

Brütt, Katharina

Working Paper

How teams overcome free-riding in strategic experimentation

Tinbergen Institute Discussion Paper, No. TI 2025-049/I

Provided in Cooperation with:

Tinbergen Institute, Amsterdam and Rotterdam

Suggested Citation: Brütt, Katharina (2025) : How teams overcome free-riding in strategic experimentation, Tinbergen Institute Discussion Paper, No. TI 2025-049/I, Tinbergen Institute, Amsterdam and Rotterdam

This Version is available at:

<https://hdl.handle.net/10419/331381>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

TI 2025-049/I
Tinbergen Institute Discussion Paper

How teams overcome free-riding in strategic experimentation

*Katharina Brütt*¹

Tinbergen Institute is the graduate school and research institute in economics of Erasmus University Rotterdam, the University of Amsterdam and Vrije Universiteit Amsterdam.

Contact: discussionpapers@tinbergen.nl

More TI discussion papers can be downloaded at <https://www.tinbergen.nl>

Tinbergen Institute has two locations:

Tinbergen Institute Amsterdam
Gustav Mahlerplein 117
1082 MS Amsterdam
The Netherlands
Tel.: +31(0)20 598 4580

Tinbergen Institute Rotterdam
Burg. Oudlaan 50
3062 PA Rotterdam
The Netherlands
Tel.: +31(0)10 408 8900

How teams overcome free-riding in strategic experimentation

Katharina Brütt*

Vrije Universiteit Amsterdam & Tinbergen Institute

August 26, 2025

Abstract

Experimentation is at the core of innovation. This project studies collaborative experimentation in teams, focusing on the inherent two-dimensional free-riding problem induced by payoff and informational externalities. The discouraging force of observing others' unsuccessful experimentation and the attempt to keep one's team members optimistic theoretically result in inefficiently low experimentation in teams. In a laboratory experiment, I study how distinct elements of the experimentation environment affect this strategic experimentation. I vary (i) the observability of experimentation and (ii) whether agents work on joint or separate projects. Teams largely overcome the free-riding problem. Contrary to theoretical predictions, both the observability of experimentation and experimenting jointly increase experimentation levels. There is no lack of sophistication in updating beliefs that drives this, neither do subjects disregard their experimentation's effect on others. Instead, the data can be best explained by joint, observable experimentation creating incentives to 'lead by example' and setting norms of high experimentation.

JEL Classification: D81, D83, D91

Keywords: Strategic experimentation; Free-riding; Informational externality; Experiment.

*k.bruett@vu.nl

This research was funded by the Research Priority Area Behavioral Economics of University of Amsterdam and by the NWO through Research Talent Grant 406.18.501. The study is pre-registered at the AEA RCT Registry with ID AEARCTR-0005503 and received ethical approval by the IRB of the University of Amsterdam. I would like to thank my advisors Arthur Schram and Joep Sonnemans for their support and advice in all stages of this project. Furthermore, I would like to thank Rebecca Morton, David Levine, and Marie Claire Villeval for valuable comments on the experimental design and draft, and conference attendants at the EYAE annual meeting, the FUR conference, the EEA-ESEM congress, the European ESA meeting, the IAREP/SABE conference, and the CCC-Meeting (University of East Anglia) for their helpful suggestions.

1 Introduction

Innovation is key to overcoming today’s most pressing issues, from fighting climate change to tackling the COVID-19 pandemic. Estimates by the International Energy Agency suggest that almost half of the required emission reductions on the road to net-zero emissions by 2050 will require technologies that are not yet available (Bouckaert et al., 2021). During the height of the COVID-19 pandemic, scientific teams, companies, governmental agencies, and civil organisations worldwide worked on creating effective vaccines to combat COVID-19, helping the economy, improving communal lives in times of social distancing, and keeping people healthy (see e.g. Kretchmer, 2020).

Experimentation, that is testing new technologies with uncertain outcomes, is central to innovation (Thomke, 2003). Innovative team projects often require a process of trial and error with the risk of spending time and resources on a doomed project. Consider a group of employees working on a business innovation or researchers developing a new technology. As an illustrating example from the medical domain, think of a team in a pharmaceutical company attempting to develop a new vaccine (i.e., the team is *experimenting*). Not all attempts will lead to success, and some will entail a waste of resources, which could have been spent on other projects with a certain reward. Recent estimates suggest that only 13.8% of all drug development programs resulted in FDA approval (Wong et al., 2019).

This project studies how the experimentation environment can be designed to encourage experimentation in teams. Large innovations usually require a team of individuals to experiment together, such as teams within companies or scientific collaborations. In the last decades, such collaborations increased substantially (Dong et al., 2017). If teams are experimenting, individual contributions to a project may provide a public good to all team members. In the illustrating example, if one team member observes that a certain vaccine is effective (i.e., there is a *breakthrough*), all individuals in the team may benefit from their team’s success (i.e., the breakthrough is a *public good*). Crucially, this public good entails two dimensions: The information created from observing successful experimentation, and the payoffs benefiting all team members.

This can result in a two-dimensional free-riding problem. First, there is the well-known moral hazard in teams problem, going back to Holmstrom (1982). Agents prefer that other team members invest their resources, such as time or individual budgets, in a project compared to investing themselves if the payoff of a breakthrough is shared among all team members. This causes lower experimenta-

tion investment than socially optimal. In addition, informational spillovers change the incentives for experimentation. Acquiring information about a project's quality yourself, such as conducting a trial of the effectiveness of a vaccine, is costly, so individuals prefer to use the information generated by others. Hence, if the information created through experimentation is public, there is an informational externality that results in a free-riding problem. This also leads to inefficiently low experimentation efforts (Bonatti and Hörner, 2011). In this study, I investigate in the laboratory which types of experimentation environments allow agents to overcome these free-rider problems.

Informational spillovers pose challenges for individuals experimenting in a team. Individuals need to anticipate how their fellow team members will react to the information generated by their own experimentation efforts. At the same time, they also have to carefully consider what other team members' actions reveal about the quality of the project they are engaged with and respond to this information. In this paper, I will empirically examine how individuals handle these challenges and whether teams can overcome free-riding problems when experimenting with projects of uncertain quality. I will focus on the dimensions of experimentation that are specifically relevant to teams.

The first crucial aspect of experimentation in teams is that how agents learn from the experimentation of team members depends on whether actions are observable. Distinct settings vary in the observability of experimentation effort. Some teams may find it easier than others to observe their team member's input in a group project. These settings are theoretically well understood, but less so empirically. It is thus important to gain a better understanding of how the provision of information on experimentation efforts changes behaviour.

At the same time, there are differences in the extent to which one team member's success is predictive of another team member's likelihood of success. This is the next aspect under consideration. Specifically, I will study settings where agents either work on separate, independent projects or jointly work on one project. In the former case, the success of one team member does not provide any information for others. Going back to the leading example, the individual team members may explore distinct technologies to develop a vaccine to combat a certain disease. The latter case, with one joint project, is the polar opposite. Here, the experimentation of all team members is equally informative, and informational spillovers are a natural part of the environment. In the example, joint experimentation translates to all team members relying on the same technology when developing a vaccine.

The experimental design builds on a simple theoretical model, close to an ex-

ample in Bonatti and Hörner (2011). I employ a two-stage variant with two agents to focus on how individuals utilize information provided by others and how they take into account the information they themselves generate. A breakthrough is a public good and results in a positive payoff for all team members. In this model, a breakthrough reveals that the project is of high quality. If no breakthrough occurs, agents can continue experimenting, but should realize that the project is now less likely of high quality.

The model allows for two observations when considering joint experimentation with a common project. First, the optimal experimentation effort will be inefficiently low, because agents do not internalise the positive externalities they have on others when choosing their effort. Second, an agent's current effort choice and the other agent's future effort choice are strategic substitutes. This implies that, if effort is observable, high effort levels are unattractive, since the other agent will interpret the fact that no breakthrough occurred as a strong negative signal that makes her more pessimistic about the quality of the project, reducing future effort provision. As agents anticipate this force when experimentation early on, they are discouraged from choosing high experimentation levels, the so-called discouragement effect. The observability of experimentation thus increases inefficiencies.

Contrary to this theoretical channel, several behavioural factors would suggest that the observability of experimentation will not result in lower levels of experimentation. This paper will systematically study these factors. First, the discouragement effect hinges on agents updating their beliefs in response to the information created by others, which agents may do insufficiently. Myopic behaviour can lead to agents disregarding the effect of their early experimentation on a later stage. Second, conditional cooperation or reciprocity can result in agents encouraging each others' experimentation, and punishing low experimentation, if effort is observable. Last, the observability of experimentation efforts may allow an agent to 'lead by example', signalling the belief that experimentation is lucrative.

To study the mechanisms that drive strategic experimentation, I contrast the setting of joint experimentation to one of separate experimentation. If team members experiment with separate projects, there is no informational externality. Therefore, the discouragement effect disappears. With separate experimentation, the theoretical predictions flip. Since agents cannot discourage each other from experimenting, a new effect dominates that is particular to a setting with separate projects. Successful innovation only requires one breakthrough in one project. Agents should therefore respond positively to others' high experimentation, as a more pessimistic partner reduces the likelihood of several simultaneous, and there-

fore inefficient, breakthroughs.

The experiment employs a 2-by-2 between-subject design closely following the theoretical setup, varying the observability of experimentation effort and whether the team members experiment with joint or separate projects. In a two-stage setting, participants' first-stage experimentation allows them to update their beliefs about the project's quality. The updated beliefs permit them to make a more informed second-stage experimentation decision. The second stage is the final stage. Therefore, decisions and outcomes from this stage do not entail any informational value. Decisions in this stage therefore inform behaviour without informational externalities. Furthermore, I elicit beliefs about the quality of the project and the partner's effort provision to disentangle different drivers of effort provision.

The experiment provides two key insights. First, in stark contrast to the theoretical prediction, the observability of experimentation reduces joint experimentation efforts. Instead, observability increases joint experimentation levels. This cannot be traced back to a lack of sophistication in belief updating. Qualitatively, subjects respond to experimentation as predicted, both when updating their beliefs and when choosing their future actions. Beliefs are updated conservatively, but in the expected direction. In addition, second-order beliefs are consistent with subjects even anticipating this response from others. Thus, there is evidence of a discouragement effect. Nevertheless, agents do not reduce their early experimentation if this is observable as a response. Hence, agents behave partially myopic. At the same time, there is no convincing evidence pointing at reciprocity as a driver of higher experimentation levels if these are observable. Other than through a change in beliefs, first-stage experimentation does not impact the partner's second-stage experimentation, which would have indicated conditionally cooperative behaviour. Second, experimentation efforts are considerably higher if individuals experiment with a joint project than with separate projects. There is no detectable difference in the response to observable effort compared to the case of joint experimentation. Therefore, the advantage of observable experimentation exists irrespective of the presence of an informational externality.

The two factors that increase experimentation are thus 1) experimenting jointly and 2) experimentation being observable. These factors share that they increase the salience of group membership. If agents observe their partner's action, they are made aware that they are not working on their own. Similarly, if agents work on the same project, this shares more noticeable elements of team productions. The observed patterns are, moreover, in line with agents creating norms of high experimentation and agents leading by example to foster such norms. Agents

respond positively to high experimentation by others in later rounds, and the variance of experimentation levels is reduced if these are observable.

The remainder of this paper proceeds as follows: Section 2 will give a brief overview of the related literature. Section 3 outlines the theoretical model that underlies the experimental design and gives its predictions, Section 4 provides the experimental design. Finally, Section 5 discusses the experimental results, and Section 6 concludes.

2 Related literature

The first bandit models of experimentation go back to Bolton and Harris (1999). Hörner and Skrzypacz (2017) review the core models in the strategic experimentation literature. These models focus on the trade-off between experimentation and exploitation in a continuous-time setting. In models of strategic experimentation, several agents face the same slot machines, so-called bandits, with uncertain payoffs. Players can learn about the underlying payoff processes by observing the outcomes of their own and others' experimentation. Arriving news comes either as breakthroughs (Keller et al., 2005), or as breakdowns (Keller and Rady, 2015). Breakthroughs have positive, breakdowns negative payoff consequences. With some exceptions (e.g. Keller and Rady, 2010), both usually provide conclusive evidence about the payoff process, which I will also focus on.¹

The study by Bonatti and Hörner (2011) is closest to the setup studied in this paper. Bonatti and Hörner (2011) introduce breakthroughs that provide a public good. In this environment, informational and payoff externalities co-exist, creating the two-dimensional free-riding problem discussed in the introduction. Theoretically, this setup induces both free-riding and delay of experimentation. Furthermore, monitoring the other agent by observing their experimentation does not reduce delay or free-riding, as this would imply that agents discourage each other from experimenting. I will provide an experimental test of this conclusion. Adding to this, I will theoretically and experimentally contrast the case of joint experimentation, studied in Bonatti and Hörner (2011), to the case where team members experiment with separate projects.

So far, experimental tests of the theoretical predictions are scarce and primarily focus on individual, not strategic experimentation. In the laboratory, agents frequently undervalue experimentation when facing individual bandit problems

¹For a discussion and comparison of the theoretical properties of these models, see Hörner and Skrzypacz (2017).

(Meyer and Shi, 1995), which can be driven by risk aversion (Hudja and Woods, 2021) or ambiguity attitudes (Anderson, 2012).² The theoretical predictions of the strategic experimentation literature have not been widely tested, though there are some notable exceptions. In a test of the model of Keller et al. (2005), there is substantial free-riding on others' experimentation (Hoelzemann and Klein, 2021). In this setting, experimentation in groups can be sustained at more pessimistic beliefs than theoretically predicted (Kwon, 2020). This implies that groups generate more information than individuals, which is in contrast to theoretical predictions. However, free-riding on others' information provision can emerge as well, and in a simpler two-stage setup as employed in this paper, participants also under-experiment compared to the theoretical predictions (Boyce et al., 2016).

There is so far no evidence on how distinct elements of strategic experimentation, such as the existence of an informational externality or whether experimentation is observable, impact experimentation efforts. In contrast to existing studies, I test the comparative statics of how behaviour depends on the observability of experimentation effort in a setting where payoff externalities exist. I do not aim at giving insights into the dynamics of behaviour in a continuous-time setting; instead, I employ a simpler discrete-time setting and focus on the determinants of experimentation. Furthermore, I am interested in how potential biases in belief formation drive experimentation, which could not be clearly studied in earlier work, where belief updating was trivial (Kwon, 2020).

There are recent experimental and theoretical papers that look at collaborative search. Search differs from strategic experimentation and the setup that is studied in this paper, because agents are not exploring the merits of one particular policy or technology but explore a set of such items. The overarching questions, nevertheless, are similar. Both collaborative search and strategic experimentation look at how teams can innovate. In collaborative search, however, agents encourage each others' search, as a breakthrough becomes more likely when more projects have been examined. In a setting of collaborative search with payoff externalities, imperfect optimisation and other-regarding preferences influence agents' experimentation (von Essen et al., 2020).

This research also relates to the public good literature. Experimenting increases the probability of a breakthrough, which constitutes a public good. Therefore, findings in the public goods literature could help us understand the be-

²There is mixed evidence on whether participants respond to parameter changes, such as changes of the discount rate and prior beliefs, in their experimentation efforts (Banks et al., 1997; Hudja and Woods, 2021).

havioural drivers of experimentation. In the experimental public good literature, many people are conditional cooperators and match contributions by others (see e.g. Fischbacher et al., 2001; Kocher et al., 2008; Thöni and Volk, 2018; Croson et al., 2005). The presence of conditional cooperators would imply that agents experiment more if they see others experiment more as well.

In public good games, the salience of group membership increases the weight that agents put on payoffs for their group members (Charness et al., 2007; Sutter, 2009). Changing the observability of others' actions and whether partners work on the same project will likely also increase the salience of group membership and could therefore increase experimentation. Similarly, observable experimentation may provide incentives to 'lead by example', as observed in public good experiments (Vesterlund, 2003; Potters et al., 2005, 2007; Güth et al., 2007; Levati et al., 2007). If agents observe each others' experimentation, it may prove beneficial to set high levels of experimentation to encourage future experimentation by others, either by signalling the profitability of experimentation, or by creating norms of high experimentation.

3 Theoretical framework

In this section, I will introduce the theoretical model underlying the experimental design, illustrating the drivers of strategic experimentation. The theoretical framework builds on a variant of a two-stage model from Bonatti and Hörner (2011). Applying a simpler setting allows me to focus on how individuals create information and utilize the information provided by others and by themselves. I study a model where a breakthrough is a public good. Therefore, everyone in a team receives a positive payoff if a team member achieves a breakthrough. News is always good, as breakthroughs reveal that the project is of high quality. If no breakthrough occurs in the first stage, agents can continue experimenting. In this two-stage model, each stage facilitates the analysis of a distinct element of strategic experimentation. First-stage experimentation captures that agents generate new information and that their experimentation entails an informational externality. This is the core element of experimentation. Second-stage experimentation captures the response to the information previously created. As the second stage is the final stage, no informational value can be generated. Therefore, there is also no informational externality.

First, I consider a setting where teams experiment jointly with one project. In a second step, I adapt this model to encompass teams in which team members

experiment with separate projects to achieve a breakthrough. For both cases, I will differentiate between a setting where experimentation efforts are observable to the other team member and a setting where these efforts are not observable.

3.1 Experimenting with a joint project

This two-stage model studies joint experimentation. There are two agents $i = 1, 2$ who can choose to invest experimentation effort $e_{i,t} \in [0, 1]$ in two stages $t = 1, 2$ in a joint project with unknown quality. Doing so entails a private cost of effort of $c(e_{i,t}) = 2e_{i,t}^2$. Both agents receive a payoff of $Y = 13$ from the project if a breakthrough occurs. A breakthrough terminates the project. Whether a breakthrough occurs depends on the quality of the project, which can be high or low, and on the effort the two agents invest in that project. The common prior that the project is of high quality is $p = 0.5$. Conditional on the project being of high quality, the probability that a breakthrough occurs in stage t is given by $\frac{e_{i,t} + e_{-i,t}}{2}$, which is increasing in the effort invested by both agents. If the project is of low quality, there will never be a breakthrough.

The experimentation effort by the two team members is either observable or not. I first consider the case in which agents observe their team member's level of first-stage experimentation before choosing their second-stage experimentation levels. In the second stage of the experiment, agents maximise the following expected utility:

$$EU_{i,2} = \rho(e_{i,1}, e_{-i,1}) \left(\frac{e_{i,2} + e_{-i,2}}{2} \right) Y - c(e_{i,2}) \quad (1)$$

This stage is only reached if there was no breakthrough. $\rho(e_{i,1}, e_{-i,1})$ is the posterior belief that the project is of high quality. By Bayes' rule, this is given by

$$\rho(e_{i,1}, e_{-i,1}) = \frac{p \left(1 - \frac{e_{i,1} + e_{-i,1}}{2} \right)}{1 - p \left(\frac{e_{i,1} + e_{-i,1}}{2} \right)} \leq p$$

Here, realise that the posterior is decreasing both in the agent's own and in their partner's first-stage experimentation:

$$\frac{\partial \rho(e_{i,1}, e_{-i,1})}{\partial e_{-i,1}} \leq 0 \quad \text{and} \quad \frac{\partial \rho(e_{i,1}, e_{-i,1})}{\partial e_{i,1}} \leq 0, \quad \forall p \in [0, 1]$$

Intuitively, if there is no breakthrough but experimentation efforts are high, it is less likely that the project is of high quality. If the project were of high quality, a

breakthrough would have been likely.

In the first stage, agents maximise the expected utility over the two stages, taking into account how first-stage experimentation affects second-stage experimentation. The first-stage expected utility as a function of the second-stage expected utility $EU_{i,2}$ is given by

$$EU_{i,1} = p \left(\frac{e_{i,1} + e_{-i,1}}{2} \right) Y - c(e_{i,1}) + \left(1 - p \left(\frac{e_{i,1} + e_{-i,1}}{2} \right) \right) \times EU_{i,2}$$

I will consider experimentation behaviour in the pure-strategy symmetric Perfect Bayesian Nash equilibrium (PBE). Similar to Bonatti and Hörner (2011), this setup allows us to make the following observations:

Lemma 1. *An agent's first- and second-stage experimentation efforts are strategic substitutes.*

An agent's second-stage effort is an increasing function of the posterior belief of the project's quality $\rho(e_{i,1}, e_{-i,1})$. Second-stage experimentation promises to pay off more likely if agents are optimistic about the project's quality. The posterior $\rho(e_{i,1}, e_{-i,1})$ is decreasing in the first-stage experimentation. Therefore, the second-stage experimentation decreases in the first-stage experimentation. Intuitively, agents become more pessimistic about the project's quality if they exerted high experimentation efforts in the first stage but do not observe a breakthrough. As a consequence, they exert less effort in the second stage.

The same underlying reasoning applies to the strategic substitutability of own experimentation and the partner's experimentation across stages.

Lemma 2. *An agent's and her partner's first- and second-stage experimentation efforts are strategic substitutes.*

As the posterior belief of the project's quality $\rho(e_{i,1}, e_{-i,1})$ is also decreasing in the partner's first-stage experimentation, agents also grow increasingly pessimistic about the project's quality the higher is their partner's first-stage experimentation.

From the two preceding lemmas, the main proposition that this model allows for follows. This proposition concerns behaviour in the PBE of the game, comparing the setting that was outlined with observable effort to a setting with unobservable effort.

Proposition 1. *The first-stage experimentation effort is higher if experimentation effort is unobservable.*

As a breakthrough provides a public good, agents benefit from their partner's experimentation. Therefore, they want to encourage their partner to experiment in the future. Hence, agents will take into account how their action changes the posterior belief and thereby the second-stage experimentation of their partner. An agent's partner's second-stage experimentation effort is decreasing in that agent's first-stage experimentation, see Lemma 2. Thus, if agents observe each others' experimentation, every agent has an incentive to decrease their experimentation in the first stage to encourage future experimentation of their partner.

If, however, experimentation effort is not observable, an agent i only forms a belief, $\hat{e}_{-i,1}$, about their partner's first stage experimentation. The posterior that enters the second-stage expected utility corresponding to Eq. 1 is then a function of the beliefs about the partner's experimentation $\hat{e}_{-i,1}$ and not of $e_{-i,1}$. Importantly, conditional on there being no breakthrough, the actual level of experimentation effort $e_{-i,1}$ has no influence on $\hat{e}_{-i,1}$ if it is unobservable. Agents know that their partner will likewise only form a belief, $\hat{e}_{-i,1}$, about their first-stage experimentation. So compared to the equilibrium level of experimentation if effort is observable, agents can deviate to a higher experimentation level if this is unobservable. This increases the chance of a breakthrough in the first period without making the partner more pessimistic.

Proposition 2. *The second-stage experimentation effort is higher if experimentation effort is observable.*

This is a direct implication of Proposition 1. First-stage experimentation is higher if unobservable. Given that beliefs are correct in the PBE, agents in the second stage are more pessimistic about the project's quality if first-stage experimentation is not observable. Agents respond to this by exerting less experimentation effort in the second stage if this was unobservable in the first stage.

All proofs for this section are presented in Appendix A.

3.2 Experimenting with separate projects

Next, I turn to a setting where agents work on separate projects. As in Section 3.1, two agents $i = 1, 2$ can choose to invest experimentation effort $e_{i,t} \in [0, 1]$ in stages $t = 1, 2$ in a project with unknown quality. Doing so again entails a private cost of $c(e_{i,t}) = 2e_{i,t}^2$.

The crucial difference to joint experimentation is that agents work on two separate projects. Each project is independently of high quality with $p = 0.5$. Both

agents receive a payoff of $Y = 13$ if there is a breakthrough in at least one of the two projects. The probability of a breakthrough in agent i 's project, conditional on her project being of high quality, is $\frac{e_{i,t}}{2}$. This probability only depends on this agent's own experimentation effort. In this setup, an agent's experimentation has the same marginal impact on the probability of a breakthrough in their own project as on the joint project in Section 3.1. Agents again either observe or do not observe their partner's experimentation effort.

If agents observe their partner's experimentation, they maximise the following expected utility in the second stage:

$$EU_{i,2} = \left(\rho(e_{i,1}) \frac{e_{i,2}}{2} + \rho(e_{-i,1}) \frac{e_{-i,2}}{2} - \rho(e_{i,1}) \rho(e_{-i,1}) \frac{e_{i,2} e_{-i,2}}{4} \right) Y - c(e_{i,2})$$

Experimentation efforts by one agent are not informative about the quality of the other agent's project, since higher experimentation in one project only makes a breakthrough in that one project more likely. Breakthroughs still represent a public good, because both agents receive a payoff of $Y = 13$ if at least one of them achieves a breakthrough in their project. The posterior belief that the project is of high quality $\rho(e_{i,1})$ is, by Bayes' rule, here given by

$$\rho(e_{i,1}) = \frac{p \left(1 - \frac{e_{i,1}}{2}\right)}{1 - p \left(\frac{e_{i,1}}{2}\right)}$$

with

$$\frac{\partial \rho(e_{i,1})}{\partial e_{i,1}} \leq 0, \quad \forall p \in [0, 1]$$

$\rho(e_{i,1})$ only depends on the agent's own experimentation, as there is no informational externality.³ This implies that an agent's second-stage experimentation effort cannot be affected by a change in beliefs about their own project's quality that results from their partners' first-stage experimentation.

In the first stage, agents again consider how their experimentation will affect second-stage experimentation. They maximise:⁴

$$EU_{i,1} = \left(p \frac{e_{i,1} + e_{-i,1}}{2} - p^2 \frac{e_{i,1} e_{-i,1}}{4} \right) Y - c(e_{i,1}) + \left(1 - p \frac{e_{i,1}}{2} \right) \left(1 - p \frac{e_{-i,1}}{2} \right) EU_{i,2}$$

Compared to joint experimentation, the strategic interaction of the two agents

³Note that in contrast to joint experimentation, full first-stage experimentation ($e_{i,1} = e_{-i,1} = 1$) therefore will also not resolve all uncertainty if there is no breakthrough.

⁴Note that $\left(1 - p \frac{e_{i,1}}{2}\right) \left(1 - p \frac{e_{-i,1}}{2}\right)$ is the probability of no breakthrough in the first stage.

across periods is now determined through a new channel. Agents know they receive a payoff of Y if there is at least one breakthrough. Given that agents work on two separate projects, an agent's incentive to experiment depends on how likely there is a breakthrough in their partner's project, as only one breakthrough is needed. This introduces an element of strategic substitutability between actions within a stage.

Lemma 3. *An agent's second- and her partner's second-stage experimentation efforts as well as an agent's first- and her partner's first-stage experimentation efforts are strategic substitutes.*

Within the second stage, an agent's incentive to experiment decreases in the other agent's experimentation effort, since only one breakthrough is needed to receive Y . There is no benefit in experimenting if the partner achieves a breakthrough, the likelihood of which is increasing in the partner's experimentation effort. The same applies in the first stage.

This strategic substitutability of experimentation within a stage drives the following result concerning experimentation across stages:

Lemma 4. *An agent's second-stage experimentation increases in their partner's first-stage experimentation.*

As second-stage experimentation of partners are strategic substitutes, an increase in the partner's posterior belief about their project's quality $\rho(e_{-i,1})$ decreases an agent's incentive to experiment, and vice versa. The mechanism behind this is that high experimentation by the partner in the first stage will discourage the partner's experimentation in the second stage. This is the case, as the partner's posterior about her project's quality decreases in her own first-stage experimentation, which decreases her second-stage experimentation incentives. As within a stage experimentation levels are strategic substitutes, an agent's second-stage experimentation increasing in their partner's first-stage experimentation.

This mechanism operates through a change in the beliefs about the quality of an agent's project associated with changes in that agent's experimentation. An agent's action does not affect the partner's posterior of their own project's quality. Therefore, there exists no informational externality.

Proposition 3. *The first-stage experimentation effort is higher if experimentation effort is observable.*

Since an agent's second-stage effort is increasing in their partner's first-stage experimentation, see Lemma 4, and a breakthrough constitutes a public good, the

observability of experimentation effort induces higher experimentation levels. The reverse logic from Proposition 1 comes into play here. Now, with unobservable effort the two agents cannot encourage their partner to increase experimentation in the second stage. Therefore, incentives to experiment are higher if this is observable.

Interestingly, Proposition 3 shows that the observability of experimentation effort has the opposite directional effect if partners experiment separately compared to when they experiment jointly (see Proposition 1). A combination of two factors drives this as we move from a setting of joint experimentation to separate experimentation: First, the possibility that experimentation is futile if the partner achieves a breakthrough and second, the lack of an informational externality.

Proposition 4. *The second-stage experimentation effort is higher if experimentation effort is unobservable.*

This is again a direct consequence of Proposition 3. In the PBE, beliefs about first-stage experimentation are correct. Agents are now more pessimistic in the second stage if experimentation is observable, because experimentation is higher in the first stage.

All proofs for this section are presented in Appendix B.

3.3 Predictions for experimentation efforts

The experimental parameters are chosen to provide large theoretical treatment differences in the two treatments with joint experimentation, while making sure optimal effort is sufficiently far from 0% and 100% to avoid boundary effects. The theoretically predicted experimentation effort levels for this set of parameters are presented in the top rows of Table 1 (first stage) and Table 2 (second stage). The efficient experimentation levels are in the bottom rows of Table 1 (first stage) and Table 2 (second stage).

		1st stage	
		Unobservable	Observable
Equilibrium levels	Joint	34%	10%
	Separate	47%	50%
Efficient levels	Joint	100%	100%
	Separate	74%	74%

Notes: The top two rows present first-stage experimentation levels in the PBE for the chosen parameters by treatment. The bottom two rows present the efficient first-stage experimentation levels for the chosen parameters by treatment.

Table 1: Theoretical treatment predictions for the first stage

		2nd stage	
		Unobservable	Observable
Equilibrium levels	Joint	65%	77%
	Separate	61%	61%
Efficient levels	Joint	0%	0%
	Separate	100%	100%

Notes: The top two rows present second-stage experimentation levels in the PBE for the chosen parameters by treatment. The bottom two rows present the efficient second-stage experimentation levels for the chosen parameters by treatment.

Table 2: Theoretical treatment predictions for the second stage

4 Experimental design

The experimental design closely follows the two theoretical models described in Section 3.1 and Section 3.2. The study was pre-registered at the AEA RCT Registry (Brütt, 2020). The experiment employs four treatments. I vary in a between-subject 2-by-2 design the observability of experimentation effort and whether experimentation is joint or separate.

In all treatments, the subjects play the experimentation game repeatedly. For each of these games, two participants are randomly paired to be in a ‘team’. Each team member has to choose how much of their individual budget of €2 to invest in two stages of the experimentation game. They can invest between 0% and 100% of their budget in each stage.⁵

⁵This excludes the possibility of negative payoffs.

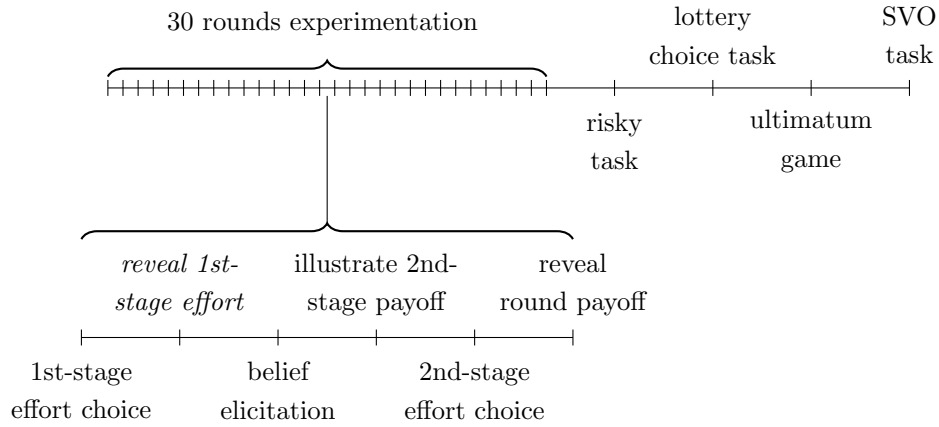


Figure 1: Outline of the experiment

4.1 Treatments

The four treatments differ along two dimensions. First, the experiment varies whether experimentation is joint. In the joint case, the two paired subjects work on one project and can achieve a ‘breakthrough’ depending on the level of joint experimentation. A breakthrough reveals the project’s quality and guarantees a payoff for all team members. If experimentation is separate, subjects work on two distinct projects with independently drawn quality. Their individual experimentation determines the likelihood of a breakthrough in their individual project. The incentives for treatments with joint experimentation are as outlined in Section 3.1, for separate experimentation as discussed in Section 3.2. In all treatments, a breakthrough results in a payoff of €13 for both agents.

Second, the treatments differ in the observability of experimentation effort. In treatments with observable experimentation investments, the participants are informed of their team member’s investment level after the first stage, before making their own second-stage investment choice. In the treatments without observable experimentation, participants only know how much they invested themselves in the first stage before moving to the second stage.

4.2 Experimental timeline

Figure 1 illustrates the timing of the experiment. All subjects face 30 rounds of the experimentation game. Each round of the experimentation game starts with the first investment stage. After the first stage, a set of beliefs is elicited, see Section 4.4. In treatments with observable experimentation, the participants afterwards learn their partner’s first-stage experimentation investment. Furthermore, the

participants receive support for the second stage, see Section 4.3.

Next, the participants make their second-stage investment decision for the case that there was no breakthrough in the first stage, using the strategy method. The strategy method ensures that I collect observations of the second-stage investment even if the project has been terminated due to a breakthrough.⁶ At the end of a round, the subjects receive feedback. This feedback includes their payoff of the round, whether a breakthrough was achieved, their investment and their partner's investment if this was observable. Afterwards, subjects are re-matched to a new partner, within matching groups of six.

The aim of repeating the game is to facilitate subjects' learning. While the game is rather complex at first, repeating it allows subjects to understand how their behaviour in the first stage can influence beliefs and behaviour in the second stage. After 30 rounds of the experimentation game, subjects go through four separate control tasks that serve to identify potential drivers of experimentation. These tasks are administered after the experimentation game to avoid any impact on the experimentation game itself.

The first control task is a decision under uncertainty. This task is closest to the experimentation game. The subjects face a risky choice to invest in a project. The project can be of high or low quality and the subjects have to choose how much of an endowment to invest. The parameters are the same as in the experimentation game. Being a one-shot game, it excludes learning and informational externalities to focus on the uncertainty about the project's quality in experimentation. There are also no payoff externalities. This task is added as studies such as Banks et al. (1997) do not document any explanatory power of risk aversion measured by standard risk aversion elicitation, despite uncertainty being central to experimentation. In contrast, Hudja and Woods (2021) find that risk aversion does entail explanatory power if measured in a setting resembling the actual task more closely. This is in line with the recent findings of Charness et al. (2020), showing that risk preference elicitation remains predictive only in closely related frameworks. I complement this with a standard lottery choice task to obtain a standard measure of the subject's risk attitude (Holt and Laury, 2002). Next, the subjects play an ultimatum game (Güth et al., 1982). The second stage employs the strategy method, where subjects indicate their lowest acceptable offer. Finally, I measure the subjects' social value orientation using the

⁶This in particular guarantees that there are also sufficiently many observations even for high experimentation levels. There would be fewer of these observations otherwise, due to the higher likelihood of a breakthrough with high experimentation investment.

ring test (Liebrand, 1984).

4.3 Decision support

The subjects need to understand the consequences of their own actions and their team members' actions on their payoffs. To facilitate this understanding, the subjects are offered a graphical interface that shows the possible payoff consequences of their actions in the second stage. The subjects can enter multiple values of possible second-stage effort levels by their team member and various beliefs they may have about the probability that the project is of high quality. Given these variables, the tool shows the expected payoffs for each possible effort level by the subject. The graph clarifies the consequences of a certain experimentation level for both the own and the other's expected payoff; it does not encourage subjects to choose any specific level.⁷

To ensure that the decision support does not push the subjects to only consider their own payoff, the graph also shows the payoff consequences of choosing a certain experimentation level for the partner. This avoids limiting subjects to maximising the own expected payoffs. If they wish, they can consider other outcome dimensions, such as overall payoffs or inequalities between payoffs, which are equally salient in the graph. Furthermore, calculators are available in both stages of the experimentation game. These allow the subjects to calculate the costs of investing and the probability of a breakthrough for given investment levels.

Both the graphical interface and the calculators are only shown to participants if they actively choose to reveal them. This way, the subjects can ignore the provided support if they want to. This aims at ensuring that the subjects' true preferences are elicited; payoff consequences are transparent, while the subjects only receive the information they desire. During the instructions, the participants see a video demonstrating how to use the graphical interface and the calculators.

4.4 Belief elicitation

The following beliefs are elicited after the first stage of the experimentation game:

- (i) The posterior belief about the project's quality
- (ii) The belief about the partner's posterior belief about the project's quality

⁷Screenshots of this tool are available in the instructions in Appendix D.

- (iii) The belief about the experimentation investment by the partner in the second stage
- (iv) *Only if effort is unobservable*: The belief about the experimentation investment by the partner in the first stage

I use the binarised scoring rule (BSR) introduced in Hossain and Okui (2013) to incentivise the belief elicitation. The chance of receiving a prize of €2 increases in the accuracy of the prediction. For this, a quadratic loss function is used. The BSR ensures that reporting true beliefs is incentive compatible even if the subjects are risk averse or non-expected utility maximisers. Danz et al. (2022) show that using the BSR may give rise to errors in the belief elicitation if the incentivisation is transparent. Therefore, the subjects are only informed that giving their truthful best guess will maximise the probability of receiving the prize for their prediction. Detailed information on the incentivisation is withheld from the subjects, unless requested. See Appendix E for the detailed instructions given to the participants.

4.5 Procedures

384 students participated in this study from September to November 2020, recruited at the CREED laboratory of the University of Amsterdam. The experiment included 32 sessions, each consisting of 12 subjects in two matching groups per session. The participants did not know the identity of the other participants in their session or matching group. The experiment was advertised as a three-hour experiment on economic decision making, without any further details. The experiment was computerised using PHP. The treatment assignment was randomised evenly at the session level. Upon starting the experiment, the subjects were randomly assigned to matching groups.

The experiment was conducted online due to the COVID restrictions at the time. The participants received a link for the experiment and an invitation to join a zoom session. The zoom session allows the participants to ask the experimenter any questions they may have.⁸ Given that this experiment is online, the participants are more likely to stop the experiment early. If a subject dropped out before the first round of the experiment, I substituted in a back-up player on their behalf.⁹ While the experiment was conducted online, the subject pool

⁸I guaranteed anonymity by re-naming subjects and ensured that no communication was possible between subjects by muting everyone and restricting the chat function to communication only with the experimenter.

⁹This way, the matching groups are not reduced in size and there is no loss in data for the

reflects the standard laboratory population, as the database of enlisted subjects of the CREED laboratory was used for recruitment. It was communicated that practices commonly used at the CREED laboratory, such as a no-deception policy, would also apply online.

The instructions are available in Appendix D. The understanding of these instructions was tested before the start of the experiment. Two rounds of the experimentation game and two other rounds of the belief elicitation were randomly chosen for payment. In addition, all control tasks were paid out. Earnings were on average €32.65. The average duration was 2 hours and 21 minutes.

5 Experimental results

Table 3 and Table 4 report the experimentation effort per treatment for the first and the second stage of the experimentation game, respectively. Average experimentation levels are shown by observability (left vs. right column) and by whether experimentation is joint or separate (top vs. bottom row). The results presented here are robust to only considering observations from the second half of the experiment, so not driven by inexperience. This pre-registered robustness check is provided in the Appendix in Tables 13 and 14.

	1st stage		
	Unobservable	Observable	
Joint	61.46%	70.98%	66.24%
Separate	50.43%	56.69%	53.56%
	55.92%	63.84%	

Notes: Average experimentation effort in the first stage.

Table 3: Experimentation effort per treatment in the first stage

I apply Permutation T-tests (*PmtT-test*) when studying treatment comparisons and comparisons of observed behaviour to the theoretical predictions.¹⁰ Given the lack of independence of observations within a matching group, the observations are averaged at the matching-group level. For regression analyses, I

remaining players. As the experience of these back-up players does not differ from the experience of any other participant, I will include their data in the analysis. In total, there were 17 drop-outs *before* the start of the experimentation game for whom back-ups were substituted in. In contrast, I discard the data of participants who dropped out prematurely after the first round of the experimentation game and the data of those that replaced them for the analysis. There was only one drop-out *after* the start of the experimentation game.

¹⁰These tests are more powerful than traditional non-parametric techniques such as the Wilcoxon signed-rank and Mann-Whitney U tests (Siegel and Castellan, 1981).

	2nd stage		
	Unobservable	Observable	
Joint	33.51%	26.74%	30.11%
Separate	38.53%	38.88%	38.71%
	36.03%	32.81%	

Notes: Average experimentation effort in the second stage.

Table 4: Experimentation effort per treatment in the second stage

cluster the observations at the matching-group level to account for the dependence of observations.

Section 5.1 discusses the observed behaviour in the two treatments with joint experimentation in teams, comparing observable and unobservable experimentation. Section 5.2 then compares this to the setting where individuals experiment separately, Section 5.3 explores behavioral channels to explain the observed experimentation behaviour.


5.1 The observability of experimentation in joint projects

First, I consider the case of joint experimentation. There are stark difference between the theoretically predicted behaviour and actual behaviour in the laboratory. Figure 2 illustrates this for experimentation in the first and second stage of the experimentation game. The left panel sets the first-stage experimentation in both joint treatments against the PBE predictions, the right panel does so for the second stage.

5.1.1 First-stage experimentation in joint projects

It is important to keep in mind that only first-stage experimentation includes some critical elements of experimentation. Here, participants' experimentation effort can generate new information to be used both by themselves and by their partner in the second stage. If effort is observable, participants invest 70.98% in the project in the first stage, while they invest 61.46% if this is unobservable (see Table 3). For both these values, it is evident from Figure 2 that I can reject the theoretically predicted experimentation levels of 10% if observable and 34% if unobservable in favour of higher experimentation (*PmtT-test*; both $p < 0.001$).¹¹

¹¹To myopically maximise first-stage payoffs, not considering the effect on second-stage payoffs, agents should choose an effort level of $e_{i,1} = \frac{p \times Y}{8} = 81.25\%$. This is significantly higher than the observed experimentation levels for both treatments (*PmtT-test*; both $p < 0.001$). Thus, while agents experiment more than theoretically predicted, they are not fully myopic either.



Notes: Comparison of experimentation in the first stage (left) and second stage (right) over 30 rounds of joint experimentation to the PBE predictions. Shaded regions indicate the 95% confidence interval, clustering observations on a matching group level.

Figure 2: Experimentation in treatments with joint experimentation

Result 1. *First-stage experimentation is higher than predicted in joint experimentation.*

At first glance, this is surprising, given that earlier studies discussing individual experimentation, such as Meyer and Shi (1995), tend to observe under-experimentation. This is also in contrast to Boyce et al. (2016), who study strategic experimentation without a payoff externality. However, there are several distinct features of strategic experimentation with payoff externalities that may help explain this observation, which will be discussed throughout this section.

Strikingly, first-stage experimentation is higher if it is observable than if unobservable, contrary to the theoretical predictions. There is an approximately 15% increase with observable experimentation effort (*PmtT-test*; $p = 0.058$).

Result 2. *First-stage experimentation is higher if experimentation effort is observable.*

Result 2 implies that I can clearly reject the theoretical prediction that the observability of experimentation effort decreases experimentation. This effect is

not an artefact of early rounds of experimentation, where subjects are still learning about the exact incentives they face. Instead, in the last half of the experiment experimentation is also 18% higher if observable (*PmtT-test*; $p = 0.046$). Thus, the presence of an informational externality does not decrease experimentation levels.

Several channels can drive higher levels of experimentation than predicted in the first stage if experimentation is observable. To explore why a discouragement effect may not be present, I first focus on the channels that can be identified by studying joint experimentation. First, the lack of a discouragement effect could be explained through biases in belief formation. Second, social preferences, specifically reciprocity, could account for this.

5.1.2 Belief formation in joint projects

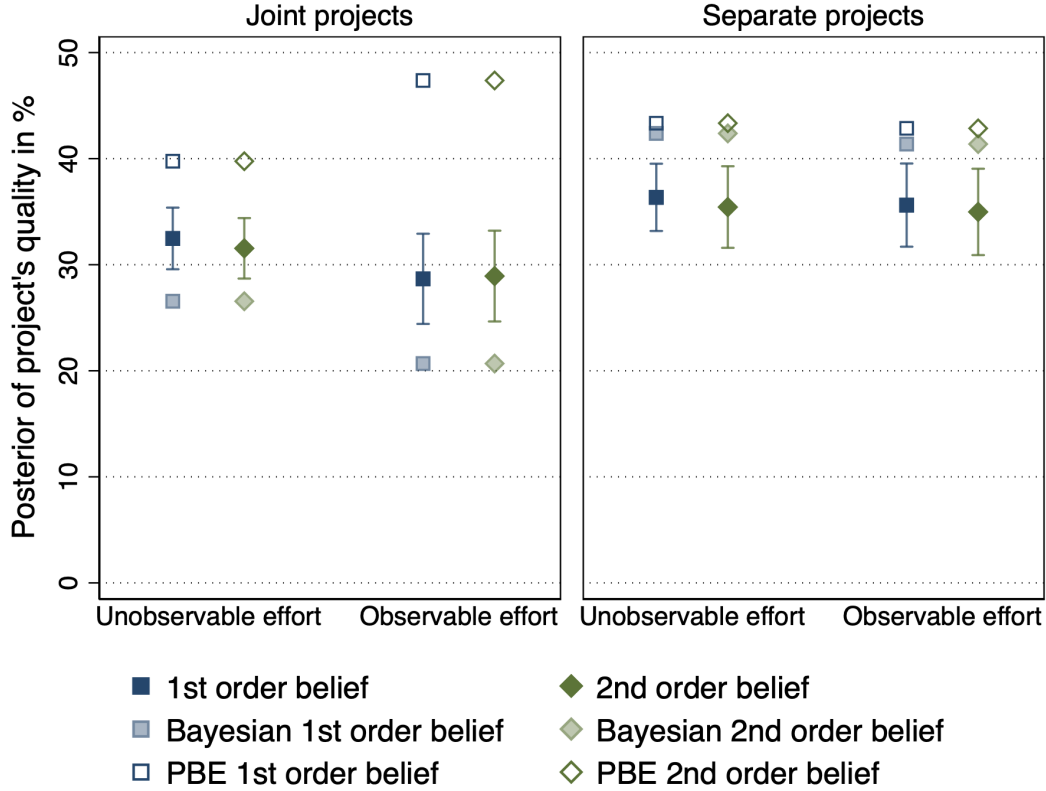
Figure 3 contrasts the participants' average beliefs with both the Bayesian posterior given their first-stage experimentation and the beliefs in the PBE for all four treatments. The left panel presents these differences for joint projects, the right for separate projects. Table 5 provides the overview of this comparison, contrasting the elicited beliefs both to the beliefs in the PBE and Bayesian posteriors.

	Unobservable	Observable
Joint	32.48% (<i>26.56%</i> , <i>39.74%</i>)	28.67% (<i>20.63%</i> , <i>47.37%</i>)
Separate	36.35% (<i>42.38%</i> , <i>43.34%</i>)	35.62% (<i>41.38%</i> , <i>42.86%</i>)

Notes: Average posteriors after first-stage experimentation. In parentheses, the table provides Bayesian posteriors in *italics* and PBE beliefs in *grey italics*.

Table 5: Posterior of the project's quality

There is no significant difference in posteriors between treatments with joint experimentation depending on whether experimentation is observable or not (*PmtT-test*; $p = 0.131$). Remember that since first-stage experimentation is higher if it is observable, agents should become more pessimistic in the observable treatment. However, the experiment may be under-powered to see this reflected in beliefs. In particular, beliefs appear to be updated similarly across treatments. Both in the case of unobservable and of observable effort, beliefs are significantly below the beliefs in the PBE (*PmtT-test*; both $p < 0.001$). This is consistent with higher-than-predicted first-stage experimentation and no evidence for biases in belief updating, as the observed experimentation efforts differ from the PBE predictions.



Notes: Comparison of own (1st order) and beliefs about partner's (2nd order) posteriors in the joint treatments (left) and the separate treatments (right) to PBE beliefs and Bayesian posteriors given the (beliefs about) first-stage experimentation. Bars indicate the 95% confidence interval, clustering observations on a matching group level.

Figure 3: Elicited beliefs vs. theoretical predictions

To establish whether biases in belief updating exist, consider the comparison between the elicited beliefs and the Bayesian posteriors. The Bayesian posteriors are calculated based on the empirical first-stage experimentation, if observable, and the subjects' beliefs about their partner's first-stage experimentation, if unobservable. Here, Figure 3 illustrates that there is a comparable difference between the elicited beliefs and Bayesian posteriors in both treatments. Beliefs are significantly more optimistic than the Bayesian posterior, both if the first-stage experimentation effort is unobservable and if it is observable (*PmtT-test*; both $p < 0.001$). This suggests that agents update their beliefs conservatively in both treatments, as frequently observed in the literature (see Benjamin (2019) for an overview). Since there is no significant difference in this measure of conservatism between treatments (*PmtT-test*; $p = 0.379$), this does not point towards a lack of first-stage experimentation being able to discourage future experimentation due

to an absent effect on beliefs.

Result 3. *Beliefs are updated conservatively, but respond to experimentation in the predicted direction.*

I will now examine more closely how a discouragement effect impacts second-stage experimentation through changes in beliefs. For a discouragement effect to exist in the treatment with observable experimentation, a first necessary condition is that first-stage experimentation affects the subjects' posterior beliefs about the project's quality. More specifically, a participant's posterior belief has to decrease in her partner's first-stage experimentation. Table 6 shows the regression results of individuals' beliefs on their own experimentation and their partner's experimentation, when the partner's experimentation is observable, or the elicited belief about the partner's experimentation if unobservable. Own and partner's first-stage experimentation indeed significantly and negatively correlate with the posterior beliefs if experimentation is joint and observable ($p = 0.060$ and $p = 0.005$, respectively; column (1) in Table 6).¹² Observing a one percentage point increase in first-stage experimentation by a subject's partner is associated with a 0.16 percentage point lower posterior about the project's quality.

This is in contrast to the case where experimentation effort is not observable. In that case, there is no correlation between beliefs about the partner's first-stage experimentation and the posterior ($p = 0.543$; column (2) in Table 6). This suggests that the subjects only react to the elements they actually observe when forming their beliefs, reflecting the inherent uncertainty about their partner's first-stage experimentation if this is unobservable.

A second crucial element of the discouragement effect is that subjects not only update their beliefs in the prescribed manner, but also expect their partners to do so. Only in this case individuals face an incentive to decrease first-stage experimentation to avoid discouraging future experimentation of their partner. Table 7 reports the correlations of the participants' beliefs about their partner's posterior with own experimentation and the partner's experimentation (for the treatment with observable experimentation) or the beliefs about the partner's experimentation (for the treatment with unobservable experimentation). For the case of observable experimentation, column (1) reveals a pattern consistent with individuals

¹²This is explained largely by between-subject variation, not within-subject variation. In Appendix C, I show that subject-level fixed effects absorb the effect of own first-stage experimentation on the posterior beliefs, suggesting that variation in first-stage experimentation and associated changes in beliefs between subject drive the effect of experimentation on beliefs.

Treatments	<i>Dependent variable:</i>			
	Posterior of project's quality			
	Joint projects		Separate projects	
	Obs	Unobs	Obs	Unobs
	(1)	(2)	(3)	(4)
Own effort	-0.08 (0.04)	-0.09 (0.03)	0.06 (0.05)	-0.03 (0.04)
Partner's effort	-0.16 (0.05)	0.04 (0.06)	0.07 (0.04)	0.01 (0.03)
Constant	46.18 (4.74)	35.51 (4.09)	28.14 (5.45)	37.41 (2.75)
Observations	2880	2850	2880	2880
Clusters	16	16	16	16
R-squared	0.055	0.012	0.026	0.002

Notes: OLS estimating effect of own and partner's first-stage experimentation effort on posterior of project's quality for all treatments. (1) and (3) use the partner's actual experimentation effort, (2) and (4) use the subject's belief about the partner's experimentation effort. Robust standard errors (in parentheses) are clustered at the matching group level.

Table 6: The effect of first-stage experimentation on posterior beliefs

correctly anticipating that their own as well as their partners' first-stage experimentation will result in their partner having more pessimistic beliefs. Elicited beliefs are in line with individuals expecting their partner to become 0.13 percentage points more pessimistic if their first-stage experimentation increases by one percentage point. This is a statistically significant correlation ($p = 0.002$). There is no such correlation if experimentation is observable ($p = 0.680$; column (2) in Table 7). This indicates that the participants anticipate the potential of discouraging their partner if they choose high experimentation levels if experimentation is observable.

5.1.3 Second-stage experimentation in joint projects

Next, consider second-stage experimentation, which should theoretically respond to the experimentation of the first stage. The informational spillovers should affect behaviour if experimentation is observable. Figure 2 illustrates that experimentation is lower in the second stage if this is observable, albeit insignificantly so ($PmtT$ -test; $p = 0.116$). Compared to the theoretical predictions, Figure 2 shows that effort is 48% lower than theoretically predicted if unobservable and 65% lower if observable ($PmtT$ -test; both $p < 0.001$). Result 4 summarises this.

Treatments	<i>Dependent variable:</i>			
	Belief about partner's posterior of project's quality			
	Joint projects		Separate projects	
	Obs	Unobs	Obs	Unobs
	(1)	(2)	(3)	(4)
Own effort	-0.13 (0.04)	-0.01 (0.04)	0.05 (0.04)	-0.04 (0.05)
Partner's effort	-0.13 (0.05)	0.01 (0.07)	0.13 (0.05)	0.01 (0.05)
Constant	47.90 (4.50)	32.10 (3.84)	24.84 (5.13)	36.66 (2.80)
Observations	2880	2850	2880	2880
Clusters	16	16	16	16
R-squared	0.057	0.000	0.050	0.003

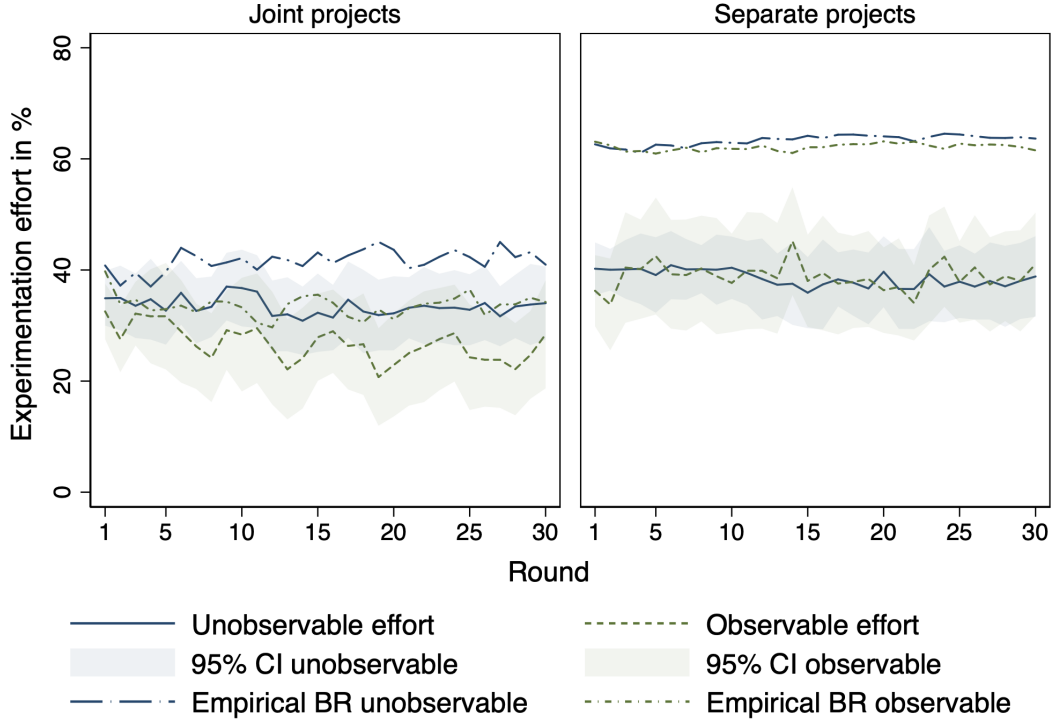
Notes: OLS estimating effect of own and partner's first-stage experimentation effort on beliefs about partner's posterior of project's quality for all treatments. (1) and (3) use the partner's actual experimentation effort, (2) and (4) use the subject's belief about the partner's experimentation effort. Robust standard errors (in parentheses) are clustered at the matching group level.

Table 7: The effect of first-stage experimentation on beliefs about partner's posterior

Result 4. *Second-stage experimentation is significantly lower than theoretically predicted.*

The lower second-stage experimentation in the treatment with observable experimentation is consistent with higher first-stage experimentation. Given first-stage experimentation and the resulting Bayesian posterior, I calculate the empirical best response for each individual in each round of experimentation. Figure 4 plots the empirical best response for each treatment. The left panel compares the empirical best response to actual second-stage experimentation for the treatments with joint experimentation. Second-stage experimentation is significantly lower than the empirical best response (*PmtT-test*; $p = 0.014$). The degree of deviation from the best response is indistinguishable between the two treatments (*PmtT-test*; $p = 0.793$).

The low second-stage experimentation is a first indicator that agents respond to information previously generated, and potentially to informational spillovers. For the discouragement effect to induce lower experimentation through a change in beliefs, second-stage experimentation must be responsive to a change in posteriors when experimentation is observable. If this were not the case, agents would have no reason to fear discouragement when deciding on first-stage experimentation,



Notes: Comparison of experimentation in the second stage in treatments with joint experimentation (left) and separate experimentation (right) over 30 rounds of experimentation to the empirical best response. Shaded regions indicate the 95% confidence interval, clustering observations on a matching group level.

Figure 4: Second stage experimentation vs. empirical best responses

knowing that the potential pessimism of their partner does not manifest itself in different actions. Table 8 gives the results of regressing second-stage experimentation on the posteriors of the project's quality. Second-stage experimentation responds significantly to the posterior in the predicted direction in all treatments. In particular, individuals invest 0.66 percentage points less experimentation effort if they are one percentage point more pessimistic when effort is observable and agents experiment jointly ($p < 0.001$), see column (1) in Table 8.

This effect is not entirely driven by variation between subjects. Instead, there is within-subject variation in beliefs across rounds that affects second-stage experimentation effort. To see this, consider the case where subject fixed effects are included in the regression, see column (5) in Table 8. Including subject fixed effects controls for different levels of experimentation and beliefs across subjects, which implies that the remaining effect (0.42 percentage points) on second-stage experimentation stems from variation in subjects' beliefs across rounds.

The final element required for a change in beliefs to yield a discouragement

Treatments	<i>Dependent variable:</i>							
	Second-stage experimentation effort							
	Joint projects		Separate projects		Joint projects		Separate projects	
	Obs	Unobs	Obs	Unobs	Obs	Unobs	Obs	Unobs
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Posterior	0.66 (0.04)	0.54 (0.11)	0.64 (0.13)	0.39 (0.10)	0.42 (0.05)	0.22 (0.03)	0.32 (0.09)	0.18 (0.06)
Constant	7.80 (2.56)	15.85 (3.23)	16.22 (5.79)	24.50 (2.88)	14.58 (1.56)	26.48 (1.11)	27.32 (3.10)	31.98 (2.29)
Fixed effects	✗	✗	✗	✗	✓	✓	✓	✓
Observations	2880	2850	2880	2880	2880	2850	2880	2880
Clusters	16	16	16	16	16	16	16	16
R-squared	0.232	0.138	0.126	0.056	0.510	0.710	0.618	0.660

Notes: OLS estimating effect of posterior about project's quality on second-stage experimentation effort for all treatments. (1)-(4) do not include subject fixed effects, (5)-(8) include subject fixed effects. Robust standard errors (in parentheses) are clustered at the matching group level.

Table 8: The effect of posterior beliefs on second-stage experimentation

effect is that the agents also anticipate that their partners respond to changes in their posterior. Individuals might not decrease experimentation levels as a response to the observability of experimentation because they fail to realise that their partner's induced pessimism will make her or him experiment less in the second stage.

As a measure, take the correlation between a subject's beliefs about their partner's posterior and their beliefs about the partner's second-stage experimentation. Table 9 reports such regression results for all treatments. The belief measures support that subjects expect their partners to respond to their posteriors, with an estimated reduction of 0.7 percentage points in the beliefs about the partner's second-stage experimentation resulting from a one percentage point change in the beliefs about the partner's posterior ($p < 0.001$; column (1) in Table 9). The strength of the correlation is thus similar to that of the subject's own beliefs and second-stage experimentation. The fixed-effect regression in column (5) again reveals that there is a substantial within-subject response.

Collectively, the belief evidence underlines that a discouragement effect to decrease experimentation incentives is present. This gives Result 5.

Result 5. *If experimentation is observable, high first-stage experimentation discourages high second-stage experimentation through a change in beliefs.*

We can therefore conclude that Result 2 (that establishes that observability in-

Treatments	<i>Dependent variable:</i>							
	Beliefs about partner's second-stage experimentation effort							
	Joint projects		Separate projects		Joint projects		Separate projects	
	Obs	Unobs	Obs	Unobs	Obs	Unobs	Obs	Unobs
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Posterior	0.70 (0.04)	0.71 (0.07)	0.70 (0.10)	0.37 (0.08)	0.58 (0.05)	0.44 (0.06)	0.46 (0.10)	0.26 (0.06)
Constant	7.21 (2.47)	10.22 (3.00)	14.60 (4.43)	22.17 (2.60)	10.51 (1.47)	18.67 (1.75)	22.97 (3.47)	26.16 (2.26)
Fixed effects	X	X	X	X	✓	✓	✓	✓
Observations	2880	2850	2880	2880	2880	2850	2880	2880
Clusters	16	16	16	16	16	16	16	16
R-squared	0.309	0.268	0.191	0.076	0.527	0.685	0.504	0.618

Notes: OLS estimating effect of belief about partner's posterior about project's quality on second-stage experimentation effort by partner for all treatments. (1)-(4) do not include subject fixed effects, (5)-(8) include subject fixed effects. Robust standard errors (in parentheses) are clustered at the matching group level.

Table 9: The effect of beliefs about partner's posterior on beliefs about partner's second-stage experimentation

creases first-round experimentation) is not a consequence of a lack of sophistication in belief updating or due to neglecting the impact of own first-stage experimentation on the partner's future behaviour. For an alternative explanation, I now turn to whether conditional cooperation is better suited to account for the observed experimentation patterns.

5.1.4 Reciprocal behaviour

Social preferences, specifically reciprocal behaviour, could provide an explanation for higher experimentation levels if these are observable. Given the positive payoff externality, subjects may reward observing high first-stage experimentation with high second-stage experimentation. To test this, I employ a two-step procedure. This separates the direct effect of first-stage on second-stage experimentation from the indirect effect through a change in beliefs.

In the first step, I regress second-stage experimentation $E_{i,t}^2$ in round t of individual i on posterior beliefs $\rho_{i,t}$.

$$E_{i,t}^2 = \beta \rho_{i,t} + \epsilon_{i,t}$$

In the second step, the residuals of the first regression $\hat{\epsilon}_{i,t}$ are regressed on i 's

partner's ($-i$) first-stage experimentation $E_{-i,t}^1$ in round t , $Observable_i$, indicating whether i is in the observable treatment, and the interaction of these two variables. When experimentation is unobservable, $E_{-i,t}^1$ is given by i 's belief of her partner's first-stage experimentation.

$$\hat{e}_{i,t} = \gamma_1 E_{-i,t} + \gamma_2 Observable_i + \gamma_3 E_{-i,t}^1 \times Observable_i + u_{i,t}$$

This only captures the direct effect of an individual's first-stage experimentation on their partner's second-stage experimentation, which is unrelated to how beliefs are affected. Table 10, column (1) provides the results of this second-stage estimation.

	<i>Dependent variable:</i>	
	Residuals of 1st-stage regression	
	(1)	(2)
Observable	13.94 (6.78)	8.46 (6.33)
Partner's effort	0.21 (0.09)	
Observable \times Partner's effort	-0.24 (0.10)	
Own effort		0.07 (0.07)
Observable \times Own effort		-0.13 (0.08)
Constant	-11.97 (5.57)	-4.07 (5.53)
Observations	5730	5730
Clusters	32	32
R-squared	0.020	0.007

Notes: OLS estimating difference-in-difference in the effect of first-stage experimentation between treatments on residuals of regression of second-stage experimentation on posteriors. In (1), partner's first-stage and own second-stage experimentation is used. In (2), own first stage and beliefs about partner's second stage experimentation are used. Elicited beliefs are used for observations from unobservable treatment. Robust standard errors (in parentheses) are clustered at the matching group level.

Table 10: The effect of (beliefs about) first-stage on second-stage experimentation

Subjects' second-stage experimentation responds significantly to (beliefs of) their partner's first-stage experimentation ($p = 0.048$). However, this effect is entirely driven by the treatment where experimentation is unobservable. The negative interaction effect of the same approximate size ($p = 0.022$) implies that in the treatment with observable experimentation, there is no correlation between

the partner’s first-stage experimentation and their second-stage experimentation beyond the effect driven by a change in posteriors about the project’s quality.

The correlation of first- and second-stage experimentation if experimentation is unobservable, controlling for belief effects, does not suggest a reciprocal motive. More likely, subjects with high experimentation levels expect others to experiment more as well. This is consistent with the fact that there is also a positive correlation between their own first-stage experimentation and their second-stage experimentation ($p = 0.017$), also controlling for belief effects using the two-step estimation.¹³ Thus, since there is no effect in the observable treatment, there is no evidence of subjects punishing or rewarding their partner’s experimentation by increasing their own experimentation.

While there is no reciprocal behaviour, it is conceivable that subjects still expect their partners to reciprocate high experimentation and thus face an incentive to increase first-stage experimentation. If this is the case, beliefs about the partner’s second-stage experimentation should increase in own first-stage experimentation. The same two-step procedure is employed as there again exists a belief channel through which first-stage experimentation can affect beliefs. The results of the second-stage regression are presented in Table 10, column (2). No significant correlation between own first-stage experimentation and the beliefs about the partner’s second-stage experimentation exists ($p = 0.364$).

In line with the preceding analysis, I show in Appendix C Table 18 that there is no differential correlation between the elicited measure of negative reciprocity from the ultimatum game and second-stage experimentation depending on whether experimentation is observable or not. There is also no significant difference in the correlation between proposer behaviour in the ultimatum game and first-stage experimentation by treatment. Reciprocity, therefore, does not seem to be driving the fact that first-stage experimentation is higher if effort is observable.

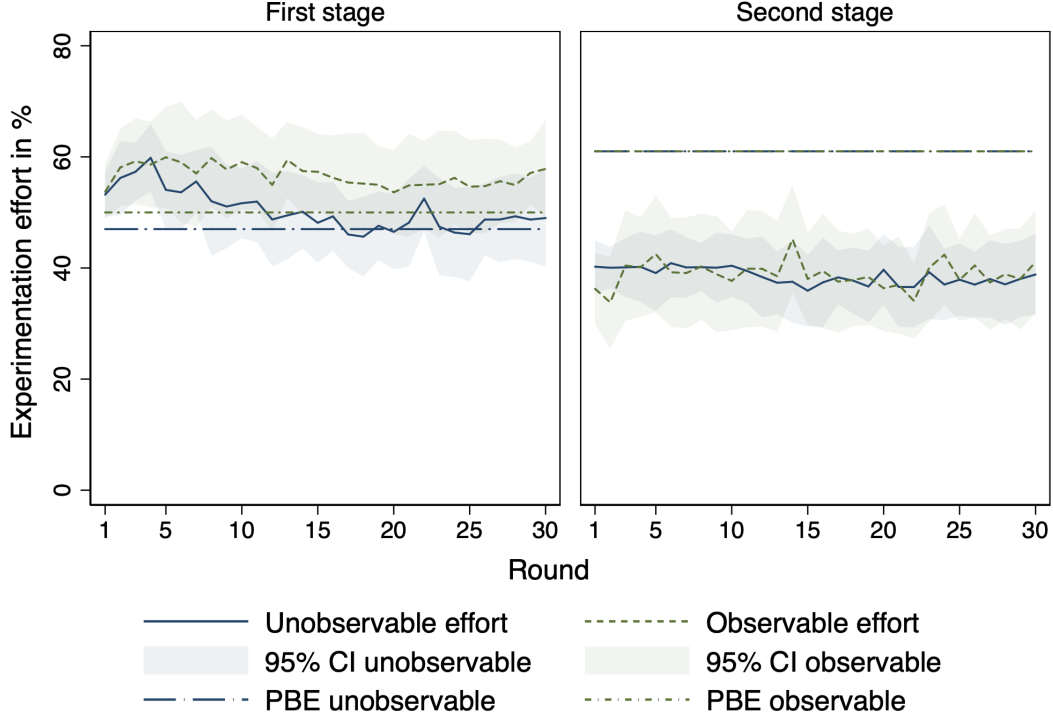
Result 6. *There is no evidence that (expected) reciprocity drives first-stage experimentation if this is observable.*

5.2 Separate compared to joint experimentation

A comparison between joint and separate experimentation sheds further light on how the determinants of experimentation, contrasting settings with and without informational externalities. Figure 5 displays first- and second-stage experimentation compared to the PBE experimentation levels in the two treatments where

¹³See Appendix C, Table 17 for the regression results.

individuals work on separate projects. In Table 11, first-stage experimentation levels are regressed on treatment indicators for observable and joint experimentation and their interaction, Table 12 reports this for second-stage experimentation.



Notes: Comparison of experimentation in the first stage (left) and second stage (right) over 30 rounds of separate experimentation to the PBE predictions. Shaded regions indicate the 95% confidence interval, clustering observations on a matching group level.

Figure 5: Experimentation in treatments with separate experimentation

First, both for observable and for unobservable experimentation, experimentation is significantly lower in the first stage if agents work on separate projects compared to joint experimentation, see column (2) in Table 11. For unobservable experimentation, experimentation is 22% higher with joint than with separate experimentation, for observable experimentation 25% higher ($PmtT$ -test; $p = 0.028$ and $p = 0.002$, respectively). Joint experimentation clearly has a positive effect on experimentation levels. This gives Result 7.

Result 7. *First-stage experimentation is higher if agents experiment jointly.*

This is in contrast to the theoretical predictions, as the lack of an informational externality implies that significantly higher experimentation levels are expected with separate experimentation if experimentation is observable. Instead,

this finding is in line with comparative statics predictions that follow from agents aiming for efficient experimentation levels, maximizing their joint payoffs. Here, separate experimentation leads to lower experimentation in the first stage, as high first-stage experimentation increases the probability of two breakthroughs, which is inefficient. With joint experimentation, two breakthroughs are not possible. Investing fully in the first stage is efficient, thereby resolving all uncertainty.

To test whether the observability of experimentation effort has a distinct effect depending on whether experimentation is joint, consider the interactions in Table 11 and Table 12. Both for the first and for the second stage, there is no statistically significant differential effect of experimentation observability on experimentation effort depending on whether experimentation is joint or separate ($p = 0.606$ in column (3), Table 11 and $p = 0.236$ in column (3), Table 12, respectively). This is against the theoretical predictions; the observability of experimentation effort is predicted to increase experimentation if separate, but decrease it if experimentation is joint. Instead, the observability of experimentation overall increases experimentation levels in the first stage ($p = 0.029$), see column (1) in Table 11. Hence, the presence of an informational externality does not have a differential impact on experimentation efforts if this information is observable or not. Thus, an environment of observable experimentation encourages experimentation, irrespective of whether the group members work on a joint or separate projects.

Result 8. *Observable experimentation increases first-stage experimentation, independent of whether experimentation is joint or separate.*

Recall that for separate experimentation, marginally higher first-stage experimentation is expected if this is observable, because this can encourage future experimentation. Considering second-stage experimentation if this is separate, Figure 4 shows that experimentation is clearly lower than the empirical best response, both for observable and unobservable experimentation ($PmtT$ -test; both $p < 0.001$). This is inconsistent with an encouraging force of higher first-stage experimentation. In line with this, the deviation from the best response is significantly larger when experimentation is joint ($PmtT$ -test; $p < 0.001$). Experimenting jointly has a positive effect on experimentation levels in both stages of the game, but this effect is not larger when it is predicted to be.

Compared to the treatments with joint experimentation, there is a stark contrast in how beliefs are updated in projects with separate experimentation. With separate experimentation, individuals' posterior beliefs do not respond to their own first-stage experimentation ($p = 0.309$ if experimentation is observable and

<i>Dependent variable:</i>						
First-stage experimentation						
	(1)	(2)	(3)	(4)	(5)	(6)
Observable	7.92 (3.53)		6.26 (4.14)	7.97 (3.44)		6.95 (3.85)
Joint		12.68 (3.30)	11.02 (4.68)		12.26 (3.28)	11.23 (4.57)
Joint \times Observable			3.26 (6.30)			2.06 (6.04)
Constant	55.92 (2.53)	53.56 (2.14)	50.43 (2.60)	33.31 (8.87)	31.12 (8.13)	26.92 (8.34)
Controls	X	X	X	✓	✓	✓
Observations	11490	11490	11490	11010	11010	11010
Clusters	64	64	64	64	64	64
R-squared	0.018	0.045	0.064	0.063	0.087	0.105

Notes: OLS estimating effect of joint experimentation, observability of experimentation and the interaction on first-stage experimentation. (1)-(3) do not include controls variables for individual characteristics, (4)-(6) do. Robust standard errors (in parentheses) are clustered at the matching group level.

Table 11: First-stage experimentation in all treatments

$p = 0.465$ if experimentation is unobservable, column (3) and (4), Table 6). As shown in Figure 3, these beliefs are more pessimistic than the Bayesian benchmark, independent of whether experimentation is observable (*PmtT-test*; $p = 0.002$ if observable and $p < 0.001$ if unobservable). The participant's beliefs are consistent with higher experimentation levels than the ones observed. Given these pessimistic beliefs, the low second-stage experimentation levels in both treatments with separate experimentation are not surprising, significantly below the best response that is based on the Bayesian posteriors (*PmtT-test*; both $p < 0.001$), see Figure 4.

5.3 Norms of high experimentation and leading by example

An intuitive explanation for higher experimentation levels with joint, observable experimentation is that both these aspects foster a stronger sense of group membership and allow teams to establish norms of high experimentation. This would be comparable to the observability of individual contributions to a public good increasing such contributions (see e.g. Andreoni and Petrie, 2004), even without punishment. While this experiment does not include direct elicitations of norms, two pieces of evidence support this argument. First, the variance in first-

<i>Dependent variable:</i>						
Second-stage experimentation						
	(1)	(2)	(3)	(4)	(5)	(6)
Observable	-3.22 (3.20)		0.35 (4.41)	-2.13 (3.17)		2.08 (4.34)
Joint		-8.60 (3.04)	-5.03 (3.34)		-8.62 (3.04)	-4.36 (3.26)
Joint \times Observable			-7.12 (5.95)			-8.41 (5.83)
Constant	36.03 (1.73)	38.71 (2.20)	38.53 (2.37)	13.58 (7.39)	17.15 (7.39)	16.37 (7.28)
Controls	X	X	X	✓	✓	✓
Observations	11490	11490	11490	11010	11010	11010
Clusters	64	64	64	64	64	64
R-squared	0.003	0.020	0.026	0.021	0.039	0.045

Notes: OLS estimating effect of joint experimentation, observability of experimentation and the interaction on second-stage experimentation. (1)-(3) do not include controls variables for individual characteristics, (4)-(6) do. Robust standard errors (in parentheses) are clustered at the matching group level.

Table 12: Second round experimentation in all treatments

stage experimentation is lower if experimentation levels are observable (*PmtT-test*; $p = 0.030$). Measured as the variance of decisions within a matching group, the lower variance in matching groups that are exposed to observable experimentation suggests that these groups coordinate on effort levels.¹⁴

Second, agents adapt their experimentation efforts to previously observed experimentation. While Section 5.1 demonstrates that agents do not reciprocate across the two stages of the experimentation game, individuals exert higher experimentation efforts if they have observed high experimentation in earlier rounds ($p = 0.059$; column (2) in Table 19). Furthermore, if experimentation is not observable, the participants' beliefs about their team member's experimentation are (correctly) below the experimentation levels that participants in the treatments with observable experimentation experience (*PmtT-test*; $p = 0.002$)¹⁵. This hampers the successful coordination on high experimentation levels.

The positive response to the partner's experimentation incentivises agents to

¹⁴Interestingly, there is not only a lower variance in groups with observable experimentation in late periods of the game, when groups have frequently observed each others' experimentation, but also in the first half of the experiment (*PmtT-test*; $p = 0.009$ in the first 15 periods). Thus, it does not seem necessary for coordination that agents see experimentation efforts frequently.

¹⁵Beliefs about the partner's first-stage experimentation do not significantly differ from the actual experimentation levels (*PmtT-test*; $p = 0.346$).

‘lead by example’. As in the leading-by-example literature, high experimentation efforts, if observable, can induce high experimentation levels in future periods through two channels. First, leading by example has a signalling value (Potters et al., 2007). High experimentation can signal the (private) belief that investment in the project is lucrative. While there is no asymmetric information, communicating private beliefs about whether experimentation is fruitful can be informative in a complex setting when the participants are unsure of their optimal actions. In line with this, in the setting of separate experimentation, individuals update more positively about the project’s quality if they observe high experimentation efforts by their partner ($p = 0.055$; column (3) in Table 6). As the projects are independent, the partner’s experimentation does not reveal any information about the project’s objective quality. However, it could signal that the partner believes the project is a worthy investment, a helpful signal in a complex environment. This only works with observable experimentation and may, therefore, explain why experimentation levels are higher in this case. Second, leading by example can result in reciprocal behaviour in later rounds (Meidinger and Villeval, 2002). As discussed, there is not observable reciprocal behaviour within one round of experimentation, but there is evidence of agents responding to earlier experimentation by their partners in later rounds, in particular, if this is observable.

6 Conclusion

This paper studies the two-dimensional free-riding problems inherent in strategic experimentation of teams, examining the type of environments that foster successful experimentation. I consider two dimensions of the experimentation environment: the observability of experimentation, and whether agents work on one joint project or on two separate projects. The observability of experimentation efforts is predicted to decrease experimentation levels when agents experiment with a joint project; this is driven by the presence of an informational externality. Agents are predicted to discourage each other from experimenting if they observe each others’ experimentation levels but do not observe a breakthrough. With two separate projects, however, the predictions flip, as there is no informational externality, and the potential of a breakthrough in another project implies that agents want to avoid futile experimentation that results in two breakthroughs.

This study employs an experiment to test these theoretical predictions and identify behavioural drivers of experimentation in teams. Strikingly, teams are capable of largely overcoming the free-riding problem that lies at the core of strategic

experimentation. In contrast to the prevalent finding in laboratory experiments that experimentation is undervalued, I find that teams experiment more than predicted. Even though their individuals form beliefs as if they grasp the discouraging effect of their experimentation with joint projects, experimentation is higher if it is observable and agents experiment more with a joint project. This is not a result of agents punishing or rewarding certain experimentation behaviour. Instead, agents can coordinate on higher effort levels if experimentation is observable. The findings are in line with agents choosing to lead by example if their team member can observe their experimentation. Moreover, the higher experimentation with joint projects suggests that agents aim for not purely individually-optimal experimentation, but instead consider efficient experimentation levels. With joint experimentation, a full resolution of uncertainty is possible and efficient in the first stage, while the possibility of having two breakthroughs with separate projects implies lower efficient experimentation levels.

To conclude, this paper both speaks to the theoretical and behavioral literature on experimentation and public good provision. First, this paper's findings speak to the theoretical literature on strategic experimentation. One way to interpret the findings in this paper is that the Markov refinement commonly used in the theoretical literature is not convincing in practice. For instance, 'leading by example' is an inherently non-Markovian explanation, as leading by example relies on individuals conditioning their actions on the history of previous experimentation decisions, not only on the information contained in the posterior about the project's quality. However, the experimental findings suggest that exactly this history of experimentation decisions may be a relevant driver of why the theoretically established rankings do not survive in the laboratory.

Second, this paper shows that observe that there are mechanisms in place that help teams overcome the theoretical hurdles to experimentation. Teams are able to innovate even in settings where it is in every team members' material interest to decrease their experimentation, as this will discourage others from experimenting in the future. Having teams active in innovative processes (as opposed to individuals) will likely not create excessive free-riding and a lack of information discovery, but instead might induce team members to work harder for their fellow team members, giving rise to more innovation. Instead of discouraging team members, informational externalities may even signal high hopes for the project's success, encouraging high experimentation.

References

- ANDERSON, C. M. (2012): “Ambiguity aversion in multi-armed bandit problems,” *Theory and Decision*, 72, 15–33.
- ANDREONI, J., AND R. PETRIE (2004): “Public goods experiments without confidentiality: A glimpse into fund-raising,” *Journal of Public Economics*, 88, 1605–1623.
- BABCOCK, L., M. P. RECALDE, L. VESTERLUND, AND L. WEINGART (2017): “Gender differences in accepting and receiving requests for tasks with low promotability,” *American Economic Review*, 107, 714–47.
- BANKS, J., M. OLSON, AND D. PORTER (1997): “An experimental analysis of the bandit problem,” *Economic Theory*, 10, 55–77.
- BENJAMIN, D. J. (2019): “Errors in probabilistic reasoning and judgment biases,” *Handbook of Behavioral Economics: Applications and Foundations* 1, 2, 69–186.
- BOLTON, P., AND C. HARRIS (1999): “Strategic experimentation,” *Econometrica*, 67, 349–374.
- BONATTI, A., AND J. HÖRNER (2011): “Collaborating,” *American Economic Review*, 101, 632–663.
- BOUCKAERT, S., A. F. PALES, C. MCGLADE, U. REMME, B. WANNER, L. VARRO, D. D’AMBROSIO, AND T. SPENCER (2021): “Net Zero by 2050: A Roadmap for the Global Energy Sector,” Flagship report, International Energy Agency.
- BOYCE, J. R., D. M. BRUNER, AND M. MCKEE (2016): “Strategic experimentation in the lab,” *Managerial and Decision Economics*, 37, 375–391.
- BRÜTT, K. (2020): “Collaborating in strategic experimentation,” Pre-registration 5503, AEA RCT Registry.
- CHARNESS, G., T. GARCIA, T. OFFERMAN, AND M. C. VILLEVAL (2020): “Do measures of risk attitude in the laboratory predict behavior under risk in and outside of the laboratory?” *Journal of Risk and Uncertainty*, 60, 99–123.
- CHARNESS, G., L. RIGOTTI, AND A. RUSTICHINI (2007): “Individual behavior and group membership,” *American Economic Review*, 97, 1340–1352.

- CROSON, R., E. FATAS, AND T. NEUGEBAUER (2005): “Reciprocity, matching and conditional cooperation in two public goods games,” *Economics Letters*, 87, 95–101.
- DANZ, D., L. VESTERLUND, AND A. J. WILSON (2022): “Belief Elicitation and Behavioral Incentive Compatibility,” *American Economic Review*, 9.
- DONG, Y., H. MA, Z. SHEN, AND K. WANG (2017): “A century of science: Globalization of scientific collaborations, citations, and innovations,” in *Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery and data mining*, 1437–1446.
- VON ESSEN, E., M. HUYSENTRUYT, AND T. MIETTINEN (2020): “Exploration in teams and the encouragement effect: Theory and experimental evidence,” *Management Science*, 66, 5861–5885.
- FISCHBACHER, U., S. GÄCHTER, AND E. FEHR (2001): “Are people conditionally cooperative? Evidence from a public goods experiment,” *Economics Letters*, 71, 397–404.
- GÜTH, W., M. V. LEVATI, M. SUTTER, AND E. VAN DER HEIJDEN (2007): “Leading by example with and without exclusion power in voluntary contribution experiments,” *Journal of Public Economics*, 91, 1023–1042.
- GÜTH, W., R. SCHMITTBERGER, AND B. SCHWARZE (1982): “An experimental analysis of ultimatum bargaining,” *Journal of Economic Behavior & Organization*, 3, 367–388.
- HOELZEMANN, J., AND N. KLEIN (2021): “Bandits in the Lab,” *Quantitative Economics*, 12, 1021–1051.
- HOLMSTROM, B. (1982): “Moral hazard in teams,” *The Bell Journal of Economics*, 13, 324–340.
- HOLT, C. A., AND S. K. LAURY (2002): “Risk aversion and incentive effects,” *American Economic Review*, 92, 1644–1655.
- HÖRNER, J., AND A. SKRZYPACZ (2017): “Learning, experimentation and information design,” *Advances in Economics and Econometrics*, 1, 63–98.
- HOSSAIN, T., AND R. OKUI (2013): “The binarized scoring rule,” *Review of Economic Studies*, 80, 984–1001.

- HUDJA, S., AND D. WOODS (2021): “Exploration Versus Exploitation: A Laboratory Test of the Single-Agent Exponential Bandit Model,” Available at SSRN 484498.
- KELLER, G., AND S. RADY (2010): “Strategic experimentation with Poisson bandits,” *Theoretical Economics*, 5, 275–311.
- (2015): “Breakdowns,” *Theoretical Economics*, 10, 175–202.
- KELLER, G., S. RADY, AND M. CRIPPS (2005): “Strategic experimentation with exponential bandits,” *Econometrica*, 73, 39–68.
- KOCHER, M. G., T. CHERRY, S. KROLL, R. J. NETZER, AND M. SUTTER (2008): “Conditional cooperation on three continents,” *Economics Letters*, 101, 175–178.
- KRETCHMER, H. (2020): “From dining pods to see-through masks: 6 ways innovations are helping in the pandemic,” Online article, World Economic Forum.
- KWON, O. (2020): “Strategic Experimentation with Uniform Bandit: An Experimental Study,” Working paper.
- LEVATI, M. V., M. SUTTER, AND E. VAN DER HEIJDEN (2007): “Leading by example in a public goods experiment with heterogeneity and incomplete information,” *Journal of Conflict Resolution*, 51, 793–818.
- LIEBRAND, W. B. (1984): “The effect of social motives, communication and group size on behaviour in an N-person multi-stage mixed-motive game,” *European Journal of Social Psychology*, 14, 239–264.
- MEIDINGER, C., AND M. C. VILLEVAL (2002): “Leadership in teams: signaling or reciprocating?” Working paper 02-13, GATE.
- MEYER, R. J., AND Y. SHI (1995): “Sequential choice under ambiguity: Intuitive solutions to the armed-bandit problem,” *Management Science*, 41, 817–834.
- POTTERS, J., M. SEFTON, AND L. VESTERLUND (2005): “After you—endogenous sequencing in voluntary contribution games,” *Journal of Public Economics*, 89, 1399–1419.
- (2007): “Leading-by-example and signaling in voluntary contribution games: An experimental study,” *Economic Theory*, 33, 169–182.

- SIEGEL, S., AND N. J. CASTELLAN (1981): *Nonparametric Statistics for the Behavioral Sciences*, New York: McGraw-Hill.
- SUTTER, M. (2009): “Individual behavior and group membership: Comment,” *American Economic Review*, 99, 2247–2257.
- THOMKE, S. H. (2003): *Experimentation matters: Unlocking the potential of new technologies for innovation*: Harvard Business Press.
- THÖNI, C., AND S. VOLK (2018): “Conditional cooperation: Review and refinement,” *Economics Letters*, 171, 37–40.
- VESTERLUND, L. (2003): “The informational value of sequential fundraising,” *Journal of Public Economics*, 87, 627–657.
- WONG, C. H., K. W. SIAH, AND A. W. LO (2019): “Estimation of clinical trial success rates and related parameters,” *Biostatistics*, 20, 273–286.

Supplemental appendix

A Proofs for Section 3.1

The results of this Section apply for parameter regions with internal solutions in both stages of the experimentation game, as used in the experiment. For all proofs, I will consider the case of observable experimentation (*Obs*) unless otherwise noted.

Lemma 1. *An agent's first- and second-stage experimentation efforts are strategic substitutes.*

To see that an agent's first and second stage experimentation are strategic substitutes for large enough Y , take the cross derivatives of the expected utility:

$$\frac{\partial^2 EU_{i,1}}{\partial e_{i,2} \partial e_{i,1}} = \frac{\partial^2 EU_{i,1}}{\partial e_{i,1} \partial e_{i,2}} = \frac{1}{4} (2pc'(e_{i,2}) - pY)$$

An agent's first- and second-stage experimentation are therefore strategic substitutes if and only if $\frac{1}{4} (2pc'(e_{i,2}) - pY) < 0$. For this to hold, it is a sufficient condition that $c'(e_{i,2}) < \frac{Y}{2}$, which was assumed. \square

Lemma 2. *An agent's and her partner's first- and second-stage experimentation efforts are strategic substitutes.*

To see that an agent's first- and her partner's second-stage experimentation are strategic substitutes, consider the following cross derivatives:

$$\frac{\partial^2 EU_{i,1}}{\partial e_{-i,2} \partial e_{i,1}} = \frac{\partial^2 EU_{i,1}}{\partial e_{-i,2} \partial e_{i,1}} = -\frac{pY}{4}$$

As $-\frac{pY}{4} < 0$, I can conclude that an agent's first and her partner's second stage experimentation are strategic substitutes. \square

Proposition 1. *The first-stage experimentation effort is higher if experimentation effort is unobservable.*

Consider the case where effort is observable. In the PBE, an agent will choose first-stage experimentation such that the marginal benefits from experimentation equal the marginal costs from experimentation, given the other player's response and their beliefs. Part of the costs of increasing experimentation are that the partner will decrease second-stage experimentation, because of the strategic substitutability.

Formally, this means that first-stage experimentation is chosen according to the following first order condition:

$$\frac{\partial EU_{i,1}}{\partial e_{i,1}} = \frac{pY}{2} - c'(e_{i,1}) + \left(1 - p \left(\frac{e_{i,1} + e_{-i,1}}{2}\right)\right) \frac{\partial EU_{i,2}}{\partial e_{i,1}} - \frac{p}{2} EU_{i,2} = 0$$

Next consider the case where effort is not observable. If effort is unobservable, this implies that agents base their decisions on their beliefs about their partner's first-stage effort $\hat{e}_{-i,1}$ instead of the actual partner's effort $e_{-i,1}$ when deciding on optimal effort in the second stage. Simultaneously, their partner will base their decisions on their beliefs about their partner's first-stage effort $\hat{e}_{i,1}$ instead of the actual effort $e_{i,1}$. In the above expression, this may affect experimentation through changes in the terms $\frac{\partial EU_{i,2}}{\partial e_{i,1}}$ and $EU_{i,2}$.

Assume agents do not observe their partner's experimentation but exert experimentation efforts that correspond to the effort levels in the PBE with observable effort. In this case, there is a profitable deviation to exert more effort in the first stage. To see this, consider how the two terms $\frac{\partial EU_{i,2}}{\partial e_{i,1}}$ and $EU_{i,2}$ compare for the two cases.

In the observable case:

$$\begin{aligned} \frac{\partial EU_{i,2}}{\partial e_{i,1}} = & \frac{\partial \rho(e_{i,1}, e_{-i,1})}{\partial e_{i,1}} \left(\frac{e_{i,2} + e_{-i,2}}{2}\right) Y + \rho(e_{i,1}, e_{-i,1}) \frac{\partial e_{i,2}(e_{i,1}, e_{-i,1})}{\partial e_{i,1}} \frac{Y}{2} + \\ & \rho(e_{i,1}, e_{-i,1}) \frac{\partial e_{-i,2}(e_{i,1}, e_{-i,1})}{\partial e_{i,1}} \frac{Y}{2} - c'(e_{i,2}) \frac{\partial e_{i,2}(e_{i,1}, e_{-i,1})}{\partial e_{i,1}} \end{aligned}$$

Now, consider how this term changes if experimentation is unobservable, while holding constant that agents exert experimentation efforts that correspond to the effort levels in the PBE of observable experimentation. In the unobservable case, $\rho(e_{i,1}, e_{-i,1}) \frac{\partial e_{i,2}(e_{i,1}, e_{-i,1})}{\partial e_{i,1}} \frac{Y}{2}$ is zero, as the unobservability of $e_{i,1}$ implies $\frac{\partial e_{i,2}(\hat{e}_{i,1}, e_{-i,1})}{\partial e_{i,1}} = 0$. As $\rho(e_{i,1}, e_{-i,1}) \frac{\partial e_{-i,2}(e_{i,1}, e_{-i,1})}{\partial e_{i,1}} \frac{Y}{2} < 0$ when effort is observable, see Lemma 2, $\frac{\partial EU_{i,2}}{\partial e_{i,1}}^{Obs} < \frac{\partial EU_{i,2}}{\partial e_{i,1}}^{Unobs}$ for a given $e_{i,1}$. Therefore, effort observability decreases incentives to experiment through a change in $\frac{\partial EU_{i,2}}{\partial e_{i,1}}$.

In the case that agents exert experimentation efforts that correspond to the effort levels in the PBE with observable experimentation, $EU_{i,2}$ is the same between the cases with observable and unobservable effort, as posterior beliefs $\rho(e_{i,1}, e_{-i,1})$ will be the same. Thus, taken together, there is a profitable deviation to experiment more with unobservable experimentation efforts. Hence, if there is an interior solution, first-stage experimentation efforts in the PBE with unobservable experimentation ($e_{i,1}^{Unobs}$) are higher than in the PBE with observable experimentation

$(e_{i,1}^{Obs})$. □

Proposition 2. *The second-stage experimentation effort is higher if experimentation effort is observable.*

The PBE requires, by sequential rationality, that the agent will maximise her expected utility given her beliefs. Thus, in the case where effort is observable, an agent will choose second-stage effort according to the following condition:

$$\frac{\rho(e_{i,1}, e_{-i,1})}{2} Y = c'(e_{i,2}) \quad (2)$$

If effort is unobservable, the agent will choose second-stage effort according to the following condition:

$$\frac{\rho(e_{i,1}, \hat{e}_{-i,1})}{2} Y = c'(e_{i,2}) \quad (3)$$

In the PBE, agents have correct beliefs about their partner's first-stage experimentation, implying that $\hat{e}_{-i,1} = e_{-i,1}$. Given that both $e_{i,1}$ and $e_{-i,1}$ are higher with unobservable than with observable effort, see Proposition 1, this gives

$$\rho(e_{i,1}^{Unobs}, \hat{e}_{-i,1}^{Unobs}) < \rho(e_{i,1}^{Obs}, e_{-i,1}^{Obs})$$

As $c''(e_{i,2}) > 0$, $e_{i,2}^{Unobs} < e_{i,2}^{Obs}$ has to hold such that Eqs. 2 and 3 are both satisfied. □

B Proofs for Section 3.2

The results in this Section apply for the parameters chosen in the experiment. For all proofs, I will consider the case of observable experimentation (*Obs*) unless otherwise noted.

Lemma 3. *An agent's second- and her partner's second-stage experimentation efforts as well as an agent's first- and her partner's first-stage experimentation efforts are strategic substitutes.*

To see that an agent's second and her partner's second-stage experimentation are strategic substitutes, take the cross derivatives of the expected utility:

$$\frac{\partial^2 EU_{i,1}}{\partial e_{i,2} \partial e_{-i,2}} = \frac{\partial^2 EU_{i,1}}{\partial e_{-i,2} \partial e_{i,2}} = -\frac{\rho(e_{i,1})\rho(e_{-i,1})Y}{4} \left(1 - \frac{p \times e_{i,1}}{2}\right) \left(1 - \frac{p \times e_{-i,1}}{2}\right) < 0$$

For the strategic interaction of an agent's first and her partner's first-stage experimentation, we will see that this depends now on how first-stage experimentation makes experimentation in the second stage more or less attractive through changing the likelihood of a breakthrough in one project, but also the likelihood of two simultaneous breakthroughs. Consider the cross derivative of the expected utility $\frac{\partial^2 EU_1}{\partial e_{i,1} \partial e_{-i,1}}$:

$$\begin{aligned} \frac{\partial^2 EU_{i,1}}{\partial e_{i,1} \partial e_{-i,1}} = & \frac{1}{4} \left(p^2 \times (EU_{i,2} - Y) + \frac{\partial^2 EU_{i,2}}{\partial e_{i,1} \partial e_{-i,1}} (2 - p \times e_{i,2}) (2 - p \times e_{-i,2}) \right. \\ & \left. - p \times \left(\frac{\partial EU_{i,2}}{\partial e_{i,1}} \times (2 - p \times e_{i,1}) - \frac{\partial EU_{i,2}}{\partial e_{-i,1}} \times (2 - p \times e_{-i,1}) \right) \right) \end{aligned}$$

Here, the cross derivative depends on how first-stage experimentation effects incentives for second-stage experimentation. Given the parameters chosen and an agent's best response to first-stage experimentation levels in the second stage, see Lemma 4, this gives $\frac{\partial^2 EU_{i,1}}{\partial e_{i,1} \partial e_{-i,1}} < 0$, and an agent's first- and her partner's first-stage experimentation are strategic substitutes. \square

Lemma 4. *An agent's second-stage experimentation increases in their partner's first-stage experimentation.*

Consider the following cross derivatives of the expected utility:

$$\frac{\partial^2 EU_{i,2}}{\partial e_{i,2} \partial e_{-i,1}} = -\frac{\rho(e_{i,1}) \times Y}{4} \left(\frac{\partial \rho(e_{-i,1})}{\partial e_{-i,1}} e_{-i,2} + \frac{\partial e_{-i,2}}{\partial e_{-i,1}} \rho(e_{-i,1}) \right)$$

We see here that this cross derivative also depends on how the partner's experimentation responds to their first-stage experimentation $\left(\frac{\partial e_{-i,2}}{\partial e_{-i,1}} \right)$, as the partner's experimentation influences the probability of a joint breakthrough. So consider an agent's best response to the observed first-stage experimentation levels. In the second stage, an agent chooses the optimal experimentation level according to the following first order condition:

$$\frac{\partial EU_{i,2}}{\partial e_{i,2}} = \left(\frac{\rho(e_{i,1})}{2} - \frac{\rho(e_{-i,1}) \rho(e_{i,1})}{4} e_{-i,2} \right) Y - c'(e_{i,2}) \stackrel{!}{=} 0$$

For the parameters chosen in the experiment, using that, by symmetry, $e_{i,2} = e_{-i,2}$ in the PBE, this implies that second-stage experimentation in the PBE is given by

$$\begin{aligned} e_{i,2}^* = & \frac{(0.9e_{i,1}^2 - 22.71e_{-i,1} + 61)e_{i,1}^2 + (136.28e_{-i,1} - 5.38e_{-i,1}^2 - 365.8)e_{i,1} + 7.17e_{-i,1}^2 - 181.70e_{-i,1} + 487.72}{(e_{-i,1}^2 - 15.77e_{-i,1} + 39.31)e_{i,1}^2 + (157.24e_{-i,1} - 15.77e_{-i,1}^2 - 345.56)e_{i,1} + 39.31e_{-i,1}^2 - 345.56e_{-i,1} + 722.21} \end{aligned}$$

The derivative of $e_{i,2}^*$ defined above w.r.t. $e_{-i,1}$ is positive, so second-stage experimentation is increasing in the partner's first-stage experimentation, if observable. \square

Proposition 3. *The first-stage experimentation effort is higher if experimentation effort is observable.*

Consider first that effort is observable. In the PBE, an agent will choose first-stage experimentation such that, given the other player's strategy and the agent's beliefs, the marginal benefits from experimentation equal the marginal costs from experimentation:

$$\begin{aligned} \frac{\partial EU_{i,1}}{\partial e_{i,1}} &= \frac{2 - pe_{-i,1}}{4} \times pY - c'(e_{i,1}) + \left(1 - p\frac{e_{-i,1}}{2}\right) \times \\ &\quad \left(\left(1 - p\frac{e_{i,1}}{2}\right) \frac{\partial EU_{i,2}}{\partial e_{i,1}} - \frac{p}{2} EU_{i,2}\right) \stackrel{!}{=} 0 \end{aligned} \quad (4)$$

Consider again the case where effort is not observable. Agents base their decisions on their beliefs about their partner's first-stage effort $\hat{e}_{-i,1}$ instead of the actual partner's effort $e_{-i,1}$ in the second stage. Their partner will base their decisions on their beliefs about their partner's first-stage effort $\hat{e}_{i,1}$. In Eq. 4, this affects $\frac{\partial EU_{i,2}}{\partial e_{i,1}}$ and $EU_{i,2}$. I will now consider how these terms depend on the observability of experimentation effort $e_{i,1}$. With effort observability, $\frac{\partial EU_{i,2}}{\partial e_{i,1}}$ is given by

$$\begin{aligned} \frac{\partial EU_{i,2}}{\partial e_{i,1}} &= \frac{\partial \rho(e_{i,1})}{\partial e_{i,1}} Y \left(\frac{e_{i,2}(e_{i,1}, e_{-i,1})}{2} - \frac{\rho(e_{-i,1}) e_{i,2}(e_{i,1}, e_{-i,1}) e_{-i,2}(e_{i,1}, e_{-i,1})}{4} \right) + \\ &\quad \frac{\rho(e_{i,1})}{2} Y \left(\frac{\partial e_{i,2}(e_{i,1}, e_{-i,1})}{\partial e_{i,1}} - \frac{\rho(e_{-i,1}) e_{-i,2}(e_{i,1}, e_{-i,1})}{2} \frac{\partial e_{i,2}(e_{i,1}, e_{-i,1})}{\partial e_{i,1}} \right) + \\ &\quad \frac{\rho(e_{-i,1})}{2} Y \left(\frac{\partial e_{-i,2}(e_{i,1}, e_{-i,1})}{\partial e_{i,1}} - \frac{\rho(e_{i,1}) e_{i,2}(e_{i,1}, e_{-i,1})}{2} \frac{\partial e_{-i,2}(e_{i,1}, e_{-i,1})}{\partial e_{i,1}} \right) - \\ &\quad c'(e_{i,2}(e_{i,1}, e_{-i,1})) \frac{\partial e_{i,2}(e_{i,1}, e_{-i,1})}{\partial e_{i,1}} \end{aligned}$$

Assume now that agents exert experimentation efforts that correspond to the effort levels in the PBE with unobservable effort while $e_{i,1}$ and $e_{-i,1}$ are unobservable. In the unobservable case, $\frac{\partial EU_{i,2}}{\partial e_{i,1}}$ differs from the expression above, as

$$\frac{\partial e_{-i,2}(\hat{e}_{i,1}, e_{-i,1})}{\partial e_{i,1}} - \frac{\rho(e_{i,1}) e_{i,2}(\hat{e}_{i,1}, e_{-i,1})}{2} \frac{\partial e_{-i,2}(\hat{e}_{i,1}, e_{-i,1})}{\partial e_{i,1}} = 0$$

given that $\frac{\partial e_{-i,2}(\hat{e}_{i,1}, e_{-i,1})}{\partial e_{i,1}} = 0$ with unobservable effort.

If agents exert experimentation effort corresponding to the PBE with unobservable effort, this creates a profitable deviation to exert more effort in the first stage if effort is observable. Recall from Lemma 4 that an agent's second-stage experimentation is increasing in their partner's first-stage experimentation if observable. Furthermore, for a given experimentation level, an agent's second-stage expected utility increases in their partner's second-stage experimentation:

$$\frac{\partial EU_{i,2}}{\partial e_{-i,2}} = \left(\frac{\rho(e_{-i,1})}{2} - \frac{\rho(e_{-i,1})\rho(e_{i,1})}{4} e_{i,2} \right) Y > 0$$

Therefore, effort observability increases incentives to experiment through a change in $\frac{\partial EU_{i,2}}{\partial e_{i,1}}$. If agents exert experimentation efforts that correspond to the effort levels in the PBE with unobservable effort, $EU_{i,2}$ is again constant between the cases with observable and unobservable effort, as posterior beliefs $\rho(e_{i,1})$ and $\rho(e_{-i,1})$ will be the same. Hence, there is a profitable deviation to experiment more with observable experimentation.

Thus, first-stage experimentation efforts in the PBE with observable experimentation are higher than in the PBE with unobservable experimentation. \square

Proposition 4. *The second-stage experimentation effort is higher if experimentation effort is unobservable.*

In the PBE the agent will maximise her expected utility given her beliefs (sequential rationality). In the case where effort is observable, an agent will choose second-stage effort according to the following condition:

$$\left(\frac{\rho(e_{i,1})}{2} - \rho(e_{i,1})\rho(e_{-i,1})\frac{e_{-i,2}}{4} \right) Y = c'(e_{i,2}) \quad (5)$$

If effort is unobservable, second-stage effort will be chosen such that:

$$\left(\frac{\rho(e_{i,1})}{2} - \rho(e_{i,1})\rho(\hat{e}_{-i,1})\frac{e_{-i,2}}{4} \right) Y = c'(e_{i,2}) \quad (6)$$

In the PBE, $\hat{e}_{-i,1} = e_{-i,1}$. As $e_{-i,1}$ and $e_{i,1}$ are higher with observable than with unobservable effort, see Proposition 3:

$$\rho(e_{i,1}^{Obs}) < \rho(e_{i,1}^{Unobs}) \quad \wedge \quad \rho(e_{-i,1}^{Obs}) < \rho(\hat{e}_{-i,1}^{Unobs})$$

With $c''(e_{i,2}) > 0$, this implies in the symmetric PBE where $\rho(e_{i,1}^{Obs}) = \rho(e_{-i,1}^{Obs})$ and $\rho(e_{i,1}^{Unobs}) = \rho(\hat{e}_{-i,1}^{Unobs})$ that $e_{i,2}^{Unobs} > e_{i,2}^{Obs}$ such that Eqs. 5 and 6 hold simultaneously. \square

C Additional analysis

Tables 13 and 14 reproduce the results from Tables 3 and 4 in the main text, only including observations from the last 15 rounds of the experimentation game.

	1st stage		
	Unobservable	Observable	
Joint	60.08%	71.00%	65.48%
Separate	48.00%	55.43%	51.62%
	53.99%	63.21%	

Notes: Average experimentation effort in the first stage for experimentation after round 15.

Table 13: Experimentation effort per treatment in the first stage in the second half of the experiment

	2nd stage		
	Unobservable	Observable	
Joint	33.04%	25.34%	29.15%
Separate	37.73%	38.45%	38.09%
	35.35%	31.90%	

Notes: Average experimentation effort in the second stage for experimentation after round 15.

Table 14: Experimentation effort per treatment in the second stage in the second half of the experiment

Table 15 and Table 16 reproduce Table 6 and Table 7 from the main text, respectively, but include subject-level fixed effects. This shows that the fixed effects absorb (parts of) the observed effect of own experimentation on posterior beliefs in treatments with joint experimentation, see columns (1) and (2). From this, I can conclude that the variation in experimentation that results in variation of posterior beliefs is mainly between-subject variation in first-stage experimentation. Table 17 provides estimates of the two-step regression in which residuals are regressed on own first-stage experimentation, controlling for an effect through beliefs. Table 18 shows the correlation between measures of reciprocity and experimentation behaviour and how this depends on the effort observability. Table 19 shows how participants' first-stage experimentation correlates with their partners' last-round's experimentation, again depending on the effort observability.

Treatments	<i>Dependent variable: Posterior of project's quality</i>			
	Joint projects		Separate projects	
	Observable	Unobservable	Observable	Unobservable
	(1)	(2)	(3)	(4)
Own effort	0.05 (0.06)	-0.05 (0.05)	0.05 (0.04)	-0.02 (0.04)
Partner's effort	-0.11 (0.05)	-0.04 (0.05)	0.02 (0.03)	0.03 (0.03)
Constant	33.19 (6.93)	38.36 (4.03)	31.63 (2.79)	35.65 (2.09)
Fixed effects	✓	✓	✓	✓
Observations	2880	2850	2880	2880
Clusters	16	16	16	16
R-squared	0.535	0.593	0.517	0.563

Notes: OLS estimating effect of own and partner's first-stage experimentation effort on posterior of project's quality. (1) and (3) use the partner's actual experimentation effort, (2) and (4) use the subject's belief about the partner's experimentation effort. Robust standard errors (in parentheses) are clustered at the matching group level. Individual fixed-effects included.

Table 15: Effect of first-stage experimentation on beliefs about partner's posterior

Treatments	<i>Dependent variable: Belief about partner's posterior</i>			
	Joint projects		Separate projects	
	Observable	Unobservable	Observable	Unobservable
	(1)	(2)	(3)	(4)
Own effort	-0.00 (0.05)	0.02 (0.03)	0.04 (0.02)	-0.02 (0.02)
Partner's effort	-0.09 (0.06)	-0.07 (0.06)	0.08 (0.04)	0.03 (0.05)
Constant	35.50 (6.57)	34.69 (4.35)	28.36 (2.48)	34.67 (2.34)
Fixed effects	✓	✓	✓	✓
Observations	2880	2850	2880	2880
Clusters	16	16	16	16
R-squared	0.502	0.588	0.503	0.598

Notes: OLS estimating effect of own and partner's first-stage experimentation effort on beliefs about partner's posterior of project's quality. (1) and (3) use the partner's actual experimentation effort, (2) and (4) use the subject's belief about the partner's experimentation effort. Robust standard errors (in parentheses) are clustered at the matching group level.

Table 16: Effect of first-stage experimentation on beliefs about partner's posterior

<i>Dependent variable:</i> Residuals of 1st-stage regression	
Observable	18.15 (6.06)
Own effort	0.20 (0.07)
Observable \times Own effort	-0.28 (0.09)
Constant	-12.11 (4.62)
Observations	5730
Clusters	32
R-squared	0.030

Notes: OLS estimating difference-in-difference in the effect of own first-stage experimentation between treatments on residuals of regression of second-stage experimentation on posteriors. Elicited beliefs are used for observations from unobservable treatment. Robust standard errors (in parentheses) are clustered at the matching group level.

Table 17: Effect of own first-stage experimentation on second-stage experimentation

Treatments	<i>Dependent variable: Experimentation effort</i>			
	Joint projects		Separate projects	
	(1)	(2)	(3)	(4)
Negative reciprocity	6.49 (4.25)		3.81 (5.01)	
Observable	4.34 (9.86)	10.93 (11.64)	2.37 (9.55)	-10.88 (10.03)
Negative reciprocity \times Observable	-9.64 (6.47)		-1.17 (6.19)	
Offer		4.31 (4.62)		-2.37 (3.91)
Offer \times Observable		-1.00 (5.94)		11.96 (5.70)
Constant	25.87 (6.65)	55.28 (9.28)	33.47 (6.83)	54.06 (6.95)
Observations	5730	5730	5760	5760
Clusters	32	32	32	32
R-squared	0.023	0.031	0.004	0.037

Notes: OLS estimating difference in the correlation of behaviour in the ultimatum game and experimentation. ‘Negative reciprocity’ refers to the minimum acceptable offer in the ultimatum game, ‘Offer’ to the amount offered in the ultimatum game. Columns (1) and (3) use second-stage experimentation as the outcome variable, columns (2) and (4) first-stage experimentation. Robust standard errors (in parentheses) are clustered at the matching group level.

Table 18: Correlation between behavior in ultimatum game and experimentation

	<i>Dependent variable: Experimentation effort</i>	
	(1)	(2)
Partner's experimentation $t - 1$	0.19 (0.05)	0.04 (0.01)
Observable	2.33 (5.83)	
Observable \times Partner's experimentation $t - 1$	0.07 (0.08)	0.04 (0.02)
Constant	45.49 (3.73)	56.29 (0.55)
Fixed effects	X	✓
Observations	11107	11107
Clusters	64	64
R-squared	0.065	0.007

Notes: OLS estimating the differential correlation of partner's first-stage experimentation in the last round on own experimentation depending on whether this was observable. Fixed effects refer to individual-level fixed effects. Robust standard errors (in parentheses) are clustered at the matching group level.

Table 19: Correlation of experimentation with partner's last-round experimentation

D Experimental instructions

Instructions part 1

The instructions are simple, and if you follow them carefully, you might earn a considerable amount of money. Your earnings will depend on your decisions and may depend on other participants' decisions.

This experiment consists of two parts. First, we are going to explain part 1 of the experiment to you. After making decisions in part 1, the next part will be explained to you.

In the first part of this experiment, you will repeatedly play a game with changing partners that consists of multiple stages. You will play 30 rounds of this game. Each round you make two decisions. Both your choices and your partner's choices will affect your payoffs.

The task

For each decision you make in each round that you play this game, you receive a budget of €2. Your main task is to decide what share to invest in a project. In each round of the experiment, you will have two opportunities to do so. We will call the percentage share you invest in the project $x\%$.

For investing in this project, you will be charged costs. Costs are higher if you invest a higher share. The higher your investment is, the costlier it becomes to further increase your investment.

More precisely, you can invest between 0% and 100%. If you invest $x\%$, $€2 \left(\frac{x}{100}\right)^2$ will be subtracted from your budget.

Examples:

If you invest 0%, the costs are $2 \times \left(\frac{0}{100}\right)^2 = 0€$.

If you invest 30%, the costs are $2 \times \left(\frac{30}{100}\right)^2 = 0.18€$.

If you invest 60%, the costs are $2 \times \left(\frac{60}{100}\right)^2 = 0.72€$.

If you invest 100%, the costs are $2 \times \left(\frac{100}{100}\right)^2 = 2€$.

These costs are subtracted from your budget.

Breakthroughs

The project you can invest in is of high or of low quality. You do not know the quality of the project. With 50% probability, the project is of high quality. With

50% probability, the project is of low quality. This means that if you would face 100 of these projects, you can expect about 50 of these to be high-quality projects.

[*Joint*: You and your partner both invest jointly in the same project. This means that if your partner is investing in a project of high quality, so are you. Similarly, if your partner is investing in a project of low quality, so are you.]

[*Separate*: You and your partner invest in separate projects. This means that if your partner is investing in a project of high quality, the project you invest in is not necessarily of high quality, too. Similarly, if your partner is investing in a project of low quality, the project you invest in is not necessarily of low quality, too.]

[*Joint*: If your project has a breakthrough, you and your partner each receive a payoff of €13 from the project's breakthrough.] [*Separate*: Both you and your partner receive a payoff of €13 each if there is at least one breakthrough in a project. This breakthrough can be in your project or in your partner's project. If both projects have a breakthrough, you also each receive €13.] Only high-quality projects can have a breakthrough. Low-quality projects can never have a breakthrough. This means that you will never receive a payoff of €13 from a low-quality project.


Next to the project's quality, whether there is a breakthrough also depends on how much [*Joint*: you and your partner invest in your joint project.] [*Separate*: is invested in each of the projects. The more either of you invests in his or her project, the more likely that project has a breakthrough.] If you face a high-quality project, the probability of a breakthrough increases with the share [*Joint*: you and your partner together invest.][*Separate*: that is invested in the project.] More specifically, the probability of a breakthrough is [*Joint*: the average of your and your partner's investment share.][*Separate*: half of the investment share.] If you invest $x\%$ [*Joint*: and your partner invests $y\%$,] the probability of a breakthrough [*Joint*: is thus $\frac{x+y}{2}\%$][*Separate*: in your project is $\frac{x}{2}\%$] for high-quality projects. [*Separate*: If your partner invests $y\%$, the probability of a breakthrough in your partner's project is $\frac{y}{2}\%$ if it is a high-quality project.]

If [*Joint*: you are] [*Separate*: someone is] facing a high-quality project and [*Joint*: both you and your partner invest][*Separate*: that person invests] a share of 100% in [*Joint*: this] [*Separate*: his or her] project, [*Joint*: you will certainly have a breakthrough and will both receive €13.] [*Separate*: there will be a breakthrough with a probability of 50%.] On the other hand, if [*Joint*: both you and your partner invest nothing in this project, you will never have a breakthrough,] [*Separate*: someone invests nothing in his or her project, there will never be a


breakthrough], no matter whether the project is of high quality or not.

You can determine the likelihood of a breakthrough in a project as follows.

[*Joint:*



] [*Separate:*



] After you and your partner have made your investment choices, the computer will determine whether there actually is a breakthrough [*Separate:* for both projects]. To determine this, the computer will use the breakthrough probability.

Examples:

Let's say that you invest 20% [*Separate:* in your project] and your partner invests 64% in [*Joint:* the] [*Separate:* his or her] project. [*Joint:* If the project is of low quality, there will not be a breakthrough. If the project is of high quality, the probability of a breakthrough is $\frac{20\%+64\%}{2} = 42\%$.] [*Separate:* If your project is of low quality, there will not be a breakthrough for your project. If your partner's project is of low quality, there will not be a breakthrough for his or her project. If your project is of high quality, the probability of a breakthrough in this project is $\frac{20}{2} = 10\%$. If your partner's project is of high quality, the probability of a breakthrough in this project is $\frac{64}{2} = 32\%$]

There can only be a breakthrough after the first or after the second investment decision. If [*Joint:* the] [*Separate:* any] project has a breakthrough after your first investment decision, you cannot invest anymore. You will receive €13 plus the budget of your second investment decision.

Within one round, the project you are investing in does not change. If the project is of high quality for your first investment decision, it will also be of high quality for your second investment decision in this round. Similarly, if the project is of low quality for your first investment decision, it will also be of low quality for your second investment decision. This means that the first investments and results of the first investments may contain information relevant to your second decisions.

In each new round you will face a new project. While in each round you face a project of high quality with 50%, the project's actual quality in one round does not say anything about the project's quality in any other round.

Your decisions

After being matched with a partner for a round, you will be asked to make three types of decisions. First, an investment decision, second, predictions about the project's quality and your partner's choices, and third, another investment decision. We now describe each of these three decisions in more detail.

First investment decision

For your first investment decision, you decide which share you want to invest and then submit your decision. You can use an on-screen calculator that will allow you to calculate your costs for any given investment and the probability of a breakthrough. You will see this interface on a later screen.

If there is a breakthrough after the first investment decisions, you will receive the payoff minus your costs from the first investment decision added to your budget. The breakthrough terminates this round. The second investment decisions are in this case not relevant.

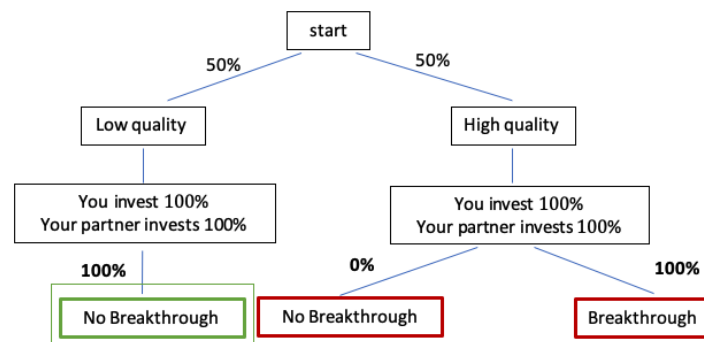
Prediction task

After the first investment decision, you are asked to state your beliefs about the project's quality, your partner's investment as well as your partner's beliefs about

the project's quality. You will also be paid according to your performance in this task. This task will be explained in more detail on a later screen.


Second investment decision

After your first investment decision, we will ask you how much you would want to invest in the project **if there was no breakthrough after the first investment decisions**. Your second investment decision will then only be implemented in case there was indeed no breakthrough after the first investment decisions. This means that you should decide how much you think is best to invest in a project where there has not yet been a breakthrough. For the case that there was a breakthrough after the first investment decisions, you receive your €2 second-period budget added to your payoff. So, after a breakthrough you will still be asked to make the second investment decision, but this will only be relevant for your payoffs if there indeed was no breakthrough! Decide as if there was no breakthrough so far. *[Joint:* Realize that if you would know for sure that both you and your partner had invested your entire budgets in the project in the first investment decision, while there was no breakthrough, then, the project cannot be of high quality. This is why: the probability of a breakthrough if the project is of high quality and you both invest 100% is given by $\frac{100\%+100\%}{2} = 100\%$. This means that if you observe no breakthrough, you are for sure in the far-left green branch of the tree below. You would have certainly seen a breakthrough if the project were of high quality.

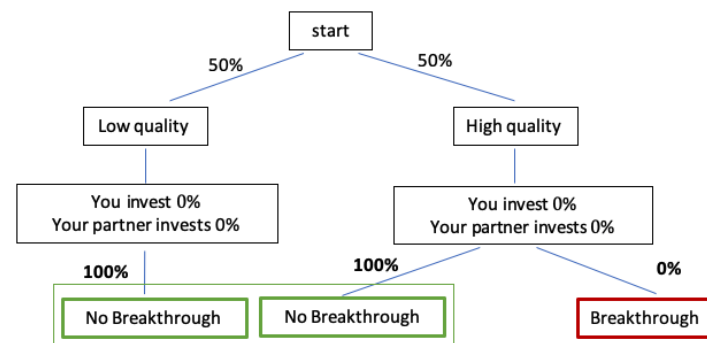


]*[Separate:* Realize that if someone had invested his or her entire budget in his or her project in the first investment decision, while there was no breakthrough, then it is twice as likely to have a low-quality project than a high-quality project. This is why: You see below that we must be in one of the two left branches of the tree, within the green box. The project could be of low quality, then

we would observe no breakthrough with a probability of 100% (far left branch). Alternatively, the project is of high quality, but there was no breakthrough (middle branch). This is only half as likely, as if you invest 100% and the project is of high quality, the probability of no breakthrough is only $\frac{100\%}{2} = 50\%$.

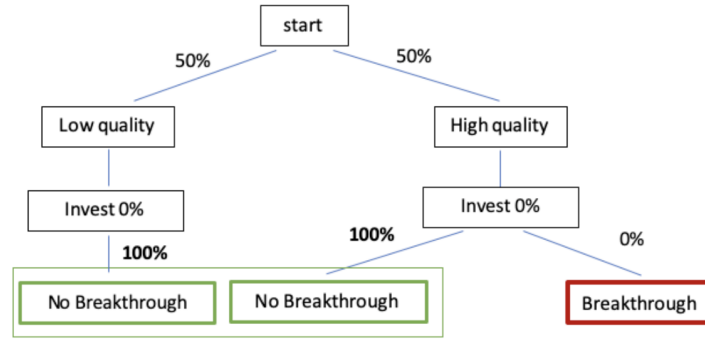


] If, in contrast, [*Joint*: both you and your partner invest] [*Separate*: someone invests] nothing in the project in the first investment decision, then you cannot learn anything new about the project. The probability of a breakthrough is 0%. [*Joint*: So if there is no breakthrough, it is equally likely to be in one of the two green branches below, the far left or the middle one.] The probability that the project is of high quality is in this case still 50%. [*Separate*: Below you see that no matter whether the project is of high or low quality, the probability of no breakthrough is always 100%. It is then equally likely to be in the left low-quality branch or in the middle high-quality branch.] [*Joint*:



] [*Separate*:

] We will show you a graph to illustrate your expected payoff from making specific investments. As before, you can also use on-screen calculators that will allow you to calculate your costs for any given investment and the probability of a breakthrough. Now, this depends on your beliefs about the probability that the project is of high quality. This will be illustrated on a later screen.



Feedback

After your first investment decision, you will [*Unobservable*: not] see which share your partner invested. Your partner will also [*Unobservable*: not] see which share you invested. After the second investment decision, you will see whether there was a breakthrough and how much your payoff from this round is.

Your partner

Your partner is anonymous and so are you. Your partner is the same for both investment decisions. You face the same decision situation. After each round of the experiment, you will be randomly assigned to a new partner. We ensure that you are never linked to the same partner for two rounds in a row. Also, your actions in any round have no influence on anything that happens in other rounds and are not known to your partners in following rounds.

Payoffs from this task

To summarize, your payoffs from each investment decision are the following:

- If you invest $x_{1st}\%$ in this project in the 1st investment decision:
 - If you [*Separate*:, your partner or both of you] achieve a breakthrough: $2 + 13 - 2 \times \left(\frac{x_{1st}}{100}\right)^2$
 - If [*Joint*: you do not achieve] [*Separate*:, no one achieves] a breakthrough: $2 - 2 \times \left(\frac{x_{1st}}{100}\right)^2$
- If you invest $x_{2nd}\%$ in this project in the 1st investment decision:
 - If there was no breakthrough after the first investment decision:
 - * If you [*Separate*:, your partner or both of you] achieve a breakthrough: $2 + 13 - 2 \times \left(\frac{x_{2nd}}{100}\right)^2$

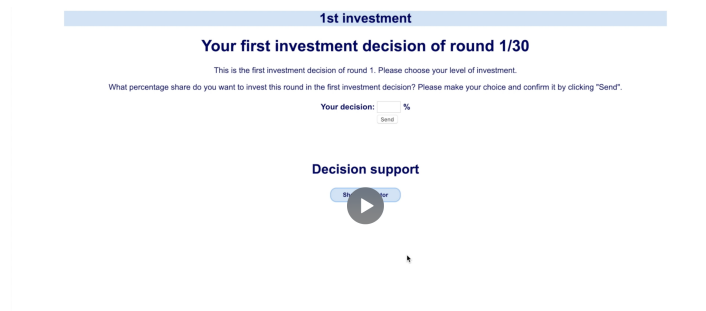
- * If [*Joint*: you do not achieve] [*Separate*: no one achieves] a breakthrough: $2 - 2 \times \left(\frac{x_{2nd}}{100}\right)^2$
- If there was a breakthrough after the first investment decision:
 - * 2

At the end of the experiment, **two rounds** of this investment game will be randomly chosen by the computer for payment. Every round is equally likely to be chosen for payment, so your actions in a round have no influence on whether that round will be paid out. Aside from the payment from the investment game, two different rounds will be chosen from which the prediction task will be paid. We will explain on another screen how payment is determined for the prediction task.

Illustration

First investment decision

For your first investment decision, you make a choice about which share to invest in the project. Please see this video illustrating how to make your first investment decision. You can pause the video at any moment, re-watch it as many times as you like and put it on full-screen if you prefer this.



Please try using the calculator below. [*Joint*:

Decision support

Calculator

Your investment: %

Your partner's investment: %

Probability that the project is of high quality: 50%

Breakthrough probability: %

Your cost of investment: cents

] [*Separate*:

Decision support

Calculator

Your investment: %

Your partner's investment: %

Probability that your project is of high quality: 50%

Probability that your partner's project is of high quality: 50%

Breakthrough probability of your project: %

Breakthrough probability of your partner's project: %

Your cost of investment: cents

]

Second investment decision

For the second investment decision, you again make a choice about which share of your budget to invest in case the project did not have a breakthrough after the first investment decisions.

On your decision screen, you will have the opportunity to see a graph of your expected payoff from investing a certain share. Expected payoff means that this is not a certain payoff from investing this share, but that this is what you are going to receive in expectation. If you would do this investment frequently, on average you would get the expected payoff. The realized payoff from investing a share x will always be either $13 - 2 * (\frac{x}{100})^2$ (if there is a breakthrough) or $-2 * (\frac{x}{100})^2$ (if there is no breakthrough), which will be added to your budget of €2.

Your expected payoff depends on [*Joint*: three] [*Separate*: four] factors:

1. How likely [*Joint*: the] [*Separate*: your] project is of high quality: Your expected payoff is higher if [*Joint*: the] [*Separate*: your] project is more likely of high quality, as only [*Joint*: high-quality projects can result in a breakthrough and thus in a payoff of €13 for you and your partner.][*Separate*: then your project can have a breakthrough.]
2. [*Separate*: How likely your partner's project is of high quality: Your expected payoff is higher if your partner's project is more likely of high quality, as only then your partner's project can have a breakthrough.]
- 2/3 The share you invest: This also increases the probability of a breakthrough [*Separate*: in your project] if the project is of high quality.
- 3/4 The share your partner invests: If he or she invests more, this increases the probability of a breakthrough [*Joint*: if the project is of high quality.][*Separate*: in his or her project if this project is of high quality.]

Please see this video illustrating how to make your second investment decision. You can pause the video at any moment, re-watch it as many times as you like and put it on full-screen if you prefer this.

2nd investment

Your second investment decision round 1/30

No breakthrough after first decision

This is the second investment decision of round 1. Please choose your level of investment.

In the first investment decision, you invested $x_1\%$.

If there was no breakthrough after the first investment decision, what percentage share do you want to invest this round in the second investment decision? Please make your choice and confirm it by clicking "Send".

Your decision: %

Send

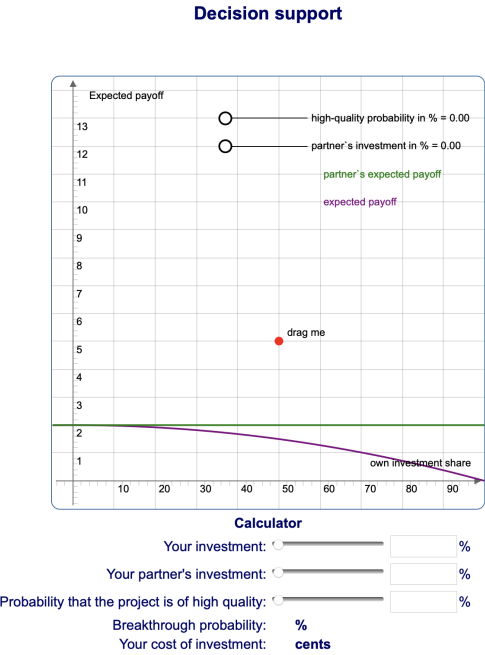
Decision support

Show calculator and graph

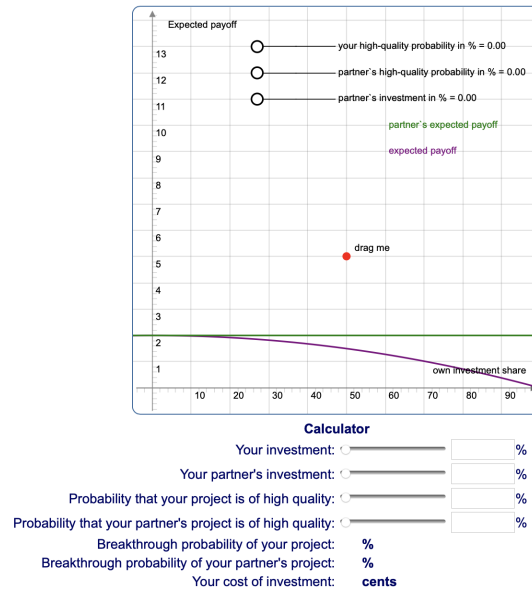
Breakthrough after first decision

If there was a breakthrough after the first decision, you receive your budget of €2 added to your payoff.

Please try using the calculator and the graph below. [Separate:



Decision support



]

Summary

- You and your partner can [*Joint*: invest in a common project] [*Separate*: each invest in a project]
- [*Joint*: This] [*Separate*: A] project is of high quality with probability 50% and of low quality with probability 50%
- [*Joint*: You don't know whether the project is of high or of low quality] [*Separate*: You know neither whether the project you face nor whether the project your partner faces is of high or of low quality]
- You and your partner invest in [*Joint*: the same project with the same quality] [*Separate*: separate projects which can have different qualities]
- You have a budget of €2 for each investment decision
- Investing a share of $x\%$ in the project costs $2 \times \left(\frac{x}{100}\right)^2$
- You and your partner receive a payoff of €13 if there is [*Joint*: a breakthrough] [*Separate*: at least one breakthrough in one of the two projects.]
- If the project is of high quality, the probability of a breakthrough [*Joint*: is given by the average of your investment share $x\%$ and your partner's

investment share $y\%$, $\frac{x+y}{2}\%$]. [*Separate:* in this project is given by half the share of $x\%$ which you or your partner invested in that project: $\frac{x}{2}\%$]

- If [*Joint:* the] [*Separate:* a] project is of low quality, no breakthrough is possible
- Breakthroughs end the project
- If there was no breakthrough after the first investment decision, you and your partner can invest again in [*Joint:* the project] [*Separate:* your projects]
- After the first investment decision, you will [*Unobservable:* not] see which share your partner invested. Your partner does [*Unobservable:* not] see the share you invested.
- Your second investment decision is for the case that there was no breakthrough after the first investment decision
- Your partner and the project are the same for the first and second investment decision, but change every round

E Experimental instructions of belief elicitation¹⁶

Your predictions

In each round, you will make several guesses after the first decision of the investment game.

1. [*Unobservable:* What share (in percent) did your partner invest in the first investment decision of this round?]
2. What is the probability (in percent) that [*Joint:* the] [*Separate:* your] project is of high quality if there was no breakthrough after the first investment decision?
3. What does your partner think is the probability (in percent) that [*Joint:* the] [*Separate:* his or her] project is of high quality if there was no breakthrough after the first investment decision?

¹⁶These instructions build on and are in parts taken from instructions in Babcock et al. (2017)

4. What share (in percent) will your partner invest in the second investment decision of this round if there was no breakthrough after the first investment decision?

Before you give your guesses, we will remind you of the share you invested in the first investment decision.

Your guess will secure you a payment of either €2 or €0 for each guess. If you win, you receive €2. If you lose, you instead receive €0 for your guess. Your payoffs from this task are such that you maximize the probability of receiving a prize of €2 by stating your best guess for each question.

Except for guess #1, you can see that you are asked about your beliefs in case there was no breakthrough after the first investment decision. Therefore, the computer randomly picks two rounds in which there was no breakthrough after the first investment decision for payment. From these two randomly selected rounds, you have the chance to win €2 for each of the guesses depending on your answer. So you can earn up to €[*Unobservable*: 16] [*Observable*: 12] from your guesses.

To determine your probability of winning the prize for each guess, we will compare your guess to what actually happens. We designed the payment rule such that you can secure the largest chance of winning the prize by reporting your most-accurate guess. Below, you can read more about how we determine whether you win the prize. To maximize your chances of winning the prize, it is not necessary that you understand how this works. While the mechanism may look complicated, what it means for you is simple: you have the highest chance of winning €2 if you report your best guess for each question.

Click here for more information on the mechanism

For each question, we will use your guess to calculate a chance-to-win. How we do this is explained below. We use this chance-to-win to determine whether you win €2. The computer generates a random number between 1 and 100 separately for each question. Each of the numbers is equally likely. You win €2 if this random number equals or falls below your chance-to-win, and you earn €0 if the random number exceeds your chance-to-win.

To maximize your earnings, you should submit a guess that secures a high chance-to-win for the events you think are most likely, and a low chance-to-win for the events that you think are least likely. If you, for instance, believe that it is very likely that the project is of high quality, you should submit a guess that secures a high chance-to-win for the case that the project is of high quality.

To secure that it is in your best interest to enter your best guess, we use the following procedure to calculate your chance-you-win for guess #2: Suppose you submitted a guess of p_1 that [*Joint: the*] [*Separate: your*] project is of high quality. Then your chance-to-win depends on whether the realized quality of [*Joint: the*] [*Separate: your*] project is high or low. If the project is of high quality, your chance-to-win is given by the equation:

$$\text{Chance} - \text{to} - \text{win} : \left(1 - \left(1 - \frac{p_1}{100}\right)^2\right) \times 100$$

If the project is of low quality, your chance-to-win would be given by the equation:

$$\text{Chance} - \text{to} - \text{win} : \left(1 - \left(\frac{p_1}{100}\right)^2\right) \times 100$$

This means that you have the highest probability of earning the €2 if your guess p_1 is what you believe is the probability that [*Joint: the*] [*Separate: your*] project is of high quality.

Example

Let's say that your best guess of the probability that [*Joint: the*] [*Separate: your*] project is of high quality is 40%. If you state this truthfully, then your chance-to-win is $\left(1 - \left(1 - \frac{40}{100}\right)^2\right) \times 100 = 64$ if [*Joint: the*] [*Separate: your*] project is of high quality and $\left(1 - \left(\frac{40}{100}\right)^2\right) \times 100 = 84$ if [*Joint: the*] [*Separate: your*] project is of low quality. As your best guess of the probability that the project is of high quality is 40%, the probability that you receive the prize is then $64 \cdot 0.4 + 84 \cdot 0.6 = 76\%$. If, for instance you would untruthfully state that your best guess of the probability is 70%, the probability of receiving the prize is lower: Your chance-to-win is $\left(1 - \left(1 - \frac{70}{100}\right)^2\right) \times 100 = 91$ if [*Joint: the*] [*Separate: your*] project is of high quality and $\left(1 - \left(\frac{70}{100}\right)^2\right) \times 100 = 64$ if [*Joint: the*] [*Separate: your*] project is of low quality. As your best guess is 40%, the probability that you receive the prize is then $91 \cdot 0.4 + 64 \cdot 0.6 = 67\%$, which is lower.

For the remaining questions, we use the following procedure: Suppose that you submitted a guess that your partner's belief about the probability that [*Joint: the*] [*Separate: his or her*] project is of high quality is $x\%$ (guess #3) or that your partner is going to invest [*Unobservable: (or invested)*] $x\%$ in the second investment decision of the round (guess #4 [*Unobservable: and guess #1*]). Then your chance-to-win depends on what your partner actually believes is the probability

of [*Joint:* the] [*Separate:* his or her] project being of high quality or on how much he or she actually invested in that second investment decision. Let's call either of these percentages $y\%$. Your chance-to-win will then be given by the equation:

$$Chance - to - win : \left(1 - \left(\frac{y}{100} - \frac{x}{100}\right)^2\right) \times 100$$

This means that you have the highest probability of earning the €2 if your guess x is what you believe is your partner's belief (guess #3) or what you believe he or she will invest (or invested) (guess #4 [*Unobservable:* and guess #1]).