# Bandits in the Lab 

Johannes Hoelzemann ${ }^{\dagger} \quad$ Nicolas Klein ${ }^{\ddagger}$

This version: January 26, 2018


#### Abstract

We test Keller, Rady, Cripps' (2005) game of strategic experimentation with exponential bandits in the laboratory. We find strong support for the prediction of free-riding because of strategic concerns. We also find strong evidence for behavior that is characteristic of Markov perfect equilibrium: non-cutoff behavior, lonely pioneers and frequent switches of action.


Keywords: Strategic Experimentation, Exponential Bandits, Learning, Dynamic Games, Markov Perfect Equilibrium, Continuous Time, Laboratory Experiments, Eye Tracking.

JEL Classification Numbers: C73, C92, D83, O32

[^0]
## 1 Introduction

Innovation and social learning are often the work of pioneers, who, by bearing the costs of experimenting with a new approach, create informational spill-overs for others. Whether we consider R\&D, resource exploration, or the testing of a new drug, the information produced by a relatively small set of agents benefits a much larger group of agents. Indeed, R\&D is universally recognized as an important factor of economic growth (Romer, 1990; Grossman \& Helpman, 1991). An economy's productivity level depends on innovation, which is driven by knowledge emerging from cumulative R\&D experience as well as an economy's overall knowledge stock (Griliches, 1988; Coe \& Helpman, 1995). It is thus important for economists to analyze pioneers' incentives for information production in the presence of informational spill-overs.

Multi-armed bandit models have become canonical in economics to study information producers' dynamic trade-offs. At each point in time, a decision maker either optimally exploits the information he already has, or he decides to invest in exploration in order to make better future decisions. Until fairly recently, the literature focussed on the trade-off of an individual decision maker acting in isolation. Bolton \& Harris (1999) and Keller, Rady, Cripps (2005; subsequently: KRC) have extended the individual choice problem to a multi-player continuous-time framework. There now appears a strategic component to the information-acquisition problem, in that other players now also benefit from the information acquired at a cost by a given player. To make the problem tractable, these papers are focussing on the choice between a safe arm, yielding a known payoff, and a risky arm, which yields payoffs following a stochastic process. The time-invariant quality of this risky arm can be good or bad. If it is good (bad), it dominates (is dominated by) the safe arm. Whether the risky arm is good or bad is initially unknown and can only be found out by trying it out over time. Trying it out is costly, however, as it means forgoing the safe payoff. As the quality of the risky arm is assumed to be the same across players, and players can observe each other's actions and payoffs, there is a positive informational externality associated with a player's use of the risky arm. This gives rise to a dynamic public-good problem, where the public good in question is the dynamically evolving information about agents' common state of the world.

While the game-theoretical analysis of these problems will lead to multiple equilibria, it has nonetheless yielded many sharp qualitative behavioral predictions. Yet, empirical evidence for these predictions has thus far been scarce. Indeed, the dynamic nature of the problem and the continuous-time setting underlying its theoretical analysis raise some challenges both for the collection of field data and the experimental
implementation in the laboratory. To the best of our knowledge, we are the first to implement an experimental test of continuous-time strategic-experimentation models in the laboratory. Our goal in doing so is twofold. Firstly, we want to test whether the bandit models correctly predict agents' behavior "in the model", by making our subjects face a setting closely resembling KRC's. This is of course a necessary, but by no means sufficient, condition for us to have "the right model" to approach these questions with. Secondly, we aim to shed some light on which of the multiple equilibria seem best-suited to capture actual behavior.

Our analysis relies on comparing the behavior of our experimental subjects in groups where the quality of the risky arm was known to be the same for all partners (which we call the strategic treatment) to that of groups where its quality was iid across members, the control treatment. When the quality of the risky arm is known to be the same across players, rational agents will take into account the result of their partners' experimentation when updating their beliefs. As they can learn from what others are doing, they have an incentive to induce others to behave in certain ways so they may learn from it. There is thus some strategic interaction across players, even though a player's payoffs depends only on his own action and the common state of the world, i.e., there are no payoff externalities.

Specifically, we use the simplest formalization of the continuous-time strategicexperimentation framework, KRC's exponential-bandit setup, as our theoretical benchmark. In this setting, a bad risky arm never yields any payoff, while a good risky arm gives lump-sum payoffs at the jumping times of a Poisson process. Thus, whenever the risky arm is used without a success, players gradually grow pessimistic about its quality; as soon as they observe a success, they know for sure that the risky arm is good.

KRC analyze Markov perfect equilibria (MPE) with the players' common posterior belief as a state variable. ${ }^{\square}$ While there is a continuum of MPE, all equilibria make two fundamental qualitative predictions regarding players' behavior: As information is a public good, players will produce too little of it. Furthermore, it is predicted that all players will not use a simple cutoff strategy in equilibrium. A cutoff strategy is defined by a unique threshold belief above which it prescribes risky play, while prescribing safe play below it.

Hörner, Klein, and Rady (2014; subsequently: HKR) analyze non-Markovian equilibria, i.e. perfect Bayesian equilibria (PBE) in which a player's action choice can depend on the history in more complex ways. They show that free-riding prevails in all PBE

[^1]as well. Moreover, the average-payoff maximizing PBE is strongly symmetric and has a particularly simple structure: Players play a cutoff strategy (on the path of play), applying the same cutoff as a single agent.

Our empirical tests are designed to contrast the qualitative predictions of MPE with those of the best PBE, which is a natural candidate for a focal equilibrium both because it maximizes players' average equilibrium payoffs and because it has a particularly simple structure. In a first step, we show that the informational externality indeed impacts subjects' behavior: the average experimentation intensity is lower, and subjects' payoffs are higher, in the strategic treatment. Secondly, we find strong evidence of the kind of qualitative behavior predicted by MPE as opposed to the simpler behavior predicted by the best PBE, with players' adopting more sophisticated behaviors than cutoff strategies in the strategic treatment. Indeed, players switch much more between safe and risky, and use cutoff strategies much less frequently, than they do in the control treatment. Moreover, there is a larger proportion of time during which exactly one player is playing risky in the strategic treatment. All this is fully consistent with the players' switching between the roles of pioneer and free-rider, which characterizes equilibrium play at intermediate beliefs in KRC and differentiates it e.g. from the best PBE in HKR.

Our game is of course very complicated, so that we cannot reasonably expect subjects to be able to compute equilibrium strategies. Nonetheless, we are documenting behavior that is very much in line with the sophisticated coordination required by MPE play, as opposed, e.g., to the simpler structure of the best PBE. Yet, we of course cannot conclude that our experimental subjects fully adopt equilibrium behavior, rather than simpler heuristics that make them behave in ways that are suggestive of MPE behavior.In fact, we proceed to subdivide the players' beliefs into a region where risky is the dominant action choice, a region where safe and risky are mutually best responses and one where safe is a dominant action, and find no striking qualitative differences in players' behavior across these regions. Thus, while our subjects adopt some of the qualitative aspects of equilibrium behavior, they do not seem strictly to separate these different strategic regions. Furthermore, they will often extend experimentation even below the single-agent cutoff in both the strategic and the control treatments, which leads us to conjecture that subjects are not able to compute beliefs and cutoffs precisely. Indeed, even in the region where safe is the dominant action, the average experimentation intensity is higher in the control treatment, which is inconsistent with the presence of an encouragement effect. ${ }^{[ }$

[^2]The rest of the paper is organized as follows: Section 2 reviews some of the related literature; Section 3 explains the KRC model in more detail; Section 4 sets out our experimental implementation; Section 5 discusses our findings, and Section 6 concludes. The Appendix exhibits and explains the interface our experimental subjects were using, and reproduces the instructions the subjects received.

## 2 Literature Review

The bandit problem as a stylized formalization of the trade-off between exploration and exploitation goes back to Thompson (1933) and Robbins (1952). It was subsequently analyzed, amongst others, by Bellman (1956) and Bradt, Johnson, Karlin (1956). Its first application to economics was in Rothschild (1974), who analyzed the price-setting problem of a firm facing an unknown demand function. Gittins \& Jones (1974) showed that, if arms are stochastically independent of each other and the state of only one arm can evolve at any one time, an optimal policy in the multi-armed bandit problem is given by the so-called "Gittins Index" policy. For this policy, one can consider the problem of stopping on each arm in isolation from the other arms. The value of this stopping problem is the so-called Gittins Index for this arm. Now, an optimal policy consists of, at each point in time, using the arm with the highest Gittins Index. Presman (1990) calculated the Gittins Index for the case in which the underlying stochastic process is a Poisson process. Bergemann \& Välimäki (2008) give a survey of this literature.

Bolton \& Harris $(1999,2000)$ were the first to consider the multi-player version of the two-armed bandit problem. While they assumed that the underlying stochastic process was a Brownian motion, KRC analyzed the corresponding problem with exponential processes. This model proved to be more tractable and is underlying our theoretical hypotheses. While the previous papers focussed on MPE, HKR extended the equilibrium concept beyond Markov perfect equilibrium. ${ }^{\text {S }}$
do so in the hope of producing public good news, which, in turn, makes their partners more optimistic. As their partners become more optimistic, they will be more inclined to experiment, thus providing some additional free-riding opportunities to the first player. This effect is absent in KRC, because here good news is conclusive: It resolves all uncertainty, so that, as soon as there is good news, players are not interested in free-riding any longer.
${ }^{3}$ Many variants of the multi-player bandit problem have been analyzed since. In Keller \& Rady (2010), a bad risky arm also sometimes yields a payoff. In Klein \& Rady (2011), the quality of the risky arm is negatively correlated across players. Klein (2013) introduces a second risky arm, with a quality that is negatively correlated with that of the first. In Keller \& Rady (2015), the lump-sum payoffs are

The only papers we are aware of that conduct experimental tests of bandit problems are Meyer \& Shi (1995), Banks, Olson, Porter (1997), Anderson (2001, 2012), and Gans, Knox, Crosson (2007). All these papers consider various single-agent problems. We are not aware of any previous experimental study of a strategic, multi-player, bandit problem.

## 3 The Theoretical Framework

We borrow our theoretical reference framework from KRC. There are $n \geq 1$ players, each of whom plays a bandit machine with two arms over an infinite horizon. One of the arms is safe, and yields a known flow payoff of $s>0$ whenever it is pulled. The other arm is risky and can be either good or bad. If it is bad, it never yields any payoff. If it is good, it yields a lump sum of $h>0$ at the jumping times of a Poisson process with parameter $\lambda>0$. It is assumed that $g:=\lambda h>s$. Players decide in continuous time which arm to pull. Payoffs are discounted at a rate $r>0$. If they knew the quality of the risky arm, players would have a strictly dominant strategy always to pull a good risky arm and never to pull a bad one. They are initially uncertain whether their risky arm is good or bad. Yet, the only way to acquire information about the quality of the risky arm is to use it, which is costly as it implies forgoing the safe payoff flow $s$. The $n$ players' risky arms are either all good or all bad. Players share a common prior belief $p_{0} \in(0,1)$ that their risky arms are good. Every player's actions as well as the outcomes of their actions are publicly observable; therefore, the information one player produces benefits the other players as well, creating incentives for players to free-ride on their partners' efforts. Players thus share a common posterior belief $p_{t}$ at all times $t \in \mathbb{R}_{+}$. All the parameter values and the structure of the game are common knowledge.

The common posterior beliefs are derived from the public information via Bayes' rule. As a bad risky arm never yields any payoff, the first arrival of a lump sum fully reveals the quality of all players' risky arms. Thus, if a success on one of the players' risky arms is observed at instant $\tau \geq 0$, the common posterior belief satisfies $p_{t}=1$ for all $t>\tau$. If no success has been observed up until instant $t$, the common posterior belief satisfies

$$
p_{t}=\frac{p_{0} e^{-\lambda \int_{0}^{t} \sum_{i=1}^{N} k_{i, \tau} d \tau}}{p_{0} e^{-\lambda \int_{0}^{t} \sum_{i=1}^{N} k_{i, \tau} d \tau}+1-p_{0}},
$$

costs to be minimized. Rosenberg, Solan, Vieille (2007) and Murto \& Välimäki (2011) analyze the case of privately observed payoffs, while Bonatti \& Hörner (2011) investigate the case of privately observed actions. Bergemann \& Välimäki $(1996,2000)$ consider strategic experimentation in buyer-seller settings. Hörner \& Skrzypacz (2016) give a survey of this literature.
where $k_{i, \tau}=1$ if player $i$ uses the risky arm at instant $\tau$ and $k_{i, \tau}=0$ otherwise.
KRC show in their Proposition 3.1 that, if players are maximizing the sum of their payoffs, all players $i \in\{1, \cdots, n\}$ choose $k_{i, t}=1$ if $p_{t}>p_{n}^{*}:=\frac{r s}{(r+n \lambda)(g-s)+r s}$, and $k_{i, t}=0$ otherwise. Note that $p_{n}^{*}$ is strictly decreasing in the number of players $n$. In particular, in the single-agent case ( $n=1$ ), the decision maker optimally sets $k_{1, t}=1$ if $p_{t}>p_{1}^{*}:=\frac{r s}{(r+\lambda)(g-s)+r s}$, and $k_{1, t}=0$ otherwise.

KRC go on to analyze the game of strategic information acquisition, where each player maximizes his own payoff, not taking into account that the information he produces is valuable to the other players as well. They analyze perfect Bayesian equilibria in Markov strategies (MPE), i.e., strategies where a player's action after any history can be written as a time-invariant function $k_{i}(p)$ of the common belief at that history. It is shown that, for beliefs close to 1 ( 0 ), playing risky (safe) is a dominant action; for intermediate beliefs, players' effort levels are strategic substitutes. In any MPE with a finite number of switches, all players will set $k_{i}(p)=0$ for all $p \leq p_{1}^{*}$ (see Proposition 6.1 in KRC). Moreover, it is shown that there exists no MPE in which all players play a cut-off strategy, i.e. a strategy that prescribes the use of the risky arm for beliefs above a single cutoff and that of the safe arm below. The intuition for this result is best described in the context of a two-player game. Indeed, suppose to the contrary that there existed an equilibrium in cutoff strategies. As there is a region of beliefs in which safe and risky are mutually best responses, both players cannot use the same cutoff in equilibrium; i.e., one player plays the role of pioneer, while the other one free-rides, throughout the belief region where safe and risky are mutually best responses. As he gets all his information for free in the relevant belief region, the free-rider's payoff function will be higher than the pioneer's. As a player's propensity to play risky is increasing in his own payoff, however, this would imply that the free-rider entered the region in which risky is dominant at a more pessimistic belief than the pioneer. Thus, the roles of pioneer and free-rider must switch at least once in equilibrium.

HKR extend the analysis to non-Markovian PBE. They show that on the path of play in the average-payoff maximizing PBE, all players set $k_{i}(p)=1$ for all $p>p_{1}^{*}$, and $k_{i}(p)=0$ otherwise. Thus, in stark contrast to the simple structure of the single-agent optimum or HKR's average-payoff maximizing PBE, every MPE has the property that, for intermediate beliefs, players change roles between experimenter and free-rider at least once. As a matter of fact, KRC show that, for any given number of role changes greater than, or equal to, one, there exists an MPE with that number of role changes. A behavioral prediction of MPE is thus that players change roles for intermediate beliefs
at least once. : $^{2}$

## 4 Parametrization and Experimental Design

### 4.1 Experimental Implementation

In our experimental treatments, the number of players will be $n=2$ or $n=3$. We choose the discount rate $r=1 / 120$. To implement the infinite-horizon game in the laboratory, we end the game at the first jump time of a Poisson process with parameter $r$. $\square^{\text {I }}$ We set the probability that the risky arm is good $p_{0}=1 / 2$, the safe payoff $s=10$, the lump-sum amount paid out by a good risky arm $h=2500$, and the arrival rate of lump sums on the good risky arm $\lambda=1 / 100$. Thus, $25=g>s=10$. The realizations of all random processes were simulated ahead of time. ${ }^{6}$ We generated six different sets of realizations of the random parameters, corresponding to six different games each of our subjects played. To make our findings more easily comparable, we have kept the same realizations for both the strategic and the control treatments. ${ }^{[1}$ One unit of time corresponds to a second in our experimental implementation. In keeping with the theoretical predictions, we have endeavored to implement our experimental investigation in continuous time, subject to the restrictions imposed by the available computing power. ${ }^{\text {B }}$

Subjects were randomly assigned to groups of $n=2$ or $n=3$ players. We used a between-subject design: Each group was randomly assigned either to a control treatment or to a strategic treatment, and played the six games in random order. To ensure

[^3]a balanced data-collection process, we replicated any order of the six games that was used for $k(k \in\{1, \cdots, 10\})$ groups in the strategic treatment for $k$ groups in the control treatment as well. Subjects could see their fellow group members' action choices and payoffs on their computer screens. They had to choose an action before the game started and could switch their action at any point in time by clicking on the corresponding button with their mouse. (Please see the Appendix for details and screen shots.)

All experimental sessions took place in July and August 2017 at the BizLab Experimental Research Laboratory at UNSW Sydney. All subjects were recruited from the university's subject pool and administered by the online recruitment system ORSEE (Greiner, 2015). All participants were native speakers of English. In total, 100 subjects, 46 of whom were female, participated in 60 sessions. The participants' age ranged from 18 to 35 years, with an average of 20.78 and a standard deviation of 2.43 . Because the implementation was computationally very intensive and because we wanted to collect eye-tracking data, only between 2 and 3 subjects participated at a time in each session. Upon arrival, participants were seated in front of a computer at desks which were separated by dividers to minimize potential communication. Participants received written instructions and had the opportunity to ask questions. ${ }^{\boxed{Q}}$ After the subjects had successfully completed a simple comprehension test, the eye-tracking devices were calibrated, after which the subjects started the experiment. The experiment was programmed in zTree (Fischbacher, 2007). At the end of the experiment, we collected some information on participants' demographic attributes and risk attitudes. They were then privately paid their cumulated experimental earnings from one randomly selected game in cash (with a conversion rate of $\mathrm{E} \$ 100=$ AU\$ 1) plus a show-up fee of AU\$ 5. The average earnings were AU 23.86 , with a standard deviation of AU\$ 9.95.

### 4.2 Behavioral Hypotheses

One of the main theoretical predictions of both MPE and PBE is that players use the risky arm less in a strategic setting than in a situation in which they are single players. This is because players free-ride on the information their partners are producing. Indeed, players are predicted to play safe at all beliefs $p \leq p_{1}^{*}$ in all these instances, while efficiency would require that they play risky at all beliefs $p>p_{n}^{*}$, where $p_{n}^{*}<p_{1}^{*}$. Single players and players playing the best PBE should play risky at all beliefs $p>p_{1}^{*}$, i.e., in the average-payoff maximizing PBE, players on path adopt the same cutoff behavior as a single agent. In any MPE, by contrast, since at least one player is not playing a cut-off

[^4]strategy, at least one player will play safe at some beliefs above $p_{1}^{*}$. Indeed, it is possible to derive a lower bound $p^{\ddagger} \in\left(p_{1}^{*}, p^{m}\right)$, where $p^{m}:=\frac{s}{g}$ is a myopic player's cutoff belief, such that, for all beliefs in $\left(p_{1}^{*}, p^{\ddagger}\right)$, at least one player plays safe. Indeed, as KRC show (their Equation (6), p.49), it is a best response for player $i$ to play safe if and only if his value function $u_{i}(p)$ satisfies $u_{i}(p) \leq s+K_{-i}(p) c(p)$, where $K_{-i}(p):=\sum_{j \neq i} k_{j}(p)$ is the number of players other than $i$ who play risky at belief $p$, and $c(p):=s-p g$ is a player's myopic opportunity cost for playing risky, given the belief $p$. An upper bound on a player's equilibrium value function $u_{i}$ is given by $V_{n, p_{1}^{*}}$, the value function of all players playing risky on $\left(p_{1}^{*}, 1\right]$, and safe on $\left[0, p_{1}^{*}\right]$. Thus, a lower bound $p^{\ddagger}$ is given by the unique root $V_{n, p_{1}^{*}}\left(p^{\ddagger}\right)-s-(n-1) c\left(p^{\ddagger}\right)=0$. By the same token, we can derive an upper bound $\bar{p}$ on the lowest belief at which risky is a dominant action. For this, we use the fact that the single-agent value function $V_{1}^{*}$ constitutes a lower bound on a player's equilibrium value function $u_{i}$, and find our upper bound $\bar{p}$ as the unique root $V_{1}^{*}(\bar{p})-s-(n-1) c(\bar{p})=0$.

With our numerical parameters, $p^{m}=0.4, \bar{p} \approx 0.3578(\bar{p} \approx 0.3742)$ if $n=2$ $(n=3), p^{\ddagger} \approx 0.3428\left(p^{\ddagger} \approx 0.3609\right)$ if $n=2(n=3), p_{1}^{*} \approx 0.2326, p_{2}^{*} \approx 0.1031$, and $p_{3}^{*} \approx 0.0535$. As $p_{0}=0.5>0.4=p^{m}$, players start out with a belief that makes playing risky the dominant action. If, in the strategic treatment, $n$ players were uninterruptedly playing risky and there was no breakthrough, the belief would drop to $p^{m}$ after $40.6 / n$ seconds, to our upper bound in the game with $n=2$ players ( $n=3$ players) $\bar{p}$ after $58.5 / n(51.5 / n)$ seconds, to our lower bound in the game with $n=2$ players ( $n=3$ players) $p^{\ddagger}$ after $65.0 / n(57.0 / n)$ seconds, to $p_{1}^{*}$ after $119.4 / n$ seconds, to $p_{2}^{*}$ after $216.4 / n$ seconds, and to $p_{3}^{*}$ after $287.4 / n$ seconds. For the control treatment, the same times apply with $n=1$.

### 4.2.1 Free-Riding

Let $\hat{T}$ be the time of a first breakthrough or the end of the game, whichever arrives first. In order to measure the prevalence of free-riding, we investigate the behavior of the average experimentation intensity, where, following KRC, we define the experimentation intensity at instant $t$ as $\sum_{i=1}^{n} k_{i, t}$. Note that, in the control treatment, a player conforming to the theoretical prediction will always play risky until his belief hits $p_{1}^{*}$. In the strategic treatment, at least one of them will switch to safe at a belief strictly above $p_{1}^{*}$ if they play an MPE. In the best PBE, they both play risky until the belief $p_{1}^{*}$ is reached. Furthermore, conditionally on no success arriving, beliefs will decrease faster in the strategic setting, as player $i$ 's belief also decreases in response to player $j$ 's hapless experimentation. As both effects go in the same direction, the average experimenta-
tion intensity should be lower in the strategic setting, whether players play MPE or the best PBE. We thus formulate the following

Hypothesis 4.1 The average experimentation intensity $\frac{\int_{0}^{\hat{T}} \sum_{i=1}^{n} k_{i, t} d t}{n \hat{T}}$ is significantly lower in the strategic treatment than in the control treatment.

Our theoretical understanding of the game would suggest that players free-ride for strategic reasons, i.e. they opportunistically take advantage of the information their partners provide them with. In order to test whether our subjects did in fact avail themselves of the additional information they received in the strategic treatment, we formulate our next

Hypothesis 4.2 Players' average final payoffs are higher in the strategic treatment.

### 4.2.2 MPE vs. Best PBE

As explained above, KRC predict that subjects will use cut-off strategies in the control treatment, whereas at least one player will not use a cut-off strategy in the strategic setting if MPE is played. By contrast, HKR show that cutoff behavior prevails on path in the strategic setting also if the best PBE is played. Cut-off behavior consists in a player's playing risky at the outset, and continuing to play risky until his risky arm is revealed to be good, the game ends, or he switches to the safe action, and continues to play safe until the game ends or his risky arm is revealed to be good. To investigate whether, qualitatively, the behavior predicted by MPE prevailed, we shall examine the following

Hypothesis 4.3 The frequency of cut-off behavior is significantly higher in the control treatment than in the strategic treatment.

In order further to discriminate between simple MPE and the best PBE, we measure the proportion of time (before a first breakthrough) during which exactly one of the players plays risky. Theory would predict this proportion to be nil both in the control treatment and in HKR's best PBE, while it is positive in KRC's MPE. We thus formulate the following

Hypothesis 4.4 The proportion of time before a first breakthrough during which exactly one player plays risky is higher in the strategic treatment than in the control treatment.

The non-cutoff behavior predicted by MPE moreover implies that players should switch arms more often in the strategic treatment. Yet, as noted above, learning also tends to be faster in the strategic setting, so that beliefs may more quickly reach the threshold at which the player will want to change his action. While this effect would add to making switching more prevalent in the strategic treatment, a substantially higher number of switches in the strategic treatment would provide further evidence in favor of subjects' adopting MPE behavior. Indeed, recall that players are predicted to switch action at most once in both the control treatment and the best PBE, while, for any number of role changes, there exists an MPE with that number of role changes, as KRC show. For a two-player game, this e.g. implies that one of the players must switch actions at least twice, with the other one switching once, before $p_{1}^{*}$ is reached. ${ }^{\text {[0 }}$

To control for the effect that, the longer the game goes on, the more time players have to switch actions, we define the incidence of switches as the number of a player's switches in a given game per unit of effective time, where effective time is understood as the time before the game ends or the player's risky arm is revealed to be good, whichever happens first. Thus, we shall check the following

Hypothesis 4.5 The incidence of switches is significantly higher in the strategic treatment than in the control treatment.

## 5 Experimental Results

### 5.1 Overview

Figures 1 and 2 display the evolution of players' action choices over all six games. Players' actions are described by dots, the width of which corresponds to one second of time. For each of the six games, we conducted four treatments à ten groups each, the parameters of which (i.e. their duration, the quality of the risky arm and the timing of successes on the risky arm in case it was good) we had simulated ahead of time, as explained in Section 4.

As the figures show, the duration of the games ranged from 32 seconds for Game 5

[^5]Figure 1: ACTION CHOICES BY PLAYERS OVER TIME: Games 1-3







 \% (.
 (\%
















Games 1 and 3 are shown at the top left and top right, respectively, while Game 2 is illustrated at the bottom. Players' actions are described by dots, the width of which corresponds to one second of time. Groups 1-10 correspond to the strategic treatment for two-player groups; groups 11-20 are the corresponding control treatments. Groups 21-30 played the strategic treatment for three-player groups, while groups 31-40 were the corresponding control treatments. In each group, we refer to the lowermost player as 'player 1', while 'player 2' will denote the player right above, and 'player 3 ' is the uppermost player. The x -axis represents calendar time. A red dot indicates that a player is playing risky in a given second, while a blue dot indicates that the player is paying safe. A black square indicates a success.

Figure 2: ACTION CHOICES BY PLAYERS OVER TIME: Games 4-6


Game 4 is shown at the top, and Games 5 and 6 at the bottom left and right, respectively. Players' actions are described by dots, the width of which corresponds to one second of time. Groups 1-10 correspond to the strategic treatment for two-player groups; groups 11-20 are the corresponding control treatments. Groups 21-30 played the strategic treatment for three-player groups, while groups 31-40 were the corresponding control treatments. In each group, we refer to the lowermost player as 'player 1', while 'player 2' will denote the player right above, and 'player 3 ' is the uppermost player. The $x$-axis represents calendar time. A red dot indicates that a player is playing risky in a given second, while a blue dot indicates that the player is playing safe. A black square indicates a success. 13
to 230 seconds for Game 4. As is furthermore evident from the figures, players change their behaviors over time. While often playing risky at the beginning, players seem to grow less inclined to use the risky arm the longer it has unsuccessfully been used before. This is consistent with Bayesian updating of a prior belief. As we shall discuss in more detail below, this of course does not imply that players adjust their behaviors precisely at the equilibrium cutoff beliefs. Nevertheless, as our subsequent analysis will show, the main qualitative predictions of Markov Perfect Equilibrium are borne out by the experimental evidence.

### 5.2 Average Experimentation Intensities

One of the main qualitative predictions of the theoretical analysis is that players will tend to free-ride on the experimentation provided by their partners. To test for treatment differences non-parametrically, we apply two-sided Wilcoxon rank-sum (MannWhitney) tests. Table 1 lists the mean experimentation intensity observed in our four treatments.

Under Hypothesis 4.1, players will use the risky arm less in the strategic treatment. The data provides support for this hypothesis.

Result 5.1 The average experimentation intensity $\frac{\int_{0}^{\hat{T}} \sum_{i=1}^{n} k_{i, t} d t}{n T}$ is significantly lower in the strategic treatment, as compared to the control treatment. This result holds for both $n=2$ and $n=3$.

As Table 1 reveals, the additional presence of one (two) perfectly positively correlated arms leads to lower experimentation intensities in all games. This is statistically significant for Games 1-5, but not for Game 6, in both settings with $n=2$ and $n=3$. The corresponding $p$-values in the case of $n=2$ are $0.0085,0.0347,0.0001,0.0336$, 0.0002 , and 0.1218 for Games $1-6$, respectively. In the setting with $n=3$, the average experimentation intensity is also lower in the strategic treatment ( $p$-values of 0.0001 for Games 1-5, and 0.5352 for Game 6, respectively). 1 As Figure 2 highlights, Game 6

[^6]features an early success by Player 2 after 9 seconds of exploration, as well as successes by Player 1 after 39 and 44 seconds of exploration, respectively.

Table 1: AVERAGE EXPERIMENTATION INTENSITY

$$
n=2 \quad n=3
$$

|  | Strategic Treatment |  | Control Treatment |  | Strategic Treatment |  | Control Treatment |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Game | Eff. <br> Time | Av. Exp. <br> Intensity | Eff. <br> Time | Av. Exp. <br> Intensity | Eff. <br> Time | Av. Exp. <br> Intensity | Eff. <br> Time | Av. Exp. <br> Intensity |
| 1 | 155 [0] | . 508 [.124] | 146.5 [15.1] | . 730 [.293] | 147.7 [12.6] | . 455 [.223] | 134.6 [25.7] | . 797 [.236] |
| 2 | 185.9 [.2] | . 512 [.157] | 186 [0] | . 696 [.283] | 83.6 [21.9] | . 543 [.269] | 139 [67.6] | . 833 [.218] |
| 3 | 88 [0] | . 565 [.126] | 88 [0] | . 878 [.235] | 88 [0] | . 457 [.305] | 88 [0] | . 866 [.248] |
| 4 | 230 [0] | . 519 [.134] | 230 [0] | . 678 [.239] | 230 [0] | . 383 [.245] | 230 [0] | . 728 [.243] |
| 5 | 32 [0] | . 653 [.349] | 32 [0] | . 984 [.072] | 32 [0] | . 596 [.381] | 32 [0] | . 953 [.196] |
| 6 | 14.6 [7.6] | . 810 [.314] | 30 [25.6] | . 941 [.167] | 25 [30.1] | . 800 [.336] | 56.4 [43.0] | . 857 [.250] |

As we have mentioned above, information accumulation is potentially faster in the strategic treatment. Indeed, on account of the conditionally independent Poisson processes, the information acquired within a given unit of time is multiplied by the number of players currently playing risky. Therefore, conditionally on no success arriving, players' beliefs will tend to decrease more quickly in the strategic setting, implying that more time will be spent at more pessimistic beliefs. To ensure that Result 5.1 is not solely due to this effect, we conduct our parameter tests separately by belief region. Specifically, we consider the belief regions $\left[\bar{p}, \frac{1}{2}\right]$, where risky is a dominant action, and $\left(p_{1}^{*}, p^{\ddagger}\right)$, where risky and safe are mutually best responses in MPE. ${ }^{[12}$ In the control treatment or if players were behaving according to the best PBE in the strategic setting, by contrast, all players should play risky in both regions. The following tables

[^7]Table 2: AVERAGE EXPERIMENTATION INTENSITY BY REGIONS FOR $n=2$

| Game | Strategic Treatment |  |  |  | Control Treatment |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | Obs. | Exp. <br> Intensity | Min | Max | Obs. | Exp. <br> Intensity | Min | Max |
| Panel A: R Dominant |  |  |  |  |  |  |  |  |
| 1 | 20 | . 648 [.315] | . 094 | 1 | 20 | . 835 [.304] | . 156 | 1 |
| 2 | 20 | . 723 [.291] | . 211 | 1 | 20 | . 888 [.259] | . 130 | 1 |
| 3 | 20 | . 617 [.281] | . 125 | 1 | 20 | . 906 [.235] | . 241 | 1 |
| 4 | 20 | . 732 [.323] | . 117 | 1 | 20 | . 880 [.243] | . 197 | 1 |
| 5 | 20 | . 653 [.349] | . 065 | 1 | 20 | . 984 [.072] | . 677 | 1 |
| Panel B: Mutually BR |  |  |  |  |  |  |  |  |
| 1 | 20 | . 503 [.265] | . 095 | 1 | 16 | . 758 [.343] | . 114 | 1 |
| 2 | 20 | . 445 [.184] | 0 | . 788 | 20 | . 783 [.362] | . 118 | 1 |
| 3 | 20 | . 589 [.341] | 0 | 1 | 16 | . 935 [.182] | . 381 | 1 |
| 4 | 20 | . 484 [.254] | 0 | 1 | 19 | . 765 [.343] | . 092 | 1 |

summarize our findings by belief region. As player 2 has a success after 9 seconds of using the risky arm, we omit Game 6 from these tables. We furthermore omit Game 5 from the tables for the "mutually BR" region, as this game lasts only 32 seconds, implying that the "mutually BR " region cannot be attained in the control treatment and only lasts for a few seconds in the strategic treatment, if it is attained at all. For Games $1-4$, the missing observations for the "mutually BR " region correspond to groups (in the strategic treatment) or individual players (in the control treatment) that have not reached the "mutually BR" region either on account of an early success or because they did not use the risky arm enough.

The comparison of the strategic treatment with the control treatment shows that the average experimentation intensity is substantially lower in the strategic treatment, for both belief regions. We first turn to the two-player setup and focus on the " $R$ dominant" region, where the effect is statistically significant at least at the $10 \%$-level in Games 1-5, the $p$-values of the two-sided Wilcoxon ranksum test amounting to 0.0367 , $0.0371,0.0010,0.0646$, and 0.0002 , respectively. ${ }^{[3]}$ Now, let us consider the "mutually

[^8]Table 3: AVERAGE EXPERIMENTATION INTENSITY BY REGIONS FOR $n=3$

| Game | Strategic Treatment |  |  |  | Control Treatment |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | Obs. | Exp. <br> Intensity | Min | Max | Obs. | Exp. <br> Intensity | Min | Max |
| Panel A: R Dominant |  |  |  |  |  |  |  |  |
| 1 | 30 | . 709 [.365] | 0 | 1 | 30 | . 935 [.190] | . 260 | 1 |
| 2 | 30 | . 649 [.340] | 0 | 1 | 30 | . 976 [.076] | . 689 | 1 |
| 3 | 30 | . 593 [.410] | 0 | 1 | 30 | . 906 [.245] | 0 | 1 |
| 4 | 30 | . 613 [.373] | 0 | 1 | 30 | . 889 [.246] | . 092 | 1 |
| 5 | 30 | . 596 [.381] | 0 | 1 | 30 | . 953 [.196] | 0 | 1 |
| Panel B: Mutually BR |  |  |  |  |  |  |  |  |
| 1 | 30 | . 537 [.370] | 0 | 1 | 29 | . 766 [.325] | . 113 | 1 |
| 2 | 30 | . 549 [.369] | 0 | 1 | 20 | . 673 [.340] | . 023 | 1 |
| 3 | 30 | . 482 [.398] | 0 | 1 | 25 | . 875 [.279] | . 113 | 1 |
| 4 | 30 | . 471 [.353] | 0 | 1 | 29 | . 815 [.275] | . 299 | 1 |

BR" region. Here, the contrast between the strategic and the control treatment is even more pronounced and statistically significant at least at the $5 \%$-level for all four games. The corresponding $p$-values are $0.0179,0.0032,0.0011$, and 0.0071 for Games $1-4$, respectively.

We now turn to the three-player setup, where we expect the same predictions to hold. When considering the " R dominant" region, we find the difference in average experimentation intensities between the two treatments to be highly statistically significant. The $p$-values are $0.0042,0.0001,0.0005,0.0012$, and 0.0001 for Games $1-5$, respectively. The same is true for the "Mutually BR" region, with the exception of Game 2. This is most likely due to an early success by player 3 after only 44 seconds of exploration, which accounts for the sharp decrease in the number of observations for the control treatment, which we report in Table 3. The $p$-values are $0.0161,0.2274,0.0002$, and 0.0002 for Games $1-4$, respectively.

Since we are conditioning on the belief region, these results provide strong evidence that players are free-riding because of strategic considerations. Our analysis by 0.0258 , which would give us a $5 \%$ significance level for all games.
belief region also shows that, while players tend to use the risky arm less in the "mutually $B R$ " region than in the " $R$ dominant" region, there does not appear to be any major qualitative difference between the two regions. This is true for both the strategic and control treatments. By contrast, theory would predict that, in the control treatment, players should play risky in both regions (i.e. we should observe average experimentation intensities of 1). In the strategic treatment, we should likewise observe average experimentation intensities of 1 in both regions if subjects behaved according to the best PBE, whereas, if they behaved according to MPE, the average experimentation intensity should be 1 only for the " $R$ dominant" region. Our results would suggest that our experimental subjects did not distinguish between the two regions, possibly because they were not able to update their subjective beliefs with enough precision to tell them apart. Thus, while, at least for groups of size $n=2$, the difference between the control and strategic treatments might be slightly more pronounced in the "Mutually BR" region, the analysis of average experimentation intensities is inconclusive when it comes to the comparison between MPE and the best PBE. We furthermore observe that, as far as free-riding is concerned, there do not seem to be any major differences between groups of size two and groups of size three.

### 5.3 Payoffs

Strategic interaction is predicted to arise among players as a result of (positive) informational externalities, i.e. the information produced by their partners allows players to make better decisions and hence to secure themselves higher payoffs. Thus, players' payoffs should be higher on average in the strategic treatment. Table 4 displays the average final payoffs per player across games for our four treatments. With the exception of Game 1, average final payoffs are much higher in the strategic treatment than in the control treatment, for both group sizes. This is statistically significant for Games $2-5$, but not for Games 1 and 6 . For $n=2(n=3)$, the $p$-values are 0.0105 ( 0.0013 ), 0.0001 ( 0.0001 ), 0.0172 ( 0.0001 ), and 0.0002 ( 0.0001 ) for Games $2-5$, respectively. The average-payoff difference is not statistically significant with $p$-values of 0.6747 ( 0.8819 ), and 0.2885 ( 0.1135 ) for Games 1 and 6 , respectively. Thus, with the exception of Game 1 , our subjects indeed take advantage of the positive informational externalities in the strategic treatment.

Table 4: AVERAGE FINAL PAYOFFS

| Game | Strategic Treatment |  |  |  | Control Treatment |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | Obs. | Final Payoff | Min | Max | Obs. | Final Payoff | Min | Max |
| Panel A: $n=2$ |  |  |  |  |  |  |  |  |
| 1 | 20 | 817.5 [199.232] | 320 | 1180 | 20 | 1176.5 [1007.207] | 0 | 2860 |
| 2 | 20 | 1092 [546.401] | 520 | 3090 | 20 | 577.5 [531.561] | 0 | 1620 |
| 3 | 20 | 407 [120.573] | 110 | 640 | 20 | 109.5 [209.422] | 0 | 670 |
| 4 | 20 | 1181 [309.905] | 760 | 2170 | 20 | 761 [564.763] | 0 | 1930 |
| 5 | 20 | 115 [114.133] | 0 | 310 | 20 | 5.5 [24.597] | 0 | 110 |
| 6 | 20 | 3800.5 [1285.213] | 2500 | 5190 | 20 | 3554.5 [1428.173] | 760 | 5240 |
| Panel B: $n=3$ |  |  |  |  |  |  |  |  |
| 1 | 30 | 1177.667 [745.037] | 310 | 3170 | 30 | 1488.667 [1067.507] | 0 | 3130 |
| 2 | 30 | 2110.333 [1423.033] | 0 | 5020 | 30 | 1161 [1037.302] | 0 | 2720 |
| 3 | 30 | 496.333 [266.775] | 0 | 880 | 30 | 123.333 [222.963] | 0 | 880 |
| 4 | 30 | 1465.333 [558.006] | 0 | 2300 | 30 | 641.667 [569.132] | 0 | 2090 |
| 5 | 30 | 137 [123.460] | 0 | 320 | 30 | 15.333 [62.958] | 0 | 320 |
| 6 | 30 | 3135 [1607.675] | 200 | 5660 | 30 | 2457.333 [1973.708] | 0 | 5320 |

### 5.4 Eye-Tracking Data

To study the players' information-acquisition processes further, we employ eye-tracking data obtained by using two (three) Tobii-TX300 eye trackers with a sampling rate of 300 Hz . The relative frequency of fixations corresponds to the relative importance of an information in the subject's decision-making process (Jacob \& Karn, 2003; Poole, Ball, Phillips, 2005). In our setting, eye fixations can thus provide information about the importance subjects assigned to the different payoff streams, which revealed both a player's actions and payoffs. [⿶] We define a subject's fixation intensity as the total number of his fixations on his own payoff stream, divided by the total number of all fixations (i.e. both on his own and on his partner's [partners'] payoff stream[s]) during a game before a breakthrough arrives or the game ends.

Figure 3 displays (non-representative) heatmaps to illustrate the different infor-

[^9]Figure 3: HEATMAPS OF FOUR TREATMENTS


In the top-left corner, the strategic treatment with $n=2$ is illustrated, with the corresponding control treatment represented just below. In the top-right corner, the strategic treatment with $n=3$ is displayed, while the control treatment with $n=3$ is shown at the bottom-right. All four heatmaps show the total number of fixations. The accumulated number of fixations is calculated for an entire game (Game 4 in the $n=2$ setup and Game 2 in the $n=3$ setup). Each fixation made has the same value and is indepentent of its duration. A color gradient is used to indicate the areas with more fixations (low=green to high=red).
mation acquisition behavior in our four treatments. The measure of interest is the total number of fixations. For each heatmap, the accumulated number of fixations is calculated for an entire game and the image corresponds to the last point in calendar time before the game ends. A color gradient is employed to display the areas that attained more fixations (low=green to high=red). As Figure 3 illustrates, players not only switch actions more frequently in the strategic treatment but also focus much more intensively on their partners' actions and payoffs. This is in sharp contrast to the corresponding control treatment, where players seem to focus almost exclusively on their own stream of payoffs. Indeed, a rational player should completely ignore a partner's actions and payoffs in the control treatments, as they are informative for his own problem only in the strategic setting. This observation is also consistent with the results presented below in Tables 6 and 8 . Indeed, in the control treatments, players' behavior requires less coordination, as significantly more cut-off behavior and less switching is observed.

Table 5: AVERAGE FIXATION INTENSITY

|  | $n=2$ |  |  |  | $n=3$ |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | Strategic Treatment |  | Control Treatment |  | Strategic Treatment |  | Control Treatment |  |
| Game | Eff. <br> Time | Fixation <br> Intensity | Eff. <br> Time | Fixation <br> Intensity | Eff. <br> Time | Fixation Intensity | Eff. <br> Time | Fixation Intensity |
| 1 | 155 [0] | . 619 [.078] | 146.5 [15.1] | . 870 [.089] | 147.7 [12.6] | . 384 [.118] | 134.6 [25.7] | . 710 [.158] |
| 2 | 185.9 [.2] | . 620 [.121] | 186 [0] | . 882 [.131] | 83.6 [21.9] | . 365 [.113] | 139 [67.6] | . 709 [.176] |
| 3 | 88 [0] | . 600 [.086] | 88 [0] | . 874 [.111] | 88 [0] | . 392 [.164] | 88 [0] | . 762 [.112] |
| 4 | 230 [0] | . 615 [.095] | 230 [0] | . 875 [174] | 230 [0] | . 389 [.124] | 230 [0] | . 700 [.139] |
| 5 | 32 [0] | . 633 [.139] | 32 [0] | . 876 [.149] | 32 [0] | . 383 [.151] | 32 [0] | . 745 [.199] |
| 6 | 14.6 [7.6] | . 594 [.169] | 30 [25.6] | . 814 [.112] | 25 [30.1] | . 382 [.157] | 56.4 [43.0] | . 646 [.159] |

As Table 5 shows, the average fixation intensity is significantly lower in the strategic treatment. This is statistically significant for all six games independently of the group size (all $p$-values are 0.0001 for $n=2$ and $n=3$ ). The sophisticated coordination required by the switching of roles between pioneer and free-rider, which is characteristic of MPE and which we shall analyze in detail below, seems to force players to pay a lot of attention to their partner's (partners') behavior. This provides additional evi-
dence that players behave strategically and try to learn from their partners' exploration efforts in the strategic treatments only.

### 5.5 Cut-Off Behavior

As we have pointed out above, optimality in the individual decision-making problem in our control treatment implies cut-off behavior. The best PBE also features cutoff behavior on the path of play, while KRC have shown that there does not exist an MPE in cut-off strategies. This prediction of MPE is confirmed by our experiment, where subjects often play cut-off strategies in the control treatment, while they hardly ever do so in the strategic treatment.

Result 5.2 The frequency of cut-off behavior is higher in the control treatment than in the strategic treatment. We find evidence for both $n=2$ and $n=3$.

Indeed, Table 6 shows that the frequency of cut-off behavior is much higher in the control treatment than in the strategic treatment for both groups of size $n=2$ and groups of size $n=3$. While it increases sharply in Games 5 and 6 as compared to Games 1-4 in the strategic treatments, it is still higher in the corresponding control treatments (for either given group size). In Game 5, this sharp increase is most likely due to the short duration of that game. In Game 6, it is most likely driven by the resolution of uncertainty very early in the game, with Player 2 achieving a success after exploring for 9 seconds.

### 5.6 Pioneers

In the control treatment as well as in the best PBE, players are predicted to play risky on ( $p_{1}^{*}, \frac{1}{2}$ ]; i.e., conditionally on no success arriving, players should switch from risky to safe only once, and do so at the same time, at which their beliefs reach $p_{1}^{*}$. By contrast, as KRC have shown, there is a range of beliefs containing $\left(p_{1}^{*}, p^{\ddagger}\right)$ such that safe and risky are mutually best responses in any Markov Perfect Equilibrium. In particular, there exists a range of beliefs in which just one pioneer should play risky while the other player(s) free-ride(s). The following result thus provides further evidence that MPE seems to predict the qualitative features of subjects' behavior better, while confirming the prevalence of free-riding in our strategic treatment.

Table 6: FREQUENCY OF CUT-OFF BEHAVIOR

|  | $n=2$ |  |  |  | $n=3$ |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | Strategic Treatment |  | Control Treatment |  | Strategic Treatment |  | Control Treatment |  |
| Game | Eff. <br> Time | Tot. (Rel.) Freq. | Eff. <br> Time | Tot. (Rel.) Freq. | Eff. <br> Time | Tot. (Rel.) Freq. | Eff. <br> Time | Tot. (Rel.) Freq. |
| 1 | 155 [0] | 0 (0) | 146.5 [15.1] | 15 (.75) | 147.7 [12.6] | 3 (.10) | 134.6 [25.7] | 21 (.70) |
| 2 | 185.9 [.2] | 0 (0) | 186 [0] | 15 (.75) | 83.6 [21.9] | 3 (.10) | 139 [67.6] | 22 (.73) |
| 3 | 88 [0] | 5 (.25) | 88 [0] | 17 (.85) | 88 [0] | 11 (.37) | 88 [0] | 26 (.87) |
| 4 | 230 [0] | 0 (0) | 230 [0] | 14 (.70) | 230 [0] | 6 (.20) | 230 [0] | 19 (.63) |
| 5 | 32 [0] | 17 (.85) | 32 [0] | 20 (1) | 32 [0] | 17 (.57) | 32 [0] | 29 (.97) |
| 6 | 14.6 [7.6] | 13 (.65) | 30 [25.6] | 17 (.85) | 25 [30.1] | 19 (.63) | 56.4 [43.0] | 25 (.83) |

Result 5.3 The proportion of time before a first breakthrough during which exactly one player plays risky is higher in the strategic treatment than in the control treatment.

Table 万shows the average proportion of time during which exactly one player is exploring before a first breakthrough by any player in his group. In each game, it is more than twice as large in the strategic treatment.

Table 7: PROPORTION OF TIME WITH A SINGLE PIONEER

|  | Strategic | Treatment | Control T | reatment | Strategic T | reatment | Control T | eatment |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Game | Eff. <br> Time | Single <br> Pioneer | Eff. <br> Time | Single <br> Pioneer | Eff. <br> Time | Single <br> Pioneer | Eff. <br> Time | Single <br> Pioneer |
| 1 | 155 [0] | . 724 [.156] | 146.5 [15.1] | . 284 [.258] | 147.7 [12.6] | . 670 [.178] | 134.6 [25.7] | . 097 [.156] |
| 2 | 185.9 [.2] | . 708 [.176] | 186 [0] | . 315 [.254] | 83.6 [21.9] | . 425 [.352] | 139 [67.6] | 0 [0] |
| 3 | 88 [0] | . 745 [.156] | 88 [0] | . 187 [.253] | 88 [0] | . 563 [.348] | 88 [0] | . 136 [.256] |
| 4 | 230 [0] | . 757 [.175] | 230 [0] | . 294 [.214] | 230 [0] | . 741 [.171] | 230 [0] | . 249 [.198] |
| 5 | 32 [0] | . 581 [.360] | 32 [0] | . 029 [.092] | 32 [0] | . 361 [.304] | 32 [0] | 0 [0] |
| 6 | 14.6 [7.6] | . 288 [.399] | 30 [25.6] | . 078 [.246] | 25 [30.1] | . 219 [.369] | 56.4 [43.0] | 0 [0] |

### 5.7 Switches of Actions

Cut-off behavior implies at most a single switch of action from risky to safe per player in a given game. However, players should switch roles at least once in any Markov Perfect Equilibrium. Hence, if players' behavior is predicted by MPE, we should expect significantly more switches in the strategic treatment. Recall that we have defined the incidence of switches as the number of a player's changes in action choice in a given game per unit of effective time.Recall that effective time is defined as the time elapsed before the game ends or the player's risky arm is revealed to be good, whichever happens first.

Result 5.4 The incidence of switches is significantly higher in the strategic treatment than in the control treatment. This holds for both $n=2$ and $n=3$.

Table 8 displays the average number of switches per player across games for our four treatments. ${ }^{[15}$ The incidence of switches in the strategic treatment is much higher than in the control treatment in all games except for Game 6 (all $p$-values of 0.0001 for Games 1-5 for either group size, with the exception of Game 4 in the setting with $n=3$ with $p$-value of 0.0011 ). As noted above, the early success in Game 6 reveals the risky arm to be good and thus resolves all uncertainty at the very beginning of the game. While still marginally higher incidences of switches are observed in the strategic treatment for both $n=2$ and $n=3$, this is not statistically significant ( $p$-values of 0.1751 and 0.1929 , respectively).

[^10]Table 8: AVERAGE NUMBER OF SWITCHES PER PLAYER

|  | $n=2$ |  |  |  | $n=3$ |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | Strategic Treatment |  | Control Treatment |  | Strategic Treatment |  | Control Treatment |  |
| Game | Eff. <br> Time | Switches <br> Per Pl. | Eff. <br> Time | Switches Per Pl. | Eff. <br> Time | Switches <br> Per Pl. | Eff. <br> Time | Switches Per Pl. |
| 1 | 155 [0] | 4.45 [1.85] | 146.5 [15.1] | . 90 [.91] | 147.7 [12.6] | 3.40 [1.87] | 134.6 [25.7] | 1.13 [1.68] |
| 2 | 185.9 [.2] | 4.50 [1.91] | 186 [0] | 1.35 [1.46] | 83.6 [21.9] | 2.77 [1.85] | 139 [67.6] | . 97 [1.50] |
| 3 | 88 [0] | 2.20 [1.11] | 88 [0] | . 30 [.47] | 88 [0] | 1.73 [1.46] | 88 [0] | . 47 [.82] |
| 4 | 230 [0] | 6.05 [2.04] | 230 [0] | 1.85 [1.90] | 230 [0] | 4.00 [3.02] | 230 [0] | 1.7 [1.97] |
| 5 | 32 [0] | . 60 [.50] | 32 [0] | . 05 [.22] | 32 [0] | . 70 [.84] | 32 [0] | . 03 [.18] |
| 6 | 14.6 [7.6] | . 60 [.88] | 30 [25.6] | . 30 [.80] | 25 [30.1] | . 97 [1.47] | 56.4 [43.0] | . 37 [.72] |

## 6 Conclusion

We have tested a game of strategic experimentation with bandits in the laboratory. As this involves a rather complex game, we of course cannot conclusively prove that subjects play, or aim to play, an MPE, as it is impossible to rule out other, more heuristic, forms of behavior. Yet, we believe that our results provide at least suggestive evidence for the main qualitative behavioral predictions of KRC's simple MPEs. Indeed, in a first step, we have exhibited strong evidence for strategic free-riding, as experimentation intensities are lower, and payoffs higher, in the strategic setting. Our eye-tracking data furthermore suggest that, in the strategic setting, subjects were paying keen attention to what their partners were up to. Moreover, subjects seem to attempt to coordinate in rather complex ways, as evidenced, inter alia, by the much lower incidence of cutoff behavior and the higher incidence of switches in the strategic setting. This, together with the greater prevalence of lonely pioneers, suggests that the qualitative aspects of subjects' behavior is more accurately predicted by KRC's Markov perfect equilibrium than by the average-payoff maximizing perfect Bayesian equilibrium of HKR.

KRC have also shown that there is a unique symmetric MPE in this game, which is characterized by players' using both arms at interior levels of intensity in the belief region where safe and risky are mutually best responses. In our experimental implementation, we do not allow subjects to pick experimentation intensities $k_{i} \in(0,1)$. While one could in principle imagine an experimental setup that does this (e.g. by letting subjects handle a gas pedal or a joystick), we have decided against doing so here in
order to keep the already highly complex game as simple as possible for our subjects. Yet we think that allowing for interior experimentation intensities would be an interesting robustness check to perform in future research, in order to see whether our results would continue to apply in this more complex setting, in particular as they pertain to the role changes between free-rider and pioneer that characterize the asymmetric MPEs of KRC.

Clearly, subjects do not precisely conform to equilibrium behavior. For instance, they seem loath to give up on the risky arm, extending experimentation to beliefs below the cutoff $p_{1}^{*}$. One might wonder whether they learned to play the game better over time, so that their behavior might converge to equilibrium with increasing experience. We find no evidence in our data to support this hypothesis. Recall that our subjects played the six games in random order. Subjects' behavior in the games they played later does not differ from their behavior in earlier games in any systematic way, be it in the control or the strategic treatment. We conjecture that this might be due to the subjects' inability to compute the relevant cutoffs, and to update their momentary beliefs, with sufficient precision, something they would be unlikely to learn over the course of six games. As a robustness check, one could in principle show subjects the current updated belief on their screens, in order to distinguish between the role of belief updating and that of determining the cutoffs. We have decided against doing so here, as we were concerned about prodding subjects toward certain behaviors, which would have made the interpretation of our results more difficult.

We have confined our analysis to the exponential-bandit setting of KRC. While the tractability of the exponential-bandit setting will certainly have facilitated its experimental implementation, the model does have some special features. For instance, as successes are fully revealing, there is no encouragement effect in KRC, which our experimental investigation confirms. Indeed, we can compute the average experimentation intensities in the region where safe is a dominant action, $\left[0, p_{1}^{*}\right]$, for Game 4 as well as for the two-player groups in Game 2. ${ }^{\boxed{6}]}$ Even in this region, the average experimentation intensity is lower in the strategic treatment: . 511 [.138] in the strategic treatment for Game 4 with $n=2$ vs. . 660 [.249] in the control treatment; .325 [.310] vs. 736 [.302] in Game 4 for $n=3$, and .510 [.270] vs. .743 [.262] in Game 2, where we report the standard deviation in square brackets. By contrast, if there were an encouragement effect, we should expect higher experimentation intensities in the strategic treatment for this belief region.

[^11]It might be interesting to test whether the encouragement effect can be shown in the laboratory for settings in which the theory would predict it to arise. This would be the case for instance in the Poisson setting with inconclusive breakthroughs à la Keller \& Rady (2010), or in the Brownian-motion setting of Bolton \& Harris (1999). It would also be intriguing to try and test the impact of privately observed actions or payoffs in the laboratory. We commend these questions for future research.

## References

Anderson, C. (2012): "Ambiguity Aversion in Multi-Armed Bandit Problems," Theory and Decision, 72, 15-33.

Anderson, C. (2001): "Behavioral Models of Strategies in Multi-Armed Bandit Problems," PhD thesis, California Institute of Technology.

Banks, J., M. Olson and D. Porter (1997): "An Experimental Analysis of the Bandit Problem," Economic Theory, 10, 55-77.

Bellman, R. (1956): "A Problem in the Sequential Design of Experiments," Sankhya: The Indian Journal of Statistics (1933-1960), Vol. 16, No. 3/4, 221-229.

Bergemann, D. and J. Välimäki (2008): "Bandit Problems," in: The New Palgrave Dictionary of Economics, 2nd edition ed. by S. Durlauf and L. Blume. Basingstoke and New York, Palgrave Macmillan Ltd.

Bergemann, D. and J. Välimäki (2000): "Experimentation in Markets," Review of Economic Studies, 67, 213-234.

Bergemann, D. and J. Välimäki (1996): "Learning And Strategic Pricing," Econometrica, 64, 1125-1149.

Bolton, P. and C. Harris (1999): "Strategic Experimentation," Econometrica, 67, 349374.

Bolton, P. and C. Harris (2000): "Strategic Experimentation: the Undiscounted Case," in: Incentives, Organizations and Public Economics - Papers in Honour of Sir James Mirrlees, ed. by P.J. Hammond and G.D. Myles. Oxford: Oxford University Press, 53-68.

Bonatti, A. and Hörner, J. (2011): "Collaborating," American Economic Review, 101(2), 632-663.

Bradt, R., S. Johnson and S. Karlin (1956): "On Sequential Designs for Maximizing the Sum of $n$ Observations," The Annals of Mathematical Statistics, 27, 10601074.

Calford, E. and R. Oprea (2017): "Continuity, Inertia, and Strategic Uncertainty: A Test of the Theory of Continuous Time Games," Econometrica, 85 (3), 915-935.

Coe, D. and E. Helpman (1995): "International R\&D Splillovers," European Economic Review, 39(5), 859-887.

Fischbacher, U. (2007): "z-Tree: Zurich Toolbox for Ready-Made Economic Experiments," Experimental Economics, 10(2), 171-178.

Gans, N., G. Knox and R. Croson (2007): "Simple Models of Discrete Choice and Their Performance in Bandit Experiments," Manufacturing and Service Operations Management, 9, 383--408.

Gittins, J. and D. Jones (1974): "A Dynamic Allocation Index for the Sequential Design of Experiments," in: Progress in Statistics, European Meeting of Statisticians, 1972, 1. Amsterdam: North-Holland, 241-266.

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

Griliches, Z. (1988): "Productivity Puzzles and R\&D: Another Nonexplanation," The Journal of Economic Perspectives, 2(4), 9-21.

Grossman, G. and E. Helpman (1993): Innovation and Growth in the Global Economy. MIT Press.

Hörner, J., N. Klein and S. Rady (2014): "Strongly Symmetric Equilibria in Bandit Games," Cowles Foundation Discussion Paper No. 1956.

Hörner, J. and A. Srzypacz (2016): "Learning, Experimentation and Information Design," mimeo.

Jacob, R. and K. Karn (2003): "Eye tracking in human-computer interaction and usability research: Ready to deliver the promises," Mind, 2(3), 573-605.

Keller, G. and S. Rady (2015): "Breakdowns," Theoretical Economics, 10, 175-202.
Keller, G. and S. Rady (2010): "Strategic Experimentation with Poisson Bandits," Theoretical Economics, 5, 275-311.

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

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

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

Meyer, R. and Y. Shi (1995): "Sequential Choice Under Ambiguity: Intuitive Solutions to the Armed-Bandit Problem," Management Science, 41(5), 817-834.

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

Poole, A., L. Ball and P. Phillips (2005): "In search of salience: A response-time and eye-movement analysis of bookmark recognition," People and Computers, 18, 363-378.

Presman, E.L. (1990): "Poisson Version of the Two-Armed Bandit Problem with Discounting," Theory of Probability and its Applications, 35, 307-317.

Robbins, H. (1952): "Some Aspects of the Sequential Design of Experiments," Bulletin of the American Mathematical Society, 58, 527-535.

Romer, P. (1990): "Endogenous Technological Change," Journal of Political Economy, 98(5), 71-102.

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

Rothschild, M. (1974): "A Two-Armed Bandit Theory of Market Pricing," Journal of Economic Theory, 9, 185-202.

Thompson, W. (1933): "On the Likelihood that One Unknown Probability Exceeds Another in View of the Evidence of Two Samples," Biometrika, 25, 285-294.

## Appendix

## A Screens Faced By Our Experimental Subjects

In this Appendix, we exhibit examples of the interfaces subjects saw during the game, showing the evolution of the screen over intervals of 30 seconds. In the top half (third) of his screen, a subject could see his own past actions and payoffs, while the bottom half (two thirds) of the screen showed his fellow group members' actions and payoffs. A blue (red) part of the payoff curve indicated that the player used the safe (risky) arm over the corresponding period. The x -axis represented calendar time, while the y -axis gave the player's cumulated total earnings up to each point in time. There was no prior indication of the point in time the game would end.
$n=2$ Strategic Setup: Example for Game 1



5ism

$\boxed{\infty}$
$n=2$ Control Setup: Example for Game 1




SAFE
$\square$ RISKY


5iswan
$\xrightarrow{3}$

$\omega$
$n=3$ Strategic Setup: Example for Game 2



5+

$\infty$
$n=3$ Control Setup: Example for Game 2



5-m
-

## B Instructions

The order of the instructions is as follows:

1. $n=2$ : Strategic Treatment
2. $n=2$ : Control Treatment
3. $n=3$ : Strategic Treatment
4. $n=3$ : Control Treatment

## Experiment Instructions

## Ground Rules

Welcome to the experiment. Please read the instructions carefully. The earnings you make in this experiment will be paid to you, in cash, at the end of the session.

Your earnings will be determined by your choices and the choices of other participants.
Communication between participants is not allowed. Please use only the computer to input your decisions. Please do not start or end any programs, and do not change any settings.

## How Groups are Organized

This experiment consists of six games in total. In the beginning of the first game, participants are randomly matched to pairs and the pairs stay the same in all six games. Therefore, in each game you will interact with the same participant.

## How the Timing Works

Games will last on average 120 seconds but may end at any time. The probability that the game ends is the same at each instant. Equivalently, the probability that the game ends during a given period of time depends only on the length of that period of time, and not on how long the game has already been going on. (Such processes are known as exponential processes in statistics.)

## How the Game Works

In every game, you have to decide whether you want to play the "safe" or the "risky" option. You can switch between the two options at any time and as often as you like by clicking on the safe (Blue) or risky (Red) button on the screen.

Whenever you choose the safe option, your payoff will increase for sure at the rate $\boldsymbol{E} \boldsymbol{\$ 1 0}$. That means the safe option will give you a reward of $\boldsymbol{E} \boldsymbol{\$} \mathbf{1 0}$ every second during which you use it.

When you choose the risky option, however, what you will be getting depends on the quality of that risky option. The quality of the risky option is determined by the computer once and for all at the start of each game; it never changes during the course of the game. We have programmed the computer so that the risky option will be good or bad with equal probability in each of the six games. The quality of the risky option in later games is independent of its quality in previous games. That is, in each of your six games, with probability $1 / 2$ your risky option will be good; with probability $1 / 2$ it will be bad. The same is true for your partner. Note that your risky option and that of your partner's might or might not be of the same quality.

If your risky option is good, it may give you a reward of $\boldsymbol{E} \$ \mathbf{2 5 0 0}$, but it will only ever do so if you use it. A good risky option yields such a reward after using it on average for $\mathbf{1 0 0}$ seconds. The probability that you get this reward from a good risky option during a given period of time during which you use it depends only on the length of that period of time; it does not depend on anything else, e.g. on how long the game has already been going on. Note that a good risky option may give you more than one reward of $E \$ 2500$ per game.

If your risky option is bad, it will never give you any reward.
You can switch back and forth between the risky option and the safe option at will and as many times as you like. All that matters for your chance of getting the reward is (1) the quality of the risky arm as determined by the computer before the game starts and (2) the overall amount of time you choose to spend on it.

The following graphic illustrates what you are going to see on your screen during the game. The graphs will be updated every second.


- The upper diagram always shows your actions and payoffs.
- In this example, you have started playing the risky option (highlighted in Red), then you have switched to the safe option (highlighted in Blue), then you have switched back again to the risky option, etc.
- The lower diagram always shows your partner's actions and payoffs.
- In this example, your partner has started playing the risky option and continues to do so.
- Note that, in this example, your partner's risky option was good and gave him once a reward of $\boldsymbol{E} \boldsymbol{\$} \mathbf{2 5 0 0}$.

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. In short:

Good risky option $>$ Safe option $>$ Bad risky option

Your partner is solving the exact same problem as you and has read the exact same instructions.

## Payment

In the experiment you will be making decisions that will earn you $\mathrm{E} \$$ (Experimental Dollars). At the end of the experiment, the $\mathrm{E} \$$ you earned will be converted into Australian Dollars at an exchange rate of $\mathrm{E} \$ 100=$ AU $\$ 1$, and paid out in cash. This amount will be added to your showup fee of AU\$ 5 .

After completing the experiment, the computer will randomly select one out of the six games (this will be the same game for all participants), and this game will then be used to determine your payoffs.

## Experiment Instructions

## Ground Rules

Welcome to the experiment. Please read the instructions carefully. The earnings you make in this experiment will be paid to you, in cash, at the end of the session.

Your earnings will be determined by your choices and the choices of other participants.
Communication between participants is not allowed. Please use only the computer to input your decisions. Please do not start or end any programs, and do not change any settings.

## How Groups are Organized

This experiment consists of six games in total. In the beginning of the first game, participants are randomly matched to pairs and the pairs stay the same in all six games. Therefore, in each game you will interact with the same participant.

## How the Timing Works

Games will last on average 120 seconds but may end at any time. The probability that the game ends is the same at each instant. Equivalently, the probability that the game ends during a given period of time depends only on the length of that period of time, and not on how long the game has already been going on. (Such processes are known as exponential processes in statistics.)

## How the Game Works

In every game, you have to decide whether you want to play the "safe" or the "risky" option. You can switch between the two options at any time and as often as you like by clicking on the safe (Blue) or risky (Red) button on the screen.

Whenever you choose the safe option, your payoff will increase for sure at the rate $\boldsymbol{E} \boldsymbol{\$ 1 0}$. That means the safe option will give you a reward of $\boldsymbol{E} \boldsymbol{\$} \mathbf{1 0}$ every second during which you use it.

When you choose the risky option, however, what you will be getting depends on the quality of that risky option. The quality of the risky option is determined by the computer once and for all at the start of each game; it never changes during the course of the game. We have programmed the computer so that the risky option will be good or bad with equal probability in each of the six games. The quality of the risky option in later games is independent of its quality in previous games. That is, in each of your six games, with probability $1 / 2$ your (and your partner's!) risky option will be good; with probability $1 / 2$ they will be bad. Note that your risky option and that of your partner's will always be of the same quality.

If your risky option is good, it may give you a reward of $\boldsymbol{E} \$ \mathbf{2 5 0 0}$, but it will only ever do so if you use it. A good risky option yields such a reward after using it on average for $\mathbf{1 0 0}$ seconds. The probability that you get this reward from a good risky option during a given period of time during which you use it depends only on the length of that period of time; it does not depend on anything else, e.g. on how long the game has already been going on. Note that a good risky option may give you more than one reward of $E \$ 2500$ per game.

If your risky option is bad, it will never give you any reward.
You can switch back and forth between the risky option and the safe option at will and as many times as you like. All that matters for your chance of getting the reward is (1) the quality of the risky arm as determined by the computer before the game starts and (2) the overall amount of time you choose to spend on it.

The following graphic illustrates what you are going to see on your screen during the game. The graphs will be updated every second.


- The upper diagram always shows your actions and payoffs.
- In this example, you have started playing the risky option (highlighted in Red), then you have switched to the safe option (highlighted in Blue), then you have switched back again to the risky option, etc.
- The lower diagram always shows your partner's actions and payoffs.
- In this example, your partner has started playing the risky option and continues to do so.
- Note that, in this example, your partner's risky option was good and gave him once a reward of $\boldsymbol{E} \mathbf{\$ 2 5 0 0}$. This means that your risky option was good too.

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. In short:

$$
\text { Good risky option }>\text { Safe option }>\text { Bad risky option }
$$

Your partner is solving the exact same problem as you and has read the exact same instructions. Note that by observing the behaviour of his risky option (provided he uses it) you can learn something about your risky option as well.

## Payment

In the experiment you will be making decisions that will earn you $\mathrm{E} \$$ (Experimental Dollars). At the end of the experiment, the E\$ you earned will be converted into Australian Dollars at an exchange rate of $\mathrm{E} \$ 100=\mathrm{AU} \$ 1$, and paid out in cash. This amount will be added to your showup fee of AU\$ 5 .

After completing the experiment, the computer will randomly select one out of the six games (this will be the same game for all participants), and this game will then be used to determine your payoffs.

## Experiment Instructions

## Ground Rules

Welcome to the experiment. Please read the instructions carefully. The earnings you make in this experiment will be paid to you, in cash, at the end of the session.

Your earnings will be determined by your choices and the choices of other participants.
Communication between participants is not allowed. Please use only the computer to input your decisions. Please do not start or end any programs, and do not change any settings.

## How Groups are Organized

This experiment consists of six games in total. In the beginning of the first game, participants are randomly matched to groups of three players and the groups stay the same in all six games. Therefore, in each game you will interact with the same participants.

## How the Timing Works

Games will last on average $\mathbf{1 2 0}$ seconds but may end at any time. The probability that the game ends is the same at each instant. Equivalently, the probability that the game ends during a given period of time depends only on the length of that period of time, and not on how long the game has already been going on. (Such processes are known as exponential processes in statistics.)

## How the Game Works

In every game, you have to decide whether you want to play the "safe" or the "risky" option. You can switch between the two options at any time and as often as you like by clicking on the safe (Blue) or risky (Red) button on the screen.

Whenever you choose the safe option, your payoff will increase for sure at the rate $\boldsymbol{E} \boldsymbol{\$ 1 0}$. That means the safe option will give you a reward of $\boldsymbol{E} \boldsymbol{\$ 1 0}$ every second during which you use it.

When you choose the risky option, however, what you will be getting depends on the quality of that risky option. The quality of the risky option is determined by the computer once and for all at the start of each game; it never changes during the course of the game. We have programmed the computer so that the risky option will be good or bad with equal probability in each of the six games. The quality of the risky option in later games is independent of its quality in previous games. That is, in each of your six games, with probability $1 / 2$ your risky option will be good; with probability $1 / 2$ it will be bad. The same is true for your partners. Note that your risky option and that of your partners' might or might not be of the same quality.

If your risky option is good, it may give you a reward of $\boldsymbol{E} \$ \mathbf{2 5 0 0}$, but it will only ever do so if you use it. A good risky option yields such a reward after using it on average for $\mathbf{1 0 0}$ seconds. The probability that you get this reward from a good risky option during a given period of time during which you use it depends only on the length of that period of time. It does not depend on anything else, e.g. on how long the game has already been going on. Note that a good risky option may give you more than one reward of $E \$ 2500$ per game.

If your risky option is bad, it will never give you any reward.
You can switch back and forth between the risky option and the safe option at will and as many times as you like. All that matters for your chance of getting the reward is (1) the quality of the risky arm as determined by the computer before the game starts and (2) the overall amount of time you choose to spend on it.

The following graphic illustrates what you are going to see on your screen during the game. The graphs will be updated every second.


- The upper diagram always shows your actions and payoffs.
- In this example, you have started playing the risky option (highlighted in Red), then you have switched to the safe option (highlighted in Blue), then you have switched back again to the risky option, etc.
- The lower diagram always shows your partner's actions and payoffs.
- In this example, one of your partners has started playing the risky option and continues to do so. The other partner has started and continues playing the safe option.
- Note that, in this example, at least one of your partner's risky option was good and gave him once a reward of $\boldsymbol{E} \$ 2500$.

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. In short:

$$
\text { Good risky option }>\text { Safe option }>\text { Bad risky option }
$$

Your partners are solving the exact same problem as you and have read the exact same instructions.

## Payment

In the experiment you will be making decisions that will earn you $\mathrm{E} \$$ (Experimental Dollars). At the end of the experiment, the E\$ you earned will be converted into Australian Dollars at an exchange rate of $\mathrm{E} \$ 100=\mathrm{AU} \$ 1$, and paid out in cash. This amount will be added to your showup fee of AU\$ 5 .

After completing the experiment, the computer will randomly select one out of the six games (this will be the same game for all participants), and this game will then be used to determine your payoffs.

## Experiment Instructions

## Ground Rules

Welcome to the experiment. Please read the instructions carefully. The earnings you make in this experiment will be paid to you, in cash, at the end of the session.

Your earnings will be determined by your choices and the choices of other participants.
Communication between participants is not allowed. Please use only the computer to input your decisions. Please do not start or end any programs, and do not change any settings.

## How Groups are Organized

This experiment consists of six games in total. In the beginning of the first game, participants are randomly matched to groups of three players and the groups stay the same in all six games. Therefore, in each game you will interact with the same participants.

## How the Timing Works

Games will last on average $\mathbf{1 2 0}$ seconds but may end at any time. The probability that the game ends is the same at each instant. Equivalently, the probability that the game ends during a given period of time depends only on the length of that period of time, and not on how long the game has already been going on. (Such processes are known as exponential processes in statistics.)

## How the Game Works

In every game, you have to decide whether you want to play the "safe" or the "risky" option. You can switch between the two options at any time and as often as you like by clicking on the safe (Blue) or risky (Red) button on the screen.

Whenever you choose the safe option, your payoff will increase for sure at the rate $\boldsymbol{E} \boldsymbol{\$ 1 0}$. That means the safe option will give you a reward of $\boldsymbol{E} \boldsymbol{\$ 1 0}$ every second during which you use it.

When you choose the risky option, however, what you will be getting depends on the quality of that risky option. The quality of the risky option is determined by the computer once and for all at the start of each game; it never changes during the course of the game. We have programmed the computer so that the risky option will be good or bad with equal probability in each of the six games. The quality of the risky option in later games is independent of its quality in previous games. That is, in each of your six games, with probability $1 / 2$ your (and your partners'!) risky option will be good; with probability $1 / 2$ they will be bad. Note that your and your partners' risky option will always be of the same quality.

If your risky option is good, it may give you a reward of $\boldsymbol{E} \$ \mathbf{2 5 0 0}$, but it will only ever do so if you use it. A good risky option yields such a reward after using it on average for $\mathbf{1 0 0}$ seconds. The probability that you get this reward from a good risky option during a given period of time during which you use it depends only on the length of that period of time. It does not depend on anything else, e.g. on how long the game has already been going on. Note that a good risky option may give you more than one reward of $E \$ 2500$ per game.

If your risky option is bad, it will never give you any reward.
You can switch back and forth between the risky option and the safe option at will and as many times as you like. All that matters for your chance of getting the reward is (1) the quality of the risky arm as determined by the computer before the game starts and (2) the overall amount of time you choose to spend on it.

The following graphic illustrates what you are going to see on your screen during the game. The graphs will be updated every second.


- The upper diagram always shows your actions and payoffs.
- In this example, you have started playing the risky option (highlighted in Red), then you have switched to the safe option (highlighted in Blue), then you have switched back again to the risky option, etc.
- The lower diagram always shows your partners' actions and payoffs.
- In this example, one of your partners has started playing the risky option and continues to do so. The other partner has started and continues playing the safe option.
- Note that, in this example, your partner's risky option was good and gave him once a reward of $\boldsymbol{E} \$ \mathbf{2 5 0 0}$. This means that your risky option was good too.

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. In short:

$$
\text { Good risky option }>\text { Safe option }>\text { Bad risky option }
$$

Your partners are solving the exact same problem as you and have read the exact same instructions. Note that by observing the behaviour of their risky option (provided they use it) you can learn something about your risky option as well.

## Payment

In the experiment you will be making decisions that will earn you $\mathrm{E} \$$ (Experimental Dollars). At the end of the experiment, the E\$ you earned will be converted into Australian Dollars at an exchange rate of $\mathrm{E} \$ 100=\mathrm{AU} \$ 1$, and paid out in cash. This amount will be added to your showup fee of AU\$ 5 .

After completing the experiment, the computer will randomly select one out of the six games (this will be the same game for all participants), and this game will then be used to determine your payoffs.


[^0]:    *Financial support by the UNSW Bizlab and the School of Economics at UNSW Sydney is gratefully acknowledged. Nicolas Klein also gratefully acknowledges financial support from the Fonds de Recherche du Québec - Société et Culture and the Social Sciences and Humanities Research Council of Canada. This research has been approved by the Human Research Ethics Committee of UNSW Sydney under approval number HC 17069. We thank Francisco Alvarez-Cuadrado, Michele De Nadai, Denzil Fiebig, Gigi Foster, Raphael Godefroy, Gabriele Gratton, Ben Greiner, Richard Holden, Hongyi Li, Sven Rady and Tom Wilkening for helpful comments. A substantial part of this paper took shape while both authors enjoyed the hospitality of the Institute for Markets and Strategy at WU Vienna.
    ${ }^{\dagger}$ UNSW Sydney. Mailing address: UNSW Business School, School of Economics, Anzac Parade 2052 NSW, Australia; email: j.hoelzemann@unsw.edu.au.
    ${ }^{\ddagger}$ Université de Montréal and CIREQ. Mailing address: Université de Montréal, Département de Sciences Économiques, C.P. 6128 succursale Centre-ville; Montréal, H3C 3J7, Canada. Telephone: +1-514-343-7908; email: kleinnic@yahoo.com.

[^1]:    ${ }^{1}$ These are perfect Bayesian equilibria where a player's action choice depends on the history only via the common posterior belief.

[^2]:    ${ }^{2}$ The encouragement effect has been identified by Bolton \& Harris (1999) and is not predicted to arise in the KRC setting. By virtue of this effect, players experiment more than if they were by themselves. They

[^3]:    ${ }^{4} \mathrm{KRC}$ show that there is also a unique symmetric MPE, where players use the risky (safe) arm with an interior intensity $k(p) \in(0,1)(1-k(p))$ throughout the belief region where risky and safe are mutually best responses. As we wanted to keep the decision problem as simple as possible, our subjects do not have the option of choosing interior experimentation levels. Please also see our discussion in the Conclusion.
    ${ }^{5}$ Subjects knew that the end time of the game corresponded to the first jumping time of a Poisson process with parameter $r$ but did not know the realization of this process at any time before the game ended. In particular, the time axis they saw on their computer screens gradually grew longer as time progressed, so that they could not infer the end date. Please see the Appendix for details and for the instructions the subects received.
    ${ }^{6}$ As all our stochastic processes are Lévy processes, simulating their realizations ahead of time is equivalent to simulating them as the game progresses. In order to increase the computational efficiency of the implementation, we chose to simulate them ahead of time.
    ${ }^{7}$ Details are available from the authors upon request.
    ${ }^{8}$ Thus, our implementation corresponds to the "Inertial Continuous-Time" setting in Calford \& Oprea (2017).

[^4]:    ${ }^{9}$ The instructions handed out to all participants can be found in the Appendix.

[^5]:    ${ }^{10}$ Note that if players were to play the best PBE and the game happened to stop at a time such that $p_{1}^{*}$ is only reached in the strategic treatment, we should observe exactly one switch per player in the strategic treatment and none in the control treatment. Therefore, a higher number of switches in the strategic treatment is not inconsistent with players' playing the best PBE. However, the magnitude of the effect, which we report in Section 5 , cannot be accounted for by this explanation.

[^6]:    ${ }^{11}$ Strictly speaking, the Wilcoxon ranksum test treats players' action choices as independent observations. Yet, one might of course argue that, in the strategic setting, players' action choices are not indepen-

[^7]:    dent of each other. As a robustness check, we also report $p$-values separately for Players 1, 2, and 3, respectively, and find that this does not impact our conclusions at all. Indeed, in the case of $n=2$, the corresponding $p$-values for Player 1 (Player 2) are 0.0675 ( 0.0638 ), 0.0230 ( 0.5964 ), 0.0092 ( 0.0018 ), 0.0887 ( 0.1983 ), $0.0166(0.0051)$, and $0.6574(0.0682)$ for Games 1-6, respectively. For $n=3$, the corresponding $p$-values for Player 1 (Player 2) [Player 3] are 0.0149 ( 0.0125 ) [0.0052], 0.0120 ( 0.4270 ) [0.0001], 0.0068 ( 0.0226 ) [0.0015], 0.0281 ( 0.0025 ) [0.0152], 0.0052 ( 0.0323 ) [ 0.0051 ], and 0.3514 ( 0.4868 ) [0.5347] for Games 1-6, respectively.
    ${ }^{12}$ Besides the beliefs $\left(\frac{1}{2}, 1\right]$, which can never be reached in the absence of a success, the complementary set of these beliefs thus consists of the region $\left[0, p_{1}^{*}\right]$, where safe is a dominant action, and the (small) interval of beliefs $\left[p^{\ddagger}, \bar{p}\right)$, which we have not assigned to either region. Indeed, as we explain in Section 4. we rely on conservative bounds in defining the "R dominant"and "Mutually BR" regions.

[^8]:    ${ }^{13}$ If we eliminate a single outlier in Game 4, namely player 1 in Group 13, our test yields a $p$-value of

[^9]:    ${ }^{14}$ Video recordings illustrating the use of the eye-tracking devices are available at www.johanneshoelzemann.com.

[^10]:    ${ }^{15}$ As there is rather little variation in effective time between the strategic and the control treatments for a given game, we have decided to report the average number, rather than the average incidence, of switches in Table 8, as the former may be easier to interpret.

[^11]:    ${ }^{16}$ These are the only settings in which this region is reached (and lasts for more than a few seconds) for both the strategic and the control treatments.

