

Breakdowns in the Lab^{*}

Johannes C. Hoelzemann[†]

Nicolas Klein[‡]

August 24, 2026

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 unobservable. The treatment difference is already present in the first period, before any within-game actions or outcomes have been observed, and does not attenuate as participants gain experience with the task. 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 Heinrich Graf Hardeggsche Stiftung, SSHRC, 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, Karl Schlag, and Leeat Yariv for helpful comments.

[†] 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) analyze 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 unobservability of actions 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

¹Bergemann 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.

²Wolitzky (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.

³For 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 Achim 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) have experimentally investigated strategic experimentation under good news, and Hoelzemann, Manso, Nagaraj, and Tranchero (2024) investigate 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, unobservable actions tend 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. Importantly, the treatment difference is already present in first-period choices, before participants have observed any action or outcome within the game: risky play is 72.6% with observable actions and 55.2% with unobservable actions. The effect is also persistent: it remains virtually unchanged in later games, after participants have accumulated substantial experience with the task.

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. While observed behavior departs from the equilibrium predictions in levels, the first-period treatment effect remains large after controlling for comprehension and among participants who correctly answer the most relevant comprehension questions on their first attempt.

We close the introduction by situating the paper within the related theoretical and experimental literature on strategic experimentation.

Related Literature

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) study the case of observable outcomes with unobservable actions. In particular, Keller and Rady (2015) study the bad-news setting under observable actions, while, closest to our setting, Bonatti and Hörner (2017) study a bad-news setting where actions are not observed but outcomes are.

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 provides the first experimental investigation of breakdown-learning in a strategic setting.

A related experimental literature studies how monitoring affects strategic behavior

in repeated games. Aoyagi and Fréchette (2009) and Aoyagi, Bhaskar, and Fréchette (2019) show that changes in the information available about past play affect both the strategies adopted by participants and resulting efficiency. Relatedly, Aoyagi (2014) shows theoretically that the quality and disclosure of information about adverse events can interact with strategic incentives for costly action. Our environment differs in that information is generated endogenously through costly experimentation, making the observability of actions directly relevant for incentives to produce information.

2 The Environment

We construct a deliberately simple environment designed to isolate the effect of action observability on strategic experimentation. 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, $\hat{k}_{-i,z}$ 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*

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

⁴Subgame 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.

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

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

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 predictions 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*

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

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

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*

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

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

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*

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

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*

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

Yet, $s - p_0\lambda < 0$, and, as we have shown above, $s - p_1\lambda + (1 - p_1 + p_1(1 - \lambda))s - p_1(1 - \lambda)\lambda < 0$. 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.*

We summarize the equilibrium implications and the behavioral predictions that guide the empirical analysis in Table I.

Table I: Theoretical Implications and Behavioral Predictions by Action Observability

	Observable Actions	Unobservable Actions
Equilibrium set	Two equilibria: efficient risky play until a breakdown; always safe.	Unique equilibrium: always safe.
Efficient outcome	Supportable by the threat of reverting to always safe after a deviation.	Not supportable, since deviations are not observed.
Behavioral prediction	Higher experimentation and payoffs if participants coordinate partly on the efficient equilibrium.	Lower experimentation and payoffs due to strategic free-riding.
Dynamics after no breakdown	After histories with additional risky-arm uses and no breakdown, the risky arm becomes myopically more attractive in both treatments.	

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.

Following additional risky-arm uses without a breakdown, players become more optimistic about the risky arm, so that myopic players become increasingly inclined to choose it as the game progresses. We therefore examine whether risky play increases over time.

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, all participants were required to answer 13 comprehension questions correctly before proceeding to the main part of the experiment. The experiment was programmed in oTree (Chen, Schonger, and Wickens 2016). We recruited 104 participants, with an average of 13 participants per session. Each treatment had 52 participants. Sessions lasted approximately 50 minutes. Participants received a €10 show-up fee, while one randomly selected game determined an additional variable payment that could be positive or negative. At the experimental conversion rate of 10,000 points = €1, participants earned €10.57 on average. The lowest possible total payment was €4.86, however, no participant in the experiment received the minimum payment. All payments were made in cash and in euros. The instructions and experimental interface are reproduced in the Supplemental Appendix.

3.2 Implementation

Before data collection, we generated 25 realization sets by drawing the risky-arm type and, conditional on a bad risky arm, the breakdown realizations for each player and period. Because all stochastic processes are Bernoulli, drawing these realizations in

advance is equivalent to drawing them as play unfolds. The 25 realization sets were fixed in advance and used in both treatments. These corresponded to 25 different games that each of our participants played. 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 unobservable 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.

Breakdowns were publicly observed in both treatments, only the observability of actions varied across treatments. As illustrated in Figure 1, in the treatment with observable actions, participants 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.^{6,7}

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 organize the analysis as follows. We first present summary statistics, highlighting both the average experimentation intensity, $\frac{1}{2\tau} \sum_{t=1}^{\tau} (k_{1,t} + k_{2,t})$, where τ is the period of the first breakdown in the group (with $\tau = 3$ if no breakdown occurs) and $k_{i,t} = 1$ if player i played risky in period t and $k_{i,t} = 0$ otherwise, and the overall group payoffs. Following this, we examine the aggregate experimental outcomes, with an initial focus on the distribution of experimentation intensities and group payoffs. We then assess efficiency by comparing observed behavior to the theoretical efficient so-

⁶Screenshots can be found in the Supplemental Appendix.

⁷We did not elicit beliefs or risk preferences. This was a deliberate design choice to keep the experiment focused on the strategic experimentation task. In addition, prior studies in similar bandit environments with good news did not detect a statistically significant effect of risk preferences on behavior (Hoelzemann and Klein 2021; Hoelzemann, Manso, Nagaraj, and Tranchero 2024; or Hudja and Woods 2025). We therefore interpret our history-dependence results below as revealed-choice evidence on how participants condition behavior on available histories, rather than as direct evidence on belief updating.

Example 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

Example 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

Example Round 2 of 3

You are in Game 1 of 25

EARNINGS

Round	YOUR EARNINGS	OTHER PARTICIPANT'S EARNINGS
1	\$0	\$2857

You are still matched with the same participant. Please make your choice by clicking on one of the two buttons.

SAFE

RISKY

Example Round 2 of 3

You are in Game 1 of 25

EARNINGS

Round	YOUR EARNINGS	OTHER PARTICIPANT'S EARNINGS
1	\$0	\$?

You are still matched with the same participant. Please make your choice by clicking on one of the two buttons.

SAFE

RISKY

Example Round 3 of 3

You are in Game 1 of 25

EARNINGS

Round	YOUR EARNINGS	OTHER PARTICIPANT'S EARNINGS
1	\$0	\$2857
2	\$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

Example Round 3 of 3

You are in Game 1 of 25

EARNINGS

Round	YOUR EARNINGS	OTHER PARTICIPANT'S EARNINGS
1	\$0	\$?
2	\$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

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	\$2857
2	\$2857	\$2857-\$20000=-\$17143
3	\$0	\$0

Next

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	\$?
2	\$2857	\$2857-\$20000=-\$17143
3	\$0	\$?

Next

Left: The *Observable Actions* treatment; right: The *Unobservable Actions* treatment.

Figure 1: Experimental Implementation

lution, study how behavior relates to our equilibrium predictions, and examine the dynamic evolution of behavior within games. Finally, we analyze history dependence in risky play, heterogeneity in experimentation behavior, and the robustness of the main treatment effects.

4.1 Experimentation and Payoffs

As outlined in Section 2, we anticipate that average experimentation intensity and group payoffs will be higher when actions are observable.

Initial Experimentation We first examine all 2,600 first-period decisions, without conditioning on whether a breakdown occurs later in the game. Participants choose the risky arm in 72.6% of first-period decisions when actions are observable and 55.2% when actions are unobservable, a difference of 17.4 percentage points. The difference is already present before participants observe any action or outcome within the game and is statistically significant using matching-group clustered inference. Moreover, it does not attenuate with experience: the treatment difference is 18.3 percentage points in games 1–12 and 16.6 percentage points in games 13–25. The result is also robust to clustering at the participant level; Supplemental Appendix A.1 reports both clustering schemes.


We measure experimentation intensity for each player up until the moment a first breakdown occurs to any player in the group. Figure 2 presents the observed mean experimentation intensities and average total payoffs by action observability.

The figure shows pronounced positive treatment differences for both outcomes. In OLS treatment-difference regressions with standard errors clustered at the matching-group level, observable actions increase experimentation intensity by 0.156 (s.e. 0.031, $p < 0.001$) and group payoffs by 2,813 (s.e. 319, $p < 0.001$).

Participants therefore pull the risky arm considerably more often when actions are observable, and this translates into markedly higher group payoffs.


Figure 3 illustrates the empirical distribution of experimentation intensities and group payoffs across the different treatments.

As can be seen in Figure 3, the empirical distribution functions for both experimentation intensity and group payoffs lie everywhere weakly below the corresponding distributions under unobservable actions. Thus, the sample distributions are consistent with first-order stochastic dominance of observable actions over unobservable actions for both outcomes. Formal Barrett–Donald tests of first-order stochastic dom-



The figure reports mean experimentation intensity and mean payoffs by action observability. Payoffs are shown in thousand points. Error bars indicate 95% confidence intervals based on standard errors clustered at the matching-group level.

Figure 2: Experimentation Intensity and Payoffs by Action Observability



The sample cumulative distribution functions for experimentation intensity and payoffs are shown, by action observability. The two-player group–game is the unit of observation, with 650 group–game observations per treatment (1,300 in total); see Supplemental Appendix A.6.

Figure 3: Empirical CDFs of Experimentation Intensities and Payoffs

inance do not reject dominance of observable actions over unobservable actions for either outcome, while rejecting the reverse dominance relation for both experimen-

tation intensity and payoffs.⁸

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 played. In particular, the unconditional first-period treatment difference reported above is present before any within-game actions or outcomes are observed, so it cannot be generated by differential within-game learning histories.

We also examine behavior after a breakdown, when playing safe is a dominant action. In approximately 68% of games with a breakdown, all players subsequently play safe. Equivalently, in 31.9% of games with a breakdown, at least one player chose the risky arm after a breakdown had already occurred. At the player-period level, participants chose risky after a previous breakdown in 30.0% of such observations, with a rate of 35.4% under observable actions and 24.4% under unobservable actions. We interpret these choices as post-breakdown dominance violations rather than direct evidence of misunderstanding. Among participants who answered the corresponding post-breakdown comprehension question correctly on their first attempt, the post-breakdown risky-choice rate remains 29.2%. In the restricted sample of games with optimal post-breakdown behavior, both experimentation intensity and payoffs remain higher when actions are observable, and these differences are also highly significant, with p -values of 0.001.

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


⁸For the null that observable actions first-order stochastically dominate unobservable actions, the Barrett–Donald test statistic is 0.000 for both experimentation intensity and payoffs, with bootstrap p -values of 1.000. For the reverse dominance relation, the corresponding test statistics are 5.131 for experimentation intensity and 3.384 for payoffs, with bootstrap $p < 0.001$ in both cases.

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 use the risky arm *less* than the efficient benchmark in both treatments.

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 4, 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 action observability.

Figure 4: 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 unobservable 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.

A complementary empirical-best-response exercise, reported in Supplemental Appendix A.2, confirms that observed play should not be interpreted as literal implementation of the equilibrium strategies.

4.4 Dynamic Evolution of Behavior

We now shift our focus to the dynamics of observed behavior. We are particularly interested in whether 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 5, 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 action observability.

Figure 5: Experimentation Intensities over Time

Although our equilibrium predictions do not imply gradual adjustment within a game, the risky arm becomes myopically more attractive following additional risky choices without a breakdown. We therefore examine descriptively whether risky play changes across periods. At the beginning of a game without any breakdowns, participants are indeed less likely to choose the risky arm than in later periods. To exam-

ine changes in behavior over time, we employ paired two-sided t-tests for parametric analysis and two-sided Wilcoxon signed-rank tests for non-parametric analysis, treating each group–game as a paired observation across periods. In both treatments, first-period risky play is significantly lower than both second-period and third-period risky play, while risky play does not differ significantly between periods 2 and 3. In the treatment with observable actions, paired t-tests (two-sided Wilcoxon signed-rank tests) produce p -values of $p_{12} = 0.002$ (0.002), $p_{13} = 0.039$ (0.034), and $p_{23} = 0.278$ (0.278), where $p_{\alpha\beta}$ is the p -value when comparing periods α and β . With unobservable actions, by contrast, we find p -values of $p_{12} < 0.001$ (< 0.001), $p_{13} = 0.001$ (0.001), and $p_{23} = 0.500$ (0.470), 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.

To assess whether the treatment difference is driven primarily by initial behavior or by differential evolution across rounds, we also compare treatment differences period by period. In games without breakdowns, experimentation is higher when actions are observable in every period. The treatment difference is already pronounced in period 1: experimentation intensity is 0.724 under observable actions and 0.538 under unobservable actions, a difference of 0.186. The corresponding differences are 0.165 in period 2 and 0.157 in period 3. Thus, the treatment effect is not driven mainly by divergent behavior over the course of the game; rather, a substantial gap is present from the first period and persists throughout the interaction. Because this subsection conditions on games without a subsequent breakdown, we do not use these period-specific comparisons as our primary evidence on initial behavior, the unconditional first-period analysis above does not condition on later outcomes.

We obtain the same conclusion from a linear probability model in which risky play is regressed on the observable-actions treatment indicator, period indicators, and interactions between observable actions and the period indicators. In this specification, the observable-actions coefficient captures the treatment difference in period 1, while the interaction terms capture whether the treatment difference changes in periods 2 and 3. The observable-actions coefficient is positive and statistically significant, whereas the interaction terms are small and not statistically significant. Hence, the treatment difference is already present in the first period and does not appear to be driven by differential evolution across periods.

We summarize these results as follows:

Risky Play Conditional on no breakdown having occurred, participants use the risky arm significantly more in periods 2 and 3 than in period 1 in both treatments, while risky play does not differ significantly between periods 2 and 3.

4.5 History Dependence in Risky Play

To better understand the dynamics underlying risky play, we examine whether participants condition their choices on their own and their partner's previous experimentation within a game. We estimate linear probability models in which the dependent variable is an indicator equal to one if the participant chooses the risky arm in the current period. The sample consists of period-2 and period-3 choices in histories in which no breakdown had previously occurred. The main explanatory variables are the cumulative number of previous risky choices by the participant, the cumulative number of previous risky choices by the partner, an indicator for observable actions, and a period-3 indicator. Table II reports the results.

The results show that current risky play is systematically associated with the history of play. Own previous experimentation is a strong predictor of current experimentation. Partner experimentation is also positively associated with current risky play. Because partner actions are unobserved in one treatment, however, the pooled partner-history coefficient should not be interpreted as a causal response to observed partner behavior. The observable-action treatment effect remains positive after controlling for these history variables. The negative period-3 coefficient conditional on prior experimentation is reported as a descriptive pattern rather than as a direct prediction of the equilibrium model.

The estimates also speak to whether participants respond similarly to their own and their partner's experimentation histories. The coefficient on own prior risky choices is larger than the coefficient on partner prior risky choices in all specifications. Wald tests reject equality in Columns (1)–(3) ($p = 0.004$, $p = 0.002$, and $p = 0.006$, respectively), though not in the individual fixed-effects specification ($p = 0.305$). Thus, while partner experimentation is positively associated with current risky play, participants' own experimentation histories appear to be the stronger predictor of subsequent behavior. This is natural in our setting. In the observable-actions treatment, participants observe both their own and their partner's previous actions and outcomes, so partner experimentation can directly affect subsequent choices. In the

Table II: History Dependence in Risky Play

	Dependent Variable: Current Risky Choice			
	(1) Baseline	(2) Interaction	(3) No-Breakdown	(4) Individual FE
Observable actions	0.066** (0.032)	0.037 (0.052)	0.076** (0.034)	
Own prior risky choices	0.256*** (0.048)	0.253*** (0.048)	0.265*** (0.050)	0.135*** (0.047)
Partner prior risky choices	0.063*** (0.024)	0.049* (0.026)	0.069** (0.027)	0.069*** (0.022)
Observable actions $\times$ partner prior risky choices		0.032 (0.039)		
Period 3	-0.245*** (0.027)	-0.246*** (0.028)	-0.251*** (0.028)	-0.163*** (0.028)
Individual fixed effects	No	No	No	Yes
Observations	4,106	4,106	3,280	4,106
R^2	0.160	0.161	0.177	0.309

Notes: The table reports OLS linear probability models. The dependent variable is an indicator equal to one if the participant chooses the risky arm in the current period. Columns (1), (2), and (4) use period-2 and period-3 observations in histories in which no breakdown had previously occurred in the group. Column (3) restricts the sample to games in which no breakdown occurred during any of the three periods. Own prior risky choices and partner prior risky choices are cumulative counts of risky choices before the current period. The observable-actions indicator is absorbed by the individual fixed effects in Column (4). Standard errors, clustered at the matching-group level, are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

unobservable-actions treatment, by contrast, participants do not observe the partner's action unless a breakdown occurs. Hence, learning from others is necessarily more limited: participants can learn from their own experimentation and from observed breakdowns suffered by the partner, while the absence of a partner breakdown is a weaker signal because it does not reveal whether the partner actually experimented. Since we did not elicit beliefs, we interpret these results as revealed-choice evidence on how participants condition behavior on observed and inferable histories, rather than as direct evidence on belief updating.

To make the role of the opponent's action history more transparent, we also examine the direct relation between the opponent's period-1 action and period-2 risky play. We restrict attention to histories in which no breakdown occurred in period 1 and regress a participant's period-2 risky choice on the observable-actions treatment indicator, the participant's own period-1 action, and the opponent's period-1 action. Opponent risky play in period 1 is positively associated with the participant's period-2

risky choice: the coefficient is 0.114 (s.e. 0.036, $p = 0.006$). The corresponding coefficient on own period-1 risky play is larger, 0.265 (s.e. 0.069, $p = 0.002$). Descriptively, under observable actions, the probability of choosing risky in period 2 is 0.830 after an opponent chose risky in period 1, compared to 0.602 after an opponent chose safe. Thus, when the opponent's action is observable, participants' period-2 choices are systematically associated with the opponent's first-period action. Full regression results, including participant-clustered standard errors, are reported in Supplemental Appendix Table A.4.

Finally, we examine whether participants mimic the behavior of past partners across games. Using the actual chronological order of play, we regress each participant's first-period risky choice in game order o on their own first-period risky choice and the first-period risky choice of their actual partner in game order $o - 1$, the observable-actions treatment indicator, and game order. Own behavior is strongly persistent: the coefficient on own previous first-period risky play is 0.313 (s.e. 0.030, $p < 0.001$). By contrast, the coefficient on the previous partner's first-period risky choice is essentially zero and not statistically significant (-0.011 , s.e. 0.027, $p = 0.694$). The observable-actions treatment effect remains positive (0.122, s.e. 0.049, $p = 0.025$). We therefore find no evidence that participants mimic past partners across games; the treatment difference is not driven by cross-game imitation. Supplemental Appendix Table A.5 reports the full specification and participant-clustered standard errors.

4.6 Heterogeneity in Experimentation Behavior

We also examine heterogeneity in experimentation behavior across participants. We classify participants according to their risky-choice frequency in pre-breakdown histories. Participants who choose the risky arm in at most one third of such histories are classified as persistent safe players; participants who choose the risky arm in at least two thirds of such histories are classified as persistent experimenters; the remaining participants are classified as intermediate experimenters.

Table III reports the distribution of behavioral types and their average payoffs by treatment. Observable actions substantially shift the distribution of types. With unobservable actions, 21.2% of participants are persistent safe players and 42.3% are persistent experimenters. With observable actions, only 3.8% are persistent safe players, while 71.2% are persistent experimenters. The overall type distribution differs significantly across treatments ($p = 0.004$, chi-squared test). The increase in the share of

Table III: Heterogeneity in Experimentation Behavior

	Observable Actions	Unobservable Actions	Difference
<i>Panel A: Behavioral Types</i>			
Persistent safe players	3.8%	21.2%	−17.3 pp
Intermediate experimenters	25.0%	36.5%	−11.5 pp
Persistent experimenters	71.2%	42.3%	28.8 pp
<i>Panel B: Average Payoff by Type</i>			
Persistent safe players	286	−177	463
Intermediate experimenters	2,356	764	1,592
Persistent experimenters	1,946	800	1,146

Notes: Participants are classified according to their risky-choice frequency in pre-breakdown histories. Persistent safe players choose the risky arm in at most one third of such histories; persistent experimenters choose the risky arm in at least two thirds of such histories; all remaining participants are classified as intermediate experimenters. Panel B reports average payoffs per game by behavioral type and treatment. The Difference column reports Observable Actions minus Unobservable Actions. The overall type distribution differs significantly across treatments ($p = 0.004$, chi-squared test). The treatment difference in the share of persistent experimenters is significant at $p = 0.005$, and the treatment difference in the share of persistent safe players is significant at $p = 0.015$ using Fisher exact tests.

persistent experimenters is significant ($p = 0.005$, Fisher exact test), as is the decrease in the share of persistent safe players ($p = 0.015$, Fisher exact test).

Persistent safe players also earn the lowest average payoffs in both treatments. Thus, the aggregate treatment effect is reflected in a systematic shift away from persistent safe play and toward more frequent experimentation. This suggests that the higher experimentation intensity under observable actions is not driven by a small number of extreme participants, but instead corresponds to a broader shift in the distribution of experimentation behavior. Supplemental Appendix A.4 shows that the same conclusion emerges from a sequence-based classification that does not rely on the one-third/two-thirds cutoffs: in particular, all-safe paths fall from 22.5% under unobservable actions to 6.9% under observable actions. Supplemental Appendix A.5 further shows that the lower realized payoff of persistent relative to intermediate experimenters under observable actions reflects greater exposure to breakdown risk before any breakdown has been observed, rather than greater post-breakdown experimentation.

4.7 Econometric Robustness Tests

As a further robustness test, we estimate regressions that address three econometric concerns raised by the structure of the experiment: repeated observations at the participant level, dependence within matching groups, and possible order effects across the 25 games. Specifically, we estimate linear random-effects models with random effects at the participant and matching-group levels and include a weighted learning function to control for potential learning or fatigue over the course of a session.⁹ The random-effects specification captures unobserved heterogeneity across participants and matching groups, while clustering standard errors at the matching-group level accounts for dependence induced by repeated interaction and rematching within matching groups.

Table IV: Random-Effects Estimates of Experimentation Intensity and Payoffs

	Experimentation Intensity			Individual
	All	No Breakdown	Until Breakdown	Payoffs
Intercept	0.596*** (0.044)	0.586*** (0.049)	0.603*** (0.054)	418.199* (239.186)
Observable actions	0.157*** (0.052)	0.170*** (0.054)	0.280*** (0.054)	1398.080*** (174.776)
Learning	0.034 (0.026)	0.040 (0.029)	-0.082 (0.106)	1059.966 (1077.624)
σ_ϵ	0.287	0.268	0.293	8494.653
σ_u	0.216	0.212	0.243	0
Observations	2600	1248	336	2600
Between R^2	0.113	0.127	0.199	0.189

Notes: The table reports random-effects estimates. The dependent variable in Columns (1)–(3) is experimentation intensity; the dependent variable in Column (4) is individual payoffs. Observable actions is an indicator equal to one in the observable-actions treatment and zero in the unobservable-actions treatment. Learning is defined as $g_o = 1/o$, where o denotes the order in which the participant encountered the game. All regressions include random effects at the participant and matching-group level. Standard errors, clustered at the matching-group level, are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

In particular, we regress experimentation intensity and individual payoffs on the treatment dummy *observable actions*, which is 0 for the unobservable-actions treatment and 1 for the observable-actions treatment. Recall that participants played the

⁹Our results remain qualitatively unchanged in a fixed-effects specification.

25 games in random order and any order of these games that was used for participants in the observable actions sessions was replicated for participants in the unobservable actions 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 ($o \in \{1, \dots, 25\}$) corresponds to the random order in which each participant was exposed to each game. The function $\{g_o\}$ flexibly controls for non-linear trends in behavior across successive games, such as learning-by-doing or potential fatigue. 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. Comprehension performance was very similar across treatments: participants averaged 2.62 first-attempt errors under observable actions and 3.00 under unobservable actions ($p = 0.240$). Controlling for the total number of first-attempt errors leaves the unconditional first-period treatment effect essentially unchanged at 0.175 ($p = 0.012$). The result is equally pronounced among participants who demonstrated understanding of the key features most relevant for the decision problem: the first-period treatment difference is 18.1 percentage points among those who correctly answered the post-breakdown decision question on their first attempt ($p = 0.006$), and 16.6 percentage points among those who immediately recognized the informational implication of a breakdown ($p = 0.030$). To account for the fact that behavior within matching groups is not independent, we treat each matching group as the unit of statistical independence and cluster standard errors at the matching-group level. Our specification ensures that estimated treatment effects are not spuriously driven by time-dependent dynamics or unobserved fixed components, but instead identify the effect of action observability on experimentation intensity and payoffs. This also mitigates the concern that behavior might be driven by session-specific coordination on particular strategies. In a given session, separate matching groups were assigned to the bad-news observable-actions and bad-news unobservable-actions treatments. Clustering at the matching-group level therefore accounts for dependence within the relevant interaction group, and the results are robust to including session controls. The treatment effect remains positive and statistically significant. Supplemental Appendix A.1 additionally reports participant-clustered inference for the principal treatment effects and

the history-dependence specifications, the substantive conclusions are unchanged.

Table IV lists the results from this analysis.

The coefficient on the learning variable is not statistically significant in any specification, and the estimated effect of observable actions remains positive and significant throughout. We therefore find no evidence that learning during the session drives our main results or that the main sample should be restricted to later games. We find a strong positive effect of observable actions on experimentation intensity across all games, games without breakdowns, games with breakdowns before the last period, and payoffs.

5 Concluding Remarks

This paper offers the first experimental investigation of *breakdown*-learning in a strategic setting. Observed behavior is consistent with strategic free-riding and shows that action observability substantially affects information production in a bad-news learning environment. The treatment difference appears already in first-period choices, before participants observe any within-game actions or outcomes, and remains similar after substantial experience with the task. Observable actions also generate substantially higher overall experimentation and payoffs. At the same time, observed behavior does not literally coincide with the equilibrium predictions: participants experiment substantially with unobservable actions, and we observe post-breakdown dominance violations. The first-period treatment effect, however, is essentially unchanged after controlling for comprehension and remains large among participants who demonstrate understanding of the relevant informational structure. We therefore interpret the evidence as showing that action observability materially affects experimentation in the direction predicted by the strategic-learning mechanism, rather than as evidence that participants implement the equilibrium strategies literally.

References

- ACHIM, P., AND N. KLEIN (2022): “Strategic Investment and Learning with Private Information,” *Journal of Economic Theory*, 204, 105523.
- AOYAGI, M. (2014): “Strategic Obscurity in the Forecasting of Disasters,” *Games and Economic Behavior*, 87, 485–496.
- AOYAGI, M., V. BHASKAR, AND G. R. FRÉCHETTE (2019): “The Impact of Monitoring in Infinitely Repeated Games: Perfect, Public, and Private,” *American Economic Journal: Microeconomics*, 11(1), 1–43.
- AOYAGI, M., AND G. FRÉCHETTE (2009): “Collusion as Public Monitoring Becomes Noisy: Experimental Evidence,” *Journal of Economic Theory*, 144(3), 1135–1165.
- 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 Street-light 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>.
- HUDJA, S., AND D. WOODS (2025): “Examining Experimentation: An Elicitation Approach,” *Working Paper*.
- 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.
- WOLITZKY, A. (2018): “Learning from Others’ Outcomes,” *American Economic Review*, 108(10), 2763–2801.

Breakdowns in the Lab

Supplemental Appendix: Additional Results

Johannes C. Hoelzemann Nicolas Klein

August 24, 2026

A Additional Robustness and Behavioral Diagnostics

This appendix provides additional analyses addressing inference with repeated observations, best responses to the empirical distribution of play, two auxiliary regressions discussed in the main text, and alternative descriptions of behavioral heterogeneity. Unless otherwise noted, standard errors are clustered at the matching-group level. There are 16 matching groups and 104 participants in the bad-news sample.

A.1 Alternative Clustering at the Participant Level

The main analysis clusters standard errors at the matching-group level, which is the level at which participants are repeatedly rematched. As an additional robustness check, Table A.1 reports participant-clustered standard errors for the principal individual-level treatment effects. The point estimates are unchanged. In particular, the unconditional first-period treatment effect remains 0.174 and is statistically significant under either clustering scheme.

Table A.2 repeats the history-dependence specifications from the main text using standard errors clustered at the participant level. The qualitative pattern is unchanged: own prior experimentation is the strongest predictor of current risky play, partner experimentation remains positively associated with current play, and the interaction between observability and partner experimentation remains statistically insignificant.

A.2 Best Responses to the Empirical Distribution of Play

We next ask whether the observed level of experimentation can be rationalized as a best response to the empirical frequency with which opponents experiment. This is a

Table A.1: Treatment Effects under Alternative Clustering Schemes

Outcome	Coefficient	MG-clustered s.e.	Participant-clustered s.e.	MG p	Participant p
First-period risky choice	0.174	0.067	0.054	0.020	0.002
Individual experimentation intensity	0.157	0.052	0.044	0.008	< 0.001
Individual payoff (points)	1,398	175	287	< 0.001	< 0.001

Notes: The first row uses all 2,600 first-period decisions. The second and third rows use 2,600 participant–game observations and control for the learning variable used in Table IV of the main text. “MG” denotes matching group. The coefficient is the effect of observable actions. Matching-group p -values use the finite-cluster t reference distribution; participant-clustered inference uses 104 participant clusters.

Table A.2: History Dependence with Participant-Clustered Standard Errors

	(1)	(2)	(3)	(4)
Observable actions	0.066 (0.036)	0.037 (0.051)	0.076 (0.037)	
Own prior risky choices	0.256 (0.026)	0.253 (0.026)	0.265 (0.027)	0.135 (0.027)
Partner prior risky choices	0.063 (0.017)	0.049 (0.020)	0.069 (0.018)	0.069 (0.014)
Observable $\times$ partner prior risky		0.032 (0.030)		
Period 3	-0.245 (0.024)	-0.246 (0.024)	-0.251 (0.025)	-0.163 (0.023)
Observations	4,106	4,106	3,280	4,106
Individual fixed effects	No	No	No	Yes

Notes: Participant-clustered standard errors are in parentheses. Columns (1), (2), and (4) use period-2 and period-3 decisions following histories in which no breakdown had previously occurred. Column (3) uses the restricted sample employed in the corresponding specification in the main text. The observable-actions indicator is absorbed by the individual fixed effects in Column (4).

descriptive empirical-best-response exercise rather than an equilibrium test. We treat empirical opponent experimentation frequencies as plug-in beliefs and let the focal player optimize given those frequencies. In period 3 the calculation is exact because the decision is myopic. In earlier periods we solve backwards, taking the opponent's history-contingent empirical experimentation frequency as given.

For reference, under the efficient observable-actions equilibrium the expected payoff gain from following the prescribed period-1 risky action rather than deviating to safe and inducing reversion to safe play is 55.1 points.

Table A.3 reports the first-period empirical-best-response calculation. At the empirical opponent experimentation rates, safe yields the higher expected payoff in period 1 in both treatments. The risky-minus-safe payoff difference is approximately -150 points under observable actions and -196 points under unobservable actions. A matching-group bootstrap yields 95% intervals of approximately $[-181, -121]$ points and $[-249, -135]$ points, respectively.

Table A.3: Period-1 Best Response to Empirical Opponent Play

Treatment	Opponent risky rate	$U(R) - U(S)$	Best response
Observable actions	0.726	-150	Safe
Unobservable actions	0.552	-196	Safe

Notes: The calculation uses the empirical period-1 experimentation rate of the opponent and history-specific empirical period-2 experimentation rates. The focal player optimizes in periods 2 and 3. The exercise describes best responses to the empirical distribution of play and does not impose equilibrium consistency on the opponent's strategy.

The later-period calculations illustrate why substantial experimentation can nevertheless be individually optimal after informative histories. Under observable actions in period 2, risky is optimal after any first-period history containing at least one risky choice at the empirical continuation frequencies, but not after a history in which both players chose safe. The risky-minus-safe payoff differences for first-period action histories (own, other) = (0, 0), (0, 1), (1, 0), (1, 1) are approximately -463 , 12, 22, and 384 points, respectively. Under unobservable actions, the corresponding period-2 difference is approximately -217 points after the participant chose safe in period 1 and 243 points after the participant chose risky.

In period 3, with observable actions, risky is myopically optimal once at least two risky-arm uses have been observed without a breakdown, exactly as implied by the theoretical analysis. The risky-minus-safe payoff differences after $n = 0, 1, 2, 3, 4$ successful observed risky uses are approximately -523 , -194 , 157, 516, and 869 points.

With unobservable actions, using empirical opponent experimentation frequencies to integrate over the unobserved action history gives approximate period-3 differences of -158 , 194 , and 553 points after the participant has personally made zero, one, or two previous risky choices without a breakdown.

These calculations reinforce the interpretation in the main text. The treatment comparison is strong, but observed play should not be interpreted as literal implementation of the equilibrium strategies: at empirical opponent frequencies, risky is not a first-period best response in either treatment, while it becomes a best response after sufficiently informative no-breakdown histories.

A.3 Additional Regressions Reported in the Main Text

Table A.4 reports the regression relating period-2 risky play to the participant's own and the opponent's period-1 actions. The sample is restricted to histories with no period-1 breakdown. The positive coefficient on the opponent's period-1 risky choice is robust to either matching-group or participant-level clustering.

Table A.4: Opponent's Period-1 Action and Period-2 Risky Play

	Coefficient	MG-clustered s.e.	Participant-clustered s.e.
Observable actions	0.073	0.043	0.043
Own period-1 risky choice	0.265	0.069	0.043
Opponent period-1 risky choice	0.114	0.036	0.028
Observations	2,194		
R^2	0.120		

Notes: OLS linear probability model. The dependent variable is the participant's period-2 risky choice. The sample contains period-2 decisions following a first period without a breakdown. Matching-group clustered p -values for own and opponent period-1 risky choices are 0.002 and 0.006, respectively; participant-clustered p -values are both below 0.001.

Table A.5 reports the cross-game persistence specification using the actual chronological order of play. For each participant, the first-period risky choice in game order o is regressed on the participant's own first-period choice and the first-period choice of the participant's actual partner in game order $o - 1$, together with the observable-actions indicator and game order. Own behavior is strongly persistent, whereas there is no evidence of imitation of the previous partner.

Table A.5: Cross-Game Persistence and Imitation

	Coefficient	MG-clustered s.e.	Participant-clustered s.e.
Own previous first-period risky choice	0.313	0.030	0.040
Previous partner's first-period risky choice	-0.011	0.027	0.022
Observable actions	0.122	0.049	0.038
Game order	0.0008	0.0011	0.0009
Observations	2,496		
R^2	0.127		

Notes: OLS linear probability model. The dependent variable is the participant's first-period risky choice in game order o . The chronological game-order variable is used to construct the lag. With matching-group clustering, the p -values are $p < 0.001$ for own previous risky choice, $p = 0.694$ for the previous partner's risky choice, $p = 0.025$ for observable actions, and $p = 0.464$ for game order. Participant-clustered inference gives the same substantive conclusion.

A.4 Behavioral-Pattern Decomposition

The main text classifies participants using their overall pre-breakdown experimentation frequency. As a complementary analysis that does not rely on arbitrary participant-level cutoffs, we classify each participant–game according to the sequence of actions taken before the first breakdown. An *all-risky* path consists exclusively of risky choices up to the first breakdown (or the end of the game); an *all-safe* path consists exclusively of safe choices; and the remaining categories distinguish the direction of within-game switching. We then compute each participant's share of games in each category and compare these participant-level shares across treatments.

Table A.6: Behavioral Patterns within Games

Pattern	Observable	Unobservable	Difference	MG-clustered p
All risky until breakdown/end	0.550	0.418	0.132	0.129
All safe until breakdown/end	0.069	0.225	-0.155	0.001
Safe $\rightarrow$ risky	0.165	0.173	-0.008	0.861
Risky $\rightarrow$ safe	0.118	0.092	0.025	0.352
Multiple switches	0.098	0.092	0.007	0.741
Post-breakdown safe-choice rate	0.655	0.776	-0.121	0.088

Notes: The first five rows partition participant–games according to the sequence of choices made before the first breakdown. Treatment means are the averages of participant-level shares across the 25 games. The final row reports the participant-level fraction of post-breakdown decision opportunities in which the safe arm is chosen. Inference clusters by matching group.

The main treatment difference in this decomposition is a sharp reduction in all-

safe paths: such paths account for 22.5% of games under unobservable actions but only 6.9% under observable actions. The shares of games involving within-game switching are very similar across treatments. Thus, the aggregate treatment effect is not simply an artifact of the particular one-third/two-thirds thresholds used for the participant typology in the main text.

A.5 Why Persistent Experimenters Earn Less than Intermediate Experimenters

Table A.7 decomposes the realized payoffs of intermediate and persistent experimenters in the observable-actions treatment using the classification in the main text. Persistent experimenters generate more successful risky-arm payoffs, but they are also exposed to substantially more breakdowns. Because a breakdown is much more costly than the gain from a successful risky choice, the additional realized breakdown losses more than offset the additional successful risky payoffs in this sample.

Table A.7: Payoff Decomposition by Behavioral Type, Observable Actions

	Intermediate	Persistent experimenters	Difference
Average payoff per game (points)	2,356	1,946	-410
Successful risky choices per game	1.526	2.037	0.511
Breakdowns per game	0.117	0.226	0.109
Breakdowns before any prior breakdown	0.108	0.218	0.111
Breakdowns after a prior breakdown	0.009	0.008	-0.002
Post-breakdown risky choices per game	0.126	0.163	0.037
Post-breakdown risky-choice rate	0.334	0.352	0.018

Notes: A successful risky choice pays 2,857 points, while a breakdown yields $2,857 - 20,000 = -17,143$ points. Persistent experimenters obtain approximately 0.511 more successful risky choices per game, worth about 1,459 additional points, but experience approximately 0.109 more breakdowns per game, contributing about 1,869 additional points of losses. The net difference is therefore approximately -410 points per game. The extra breakdown exposure occurs almost entirely before any breakdown has previously been observed, rather than through post-breakdown experimentation.

This decomposition clarifies that the lower realized payoff of persistent experimenters in the observable-actions treatment is not driven by a greater tendency to continue experimenting after a known breakdown. Instead, persistent experimenters simply expose themselves more often to the risky arm before a breakdown, producing both more successful risky outcomes and more breakdown losses; in the realized

sample, the latter dominate.

A.6 Unit of Observation in the Empirical CDFs

The empirical CDFs in Figure 3 of the main text use the two-player group–game as the unit of observation. There are 650 group–game observations in each treatment (1,300 in total). Experimentation intensity and group payoff are computed for the same group–game observation. Because participants are rematched within matching groups, inferential procedures account for matching-group dependence; the CDF figure itself is descriptive.

Breakdowns in the Lab

Supplemental Appendix: Experimental Interface

The Environment with Observable Actions

Johannes C. Hoelzemann Nicolas Klein

August 24, 2026

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

Example 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**

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

Example Round 2 of 3

You are in Game 1 of 25

EARNINGS

Round	YOUR EARNINGS	OTHER PARTICIPANT'S EARNINGS
1	 \$0	 \$2857

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

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

SAFE **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.

Example Round 3 of 3

You are in Game 1 of 25

EARNINGS

Round	YOUR EARNINGS	OTHER PARTICIPANT'S EARNINGS
1	\$0	\$2857
2	\$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

- 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**.

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	 \$2857
2	 \$2857	 \$2857 - \$20000 = -\$17143
3	 \$0	 \$0
Next		

- 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

Supplemental Appendix: Experimental Interface

The Environment with Unobservable Actions

Johannes C. Hoelzemann Nicolas Klein

August 24, 2026

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

Example 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.

SAFERISKY

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

Example Round 2 of 3

You are in Game 1 of 25

EARNINGS

Round	YOUR EARNINGS	OTHER PARTICIPANT'S EARNINGS
1	 \$0	 \$?

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

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

SAFE 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**.

Example Round 3 of 3

You are in Game 1 of 25

EARNINGS

Round	YOUR EARNINGS	OTHER PARTICIPANT'S EARNINGS
1	 \$0	 \$?
Round	YOUR EARNINGS	OTHER PARTICIPANT'S EARNINGS
2	 \$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

- 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	 \$?
2	 \$2857	 $\$2857 - \$20000 = -\$17143$
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.



Instructions



Quiz



Experiment



Questionnaire

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.