Breakdowns in the Lab

Breakdowns in the Lab*

Johannes C. Hoelzemann† Nicolas Klein‡

June 25, 2025

Abstract

We experimentally investigate a game of strategic experimentation in which information arrives through fully revealing, publicly observable, *breakdowns*. As predicted by theory, we find that players experiment significantly less, and payoffs are lower, when actions are hidden. We view this as evidence that behavior is systematically affected by the informational environment and consistent with strategic free-riding.

JEL Classification: C73, C92, D83, O32

Keywords: Dynamic Public-Good Problem, Strategic Experimentation, Exponential Bandits, Learning, Dynamic Games, Laboratory Experiments

*A previous version of this paper was circulated under the title "Information and the Bandit: The Good, the bad and the Ugly." Financial support by the Université Paris-Panthéon-Assas and the University of Vienna is gratefully acknowledged. The experimental part of this study was approved by the University of Vienna's VCEE (Vienna Center for Experimental Economics). We thank Ben Greiner, Stanton Hudja, Gustavo Manso, Abhishek Nagaraj, Sven Rady, and Karl Schlag for helpful comments. Corresponding author: Nicolas Klein.

†Universität Wien: www.johanneshoelzemann.com & jc.hoelzemann@gmail.com.

‡Université de Montréal: www.nicolasklein.com & kleinnic@yahoo.com.

1 Introduction

Gradual learning about the world is at the center of many economic decisions. For instance, it is well understood that learning is an important aspect in investment decisions in private assets (see, e.g., Sorensen 2008). The trade-off between exploitation vs. exploration, which, amongst others, is at the heart of venture-capital investment decisions, is canonically analyzed in the framework of the so-called bandit model (see e.g., Bergemann and Välimäki 2008). Bergemann and Hege (2005) analyzes the impact of feedback on incentives for an entrepreneur who receives financing from a venture capitalist and Krieger (2021) studies project continuation decisions where firms may resolve uncertainty through news about their own results and competitors' R&D failures.[1] In this paper, we provide experimental evidence showing that, in a strategic setting, transparency encourages investment. Rather than the relationship between investor and entrepreneur, this paper is concerned with the informational externalities between different private-asset investors. To isolate the informational dimension, we examine a setting with purely informational externalities, deviating from the contest framework employed by, e.g., Mihm and Schlapp (2019) and Halac, Kartik, and Liu (2017). In line with theoretical predictions, we find that private information hurts effort provision with an uncertain technology if news arrives in the form of lumpy breakdowns.

Games of pure informational externalities have received a lot of attention in the literature (see, e.g., Bolton and Harris 1999, Keller, Rady, and Cripps 2005 or Hörner, Klein, and Rady 2022).[2] In these games, the information produced by a given player benefits other players as well—information is a public good, and players tend to produce inefficiently little of it in equilibrium. Following Keller, Rady, and Cripps (2005), most papers in this literature have focused on *breakthrough*, so-called good-news, environments, where discontinuous events bring good news; the absence of news consequently leads to a continuous deterioration in beliefs. In many real-world applications, however, discontinuous news events take the form of *breakdowns*, i.e., bad news; after all, a staggering 90% of all start-ups fail.[3] Such failure may well come about as

[1] Bergemann and Hege (1998), Manso (2011) and Klein (2016) also analyze incentive provision in a principal-agent setting; Ederer and Manso (2013) provide experimental evidence. Levinthal (1997), March (1991), Azoulay, Graff Zivin, and Manso (2011) study innovation search and examine what drives risky exploration among innovators.

[2] Wolitzky (2018) analyzes the trade-off between a safe and a risky option in the context of a continuum of short-lived players making decisions sequentially. For an application to political economics, see Callander and Harstad (2015) who analyze policy experimentation in federal systems.

[3] For up-to-date statistics, see, e.g., Kotashev (2024) and Howarth (2025).

a result of a sudden, singular, event, such as the discovery of severe side effects stemming from a medical drug, or the catastrophic malfunctioning of some technology. Theoretically, it is well understood (see, e.g., Keller and Rady 2015, or Wagner and Klein 2022) that the mechanisms underlying the bad-news strategic-learning models differ sharply from those under good news. While Hoelzemann and Klein (2021) has experimentally investigated strategic experimentation under good news, and Hoelzemann, Manso, Nagaraj, and Tranchero (2024) investigates the role of players' information in a strategic setting with good news, we are, to the best of our knowledge, the first to experimentally investigate a bad-news strategic-experimentation setting. As predicted, behavior contrasts sharply with that in breakthrough environments, we believe our investigation to be filling an important gap in the literature.

The scant attention given to settings with breakdowns is surprising because of their economic importance beyond investment choices and the financial world: Bad-news learning processes naturally occur upon the introduction of a new technology that holds out hopes of cost savings but entails risks. Such risky technologies include new drugs and medical devices, and innovative processes such as hydraulic fracturing for oil production. Some technologies that are socially undesirable, perhaps because they impose negative externalities on other sectors, also fit in this broad class. Consider financial fraud or tax evasion when agents have incomplete information about the effectiveness of the detection technology. In all these cases, there also exist significant barriers to the flow of information, making unobservable actions a good starting point for the analysis. For example, the decision to evade taxes is private, but getting caught is typically a public event.

In this paper, we are investigating in particular the role of the observability of actions in a bad-news game of strategic experimentation with bandits. These are games of purely informational externalities, where players have an incentive to free-ride on the information produced by the other players. In a continuous-time, infinite-horizon, setting, it is theoretically known that, in a conclusive bad-news model, private information tends to be bad for welfare (Bonatti and Hörner 2017). This is because, in the absence of conclusive news, observing a player's shirking in information production makes the other player(s) more pessimistic than they would be on the equilibrium path if the conclusive bad news fails to materialize. Therefore, with conclusive bad news, players will be less prone to slack off in information production if their actions are observable, because, after observing a deviation, the other player(s) will be warier about the risky option than they would be absent a deviation. Because

the only externality in the game is the positive informational externality between players, leading to a tendency toward under-production of information in equilibrium, we should expect that making deviations unobservable ought to dampen welfare in a conclusive bad-news environment.

The main goal of this investigation is to test whether this qualitative prediction of the theory is borne out by actual behavior in a controlled laboratory environment. Empirically, we indeed find that both experimentation and payoffs are higher with observable actions, as predicted by theory. We have constructed our game in a particularly stark way so that it has the feature that the efficient solution is an equilibrium *if and only if* actions are observable. Further, participants use the risky option more frequently over time, reflecting growing optimism consistent with Bayesian updating.

In summary, the paper makes two main contributions. First, we present evidence that, in a *bad-news* setting, behavior is systematically affected by the informational environment. We find that both experimentation and payoffs are *higher* with observable actions. Second, behavior is consistent with strategic free-riding, as information is a public good and participants produce inefficiently little of it. Participants experiment, on average, too little even when the efficient solution is an equilibrium.

2 The Environment

We have endeavored to come up with the simplest possible environment in which theory would predict lower welfare with unobservable actions. For the effect to arise, we need at least three periods. This is because, in the last period, a player does not care what their opponent will do, as they have no future use for the information learned in this period. So, only in the first period do players want to alter their opponent's future behavior for strategic considerations. We therefore construct a three-period, two-player, simultaneous-move game, calibrating the parameters in such a way that the game features the strategic effects we are interested in. The efficient solution has both players using the risky option in all periods (absent a breakdown); the unique equilibrium with unobservable actions has both players never using the risky option, while either always or never playing risky are the two equilibria with observable actions.

Specifically, there are two risk-neutral players, and the game is played over three periods $t = 1, 2, 3$. In each period, players make a simultaneous choice. At the end of the period, outcomes are revealed; we vary whether a player's choice is observable. If

the safe arm is used, the payoff will be 0 for certain in that period. Using the risky arm entails a *benefit* of $s = 2857$. The risky arm is either *good* or *bad*, its type remaining constant over the three periods of the game. If it is good, its use never imposes a cost. If it is bad, it leads to a breakdown, imposing a cost of 20000, with a probability of $\lambda = 0.25$ in any period it is used. Conditionally on the risky arm's type, the draws are i.i.d. between players and across periods; there are thus no payoff externalities between the players, as only the player whose arm incurs the breakdown bears its cost. Players do not initially know if the risky arm is good or bad; they know that Nature (or the computer) makes the risky arm bad with a probability of $p_0 = 0.676$. After a breakdown is observed, the risky arm is known to be bad with probability 1. In the absence of a breakdown and n successful tries of the risky arm, Bayes' rule implies that an observer knowing this information should hold the belief $p_n = \frac{p_0(1-\lambda)^n}{p_0(1-\lambda)^n+1-p_0}$ that the risky arm is bad; i.e., observing that the risky arm has been used without a breakdown makes players increasingly *optimistic* about the quality of the risky arm. Thus, the updated posterior belief either jumps to 1 in case of a breakdown, or declines with the number of unsuccessful tries n. Arm types are i.i.d. across games. One player's risky arm is good *if and only if* the other one's is as well. In the treatment with *observable* actions, a player observes all of the other player's previous actions as well as the outcomes of these actions. In the treatment with *unobservable* actions, a player observes only if the other player has suffered a breakdown of 20000 from the risky arm or not.

More formally, the solution concept is that of pure-strategy perfect Bayesian equilibrium.[4] Sequential rationality is verified by moving backwards in time. After a breakdown has been publicly observed, the risky arm is known to be bad, so that playing safe is the dominant action. Subsequently, we thus verify sequential rationality conditionally on no breakdown having been observed. We normalize the cost of a breakdown to 1.[5] We write $p_i(t) \equiv p_n$ for player i's Bayesian belief in period $t \in \{1,2,3\}$, if $n = \sum_{z=1}^{t-1}(k_{i,z} + \hat{k}_{-i,z})$, where we write $k_{q,z} = 1$ ($k_{q,z} = 0$) if player $q \in \{i,-i\}$ has used the risky (safe) arm in period z without suffering a breakdown, and $\hat{k}_{-i,z}$ denotes the action that player i thinks that player $-i$ has taken in period z. In the case of observable actions, $\hat{k}_{-i,z} \equiv k_{-i,z}$; in the case of unobservable actions, $\hat{k}_{-i,z}$

[4] Subgame perfection has no bite as an equilibrium refinement, because the game starts with an initial move of Nature, which determines the quality of the risky arm; the game therefore admits of no proper subgames.

[5] For our parameters, this implies that $s = 2857/20000 = 0.14285$. Recall furthermore that $\lambda = 0.25$, and $p_0 = 0.676$.

is pinned down by player i's expectations, which are correct in equilibrium.

In the last period $t = 3$, players face a myopic decision problem, where playing risky is a best response *if and only if* $p(t = 3)\lambda \leq s$. For our parameters, $p_n\lambda < s$ *if and only if* $n \geq 2$.

Indeed, $p_2\lambda < s$ implies that, after a history of two tries without a breakdown, playing risky becomes a dominant action. This pins down play in all equilibrium candidates in which $k_{i,1} + k_{-i,1} = 2$.

Now, let us assume that $k_{i,1} + k_{-i,1} = 1$. Suppose that $k_{-i,2} = 1$ (so that $k_{i,3} + k_{-i,3} = 2$). In this case, $k_{i,2} = 1$ is a best response *if and only if*

$$s - p_1\lambda + \big(1 - p_1 + p_1(1-\lambda)^2\big)s - p_1\lambda(1-\lambda)^2 \geq \big(1 - p_1 + p_1(1-\lambda)\big)s - p_1\lambda(1-\lambda)$$

$$\iff p_1 \leq \frac{s}{\lambda}\frac{1}{1-(1-\lambda)(\lambda-s)},$$

which holds for our parameters. Thus, $k_{i,2} = 1$ is a best response to $k_{-i,2} = 1$.

Next, let us assume that $k_{-i,2} = 0$. If actions are observable, $k_{i,2} = 1$ will induce $k_{i,3} + k_{-i,3} = 2$, and $k_{i,2} = 0$ will induce $k_{i,3} = k_{-i,3} = 0$. Thus, $k_{i,2} = 1$ is a best response *if and only if*

$$s - p_1\lambda + (1 - p_1\lambda)s - p_1\lambda(1-\lambda) \geq 0$$

$$\iff p_1 \leq \frac{s}{\lambda}\frac{2}{2-(\lambda-s)},$$

which is not satisfied for our parameters. Thus, $k_{i,2} = 0$ is the unique best response to $k_{-i,2} = 0$ if actions are observable, inducing $k_{i,3} = k_{-i,3} = 0$. Now suppose that actions are unobservable. We have already shown that $k_{i,2} = 1$ (inducing $k_{i,3} = k_{-i,3} = 1$) cannot happen on the equilibrium path. Now, $k_{i,2} = k_{-i,2} = k_{i,3} = k_{-i,3} = 0$ can be part of an equilibrium *if and only if*

$$0 \geq s - p_1\lambda + (1 - p_1\lambda)s - p_1\lambda(1-\lambda)$$

$$\iff p_1 \geq \frac{s}{\lambda}\frac{2}{2-(\lambda-s)},$$

which is satisfied for our parameters, as we have seen.

Thus, in conclusion, after a history such that $k_{i,1} + k_{-i,1} = 1$, both $k_{i,2} = k_{-i,2} = k_{i,3} = k_{-i,3} = 0$ and $k_{i,2} = k_{-i,2} = k_{i,3} = k_{-i,3} = 1$ are compatible with equilibrium, whether actions are observable or unobservable. It is this non-uniqueness of equilibrium play after histories $k_{i,1} + k_{-i,1} = 1$ that will lead to different first-period equilibrium predic-

tions depending on whether actions are observable or unobservable, as we shall see below.

Let us turn to histories such that $k_{i,1} = k_{-i,1} = 0$. Clearly, $k_{-i,2} = k_{i,3} = k_{-i,3} = 0$, while $k_{i,2} = 1$ is not compatible with equilibrium, because $s - p_0\lambda < 0$ implies that i has an incentive to deviate to $k_{i,2} = 0$. Can $k_{i,2} = k_{-i,2} = k_{i,3} = k_{-i,3} = 1$ occur in equilibrium? If actions are observable (unobservable), a deviation by player i in $t = 2$ leads to a path of play of $k_{i,2} = k_{i,3} = k_{-i,3} = 0$, with $k_{-i,2} = 1$ ($k_{i,2} = k_{i,3} = 0$, with $k_{-i,2} = k_{-i,3} = 1$), giving the deviator i a payoff of 0 in both cases; such a deviation is therefore profitable *if and only if*

$$0 > s - p_0 + \big(1 - p_0 + p_0(1-\lambda)^2\big)s - p_0(1-\lambda)^2\lambda$$

$$\iff p_0 > \frac{s}{\lambda}\frac{2}{1 + s(2-\lambda) + (1-\lambda)^2},$$

which holds for our parameters. The only equilibrium candidate remaining is thus $k_{i,2} = k_{-i,2} = k_{i,3} = k_{-i,3} = 0$; this clearly is compatible with equilibrium as $s - p_0\lambda < 0$.

Let us now move to the first period $t = 1$, and assume that actions are observable. By our previous analysis, there are four equilibrium candidates: (1.) the utilitarian optimum, namely $k_{i,1} = k_{-i,1} = k_{i,2} = k_{-i,2} = k_{i,3} = k_{-i,3} = 1$, (2.) $k_{i,1} = k_{i,2} = k_{-i,2} = k_{i,3} = k_{-i,3} = 1$ and $k_{-i,1} = 0$, (3.) $k_{i,1} = k_{-i,1} = k_{i,2} = k_{-i,2} = k_{i,3} = k_{-i,3} = 0$, and (4.) $k_{-i,1} = k_{i,2} = k_{-i,2} = k_{i,3} = k_{-i,3} = 0$, while $k_{i,1} = 1$. Candidate (4.) can be ruled out right away, as $s - p_0\lambda < 0$, so that player i has an incentive to deviate in the first period, whether actions are observable or unobservable.

Let us turn to candidate (1.). With observable actions, a unilateral deviation by i in the first period can be "punished" with the continuation equilibrium $k_{i,2} = k_{-i,2} = k_{i,3} = k_{-i,3} = 0$, making the deviation unprofitable; indeed, in the absence of a deviation, players get the utilitarian optimum, which is strictly greater than 0, whereas, by deviating, i receives 0. For unobservable actions, however, this "punishment equilibrium" is not available, and (1.) is an equilibrium *if and only if*

$$s - p_0\lambda - 2p_0\lambda(1-\lambda)s + p_0\lambda(1-\lambda)(s+\lambda) + p_0\lambda(1-\lambda)^3(\lambda - s) \geq 0$$

$$\iff p_0\lambda\big[1 - (1-\lambda)(\lambda - s)\big(1 + (1-\lambda)^2\big)\big] \leq s,$$

which is violated for our parameters. Therefore, the utilitarian optimum (1.) is an equilibrium *if and only if* actions are observable.

Next, let us analyze candidate (2.). If actions are observable (unobservable), a

first-period deviation by player i who is supposed to play risky in that period leads to play of $k_{i,1} = k_{-i,1} = k_{i,2} = k_{-i,2} = k_{i,3} = k_{-i,3} = 0$ ($k_{i,1} = k_{-i,1} = k_{i,2} = k_{i,3} = 0$, with $k_{-i,2} = k_{-i,3} = 1$). In either case, this deviation is profitable *if and only if*

$$0 > s - p_0\lambda + (1 - p_0\lambda)s - p_0(1 - \lambda)\lambda + \big(1 - p_0 + p_0(1 - \lambda)^3\big)s - p_0(1 - \lambda)^3\lambda,$$

which holds for our parameters. Thus, candidate (2.) is eliminated, whether actions be observable or unobservable.

Finally, let us turn to candidate (3.). If actions are observable, a first-period deviation leads to play of either $k_{i,1} = 1$, with $k_{-i,1} = k_{i,2} = k_{-i,2} = k_{i,3} = k_{-i,3} = 0$ or $k_{-i,1} = 0$ with $k_{i,2} = k_{-i,2} = k_{i,3} = k_{-i,3} = 0$. In the latter case, the deviation is unprofitable as $s - p_0\lambda < 0$; in the former, it is unprofitable by the same argument as in the previous paragraph. With unobservable actions, deviating in only one period is unprofitable as $s - p_0\lambda < 0$. A deviation to $k_{i,1} = k_{i,2} = k_{i,3} = 1$ is profitable *if and only if*

$$0 < s - p_0\lambda + \big(1 - p_0 + p_0(1 - \lambda)\big)\big[s - p_1\lambda + \big(1 - p_1 + p_1(1 - \lambda)\big)s - p_1(1 - \lambda)\lambda\big].$$

Yet, $s - p_0\lambda < 0$, and, as we have shown above, $s - p_1\lambda + \big(1 - p_1 + p_1(1 - \lambda)\big)s - p_1(1 - \lambda)\lambda < 0$. We thus conclude that candidate (3.) is an equilibrium, whether actions be observable or unobservable.

We summarize our findings in the following

Proposition 1. *If actions are unobservable, players uniquely always play safe in perfect Bayesian equilibrium. This remains an equilibrium with observable actions. With observable actions, the utilitarian optimum (in which players play risky until a breakdown arrives) is an additional equilibrium, which is supported by the threat of always playing safe in case of a deviation.*

Implications for Behavior Consequently, we hypothesize that action observability matters. Our behavioral hypotheses are as follows:

- We observe efficient behavior more often with observable than with unobservable actions.
- Participants use the risky arm *more* when actions are observable.
- Participants' payoffs are *higher* when actions are observable.
- Updating of beliefs: Conditionally on no breakdown having occurred, participants use the risky arm *more* in later periods.

3 The Experiment

3.1 Organization

We conducted all experiments in the months of July to November 2023 at the University of Vienna. Participants were recruited from the Vienna Center for Experimental Economics (VCEE) subject pool using ORSEE (Greiner 2015). No one participated in more than one session. During the experiments, participants could contact an experimenter anytime for assistance. After reading the instructions, participants had to correctly answer several comprehension questions before starting the main part of the experiment. The experiment was programmed in oTree (Chen, Schonger, and Wickens 2016). We recruited 104 participants and all payments were made in cash. Participants earned on average approximately €10.57 from one randomly selected game and all payments were in cash and in Euros (with a conversion rate of 1000 = €1). The instructions and experimental interface are reproduced in the Online Appendix.

3.2 Implementation

In order to increase the computational efficiency of the implementation and to increase control, we had simulated all the relevant parameters ahead of time. As all our stochastic processes are Bernoulli processes, simulating their realizations ahead of time is equivalent to simulating them as the game progresses. These included separate processes for the quality of the risky arm and the timing of breakdowns on the risky arm in case it was bad.[6] We generated 25 different sets of realizations of the random parameters controlling the quality of the risky arm and the arrivals of the bad risky arm. These corresponded to 25 different games that each of our participants played. To make our findings more easily comparable, we have kept the same realizations for both observable and hidden actions. Participants were randomly assigned to groups of two players and randomly rematched within a matching group of six to eight participants after each game. Each participant was randomly assigned either to the treatment with observable or hidden actions, and played the 25 games in random order. To ensure a balanced data-collection process, we replicated any order of the 25 games that was used for a matching group in the treatment with observable actions for a matching group in the treatment with unobservable actions.

As illustrated in Figure 1, in the treatment with observable actions, participants

[6] Details are available upon request.

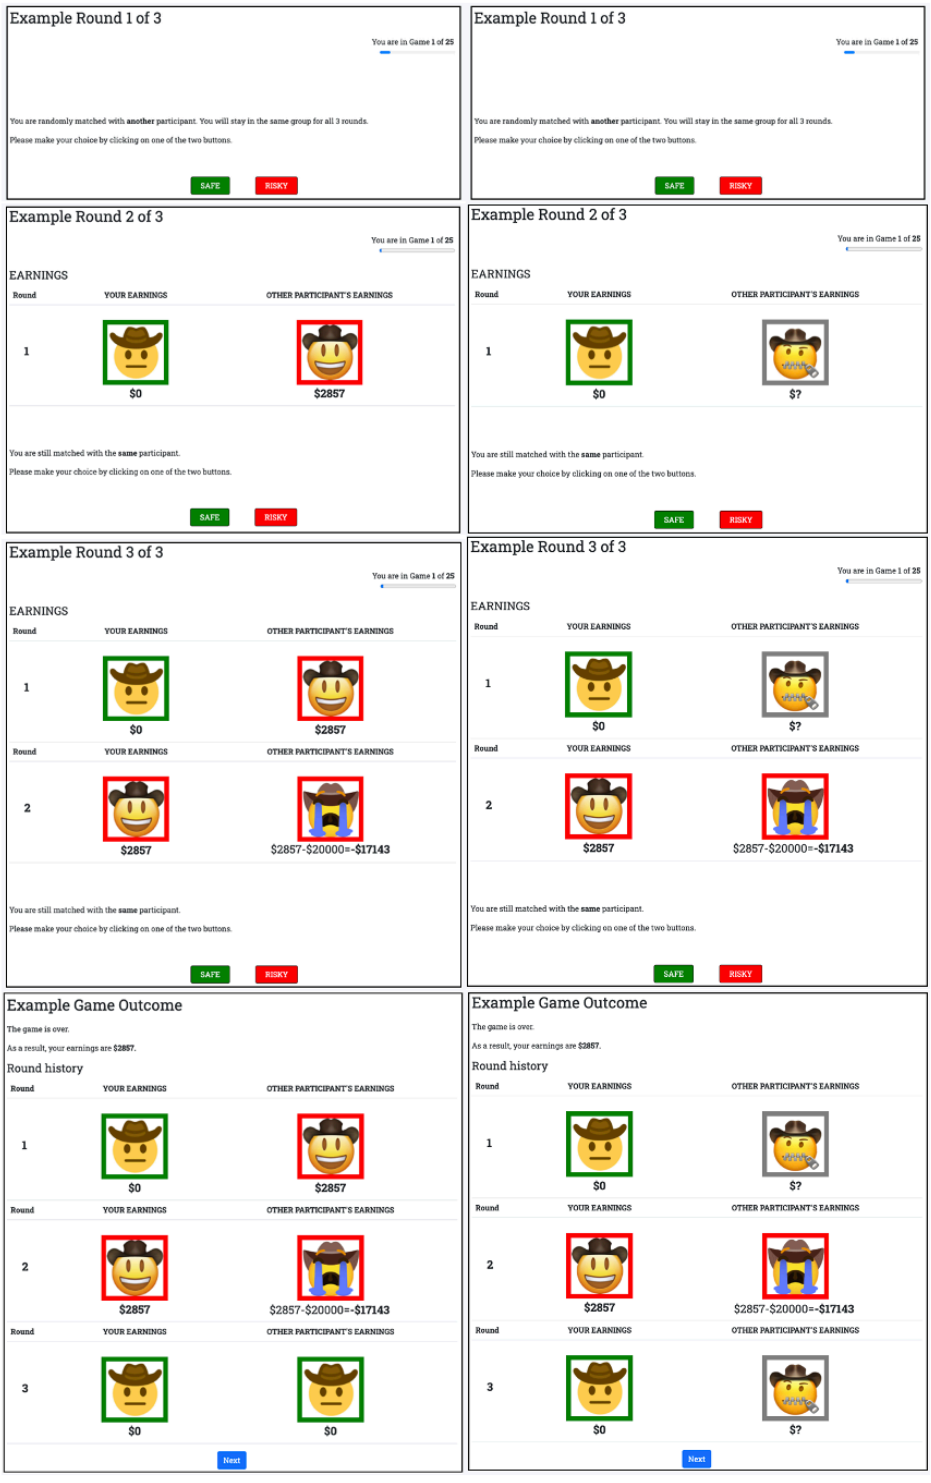

Left: The *Observable Actions* treatment with public information; right: The *Unobservable Actions* treatment with private information.

Figure 1: Experimental Implementation

could see their opponent's as well as their own past action choices and payoffs. In the treatment with unobservable actions, participants could see only if, and when, their opponent had suffered a breakdown so far as well as their own past action choices and payoffs.[7,8]

4 Findings

This section is dedicated to examining the implications for behavior as detailed in Section 2. For each of the 25 games, we implemented two treatments, with observable and unobservable actions respectively, within two-player groups, comprising 52 groups in total. These groups were randomly re-matched within a matching group after each game.

We divide our analysis into four distinct sections. Initially, we provide summary statistics, highlighting both the average intensity of experimentation, $\sum_{t=1}^{3} \frac{k_{1,t}+k_{2,t}}{2}$ (where $k_{i,t} = 1$ ($k_{i,t} = 0$) if player i played risky (safe) in period t), and the overall group payoffs. Following this, our primary analysis examines the aggregate experimental outcomes, with an initial focus on the distribution patterns of experimentation intensities and group payoffs. Subsequently, we assess efficiency by comparing observed behavior to the theoretical efficient solution. In addition, we study how behavior relates to our theoretical predictions, in particular consistency with equilibrium. The final part of our analysis examines the evolution of behavior over time, specifically how participants adjust their action choices in games where no breakdowns occur. To enhance the robustness of our findings, we include a robustness test, utilizing ordinary least squares (OLS) regressions with random effects and clustering of standard errors at the matching-group level. These results are reported in the Online Appendix. Our results show that the number and order of games previously played by participants does not significantly influence their behavior.

4.1 Experimentation and Payoffs

As outlined in Section 2, we anticipate that average experimentation intensities and group payoffs will be higher in the treatment where monitoring by others is possible.

[7] Screenshots can be found in the Online Appendix.

[8] We chose not to elicit risk preferences due to the small stakes involved and prior studies not detecting any statistically significant effect or impact on behavior in similar bandit environments with good news (Hoelzemann and Klein 2021, Hoelzemann, Manso, Nagaraj, and Tranchero 2024).

We measure experimentation intensity for each player up until the moment a first breakdown occurs to any player in the group. Table I presents the observed mean experimentation intensities and the average total payoffs, calculated using group averages across all games for both treatments. In particular, we regressed experimentation intensity and payoffs on the treatment dummy *Public*, which is 0 for the private-information (unobservable-actions) treatment and 1 for the public-information (observable-actions) treatment.

Table I: OLS Estimations

	Experimentation Intensity	Payoffs
Intercept	0.603^{***}	1159.982^{***}
	(0.045)	(255.138)
Public	0.156^{***}	2813.060^{***}
	(0.062)	(446.359)
N	1300	1300
R-squared	0.076	0.014

For all estimations, robust standard errors are clustered at the matching-group level and shown in brackets. ***Significant at the 1 percent level; **Significant at the 5 percent level; *Significant at the 10 percent level

We observe a pronounced positive impact of action observability on both experimentation intensity and payoffs. Specifically, participants pull the risky arm considerably more often when monitoring is feasible, resulting in markedly higher payoffs at the group level.

Extending our analysis beyond mere point estimates, Figure 2 illustrates the empirical distribution of experimentation intensities and group payoffs across the different treatments.

As can be seen in Figure 2, both experimentation intensities and group payoffs exhibit first-order stochastic dominance in the treatment with observable actions over the treatment where actions are unobservable. Kolmogorov-Smirnov tests for evaluating similarity of distributions yield p-values of 0.001.

We summarize these findings in the following:

Experimentation Participants use the risky arm *more* when actions are observable.

Payoffs Group payoffs are *higher* when actions are observable.

That participants should use the risky arm more with observable actions is consistent with the (larger) equilibrium set, unless one expected for some reason that the equilibrium in which everyone plays safe all the time was the only equilibrium being

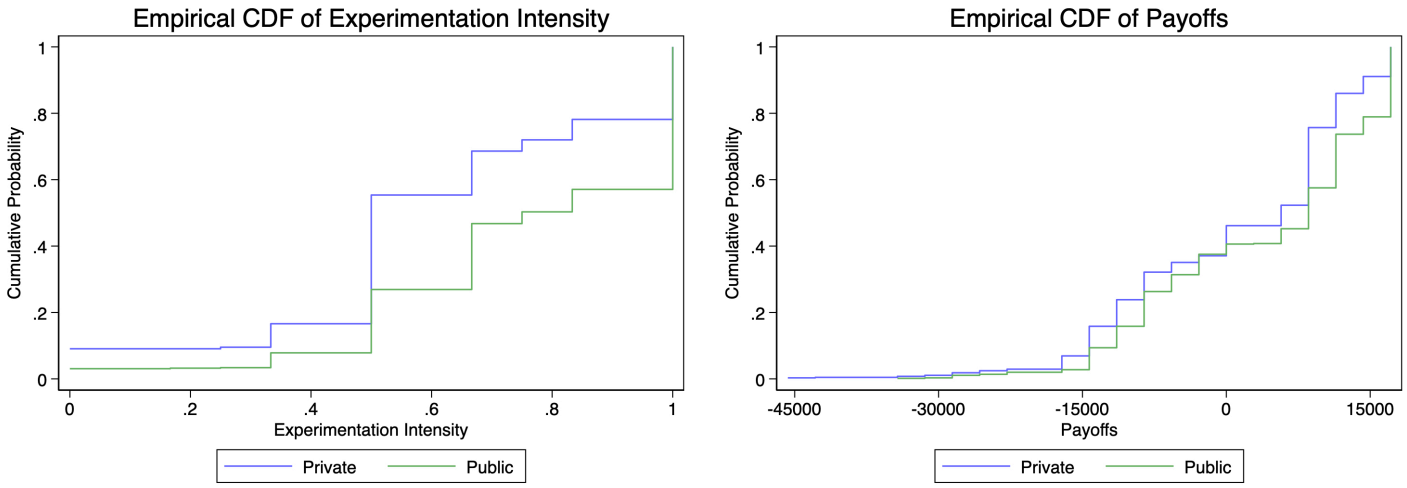

The sample cumulative distribution functions for experimentation intensity and payoffs are shown, by information condition.

Figure 2: Empirical CDFs of Experimentation Intensities and Payoffs

played. In particular, note that there is a stark difference in first-round experimentation intensities, see Figure 4, an effect that cannot be accounted for by channels other than the strategic forces related to action observability.

4.2 Efficiency Benchmark

To assess the efficiency of participants' behavior, our analysis concentrates on games where no breakdown was suffered during the three interaction periods. In the treatment where actions are observable, the average experimentation intensity stands at 0.762 (with a standard deviation of 0.253, N=312), significantly deviating from the theoretical efficient solution of 1. Out of 312 observations, 136 are in alignment with the efficient solution. Conversely, for games with unobservable actions, the average experimentation intensity with 0.592 and a standard deviation of 0.280 for N=312 is also significantly different from the efficient solution, and only 57 of 312 observations directly coincide with the efficient solution.

Efficiency Participants free-ride, i.e., they use the risky arm *less* than what would be efficient with either observable or unobservable actions.

4.3 Consistency of Behavior with Equilibrium

Here too, our analysis focuses on the 312 games where no breakdown was suffered during the three interaction periods. In Figure 3, we highlight the observed risky play for each treatment separately, providing a detailed view of how behavior is associated with equilibrium play. In particular, we plot the fraction of games in which risky was played r times divided by the number of periods t multiplied by the number of players n, i.e., $\frac{r}{t \times n}$.

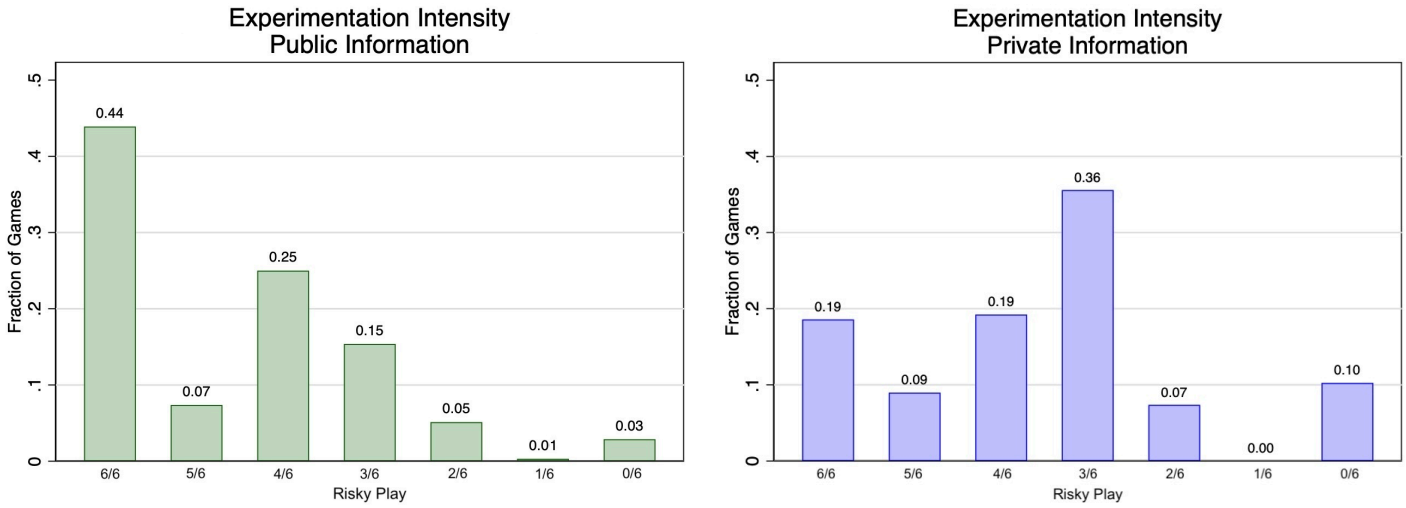

The experimentation intensity in games without breakdowns is shown, by information condition.

Figure 3: Experimentation Intensities

When actions are observable, 144 of 312 games are consistent with equilibrium; among these, the overwhelming majority, namely 136 games, coincide with the efficient solution. By contrast, with hidden actions, play that is consistent with the—now smaller—equilibrium set significantly decreases, with only 32 of 312 in line with the theoretical prediction. Unsurprisingly, the difference in equilibrium play by treatment is highly statistically significant with p-values of 0.001 for both a t-test and a two-sided Wilcoxon rank-sum test.

This leads to the following:

Equilibrium Behavior is more often consistent with equilibrium when actions are observable.

4.4 Dynamic Evolution of Behavior

We now shift our focus to the dynamics of observed behavior. We are particularly interested in whether, consistently with Bayesian updating, participants increase their use of the risky arm as the game progresses. Our attention remains on games where no breakdown is incurred throughout the three periods of interaction. In Figure 4, we graph the observed experimentation intensities for each period and treatment separately, providing a detailed view of how behavior evolves over the course of the game.

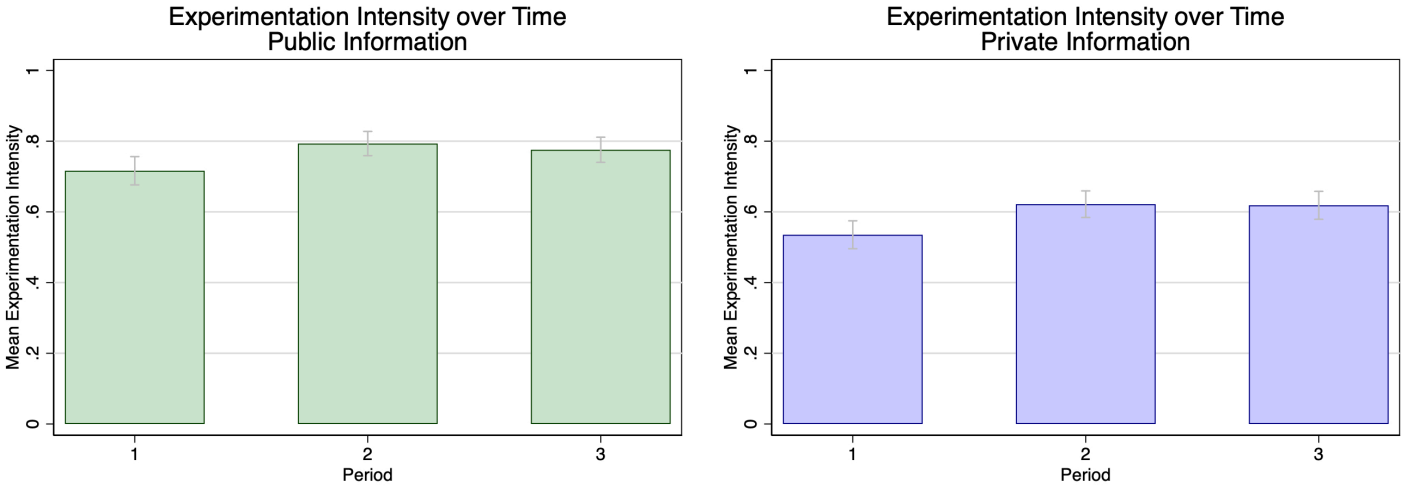

The average experimentation intensity over time in games without breakdowns is shown, by information condition.

Figure 4: Experimentation Intensities over Time

While it is unrealistic to expect our participants to calculate posterior beliefs precisely using Bayes' rule, we nevertheless anticipate that, in the absence of a breakdown, participants will increasingly use the risky arm as the game progresses, reflecting growing optimism. At the beginning of a game without any breakdowns, participants are indeed significantly less likely to choose the risky arm than in later periods. To examine changes in behavior over time, we employ two-sided t-tests for parametric analysis and two-sided Wilcoxon rank-sum (Mann-Whitney) tests for nonparametric analysis, treating group averages as independent observations. Regardless of whether actions were observable, we find that the differences in experimentation intensities across time are highly statistically significant when comparing behavior in the first period to that in either the second or the last period. In the treatment with observable actions, t-tests (two-sided Wilcoxon rank-sum tests) produce p-values of $p_{12} = 0.012$ (0.024), $p_{13} = 0.089$ (0.121), and $p_{23} = 0.411$ (0.480), where $p_{\alpha\beta}$ is the

p-value when comparing periods α and β. With unobservable actions, by contrast, we find p-values of $p_{12} = 0.001$ (0.002), $p_{13} = 0.008$ (0.007), and $p_{23} = 0.643$ (0.751), respectively.

Additionally, we also analyze mean experimentation intensities across treatments for each period. Participants engage with the risky arm more frequently when actions are observable. The differences in all periods are highly statistically significant, with p-values of 0.001 for both tests.

We summarize these results as follows:

Belief Updating Conditionally on no breakdown having occurred, participants use the risky arm *more* in later periods.

4.5 Econometric Robustness Tests

As a further robustness test and to complement our previous analyses and key elements discussed, we ran ordinary least-square regressions with random effects controlling for learning effects. In particular, we regressed experimentation intensity and individual payoffs on the treatment dummy *Public*, which is 0 for the private-information (unobservable-actions) treatment and 1 for the public-information (observable-actions) treatment. Recall that participants played the 25 games in random order and any order of these games that was used for participants in the public information sessions was replicated for participants in the private information sessions. In order to verify that participants treated the games they successively played as independent games rather than as parts of a larger super-game, we define a weighted learning function $\{g_o\} = \{1/o\}$ where o ($o \in \{1, \cdots, 25\}$) corresponds to the random order in which each participant was exposed to each game. All regressions control for trends over time using this weighted learning function. The results do not qualitatively change when we replace the learning function with a linear version such that $\{g_o\} = \{o\}$. Further, the results do not qualitatively change either when we include controls for matching groups or sessions, age, gender, field of study as well as attempts needed to correctly answer the quiz questions at the start of the experiment. To account for the fact that behavior within matching groups is not independent, we treat each matching group as our units of statistically independent observations and cluster standard errors by matching group.

Table II lists the results from this analysis.

Table II: OLS Estimations with Random Effects of Experimentation Intensity and Payoffs.

	Experimentation Intensity			*Individual*
	All	*No Breakdown*	*Until Breakdown*	*Payoffs*
Intercept	0.596^{***}	0.586^{***}	0.603^{***}	418.199^{*}
	(0.044)	(0.049)	(0.054)	(239.186)
Public	0.157^{***}	0.170^{***}	0.280^{***}	1398.080^{***}
	(0.052)	(0.054)	(0.054)	(174.776)
Learning	0.034	0.040	-0.082	1059.966
	(0.026)	(0.029)	(0.106)	(1077.624)
σ_ϵ	0.287	0.268	0.293	8494.653
σ_u	0.216	0.212	0.243	0
N	2600	1248	336	2600
(Between) R-squared	0.113	0.127	0.199	0.189

For all estimations, robust standard errors are clustered at the session level and shown in brackets. ***Significant at the 1 percent level; **Significant at the 5 percent level; *Significant at the 10 percent level

We find a strong positive effect of public information on experimentation intensity across all games, games without breakdowns, games with breakdowns before the last period, and payoffs.

5 Related Literature

Strategic Experimentation Our theoretical framework can be viewed as a model of experimentation. Initially, the literature focussed on the trade-off of an individual decision maker who acts in isolation. Bolton and Harris (1999) and Keller, Rady, and Cripps (2005) have extended the individual decision problem to a multi-player framework. Since then this literature is steadily growing. For example, Klein and Rady (2011), Klein (2013), and Hörner, Klein, and Rady (2022) study various bandit problems in which different players may choose different arms. While free-riding is a central element in these studies as well, players' actions are observable. Several studies analyzed experimentation in teams where the outcome of each player's action is unobservable while their actions are observable (Rosenberg, Solan, and Vieille 2007, Murto and Välimäki 2011, Hopenhayn and Squintani 2011), while Bonatti and Hörner (2011) studies the case of observable outcomes with unobservable actions. In particular, Keller and Rady (2015) study the bad-news setting under public information, while, closest to our setting, Bonatti and Hörner (2017) study a bad-news setting

where actions are not observed but outcomes are.

Strategic Experimentation Experiments in Good-News Environments Our paper is embedded in an emerging literature that studies behavior in experiments of experimentation with breakthrough-learning in groups; that is, multi-player settings with strategic links across players where decision-makers are *not* acting in isolation *and* informational externalities exist. We are aware of only four other experimental investigations of good-news strategic-experimentation problems with bandits. Hoelzemann and Klein (2021) experimentally implement a dynamic public-good problem where information about agents' common state of the world is dynamically evolving. Observed behavior is consistent with free-riding because of strategic concerns, and participants adopt non-cut-off behavior and frequent switches of action. Boyce, Bruner, and McKee (2016) study a setting with ambiguity concerning the type of the risky arm to test strategic free-riding in a two-player, two-period, game. Players are asymmetric in their costs in that one player was known to have lower opportunity costs for playing risky than the other, so that it was clear which player ought to play the free-rider in the first period. Von Essen, Huysentruyt, and Miettinen (2020) implement a treasure-hunt game in the laboratory, finding that the information externality can induce an encouragement effect. Hoelzemann, Manso, Nagaraj, and Tranchero (2024) study an environment where players must explore across different options with varying but uncertain payoffs. While informative signals, interpreted as data, can typically reduce uncertainty and improve welfare, in their setting it can instead decrease individual and group payoffs. When data highlights sufficiently attractive but dominated options, it can crowd-out exploration and thus lower payoffs as compared to when no data is provided. Importantly, empirical evidence from the field of genetics research provides a real-world confirmation of their framework and shows that data on genetic targets of medium promise can significantly increase the delay of valuable discoveries.

By contrast, this paper offers the first experimental investigation of *breakdown*-learning in a strategic setting. Observed behavior is markedly different from that documented in strategic environments with good news. Moreover, observed behavior is consistent with strategic free-riding; thus documenting that free-riding on information because of strategic concerns also exists in bad-news learning environments, while differing in its form of manifestation.

References

Azoulay, P., J. S. Graff Zivin, and G. Manso (2011): "Incentives and Creativity: Evidence from the Academic Life Sciences," *The RAND Journal of Economics*, 42(3), 527–554.

Bergemann, D., and U. Hege (1998): "Venture Capital Financing, Moral Hazard, and Learning," *Journal of Banking & Finance*, 22(6-8), 703–735.

——— (2005): "The Financing of Innovation: Learning and Stopping," *RAND Journal of Economics*, pp. 719–752.

Bergemann, D., and J. Välimäki (2008): "Bandit Problems," The New Palgrave Dictionary of Economics, 2nd ed. Palgrave and Macmillan Press.

Bolton, P., and C. Harris (1999): "Strategic Experimentation," *Econometrica*, 67(2), 349–374.

Bonatti, A., and J. Hörner (2011): "Collaborating," *American Economic Review*, 101(2), 632–63.

——— (2017): "Learning to Disagree in a Game of Experimentation," *Journal of Economic Theory*, 169, 234–269.

Boyce, J. R., D. M. Bruner, and M. McKee (2016): "Strategic Experimentation in the Lab," *Managerial and Decision Economics*, 37(6), 375–391.

Callander, S., and B. Harstad (2015): "Experimentation in Federal Systems," *The Quarterly Journal of Economics*, 130(2), 951–1002.

Chen, D. L., M. Schonger, and C. Wickens (2016): "oTree – An Open-Source Platform for Laboratory, Online, and Field Experiments," *Journal of Behavioral and Experimental Finance*, 9, 88–97.

Ederer, F., and G. Manso (2013): "Is Pay for Performance Detrimental to Innovation?," *Management Science*, 59(7), 1496–1513.

Greiner, B. (2015): "Subject Pool Recruitment Procedures: Organizing Experiments with ORSEE," *Journal of the Economic Science Association*, 1(1), 114–125.

Halac, M., N. Kartik, and Q. Liu (2017): "Contests for Experimentation," *Journal of Political Economy*, 125(5), 1523–1569.

Hoelzemann, J., and N. Klein (2021): "Bandits in the Lab," *Quantitative Economics*, 12(3), 1021–1051.

Hoelzemann, J., G. Manso, A. Nagaraj, and M. Tranchero (2024): "The Streetlight Effect in Data-Driven Exploration," *NBER Working Paper #32401*.

Hopenhayn, H. A., and F. Squintani (2011): "Preemption Games with Private Information," *The Review of Economic Studies*, 78(2), 667–692.

Hörner, J., N. Klein, and S. Rady (2022): "Overcoming Free-Riding in Bandit Games," *The Review of Economic Studies*, 89(4), 1948–1992.

Howarth, J. (2025): "Exploding Topics: Startup Failure Rate Statistics," in *https://explodingtopics.com/blog/startup-failure-stats*.

Keller, G., and S. Rady (2015): "Breakdowns," *Theoretical Economics*, 10(1), 175–202.

Keller, G., S. Rady, and M. Cripps (2005): "Strategic Experimentation with Exponential Bandits," *Econometrica*, 73(1), 39–68.

Klein, N. (2013): "Strategic Learning in Teams," *Games and Economic Behavior*, 82, 636–657.

——— (2016): "The Importance of Being Honest," *Theoretical Economics*, 11(3), 773–811.

Klein, N., and S. Rady (2011): "Negatively Correlated Bandits," *The Review of Economic Studies*, 78(2), 693–732.

Kotashev, K. (2024): "Startup Failure Rate: How Many Startups Fail and Why in 2024?," in *https://www.failory.com/blog/startup-failure-rate*.

Krieger, J. L. (2021): "Trials and Terminations: Learning from Competitors' R & D Failures," *Management Science*, 67(9), 5525–5548.

Levinthal, D. A. (1997): "Adaptation on Rugged Landscapes," *Management Science*, 43(7), 934–950.

Manso, G. (2011): "Motivating Innovation," *The Journal of Finance*, 66(5), 1823–1860.

March, J. G. (1991): "Exploration and Exploitation in Organizational Learning," *Organization Science*, 2(1), 71–87.

Mihm, J., and J. Schlapp (2019): "Sourcing Innovation: On Feedback in Contests," *Management Science*, 65(2), 559–576.

Murto, P., and J. Välimäki (2011): "Learning and Information Aggregation in an Exit Game," *The Review of Economic Studies*, 78(4), 1426–1461.

Rosenberg, D., E. Solan, and N. Vieille (2007): "Social Learning in One-Arm Bandit Problems," *Econometrica*, 75(6), 1591–1611.

Sorensen, M. (2008): "Learning by Investing: Evidence from Venture Capital," in *AFA 2008 New Orleans Meetings Paper*.

Von Essen, E., M. Huysentruyt, and T. Miettinen (2020): "Exploration in Teams and the Encouragement Effect: Theory and Experimental Evidence," *Management Science*, 66(12), 5861–5885.

Wagner, P. A., and N. Klein (2022): "Strategic Investment and Learning with Private Information," *Journal of Economic Theory*, 204, 105523.

Wolitzky, A. (2018): "Learning from Others' Outcomes," *American Economic Review*, 108(10), 2763–2801.

Breakdowns in the Lab

Online Appendix: Experimental Interface

The *Bad-News* Environment with *Public* Information

June 25, 2025

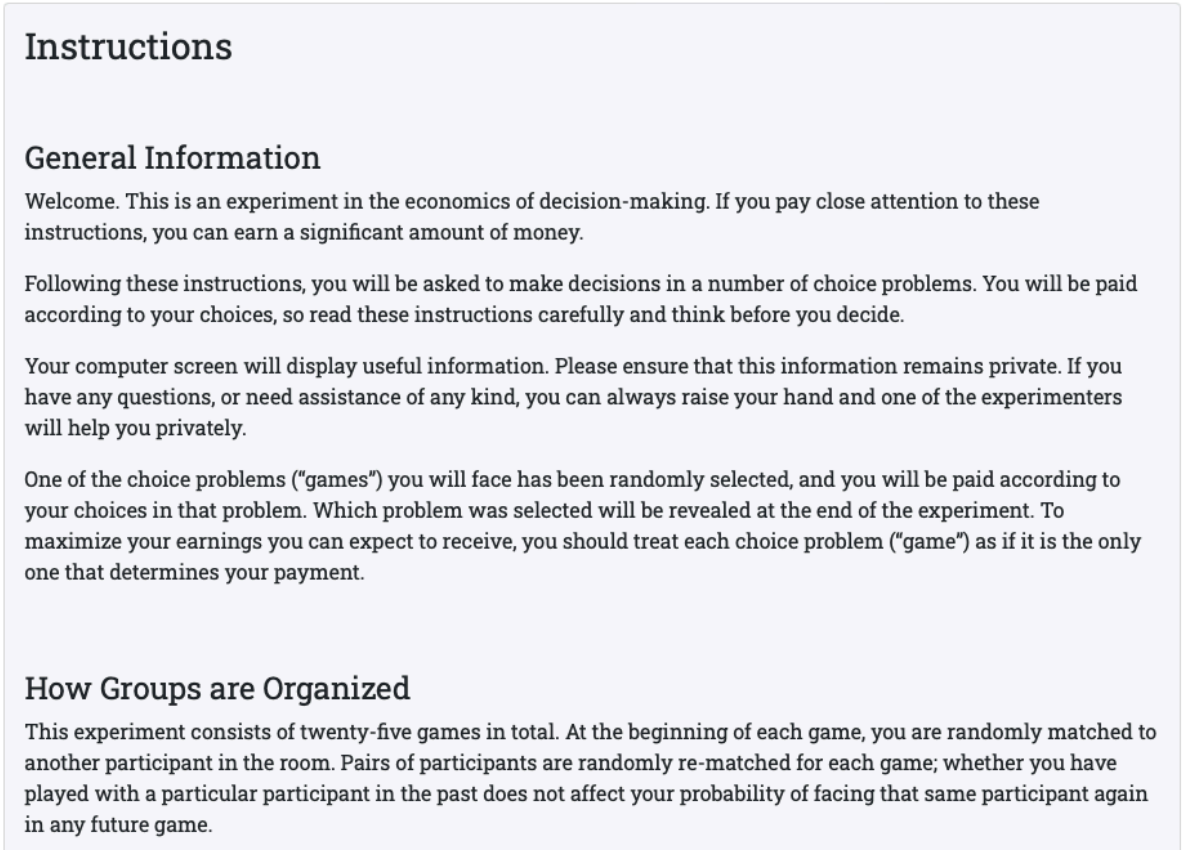

Instructions

General Information

Welcome. This is an experiment in the economics of decision-making. If you pay close attention to these instructions, you can earn a significant amount of money.

Following these instructions, you will be asked to make decisions in a number of choice problems. You will be paid according to your choices, so read these instructions carefully and think before you decide.

Your computer screen will display useful information. Please ensure that this information remains private. If you have any questions, or need assistance of any kind, you can always raise your hand and one of the experimenters will help you privately.

One of the choice problems ("games") you will face has been randomly selected, and you will be paid according to your choices in that problem. Which problem was selected will be revealed at the end of the experiment. To maximize your earnings you can expect to receive, you should treat each choice problem ("game") as if it is the only one that determines your payment.

How Groups are Organized

This experiment consists of twenty-five games in total. At the beginning of each game, you are randomly matched to another participant in the room. Pairs of participants are randomly re-matched for each game; whether you have played with a particular participant in the past does not affect your probability of facing that same participant again in any future game.

How the Game Works

Every game consists of three rounds; you will make a decision in each round. That means that you have to decide whether you want to play the "safe" (SAFE) or the "risky" (RISKY) option in each round. Your previous choices in no way constrain your subsequent ones.

When you choose the SAFE option, you will receive no earnings in the current round.

If you choose the RISKY option, you will always earn **$2857**. In addition, you may, however, lose **$20000**. The chance of this happening depends on whether your risky option is good or bad. Whether it is good or bad is determined by the computer at the start of each game.
With a 67.6% chance, your and the other participant's risky option will be **bad**;
with a 32.4% chance, your and the other participant's risky option will be **good**.
Note that your, and the other participant's, risky option will always be of the **same** type (either both are good or both are bad).

If your RISKY option is **bad**, you may **lose $20000**. This will happen **with a 25% chance** in each round you use it. That is, with a 25% chance you lose $17143, and with a 75% chance you gain $2857. The probability that you incur this loss in a given round is constant over time, and does not depend on your previous earnings or choices, nor does it depend on the other participant's previous or current earnings or choices. It only depends on the type of the risky option (which is determined before the start of the game and which remains constant throughout the game) and your choice in the current round.

If your RISKY option is **good**, you will never incur a loss and will always **gain $2857** in each round you use it.

At the end of each round, you and the other participant observe each other's choices and earnings in that round. Note that by observing the behavior of their RISKY option (provided they use it) you can learn something about your own RISKY option as well.

Your earnings for the game are the sum of the earnings you will have accumulated over the three rounds of the game.

Please Note

The parameters are chosen in such a way that, *if you knew* the RISKY option to be good, you would be best off by always choosing it. Yet, *if you knew* the RISKY option to be bad, you would be best off by always choosing the SAFE option. The other participant is solving the exact same problem as you and has read the exact same instructions.

The Following Graphics Illustrate How the Game Evolves

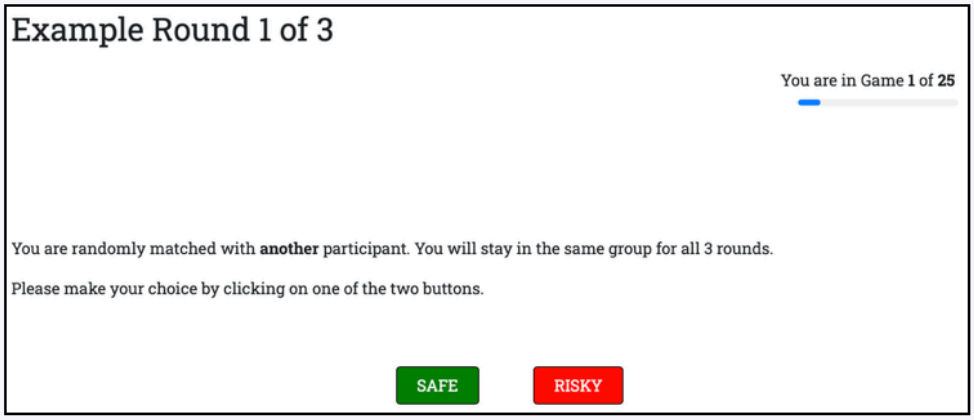

- In this example, you are in Round 1.
- You now need to choose between playing SAFE or RISKY.

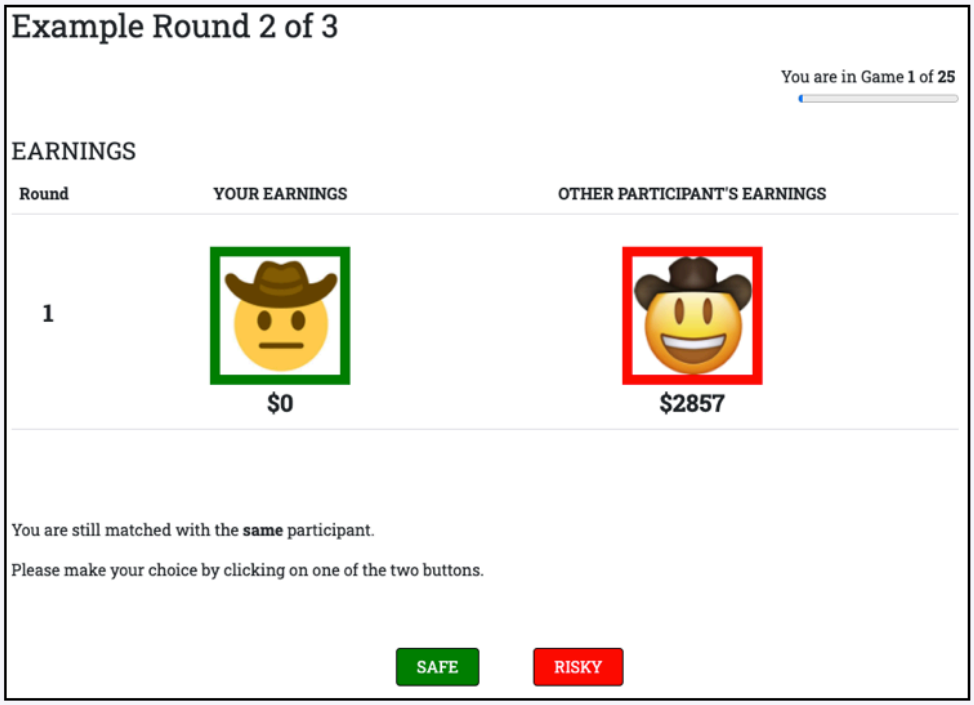

- In this example, you are in Round 2, and therefore you get to see your own and the other participant's choices and earnings in Round 1.
- The **left** column always shows **your** choices and earnings:
  - In this example, you have started playing safe (highlighted by the green frame), which does not yield a reward.
- The **right** column always shows the **other participant's** choices and earnings:
  - In this example, the other participant has played risky (highlighted by the red frame) in Round 1 and has not incurred a loss, thus earning $2857.
- Finally, you now need to choose between playing SAFE or RISKY.

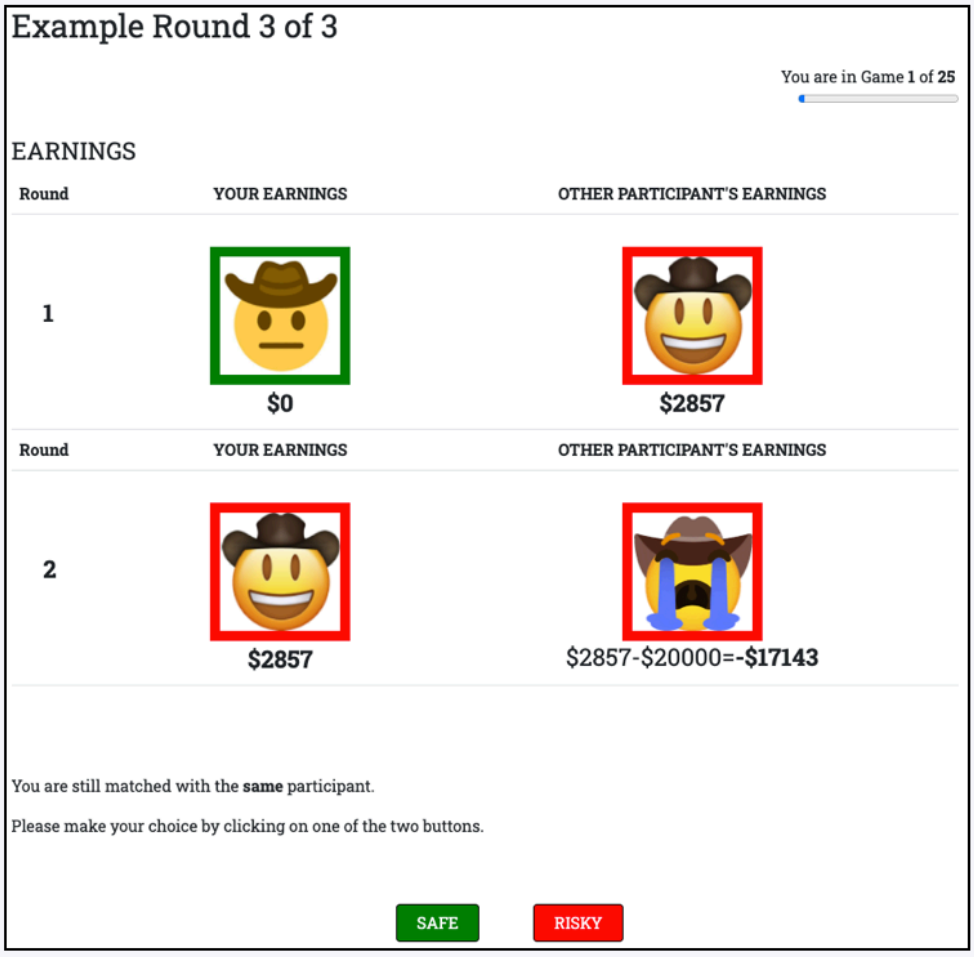

- In this example, you are in Round 3, and therefore you get to see your own and the other participant's choices and earnings in Rounds 1 and 2.
- The **left** column always shows **your** choices and earnings:
  - In this example, you have played risky (highlighted by the red frame) in Round 2 and have not incurred a loss, thus earning $2857.
- The **right** column always shows the **other participant's** choices and earnings:
  - In this example, the other participant has played risky (highlighted by the red frame) in Round 2, yielding $2857, but has also incurred a loss of $20000: 🤠

    This means that **your and the other participant's** RISKY options are **bad**.
- Finally, you now need to choose between playing SAFE or RISKY.

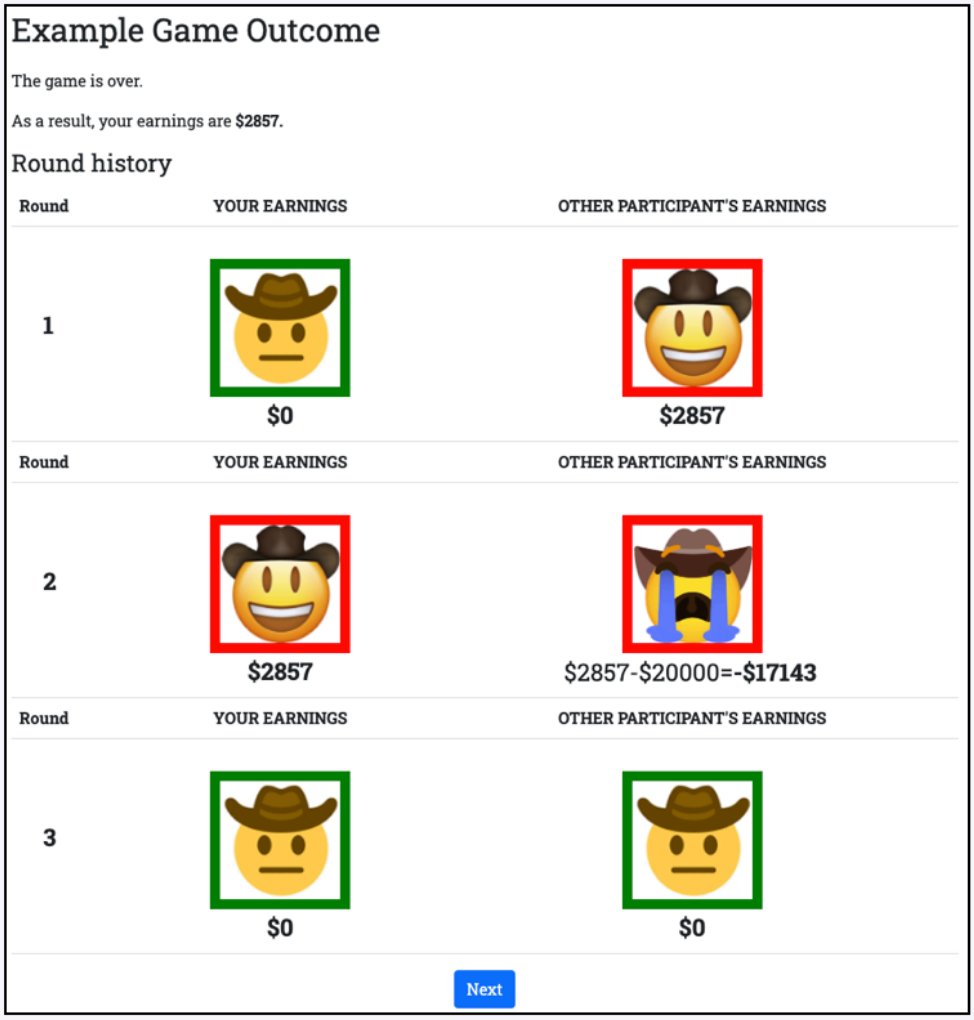

- In this example, the game is over, and therefore you get to see your own and the other participant's choices and earnings in all rounds.
- The **left** column always shows **your** choices and earnings:
  - In this example, you have played safe (highlighted by the green frame) in Round 1, which does not yield a reward. In Round 2 you have played risky (highlighted by the red frame), which has yielded a reward of $2857; and in Round 3 you have played safe (highlighted by the green frame), which does not yield a reward. Thus, your earnings in this game are $2857.
- The **right** column always shows the **other participant's** choices and earnings:
  - In this example, the other participant has played risky (highlighted by the red frame) in Round 1, which has yielded a reward of $2857. In Round 2 they have played risky (highlighted by the red frame) and incurred a loss of $17143; and in Round 3 they have played safe (highlighted by the green frame), which does not yield a reward. Thus, the other participant's loss in this game is $14286.

Payment

In the experiment you will be making decisions that will earn you $ ("Dollars"). At the end of the experiment, the $ you have earned will be converted into Euros at an exchange rate of $ 10000 = € 1, and paid out in cash. This amount will be added to your participation payment of € 10.

Before the actual experiment starts, you will be asked to answer some questions. You must answer these correctly in order to proceed to the next question.

After completing the experiment, the computer will reveal which of the twenty-five games was randomly selected to determine your payment. You will be paid according to the choices you made in that particular game and you will be paid your earnings in cash. This protocol of determining payments suggests that you should choose in each game as if it is the only game that determines your payment.

Frequently Asked Question

Is this some kind of psychology experiment with an agenda you haven't told us? Answer. No. It is an economics experiment. If we do anything deceptive or don't pay you cash as described, then you can complain to the University of Vienna Research Ethics Board and we will be in serious trouble. These instructions are meant to clarify how you earn money, and our interest is in seeing how people make decisions.

Next

This button will be activated after 300 seconds. Please take your time to read through the instructions.

You have successfully finished reading the instructions.

The quiz, consisting of 13 questions in total, follows.

Next

Quiz Time!

Q1. The computer has randomly selected one of the games to determine your payment.

A Correct

B Incorrect

Quiz Time!

Q2. In each game (not round) you will be randomly matched with another participant in this session and stay together for that game.

A Correct

B Incorrect

Quiz Time!

Q3. In all games, the RISKY option is always **bad** and always leads to a loss of $20000 each time it is used.

A Correct

B Incorrect

Quiz Time!

Q4. If the RISKY option is **bad**, then it leads to a loss of $20000 with a chance of 25%.

A Correct

B Incorrect

Quiz Time!

Q5. If you knew the RISKY option to be **bad**, you would be best off by always choosing the SAFE option

A Correct

B Incorrect

Quiz Time!

Q6. If you experienced a loss of $20000 in an earlier round, you would be best off by always choosing the _____ .

A SAFE option

B RISKY option

Quiz Time!

Q7. A **bad** RISKY option always leads to a loss of $20000 each time you use it.

A Correct

B Incorrect

Quiz Time!

Q8. A **good** RISKY option leads to a loss of $20000 with a 25% chance.

A Correct

B Incorrect

Quiz Time!

Q9. If the other participant has suffered a loss of $20000 , this means that your RISKY option is **bad**.

A Correct

B Incorrect

Quiz Time!

Q10. If you have experienced a loss of $20000 , this means that your RISKY option is **bad**.

A Correct

B Incorrect

Quiz Time!

Q11. If you choose the SAFE option in a given round, then you can never learn anything about your and the other participant's RISKY option in that round.

A Correct

B Incorrect

Quiz Time!

Q12. If both you and the other participant choose the SAFE option in a given round, then you cannot learn anything new about your and the other participant's RISKY option in that round.

A Correct

B Incorrect

Quiz Time!

Q13. The more you have observed you and/or the other participant using the RISKY option without experiencing a loss of $20000, the more inclined you should be to choose the RISKY option in the future.

A Correct

B Incorrect

You have successfully finished the quiz.

The experiment follows. When you are ready please click "Next" to start the experiment.

Next

Round 1 of 3

You are in Game **1** of **25**

You are randomly matched with **another** participant. You will stay in the same group for all 3 rounds.

Please make your choice by clicking on one of the two buttons.

SAFE RISKY

Round 2 of 3

You are in Game **1** of **25**

EARNINGS

Round	YOUR EARNINGS	OTHER PARTICIPANT'S EARNINGS
1	**$2857**	$2857-$20000=**-$17143**

You are still matched with the **same** participant.

Please make your choice by clicking on one of the two buttons.

SAFE RISKY

Round 3 of 3

You are in Game **1** of **25**

EARNINGS

Round	YOUR EARNINGS	OTHER PARTICIPANT'S EARNINGS
1	**$2857**	$2857-$20000=**-$17143**

Round	YOUR EARNINGS	OTHER PARTICIPANT'S EARNINGS
2	**$2857**	**$0**

You are still matched with the **same** participant.

Please make your choice by clicking on one of the two buttons.

SAFE RISKY

Game Outcome

The game is over.

As a result, your earnings are **$5714**.

Round history

Round	YOUR EARNINGS	OTHER PARTICIPANT'S EARNINGS
1	**$2857**	$2857-$20000=**-$17143**
Round	**YOUR EARNINGS**	**OTHER PARTICIPANT'S EARNINGS**
2	**$2857**	**$0**
Round	**YOUR EARNINGS**	**OTHER PARTICIPANT'S EARNINGS**
3	**$0**	**$0**

Next

You have successfully finished the main part of the experiment.

A brief questionnaire follows.

Your answers will be kept confidential and will not affect your earnings for today's experiment.

Please state your age:

Please state your gender:

Please state your degree and field of study:

Please briefly explain, in your own words, the rules of today's experiment:

Please briefly describe how you reached your decisions in this experiment:

Please briefly describe how, in your opinion, other participants reached their decisions in this experiment:

Next

Thank you for participating in this experiment!

Your Earnings

The randomly selected game for payment is **Game #21**. Your total earnings are therefore **$5714**.

This corresponds to **€0.57**.

In total, you earned **€10.57** in this experiment. This includes your participation payment of €10 for taking part in this experiment.

Breakdowns in the Lab

Online Appendix: Experimental Interface

The *Bad-News* Environment with *Private* Information

June 25, 2025

Instructions

General Information

Welcome. This is an experiment in the economics of decision-making. If you pay close attention to these instructions, you can earn a significant amount of money.

Following these instructions, you will be asked to make decisions in a number of choice problems. You will be paid according to your choices, so read these instructions carefully and think before you decide.

Your computer screen will display useful information. Please ensure that this information remains private. If you have any questions, or need assistance of any kind, you can always raise your hand and one of the experimenters will help you privately.

One of the choice problems ("games") you will face has been randomly selected, and you will be paid according to your choices in that problem. Which problem was selected will be revealed at the end of the experiment. To maximize your earnings you can expect to receive, you should treat each choice problem ("game") as if it is the only one that determines your payment.

How Groups are Organized

This experiment consists of twenty-five games in total. At the beginning of each game, you are randomly matched to another participant in the room. Pairs of participants are randomly re-matched for each game; whether you have played with a particular participant in the past does not affect your probability of facing that same participant again in any future game.

How the Game Works

Every game consists of three rounds; you will make a decision in each round. That means that you have to decide whether you want to play the "safe" (SAFE) or the "risky" (RISKY) option in each round. Your previous choices in no way constrain your subsequent ones.

When you choose the SAFE option, you will receive no earnings in the current round.

If you choose the RISKY option, you will always earn **$2857**. In addition, you may, however, lose **$20000**. The chance of this happening depends on whether your risky option is good or bad. Whether it is good or bad is determined by the computer at the start of each game.
With a 67.6% chance, your and the other participant's risky option will be **bad**;
with a 32.4% chance, your and the other participant's risky option will be **good**.
Note that your, and the other participant's, risky option will always be of the **same** type (either both are good or both are bad).

If your RISKY option is **bad**, you may **lose $20000**. This will happen **with a 25% chance** in each round you use it. That is, with a 25% chance you lose $17143, and with a 75% chance you gain $2857. The probability that you incur this loss in a given round is constant over time, and does not depend on your previous earnings or choices, nor does it depend on the other participant's previous or current earnings or choices. It only depends on the type of the risky option (which is determined before the start of the game and which remains constant throughout the game) and your choice in the current round.

If your RISKY option is **good**, you will never incur a loss and will always **gain $2857** in each round you use it.

At the end of each round, you observe only if the other participant incurred a loss of **$20000** from their RISKY option as well as your own choice and earnings.

Your earnings for the game are the sum of the earnings you will have accumulated over the three rounds of the game.

Please Note

The parameters are chosen in such a way that, *if you knew* the RISKY option to be good, you would be best off by always choosing it. Yet, *if you knew* the RISKY option to be bad, you would be best off by always choosing the SAFE option. The other participant is solving the exact same problem as you and has read the exact same instructions.

The Following Graphics Illustrate How the Game Evolves

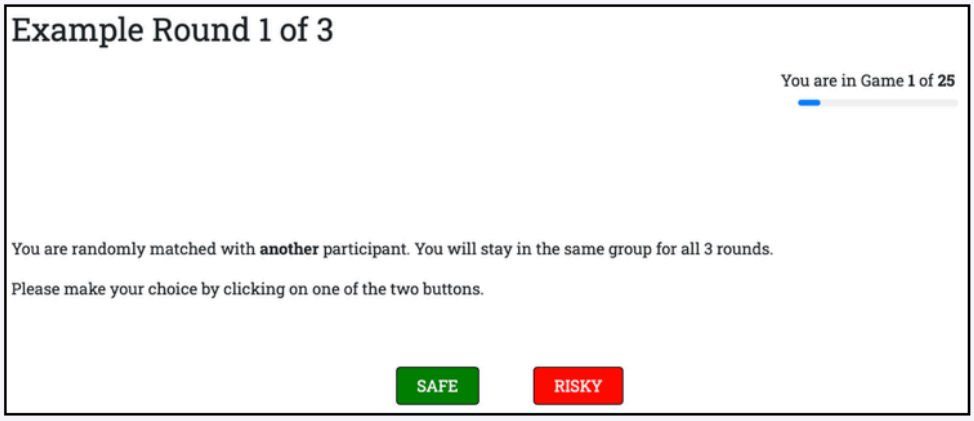

- In this example, you are in Round 1.
- You now need to choose between playing SAFE or RISKY.

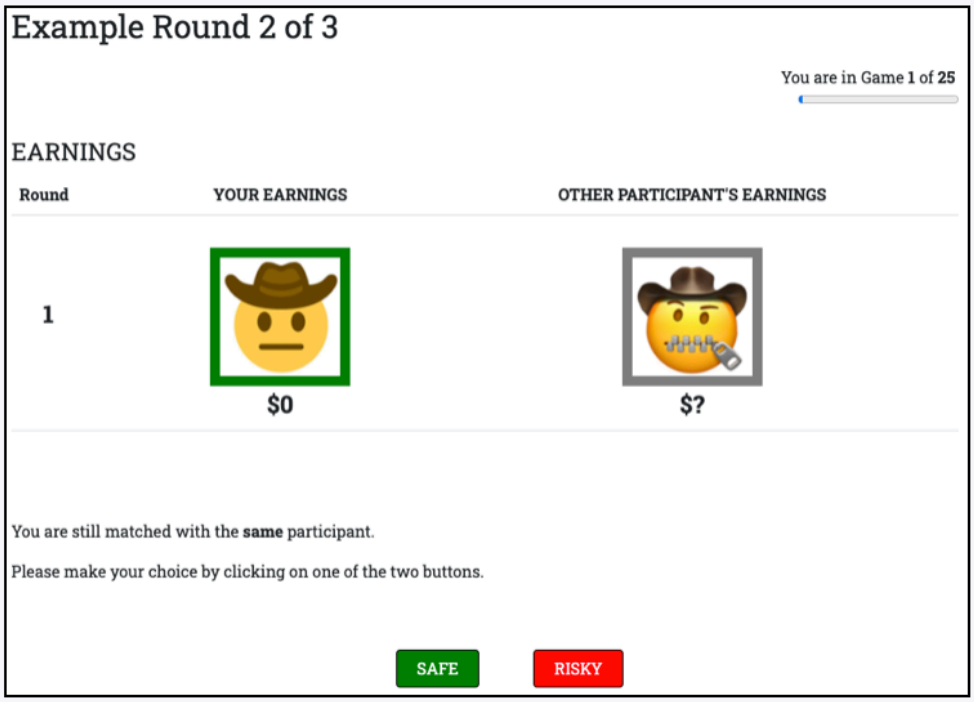

- In this example, you are in Round 2, and therefore you get to see your own choices and earnings in Round 1 and whether or not the other participant has incurred a loss of $20000.
- The **left** column always shows **your** choices and earnings:
  - In this example, you have started playing safe (highlighted by the green frame), which does not yield a reward.
- The **right** column always shows whether or not the **other participant** has incurred a loss of $20000:
  - In this example, the other participant has played either safe or risky (highlighted by the grey frame) in Round 1 and has not incurred a loss. So, their earnings in Round 1 are either $0 if they chose the safe option or $2857 if they chose the risky option.
- Finally, you now need to choose between playing SAFE or RISKY.

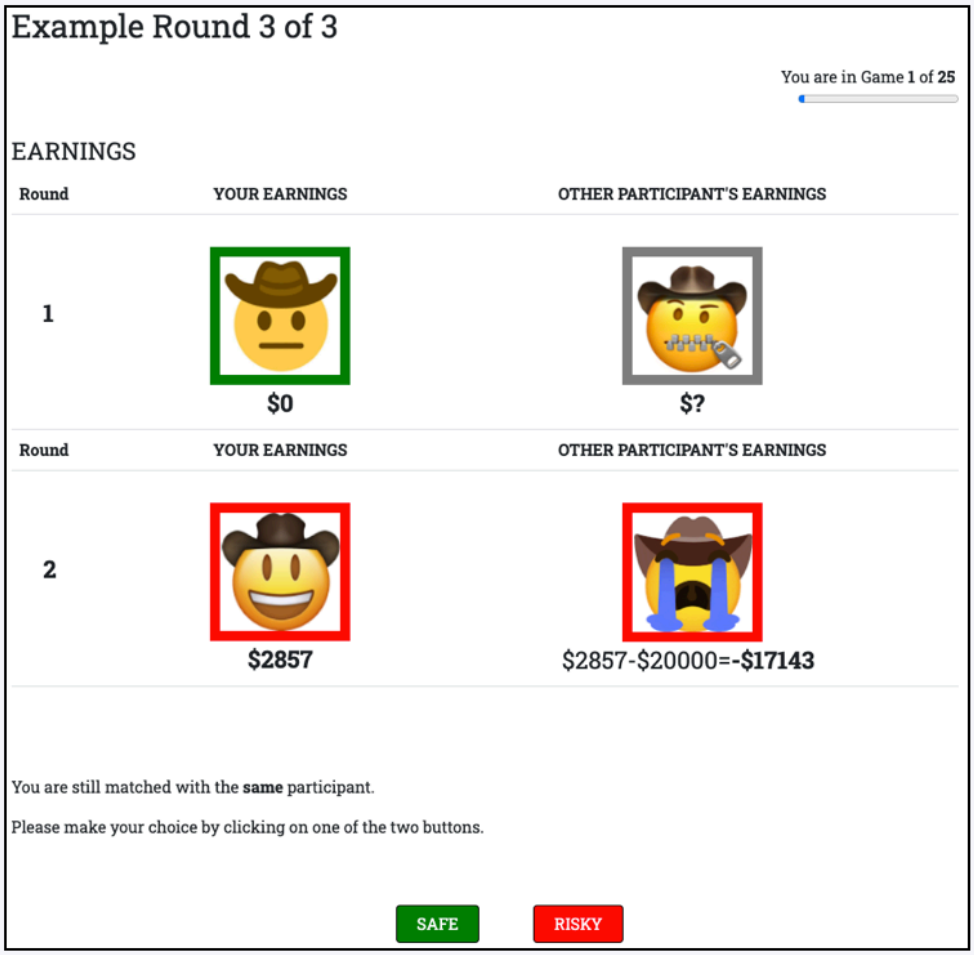

- In this example, you are in Round 3, and therefore you get to see your own choices and earnings in Rounds 1 and 2, and whether or not the other participant has incurred a loss of $20000.
- The **left** column always shows **your** choices and earnings:
  - In this example, you have played risky (highlighted by the red frame) in Round 2 and have not incurred a loss, thus earning $2857.
- The **right** column always shows whether or not the **other participant** has incurred a loss of $20000:
  - In this example, the other participant has played risky (highlighted by the red frame) in Round 2, yielding $2857, but has also incurred a loss of $20000: 😭

    This means that **your and the other participant's** RISKY options are **bad**.
- Finally, you now need to choose between playing SAFE or RISKY.

Example Game Outcome

The game is over.

As a result, your earnings are **$2857.**

Round history

Round	YOUR EARNINGS	OTHER PARTICIPANT'S EARNINGS
1	$0	$?
Round	**YOUR EARNINGS**	**OTHER PARTICIPANT'S EARNINGS**
2	$2857	$2857-$20000=**-$17143**
Round	**YOUR EARNINGS**	**OTHER PARTICIPANT'S EARNINGS**
3	$0	$?

Next

- In this example, the game is over, and therefore you get to see your own choices and earnings in all rounds, and whether or not the other participant has incurred losses of $20000.
- The **left** column always shows **your** choices and earnings:
  - In this example, you have played safe (highlighted by the green frame) in Round 1, which does not yield a reward. In Round 2 you have played risky (highlighted by the red frame), which has yielded earnings of $2857; and in Round 3 you have played safe (highlighted by the green frame), yielding no reward. Thus, your earnings in this game are $2857.
- The **right** column always shows whether or not the **other participant** has incurred a loss of $20000:
  - In this example, the other participant has played either safe or risky (highlighted by the grey frame) in Round 1 with no loss but potentially with earnings of $2857; and risky in Round 2 and incurred a loss of $17143. In Round 3, they have played either safe or risky (highlighted by the grey frame) with no loss but potentially with earnings of $2857.

Payment

In the experiment you will be making decisions that will earn you $ ("Dollars"). At the end of the experiment, the $ you have earned will be converted into Euros at an exchange rate of $ 10000 = € 1, and paid out in cash. This amount will be added to your participation payment of € 10.

Before the actual experiment starts, you will be asked to answer some questions. You must answer these correctly in order to proceed to the next question.

After completing the experiment, the computer will reveal which of the twenty-five games was randomly selected to determine your payment. You will be paid according to the choices you made in that particular game and you will be paid your earnings in cash. This protocol of determining payments suggests that you should choose in each game as if it is the only game that determines your payment.

Frequently Asked Question

Is this some kind of psychology experiment with an agenda you haven't told us? Answer. No. It is an economics experiment. If we do anything deceptive or don't pay you cash as described, then you can complain to the University of Vienna Research Ethics Board and we will be in serious trouble. These instructions are meant to clarify how you earn money, and our interest is in seeing how people make decisions.

Next

This button will be activated after 405 seconds. Please take your time to read through the instructions.

You have successfully finished reading the instructions.

The quiz, consisting of 13 questions in total, follows.

Next

Quiz Time!

Q1. The computer has randomly selected one of the games to determine your payment.

A Correct

B Incorrect

Quiz Time!

Q2. In each game (not round) you will be randomly matched with another participant in this session and stay together for that game.

A Correct

B Incorrect

Quiz Time!

Q3. In all games, the RISKY option is always **bad** and always leads to a loss of $20000 each time it is used.

A Correct

B Incorrect

Quiz Time!

Q4. If the RISKY option is **bad**, then it leads to a loss of $20000 with a chance of 25%.

A Correct

B Incorrect

Quiz Time!

Q5. If you knew the RISKY option to be **bad**, you would be best off by always choosing the SAFE option

A Correct

B Incorrect

Quiz Time!

Q6. If you experienced a loss of $20000 in an earlier round, you would be best off by always choosing the _____ .

A SAFE option

B RISKY option

Quiz Time!

Q7. A **bad** RISKY option always leads to a loss of $20000 each time you use it.

A Correct

B Incorrect

Quiz Time!

Q8. A **good** RISKY option leads to a loss of $20000 with a 25% chance.

A Correct

B Incorrect

Quiz Time!

Q9. If the other participant has suffered a loss of $20000 , this means that your RISKY option is **bad**.

A Correct

B Incorrect

Quiz Time!

Q10. If you have experienced a loss of $20000 , this means that your RISKY option is **bad**.

A Correct

B Incorrect

Quiz Time!

Q11. If you choose the SAFE option in a given round, then you can never learn anything about your and the other participant's RISKY option in that round.

A Correct

B Incorrect

Quiz Time!

Q12. If you choose the SAFE option in a given round, then you cannot learn anything new about your and the other participant's RISKY option in that round unless the other participant experiences a loss of $20000 from choosing their RISKY option.

A Correct

B Incorrect

Quiz Time!

Q13. The more you have used the RISKY option without experiencing a loss of $20000, the more inclined you should be to choose the RISKY option in the future unless the other participant has experienced a loss of $20000 from choosing their RISKY option.

A Correct

B Incorrect

You have successfully finished the quiz.

The experiment follows. When you are ready please click "Next" to start the experiment.

Next

Round 1 of 3

You are in Game **1** of **25**

You are randomly matched with **another** participant. You will stay in the same group for all 3 rounds.

Please make your choice by clicking on one of the two buttons.

SAFE RISKY

Round 2 of 3

You are in Game **1** of **25**

EARNINGS

Round	YOUR EARNINGS	OTHER PARTICIPANT'S EARNINGS
1	**$2857**	$2857-$20000=**-$17143**

You are still matched with the **same** participant.

Please make your choice by clicking on one of the two buttons.

SAFE RISKY

Round 3 of 3

You are in Game **1** of **25**

EARNINGS

Round	YOUR EARNINGS	OTHER PARTICIPANT'S EARNINGS
1	**$2857**	$2857-$20000=**-$17143**

Round	YOUR EARNINGS	OTHER PARTICIPANT'S EARNINGS
2	**$0**	**$?**

You are still matched with the **same** participant.

Please make your choice by clicking on one of the two buttons.

SAFE RISKY

Game Outcome

The game is over.

As a result, your earnings are **$5714**.

Round history

Round	YOUR EARNINGS	OTHER PARTICIPANT'S EARNINGS
1	**$2857**	$2857-$20000=**-$17143**
Round	**YOUR EARNINGS**	**OTHER PARTICIPANT'S EARNINGS**
2	**$2857**	**$0**
Round	**YOUR EARNINGS**	**OTHER PARTICIPANT'S EARNINGS**
3	**$0**	**$0**

Next

Thank you for participating in this experiment!

Your Earnings

The randomly selected game for payment is **Game #21**. Your total earnings are therefore **$2857**.

This corresponds to **€0.29**.

In total, you earned **€10.29** in this experiment. This includes your participation payment of €10 for taking part in this experiment.