, in b The out	come of ea	rs and perc ch bet will I	entages. be displayed immediately		7	0 (0%)	0 (0%)	22	0 (0%)	00
followin	g the end o	of the count	tdown.		8	0 (0%)	0 (0%)	23	0 (0%)	0(
Bet 30 c	out of 30				9	0 (0%)	0 (0%)	24	0 (0%)	0(
The out	come of the	e bet is Jan	e		10	0 (0%)	0 (0%)	25	0 (0%)	0(
You gue	essed				11	0 (0%)	0 (0%)	26	0 (0%)	0(
You are	ranked 1 or	ut of S			12	0 (0%)	0 (0%)	27	0 (0%)	0(0
_					13	0 (0%)	0 (0%)	28	0 (0%)	0(
Thereto	re, you did	not earn 10	0 points in this round.		14	0 (0%)	0 (0%)	29	0 (0%)	0(0
Press th	ie spaceba	r to contin	ue.		15	0 (0%)	0 (0%)	30	0 (0%)	0 (0
Bet #	Jane	Jill	Outcome	Ranking		Poir	its won in	round	Total	poi
30	0 (0%)	0 (0%)	Jane (your guess was incorrect)	20% (1	out of S	i) 0 po	pints		20 p	oints
29	0 (0%)	0 (0%)	Jill (your guess was incorrect)	20% (1	out of S	6) 0 pc	pints		20 p	oints
28	0 (0%)	0 (0%)	Jane (your guess was incorrect)	20% (1)	out of S	6) 0 pc	oints		20 p	oints
27	0 (0%)	0 (0%)	Jill (your guess was incorrect)	20% (1	out of S	6) 0 po	pints		20 p	oints
26	0 (0%)	0 (0%)	Jane (your guess was incorrect)	20% (1)	out of S	6) 0 po	pints		20 p	oints
25	0 (0%)	0 (0%)	Jane (your guess was incorrect)	20% (1)	out of S	6) 0 pc	pints		20 p	oints
24	0 (0%)	0 (0%)	Jill (your guess was incorrect)	20% (1	out of S	6) 0 po	pints		20 p	oints
23	0 (0%)	0 (0%)	Jane (your guess was incorrect)	20% (1)	out of S	i) 0 po	pints		20 p	oints
22	0 (0%)	0 (0%)	Jill (your guess was incorrect)	20% (1)	out of S	6) 0 po	pints		20 p	oints
21	0 (0%)	0 (0%)	Jane (your guess was incorrect)	20% (1)	out of S	6) 0 po	pints		20 points	
20	0 (0%)	0 (0%)	Jane (your guess was incorrect)	20% (1)	out of S	6) 0 po	pints		20 p	oints
19	0 (0%)	0 (0%)	Jane (your guess was incorrect)	20% (1)	out of S	6) 0 pe	pints		20 p	oints
18	0 (0%)	0 (0%)	Jill (your guess was incorrect)	20% (1)	out of S	6) 0 po	pints		20 p	oints
17	0 (0%)	0 (0%)	Jill (your guess was incorrect)	20% (1)	out of S	i) 0 po	pints		20 p	oints
16	0 (0%)	0 (0%)	Jill (your guess was incorrect)	20% (1	out of S	6) 0 pc	pints		20 p	oints
15	0 (0%)	0 (0%)	Jill (your guess was incorrect)	20% (1)	out of S	i) 0 po	pints		20 p	oints
14	0 (0%)	0 (0%)	Jane (your guess was incorrect)	20% (1	out of S	6) 0 pc	pints		20 p	oints
13	0 (0%)	0 (0%)	Jane (your guess was incorrect)	20% (1	out of S	i) 0 po	pints		20 p	oints
12	0 (0%)	0 (0%)	Jane (your guess was incorrect)	20% (1)	out of S	6) 0 pc	pints		20 p	oints
11	0 (0%)	0 (0%)	Jane (your guess was incorrect)	20% (1	out of S	6) 0 pc	ints		20 p	oints
10	0 (0%)	0 (0%)	Jill (your guess was incorrect)	20% (1)	out of S	i) 0 po	pints		20 p	oints
9	0 (0%)	0 (0%)	Jane (your guess was incorrect)	20% (1)	out of S	6) 0 pc	ints		20 p	oints
8	0 (0%)	0 (0%)	Jill (your guess was incorrect)	20% (1)	out of S	i) 0 pc	pints		20 p	oints
7	0 (0%)	0 (0%)	Jane (your guess was incorrect)	20% (1	out of	5) 0 p	pints		20 p	oints
6	0 (0%)	0 (0%)	Jane (your guess was incorrect)	20% (1	out of	5) 0 p	pints		20 p	oints
5	0 (0%)	0 (0%)	Jill (your guess was incorrect)	20% (1	out of	5) 0 p	pints		20 p	oints
4	0 (0%)	0 (0%)	Jill (your guess was incorrect)	20% (1	out of	5) 0 p	pints		20 p	oints
3	0 (0%)	0 (0%)	Jill (your guess was incorrect)	20% (1	out of	5) 0 p	pints		20 p	oints
	4 (80%)	1 (20%)	Jane (your quess was correct)	20% (1	out of !	5) 10	points		20 p	oints
2	4 (00/0)		9 9 V							

Final Results

Your final ranking is 1 out of 5 . You earned points in 2 outo of 30. In total, you earned 20 points.							
Please press the space bar to complete the experiment.							
Bet #	Jane	Jill	กษณา	דינג	בנוס	onic	
30	0 (0%)	0 (0%)	Jane (your guess was incorrect)	20% (1 out of 5)	0 points	20 point	
29	0 (0%)	0 (0%)	Jill (your guess was incorrect)	20% (1 out of 5)	0 points	20 point	
28	0 (0%)	0 (0%)	Jane (your guess was incorrect)	20% (1 out of 5)	0 points	20 point	
27	0 (0%)	0 (0%)	Jill (your guess was incorrect)	20% (1 out of 5)	0 points	20 poin	
26	0 (0%)	0 (0%)	Jane (your guess was incorrect)	20% (1 out of 5)	0 points	20 poin	
25	0 (0%)	0 (0%)	Jane (your guess was incorrect)	20% (1 out of 5)	0 points	20 poin	
24	0 (0%)	0 (0%)	Jill (your guess was incorrect)	20% (1 out of 5)	0 points	20 poin	
23	0 (0%)	0 (0%)	Jane (your guess was incorrect)	20% (1 out of 5)	0 points	20 poir	
22	0 (0%)	0 (0%)	Jill (your guess was incorrect)	20% (1 out of 5)	0 points	20 poir	
21	0 (0%)	0 (0%)	Jane (your guess was incorrect)	20% (1 out of 5)	0 points	20 poir	
20	0 (0%)	0 (0%)	Jane (your guess was incorrect)	20% (1 out of 5)	0 points	20 poir	
19	0 (0%)	0 (0%)	Jane (your guess was incorrect)	20% (1 out of 5)	0 points	20 poir	
18	0 (0%)	0 (0%)	Jill (your guess was incorrect)	20% (1 out of 5)	0 points	20 poir	
17	0 (0%)	0 (0%)	Jill (your guess was incorrect)	20% (1 out of 5)	0 points	20 poir	
16	0 (0%)	0 (0%)	Jill (your guess was incorrect)	20% (1 out of 5)	0 points	20 poir	
15	0 (0%)	0 (0%)	Jill (your guess was incorrect)	20% (1 out of 5)	0 points	20 poir	
14	0 (0%)	0 (0%)	Jane (your guess was incorrect)	20% (1 out of 5)	0 points	20 poir	
13	0 (0%)	0 (0%)	Jane (your guess was incorrect)	20% (1 out of 5)	0 points	20 poir	
12	0 (0%)	0 (0%)	Jane (your guess was incorrect)	20% (1 out of 5)	0 points	20 poir	
11	0 (0%)	0 (0%)	Jane (your guess was incorrect)	20% (1 out of 5)	0 points	20 poir	
10	0 (0%)	0 (0%)	Jill (your guess was incorrect)	20% (1 out of 5)	0 points	20 poir	
9	0 (0%)	0 (0%)	Jane (your guess was incorrect)	20% (1 out of 5)	0 points	20 poir	
8	0 (0%)	0 (0%)	Jill (your guess was incorrect)	20% (1 out of 5)	0 points	20 poir	
7	0 (0%)	0 (0%)	Jane (your guess was incorrect)	20% (1 out of 5)	0 points	20 poir	
6	0 (0%)	0 (0%)	Jane (your guess was incorrect)	20% (1 out of 5)	0 points	20 poir	
5	0 (0%)	0 (0%)	Jill (your guess was incorrect)	20% (1 out of 5)	0 points	20 poin	
4	0 (0%)	0 (0%)	Jill (your guess was incorrect)	20% (1 out of 5)	0 points	20 poin	
3	0 (0%)	0 (0%)	Jill (your guess was incorrect)	20% (1 out of 5)	0 points	20 poin	
2	4 (80%)	1 (20%)	Jane (your guess was correct)	20% (1 out of 5)	10 points	20 poin	
1	3 (60%)	2 (40%)	Jane (your guess was correct)	20% (1 out of 5)	10 points	10 poin	

Completion Screen

Thank you for your participation.

A lottery will be conducted in the coming days to determine whether you are entitled to bonus payment. If you win, you will be notified and paid through Prolific. Please click the following link to complete the study and return to Prolific: https://app.prolific.com/submissions/complete?cc=COBWOW66

Appendix B – "Coin Experiment"

1. Experiment

1.1. Experimental setting

This was, in fact, the original experiment which inspired the rest of this paper. It featured a shorter game (ten rounds only), and feedback only shown to participants at the conclusion of the game, considerably limiting both the herding effect and our ability to study the learning process. This experiment, however, was conducted in a more conventional laboratory setting, and as such provides additional validation for our paper. Moreover, the lack of feedback throughout the game, while significantly reducing both the herding and learning effects, is also likely to greatly reduce the presence of other emotional and cognitive biases that might affect our main results. Therefore, we have chosen to include the results in this appendix.

In this variant of the experiment, participants guessed the outcome of a coin toss, repeated across ten rounds. The experiment comprised two parts: Each participant engaged in two games of ten rounds. In the first game (inclusive environment), the top 90% of guessers received a reward, whereas in the second game (exclusive environment), only the top 10% of guessers received a reward. Participants were randomly assigned to either Treatment B1 (starting with the inclusive environment and then moving to the exclusive environment) or Treatment B2 (starting with the exclusive environment and then moving to the inclusive environment). All participants in a group received the same treatment. The within-participants component allowed us to expand the scope of our study and examine the effects of the transition from one environment type to another.

After each round, we displayed the previous rounds' guess distributions to the group, both in absolute numbers and percentages. However, the outcomes of the coin tosses and the final rewards were disclosed only after both games had been completed.

The experiment was conducted on Zoom, with all participants of a group present in the same session. Each participant received a unique link to join the session and was required to keep their camera active throughout. Supervisors were present to explain instructions, address questions, and ensure adherence to the protocol. Upon completion, payments were processed through Bit, a popular digital payment app in Israel.

1.2. Participants

A total of 208 participants were recruited from among the student body of the Hebrew University of Jerusalem. Participants were divided into thirteen groups of sixteen participants each, with seven groups assigned to Treatment B1 and six groups assigned to Treatment B2 (see Table B1). In terms of the structure of the experiment, our design is similar to those used by Amnon Rapoport and his associates for entry game experiments (e.g., Rapoport, Seale, & Winter, 2000), but uses includes more groups and a larger number of subjects. The participant pool was balanced in terms of gender, and all participants were required to have a proficient command of the Hebrew language. Session scheduling ensured consistency in terms of weekdays and hours to minimize potential external influences. Multiple supervisors were present throughout the sessions to provide instructions, address queries, and ensure compliance with the experiment's guidelines.

Treatment	First game	Groups	Participants
B1	Inclusive	7	112
B2	Exclusive	6	96
Total		13	208

Table B1: Between-participant group distribution

Notes: The table presents the distribution of participants into the B1 and B2 treatments. "Groups" indicates the number of groups (of sixteen participants each) assigned to the treatment, while "Participants" indicates the total number of participants assigned to each treatment.

A range of personal, demographic, and behavioral characteristics variables were collected through a personal information form filled out by participants at the end of the experiment. There were no statistical differences in these variables between the treatment groups.

2. Results

2.1. Clustering

Figure B1 shows the majority size for the two treatment groups. Majority size is here defined as the portion of participants in a given group and round who selected the majority choice. The range for this distribution is between 0.5 and 1.0 choice (i.e., an even split between "heads" and "tails" would yield a majority size of 0.5, while an all

"tails" or all "heads" group would both yield a majority size of 1.0). Panel 1 shows the distribution for all rounds. Panel 2 shows only the distributions for rounds in the first game played by each group (i.e., only the inclusive environment for groups assigned to Treatment B1, and only the exclusive environment for groups assigned to Treatment B2).

As expected, the majority size distribution of the inclusive environment clusters closer to 1.0, while the majority size distribution of the exclusive environment clusters closer to 0.5. Notably, however, the difference between the inclusive and exclusive environments appears to be stronger in the first game played.

A notable observation across both treatments, particularly in the inclusive environment, is a skewness towards "heads" in the distribution of choices. This skewness suggests a general preference among participants for selecting "heads" over "tails," a pattern that persists even in the first round of each game. This indicates that participants' initial choices are not entirely random, even in the absence of prior information about others' decisions.



Figure B1: Majority Size

Panel 1: All Rounds

Note: Panel 1 shows the results for the inclusive environment (on the left, in blue) and the exclusive environment (on the right, in red). It contains 130 observations for the inclusive environment (thirteen groups, ten rounds each), and 128 observations for the exclusive environment (thirteen groups, ten rounds each, with the exception of one group that only played eight rounds due to technical difficulties).





Note: Panel 2 shows the results for the inclusive environment (on the left, in blue) and the exclusive environment (on the right, in red). It contains seventy observations for the inclusive environment (seven groups, ten rounds each), and 58 observations for the exclusive environment (six groups, ten rounds each, with the exception of one group that only played eight rounds due to technical difficulties).

General notes for both panels: These histograms present the frequency of majority sizes for all participants in a given round. The x axis represents the choice average (0.5 represents an even split between "heads" and "tails," while 1.0 represents a majority where all participants selected "heads" or all participants selected "tails"), grouped into brackets of 0.05. The y axis represents frequency. Blue represents the distribution for the inclusive environment, while red represents the distribution for the exclusive environment.

A T test¹⁷ comparing the majority size in each round between the two treatment groups. shows larger majority sizes on average in the inclusive environment (65.68% compared to 59.93% in the exclusive environment. This difference is significant at the 1% level. Table B2 presents the marginal effect coefficients for clustering (i.e., whether a participant's guess in a given round matches the majority of the previous round). The dependent variable is an indicator variable which equals 1 if the participant's decision in a given round is identical to the majority choice in the previous round. Columns (1), (2), and (3) show the results for rounds 2-10 of all games played (two games per participant). Column (2) controls for session fixed effects (thirteen sessions)

¹⁷ A skewness and kurtosis test conducted on the distribution of majority sizes indicates a deviation from normal distribution at a 95% confidence level. Consequently, to accommodate the non-normality of the data, we opted for a T test instead of a Z test for our analysis. The T test is more appropriate in this context as it is less sensitive to deviations from normal distribution, especially in smaller sample sizes.

and participant fixed effects (sixteen participants per session, or 208 participants in total).

To control for any possible problems with the within-participants setup, Columns (4) and (5) use a purely between-participants setup, examining only the first game played by each participant (i.e., only the inclusive environment for B1 participants, and only the exclusive environment for B2 participants) and only the second game played by each participant (i.e., only the exclusive environment for B1 participants, and only the inclusive environment for B1 participants, and only the inclusive environment for B2 participants) respectively. Consistent with our predictions, the results show that overall, participants in the exclusive environment were 14-17% more likely to stay with the majority, with this finding being statistically significant at the 1% level. Both the coefficient and the explanatory power are higher for the smaller, first-game-only sample, where it rises to 33%. For the second-game-only sample, the coefficient is similarly positive, but not statistically significant and with a much lower explanatory power.

Columns (6) and (7) are intended to examine whether participants display a learning process throughout the course of the game. Hence, Column (6) tests the effect of the environment only for the first half of the game (rounds 2-5), and Column (7) tests for only the second half of the game (rounds 6-10). While both specifications show a positive clustering coefficient, in earlier rounds this coefficient is only 8%, and not statistically significant. In later rounds the clustering coefficient rises to 20% at a 1% significance level.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Rounds	2-10	2-10	2-10	2-10	2-10	2-5	6-10
Games	1&2	1&2	1&2	1 only	2 only	1&2	1&2
Inclusive dummy	0.1470*** (0.0460)	0.1437*** (0.0461)	0.1701*** (0.0482)	0.3311** (0.1647)	0.1935 (0.1608)	0.0852 (0.0698)	0.2058*** (0.0623)
Session fixed effects	No	Yes	Yes	Yes	Yes	Yes	Yes

Table B2: Clustering (probit)

Participant fixed effects	No	No	Yes	No	No	No	No
Number of observations	3,122	3,122	3,090	1,595	1,527	1,392	1,730
Pseudo R ²	0.0025	0.0048	0.0894	0.0139	0.0052	0.0057	0.0113

Notes: The table presents probit marginal effect coefficients for the treatment effect on whether participants' choice belonged to the majority in that round. The dependent variable is an indicator variable for whether a participant's choice belongs to the majority in that round. Column (1) shows the basic model, Column (2) shows the model with session fixed effects, and Column (3) shows the model with both session and participant fixed effects. Column (4) shows the coefficients only for the first game of each participant (only inclusive for Treatment B1 participants, only exclusive for Treatment B2 participants), while Column (5) shows the coefficient only for the second game of each participant (only exclusive for Treatment B1 participants, only inclusive for Treatment B2 participants). Column (6) shows the results only for decisions in rounds 2-5, while Column (7) shows the results only for decisions in rounds 6-10. The inclusive dummy equals 1 if the environment played is the inclusive environment, and 0 if it is the exclusive environment. Standard errors appear in parentheses. ** and *** denote significance at the 5% and 1% levels respectively.

2.2. Fluctuation

Figure B2 shows the fluctuation of decisions for each treatment. For each round of a game, the share of participants who modified their decision compared to the previous round (i.e., selected "heads" in round t - 1 and "tails" in round t, or vice versa)¹⁸. Panel 1 shows the distribution for all rounds. Panel 2 shows only the distributions for rounds in the first game played by each group (i.e., only the inclusive environment for groups assigned to Treatment B1, and only the exclusive environment for groups assigned to Treatment B2).

As per our prediction, it is evident that participants tend to change their decisions from one round to the next more frequently in the exclusive environment than in the inclusive environment. As with the majority size, however, this difference appears to be particularly strong in the first game played by participants, and is diminished in the second game.

¹⁸ Data for the first round of each game are not included in our analysis, as no prior round data exist for reference. Therefore, our analysis encompasses rounds 2-10 of each game only.

Figure B2: Changed Decisions





Note: Panel 1 shows the results for the inclusive environment (on the left, in blue) and the exclusive environment (on the right, in red). It contains 117 observations for the inclusive environment (thirteen groups, rounds 2-10), and 115 observations for the exclusive environment (thirteen groups, rounds 2-10, with the exception of one group that only played eight rounds due to technical difficulties).



Panel 2: First Game Played Only

Note: Panel 2 shows the results for the inclusive environment (on the left, in blue) and the exclusive environment (on the right, in red). It contains 63 observations for the inclusive environment (seven groups, rounds 2-10), and 52 observations for the exclusive environment (six groups, 2-10, with the exception of one group that only played eight rounds due to technical difficulties).

General notes for both panels: These histograms present the frequency of change rates for all participants in a given round. The x axis represents the share of participants in a group who changed their decision compared to the previous round, grouped into brackets of 0.05. The y axis represents frequency. Blue represents the distribution for the inclusive environment, while red represents the distribution for the exclusive environment.

A T test¹⁹ comparing the share of participants in each round who changed their decision compared to their decision in the preceding round (either from "heads" to "tails" or from "tails" to "heads") between the two treatments revealed results consistent with our predictions. On average, 34.51% of players in the inclusive environment change their decision from one round to the next, in the exclusive environment, this number goes up to 40.03%, with this difference being statistically significant at the 1% level. Table B3 presents the marginal effect coefficients for fluctuation (i.e., participants' decision in a given round to change their guess compared to the preceding round, from "heads" in the preceding round to "tails" in the current round, or vice versa). The dependent variable is an indicator variable which equals 1 if the participant's decision changed in a given round compared to the previous round, and 0 if the decision did not change. Columns (1), (2), and (3) show the results for rounds 2-10 of all games played (two games per participant). Column (2) controls for session fixed effects (thirteen sessions of two games each), while Column (3) controls for both session fixed effects (thirteen sessions) and participant fixed effects (sixteen participants per session, or 208 participants in total). To control for any possible problems with the within-participants setup, Columns (4) and (5) use a purely between-participants setup, examining only the first game played by each participant (i.e., only the inclusive environment for B1 participants, and only the exclusive environment for B2 participants) and only the second game played by each participant (i.e., only the exclusive environment for B1 participants, and only the inclusive environment for B2 participants). Consistent with our predictions, the results show that participants in the exclusive environment were 14-17% more likely to change their decision from one round to the next, with this finding being statistically significant at the 1% level. Both the coefficient and the explanatory power are higher for the first-game-only sample, where participants were 45% more likely to fluctuate between choices. For the second-game-only sample, the coefficient is positive but not statistically significant, and with a much lower explanatory power.

Columns (6) and (7) examine whether participants display a learning process throughout the course of the game. Column (6) tests the effect of the environment only for the first half of the game (rounds 2-5), and Column (7) tests for only the second half

¹⁹ Once again, a skewness and kurtosis test reveal that the distribution of the number of changed decisions is not normally distributed, with a confidence level of 95%. Hence, here too a T test is employed.

of the game (rounds 6-10). Unlike the majority choice shown in Table (4), in this case the coefficients differ only marginally (12% and 13% respectively). The coefficient for the later rounds indicates only a marginally stronger effect, with the difference between the two specifications not being statistically significant.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Rounds	2-10	2-10	2-10	2-10	2-10	2-5	6-10
Games	1&2	1&2	1&2	1 only	2 only	1&2	1&2
Inclusive	-0.1351***	-0.1357***	-0.1781***	-0.4583***	0.1752	-0.1224**	-0.1328**
dummy	(0.0424)	(0.0427)	(0.04892)	(0.1581)	(0.1468)	(0.0591)	(0.0574)
Session fixed effects	No	Yes	Yes	Yes	Yes	Yes	Yes
Participant fixed effects	No	No	Yes	No	No	No	No
Number of observations	3,652	3,652	3,044	2,009	2,059	2,051	2,017
Pseudo R ²	0.0021	0.0129	0.1402	0.0210	0.0001	0.0108	0.0166

Table B3: Fluctuation (probit)

Notes: The table presents probit marginal effect coefficients for the treatment effect on participants' choice to change their decision compared to the preceding round. The dependent variable is an indicator variable for choosing to change one's decision relative to the previous round. Column (1) shows the basic model, Column (2) shows the model with session fixed effects, and Column (3) shows the model with both session and participant fixed effects. Column (4) shows the coefficients only for the first game of each participant (only inclusive for Treatment B1 participants, only exclusive for Treatment B2 participants), while Column (5) shows the coefficient only for the second game of each participant (only exclusive for Treatment B1 participants, only inclusive for Treatment B2 participants). Column (6) shows the results only for decisions in rounds 2-5, while Column (7) shows the results only for decisions in rounds 6-10. The inclusive dummy equals 1 if the environment played is the inclusive environment, and 0 if it is the exclusive environment. Standard errors appear in parentheses. ** and *** denote significance at the 5% and 1% levels respectively.

Appendix C – Equilibrium Analysis of the n-person Game

To begin our analysis, we devise a theoretical model to explain how inclusive environments can generate a herding equilibrium, while exclusive environments can generate an equilibrium with considerably more anti-herding behavior. We consider two games with a set *N* of *n* players (where *n* is an odd natural number). Each player must bet on the outcome of a lottery $[x, y; \frac{1}{2}, \frac{1}{2}]$. Let *R* be the set of players who guessed correctly (/R/ = *r* and *n* – *r* guessed incorrectly). Players' payoffs in the two games are determined as follows:

For each $1 \le k \le n$ we define a game G_k as follows:

Exactly *k* players get the prize *M* according to the following criterion: if $r \ge k$, then k players are selected randomly (with equal probabilities) among those who guessed correctly, and the rest receive zero. If r < k all players in *R* receive the prize as well as additional randomly selected k - r players who made a wrong guess. More simply, in the game G_k , the *r* players who guessed correctly ranked above the n - r who guessed incorrectly. Within each of these two groups, players are ranked randomly. Only the first *k* players in the order receive the prize *M*. In the proposition below we define strategy profiles to be of identical type if one can be obtained by renaming players or strategies.

Proposition:

(a) For any $k \ge \frac{n+1}{2}$ there is a unique type of pure strategy Nash equilibrium of the game G_k , which involves all players betting on the same outcome.

(b) For any $k < \frac{n+1}{2}$ there is a unique type of pure strategy Nash equilibrium, in which exactly $\frac{n+1}{2}$ players bet on one outcome, and $\frac{n-1}{2}$ players bet on the other outcome.

(c) In addition to the pure equilibria described in (a) and (b), and for every $1 \le k \le n$, the game G_k also has a symmetric mixed equilibrium in which all players assign equal probability to the two bets.

As demonstrated in the introduction, according to the proposition when the environment is inclusive i.e., case (a) players will coordinate in equilibrium to make the same bet. In contrast in an exclusive environment i.e., case (b) they will diverge and divide themselves between the two bets with virtually equal number of players for each bet. Indeed, choosing x and y with equal probability is also an equilibrium in the inclusive environment, but as we shall show in the proof for the proposition this equilibrium is unstable.

Proof of proposition:

We first prove (a). We start by showing that if $k \ge \frac{n+1}{2}$, there exists a Nash equilibrium where all players make the same guess. Without loss of generality, assume that all players bet on x. Then each player receives the prize with probability $\frac{k}{n}$. This is true regardless of whether the bet is correct or not. A deviating player will receive the prize with probability 1 if that player's guess is correct (a correct guess occurs with probability $\frac{1}{2}$). Since $\frac{k}{n} > \frac{1}{2}$, the deviating player is worse off.

Consider now any other profile, and assume that the majority of players bet on $x: m \ge \frac{n+1}{2}$ are betting x and n - m are betting y. We will show that any player betting y would be better off deviating and choosing x instead. If y is correct, a player in the minority will earn the prize with probability 1. If wrong, she will earn the prize with positive probability only if k > m.

We assume first that k > m. In this case, she will earn the prize with probability $\frac{k-m}{n-m}$. Hence, the overall probability of earning the prize is $\frac{1}{2}(1 + \frac{k-m}{n-m})$ if k > m. Consider now a minority player who deviates from the majority. If k > m, then, if x is correct, she will earn the prize with probability 1, and if x is wrong, she will receive the prize with probability $\frac{k-(n-(m+1))}{m+1}$. Note that (n - (m + 1)) is the number of players in the minority, all of whom are correct and therefore receive the prize. Hence the overall probability of receiving the prize is $\frac{1}{2}(1 + \frac{k-(n-(m+1))}{m+1})$. Since 2m > n and k > m it follows that $\frac{k-(n-(m+1))}{m+1} > \frac{k-m}{n-m}$. Hence, if k > m, then every player of the minority would be better off deviating to the majority.

We next show that deviation is profitable also when $k \le m$. Indeed, being in the minority, a player gets the prize with probability 1 if correct. If a minority player is wrong, he gets the prize with probability zero. Hence, the probability of receiving the prize is $\frac{1}{2}$. By deviating to the majority, he may get the prize whether correct or not. If correct, he will get it with probability $\frac{k}{m+1}$, and if wrong, he will get it with probability

$$\frac{k-(n-(m+1))}{m+1}$$
. Hence, the overall probability of winning is $\frac{1}{2}\left(\frac{k-(n-(m+1))}{m+1} + \frac{k}{m+1}\right) = \frac{1}{2}\left(\frac{2k-n+m+1}{m+1}\right) > \frac{1}{2}$.

We now proceed to prove (b). Since n is an odd number, any profile of strategies generates both a majority and a minority. Let *m* be the size of the majority. We shall show that unless $m = \frac{n+1}{2}$, a majority player would be better off deviating to the minority. We start by assuming that k > n - m. Note that this is only possible if $m > \frac{n+1}{2}$. If correct, a majority player gets the prize with probability $\frac{k}{m}$. If wrong, she gets it with probability $\frac{k-n+m}{m}$. So, the overall probability of getting the prize is $\frac{1}{2}\left(\frac{2k-n+m}{m}\right)$. If a majority player deviates to the minority, then, if correct, she will get the prize with certainty. If incorrect, she will get it with probability zero. Hence the overall probability is $\frac{1}{2}$. Since $\frac{1}{2}\left(\frac{2k-n+m}{m}\right) < \frac{1}{2}$, and so a player will be made better off by deviating from the majority to the minority.

Finally, we consider the case in which $k \le n - m$. In this case, if a majority player's guess is correct, that player gets the prize with probability $\frac{k}{m}$. If incorrect, he gets the prize with probability zero, so the overall probability of winning the prize is $\frac{1}{2} \left(\frac{k}{m}\right)$. If he deviates to the minority, then, if correct, he gets the prize with probability $-\frac{k}{n-m+1}$. If incorrect, he gets the prize with probability zero. If $m > \frac{n+1}{2}$, then n - m + 1 < m, and deviating to the minority yields a higher expected payoff. If, however, $m = \frac{n+1}{2}$, then n - m + 1 = m, and the two options are identical. Hence, the unique equilibrium must involve the smallest majority size possible, $m = \frac{n+1}{2}$.

To show (c), note that if all players bet on x with probability p and y with probability (1-p), then for $p = \frac{1}{2}$, their strategies will form a Nash equilibrium. Given that all other players also choose $p = \frac{1}{2}$, each player will be indifferent about which bet to make. Meanwhile, if $p \neq \frac{1}{2}$, players have incentive to deviate. In the inclusive environment, $(k \ge \frac{n+1}{2})$ players will choose the bet that maximizes the probability of being part of the majority. In the exclusive environment $(k < \frac{n+1}{2})$ players will choose the bet that maximizes the probability of being part of the minority. Note that the symmetric equilibrium is unstable in the inclusive environment. If even a single player shifts his

or her probability from $\frac{1}{2}$ for both *x* and *y* to, for example, $\frac{1}{2} + \varepsilon$ for *x* (with an arbitrarily small ε), then the best response of all other players would be to switch to choosing *x* with probability 1.