Crowd-funding conservation (and other public goods)*

Erik Ansink[a], Mark Koetse[b], Jetske Bouma[c], Dominic Hauck[b], Daan van Soest[d]

August 17, 2017

Abstract

We assess the impact of crowd-funding design on the success of crowd-funded public goods using a lab-in-the-field experiment. Our design treatments aim to increase the efficiency of crowd-funding campaigns by decreasing possible coordination problems that may occur when potential donors are faced with a multitude of projects to contribute to. Amongst others, we explore the potential of seed contributions and the impact of the attraction effect. We implement our crowd-funding experiment using a web-based user interface with multiple threshold public goods over the course of four days, similar in style to conventional crowd-funding websites. Our results show that such alternative crowd-funding designs do not affect the total amount of contributions, but they do affect coordination. These restuls are confirmed in a follow-up experiment with actual nature conservation projects.

JEL Classification: C91, C92, H41, L31

Keywords: Crowd-funding, field experiment, threshold public goods, charitable giving, nature conservation

*We thank Jos van Ommeren, seminar participants at Vrije Universiteit Amsterdam, EAERE 2017, and the 2017 Workshop on Experimental Economics for the Environment in Bremen for comments, suggestions, and discussion. We thank Natuurmonumenten and Kantar TNS for excellent support and advice. We are grateful to the Netherlands Environmental Assessment Agency for financial support. Ansink, Bouma, and Hauck gratefully acknowledge financial support from the Netherlands Organisation for Scientific Research (NWO) Grant No. 841.12.002. Van Soest gratefully acknowledges receiving funding from the European Commission's Seventh Framework Program (FP7/2007-2013) Grant No. 613420.

[a]Department of Spatial Economics, Vrije Universiteit Amsterdam.

[b]Institute for Environmental Studies, Vrije Universiteit Amsterdam.

[c]Netherlands Environmental Assessment Agency.

[d]Department of Economics, Tilburg University.

1 Introduction

With the introduction of crowd-funding websites as platforms for crowd-based financing, they are increasingly being used as a mechanism for the provision of public goods. Crowd-funding is the practice of a fund-raiser raising capital from many people (or, donors) to fund a project through an on-line platform. Crowd-funding is new, with well-known platforms such as FundRazr, GoFundMe, Kickstarter, and IndieGoGo launched only in the late 2000s. The crowd-funding of public goods, sometimes denoted by the more encompassing term 'civic' crowd-funding (Davies, 2015; Stiver et al., 2015), has found its niche on these platforms but its popularity also initiated specialized crowd-funding websites such as Citizinvestor, IOBY, JustGiving, and GlobalGiving. A large share of civic crowd-funding concerns projects related to nature conservation and environmental protection.

Using crowd-funding as a mechanism to provide public goods may mitigate several obstacles that are normally associated with the private provision of public goods, such as the transaction costs involved in designing a contribution scheme for the public good (Cornes and Sandler, 1996). For instance, compared to alternative mechanisms, crowd-funding allows cheap matching between projects and potential donors through low search costs, low risk exposure, low demand uncertainty, and transparent monitoring of progress (Agrawal et al., 2014; Strausz, 2017). Yet, other obstacles remain, including cheap-riding and coordination problems. Cheap-riding occurs whenever agents do contribute to the public good but try to reach an equilibrium where their own relative contribution is low (Isaac et al., 1989). Coordination failure occurs when the multiplicity of projects offered on a crowd-funding platform causes an *'inefficient distribution of donations across projects'* (Corazzini et al., 2015), possibly leading to project failure and discouraging potential donors.

In order to mitigate cheap-riding, most platforms offer an all-or-nothing design. This design implements a threshold public good with full refund – known to boost donations in public goods games (Cadsby and Maynes, 1999). Yet, this design does not touch upon cheap-riding and coordination problems, which potentially have more impact on crowd-funding success. This is why most platforms also implement a pre-specified deadline and real-time information on contribution progress. Nevertheless, success rates of around 50% (for Kickstarter, see e.g. Mollick, 2014) indicate that there may be scope for improvement. Both cheap-riding and coordination problems could be further mitigated by signaling focal projects, which have been demonstrated to facilitate success in coordination games (Schelling, 1960; Mehta et al., 1994). Such focal projects can be provided by either the crowd-funding platform or the fund-raiser itself. Previous empirical work has found that fund-raisers use e.g. short deadlines and high thresholds as signals in order to gain donors' attraction (Devaraj and Patel, 2016), although many more signals are possible and actually being used (Mollick, 2014; Belleflamme et al., 2015).[1] Crowd-funding platforms may also signal focal projects, for instance by promoting such projects more prominently on their website. Identification of the impact of both types of signals (i.e. by fund-raisers and by crowd-funding platforms) is muddled by endogeneity problems. For instance, the platforms' choice of featured projects is likely to be related to expected success.

In this paper we take an alternative approach to analyze the impact of crowd-funding design, and specifically the use of focal projects, on project success. We do so in a controlled setting, which prevents any endogeneity problems. This setting is a lab-in-the-field experiment, allowing us to assess the impact of crowd-funding design on the success of crowd-funded public goods. We investigate two mechanisms. The first is the provision of seed money (also known as 'challenge grants' List and Lucking-Reiley, 2002; Rondeau and List, 2008), The second mechanism is adding a seemingly irrelevant project that serves as a decoy in order to exploit the so-called attraction effect (cf. Ariely and Wallsten, 1995). From the crowd-funding platform and fund-raisers' perspective, our contribution is that we analyse whether, and if so, how to increase the number of successful projects through the mitigation of cheap-riding and coordination failure.[2]

Our lab-in-the-field experiment extends upon earlier crowd-funding experiments in the lab (Wash and Solomon, 2014; Corazzini et al., 2015; Solomon et al., 2015) by providing a more natural field context. That is, the subjects in our experiment were sampled from a representative pool and could log in to the internet-based user interface of the game (further described in Section 2) at any moment over a four-day period – from any location – to make contribution decisions for one or more projects. This context contrasts starkly with the mentioned lab studies that lasted for approximately one hour (for multiple rounds of play) and were played in computer labs with a standard subject pool. A follow-up experiment with actual nature conservation projects, described in Section 5, brings us even closer to the field.

Some of our results, reported in Section 4, confirm earlier findings, such as coordination difficulties in general. Our analysis of the impact of seeding and seemingly irrelevant

[1]Other aspects that affect crowd-funding success but are not under control of the fund-raiser include the number of potential donors, social norms, the characteristics of early funders, and herding (cf. Bøg et al., 2012; Crosetto and Regner, 2014; Belleflamme et al., 2015). Note that there exist modest differences in donor behavior between for-profit and non-profit crowd-funding campaigns (Pitschner and Pitschner-Finn, 2014).

[2]Our methodology does not allow us, however, to assess the impact of our crowd-funding designs on substitution to and away from contributions to other projects, charities, and investments on other crowd-funding platforms or using alternative fund-raising mechanisms.

alternatives yields new findings. Our results reveal that such mechanisms do not affect the total amount of contributions, but they do affect coordination. Consistent with our expectations, seeded projects serve as focal projects, receiving relatively more contributions and displaying a higher success rate. This effect, however, only occurs when the number of available projects is high (6) and not when it is low (2). Perhaps counter-intuitively, adding a seemingly irrelevant project does not lead to an attraction effect, but rather a repulsion effect, with the decoy project diverting contributions from its target project rather than increasing them.

In general, our results demonstrate that signals do work in the context of crowd-funding and that they can be exploited in order to mitigate coordination failure and cheap-riding. These results confirm the coordination results found by Corazzini et al. (2015) in a lab setting and provide further support for crowd-funding as a promising mechanism for the private provision of public goods.[3] As such, our results are of interest to platforms, funders, and project owners of civic crowd-funding projects, many of which relate to nature conservation and environmental protection

2 Experimental design

The game that we implement is a combination of on-line real-time linear threshold public goods games with no rebate and full refund. The linear version of the public goods game is also known as the voluntary contribution mechanism (VCM) and has been widely used to study efficiency and behavioral aspects of various mechanisms and design options in the provisioning of public goods. To mimic the crowd-funding setting we choose six specific game features. First, subjects do not come to a lab but play the game on-line using a web-based user interface, similar in style to conventional crowd-funding websites. Second, we play the game in real-time over the course of four days, allowing subjects to incrementally increase their contributions to the public good, up to the end of the game (Dorsey, 1992; Kurzban et al., 2001; Goren et al., 2003). The time-span of four days was selected as a compromise between the length of conventional crowd-funding campaigns on the one side and keeping participants' attention and limiting their time investment on the other side. Third, we present multiple public goods at once, offering subjects the choice to invest in more than one good, if any (Corazzini et al., 2015). Fourth, we impose an exogenous threshold value that serves as the minimum level of aggregate contributions required to implement the public good which provides a discrete payoff bonus to each subject in

[3] Note that civic crowd-funding could either substitute or complement government-provided public goods (Davies, 2015).

the group (van de Kragt et al., 1983; Palfrey and Rosenthal, 1984; Bagnoli and Lipman, 1989). Fifth, there is no rebate of excess contributions (Marks and Croson, 1998; Spencer et al., 2009). Sixth, in line with common practice at major crowd-funding platforms we implement full refund, i.e. a money-back guarantee in case the threshold is not reached (Isaac et al., 1989; Wash and Solomon, 2014).

Our analysis is based on a between-subject design across five treatments. In our benchmark treatment, subjects have the option to contribute to two threshold public goods, presented simultaneously in a user interface further described below. Subsequent treatments differ from the benchmark treatment in terms of the number of goods and their characteristics. All treatments were played by 90 subjects, distributed over 15 groups of 6 subjects each. In total, across treatments, 450 subjects participated in the experiment.

2.1 Benchmark treatment

In the benchmark treatment, that we refer to as BEN, subjects were randomly assigned to groups of six. Upon logging in to the user interface, each subject received a one-time endowment of 34 tokens in his private account. Subjects could use this endowment over the course of the game, from Thursday 08:00 to Sunday 22:00, to contribute to one or both of two threshold public goods by investing (but not divesting), part or all of their endowment, either at once, or in smaller steps. The public goods are collective accounts, where each subject within a group can contribute to. When the threshold for a public good is reached at the game deadline (i.e. Sunday 22:00), contributions to that good are forfeited and a bonus is paid to each subject in the group regardless of their individual contributions to the public good. When the threshold for a public good is not reached at the game deadline, any contributions to that good are returned to the contributors' private accounts. Subjects' payoffs are the sum of what is left in their private accounts plus any bonus received from one or both goods.

The two goods differ in threshold and bonus and we refer to them as Good 72_{18} and Good 84_{24}. Good 72_{18} has a threshold value of 72 and a bonus of 18 while Good 84_{24} has a threshold value of 84 and a bonus of 24. Hence, one good is easier to implement since it has a lower threshold while the other is more efficient. Since $4 \times 34 < 72 + 84 < 5 \times 34$, at least 5 subjects need to contribute on average $(72+84)/5 = 31.2$ tokens in order to reach both goods' thresholds. With 6 subjects, the required average contribution drops to $(72+84)/6 = 26$ tokens.

2.2 Other treatments

In order to test the impact of various crowd-funding design choices, we implement four variations on the benchmark treatment, that differ from BEN in terms of the number of goods and their characteristics. Motivation for treatment characteristics and a discussion of expected behavior under these treatments is deferred to Section 3, while an overview of treatments is provided in Table 1.

In two treatments, SEED72 and SIA72, we focus attention on the least successful good in BEN. In Section 4 we show that 72_{18} is less successful, both in terms total contributions and in terms of thresholds' reached, than 84_{24}. We focus attention on this good in two distinct ways. In SEED72 we add 20 seed tokens to 72_{18} at the start of the game and raise its threshold by an equal amount of tokens so that its characteristics are 92_{18}. Note that by the 20 seed tokens, 92_{18} under SEED72 is formally equivalent to 72_{18} under BEN, although subjects may perceive this differently. In SIA72 we add a third 'seemingly irrelevant alternative' good to the two BEN goods, serving as a decoy to 72_{18}. Its characteristics are 72_{17} so that this good is dominated in terms of its bonus by 72_{18}.

In the next two treatments, SIX and SIXSEED78, we increase the number of available goods from two to six. Coordination failures are more likely to arise when more projects are available. In SIX, we add four goods that are dominated by the two BEN goods in terms of both their thresholds and bonuses. These four goods' characteristics are, respectively, 74_{16}, 78_{17}, 86_{22}, and 90_{23}, so that each good under SIX has a distinct threshold and bonus. In SIXSEED78, similar to SEED72, we add 20 seed tokens to the least successful good under SIX and raise its threshold by an equal amount of tokens. In Section 4 we show that this least successful good is 78_{17}. Under SIXSEED78 this good's characteristics become 98_{17}, which is formally equivalent to 78_{17} under SIX. A variation of treatment SIX with an additional seemingly irrelevant project would not make much sense, given that four of the six projects in this treatment are already dominated and relatively similar to one of the two BEN goods.

2.3 Sampling and procedures

Subjects were sampled from a survey panel hosted by Kantar TNS, a Dutch survey consultancy. All communication with subjects was done via email by the consultancy. Subjects were invited to participate about one month in advance, asking them for their availability to participate in an on-line game and log in at least once a day over the course of four specified days. From the sample, sub-samples for each treatment were formed that were

Table 1: Overview of projects in all treatments.

Treatment	**Project descriptions***
BEN	72_{18}, 84_{24}
SEED72	$\underline{92_{18}}$, 84_{24}
SIA72	72_{18}, 84_{24}, 72_{17}
SIX	72_{18}, 84_{24}, 74_{16}, 78_{17}, 86_{22}, 90_{23}
SIXSEED78	72_{18}, 84_{24}, 74_{16}, $\underline{98_{17}}$, 86_{22}, 90_{23}

* Underlined goods receive an anonymous initial contribution of 20 tokens.

representative for the Dutch population in terms of gender, age, and education. Three days prior to the experiment, subjects were reminded of the experiment and given the opportunity to cancel their participation. One day before the start of the experiment, an email was sent to all participating subjects with game instructions and a personal login URL. These game instructions were accessible from the user interface throughout the experiment. The experiment took place in three sessions, each lasting from Thursday 08:00 to Sunday 22:00 in the period of October-November 2015. Each day of the experiment before noon, a reminder including their personal login URL, was sent to all subjects, irrespective of whether they had already made logged in on that day.

Any subjects that canceled their participation prior to the start of the experiment were replaced by subjects from a standby sample. In order to assure full groups of active contributors we also allowed replacement on the first day of the game. Any subject that did not login on the first day of the experiment before 22:00 was replaced by a subject from the same standby sample. Other subjects in the group were not informed of such replacements, nor could they derive this information through the user interface.

Instructions (see Appendix) included the game objective, the rules of play, calculation of payoffs, an extensive example, and a FAQ with questions based on the evaluation of a pilot study of the benchmark treatment in October 2015. Instructions were kept as brief as possible since they had to be read on-screen. The specific wording was based on multiple rounds of testing and evaluation during the pilot study. Based on this evaluation, we framed the public goods as 'projects' and contributions as 'investments', without expecting significant impact of this framing on game behavior (Alekseev et al., 2017); we will use the terms interchangeably throughout the paper. Contact information was provided in case anything was unclear (20 subjects used this possibility).

Upon logging in via their personal device (e.g. PC, laptop, tablet, smartphone), subjects entered the user interface of the game, a web-based platform. The user interface presented the subjects under the benchmark treatment with two projects, displayed side-by-side (see

Appendix). The projects were identified by a number, described in neutral terms, and differed only in threshold and bonus. When more than two projects were presented, they were displayed with a maximum of two projects in each row. General information on the game is displayed at the top of the screen and includes (i) the number of remaining tokens in the subject's private account, (ii) the number of subjects in the group, and (iii) the game deadline: Sunday 22:00. In addition, specific feedback is presented for each project, including (i) total contributions by all subjects in the group, (ii) the remaining contribution gap to the threshold, (iii) the number of subjects in the group who have contributed, and (iv) the subject's own total contribution. This feedback was updated continuously so that subjects were able to make fully informed decisions. Subjects were logged off automatically after 20 minutes of inactivity.

Within three days after the experiment, earned payments were made via a bank transfer by the research consultancy, conditional on having logged in at least once each day and filling out a short on-line survey in the week after the experiment (both conditions were announced in the invitation as well as in game instructions). This survey contained items on risk preference (based on Holt and Laury, 2002), inequality aversion (based on Carlsson et al., 2005), attitude (based on 3 items from the survey 'Geven in Nederland 2013'[4]), game behavior, and game evaluation. The research consultancy provided us with readily available information on a wide range of demographic variables, as well as information on membership of and donations to charities.

The value of each token was EUR 0.20. Starting with 34 tokens, subjects earned on average EUR 7.96 for a limited time investment over four days.[5] We estimate this time investment from login behavior as follows. The average number of logins was 5.2 and the average number of bets was 4.1. If a bet was made, the average time between login and bet was 200 seconds. Hence, the average minimal time investment, excluding logins during which no bet was made, was equal to 4.1×200 seconds, less than 14 minutes. More realistically, extrapolating the average time between login and bet to logins during which no bet was made, the average time investment was equal to 5.2×200 seconds, more than 17 minutes. Still, this estimate of time investment excludes time spent on reading the instructions upon first login and any time spent thinking about possible strategies when not logged in.

[4]See http://www.geveninnederland.nl/.

[5]Data in this paragraph is based on the cleaned sample as described in Section 4.

3 Theory and predictions

3.1 Theoretical framework

Consider a public goods game with $g \in \{1, 2, \ldots, G\}$ public goods and $j \in \{1, 2, \ldots, J\}$ players. The game starts at time $t = 0$ and ends at time $t = T$. Each player has an endowment e_j in his private account, which he can use over the course of the game, to contribute to one or more of the G public goods available to him. Denote a contribution by player j to good g at time t by c^t_{jg}. Total contributions by all players to good g up to time t are denoted by $C^t_g = \sum_j \sum_{\tau \leq t} c^\tau_{jg}$.

When a threshold contribution τ_g is reached at time T for good g, a bonus b_g is paid out to each player. No rebate implies that excess contributions $C^T_g - \tau_g$ are wasted. Full refund implies that any contributions to a good that does not reach the threshold, i.e. $C^T_g < \tau_g$, are returned to the contributors' private accounts. The following function denotes payoffs π_j:

$$\pi_j = e_j + \sum_g \left(\begin{cases} 0 & \text{when } C^T_g < \tau_g \\ b_g - c^T_{jg} & \text{when } C^T_g \geq \tau_g \end{cases} \right). \tag{1}$$

Payoffs are equal to the sum of endowment and good-specific payoffs. Good-specific payoffs are equal to the bonus minus own contribution to that good if the threshold is met, and zero otherwise, independent of own or others' contributions. This payoff specification is simpler than the provision-point setup where, in addition to the bonus, a marginal return is paid on total contributions C^T_g. With multiple public goods, removing this additional payoff increases the importance of coordination and aligns it with standard all-or-nothing crowd-funding campaigns. Not only is coordination important to obtain bonuses, coordination is also facilitated by the full refund feature, a common feature of major crowd-funding platforms (Wash and Solomon, 2014), which weakens the assurance problem to the cheap-rider problem (Isaac et al., 1989). Under full refund, the only cost of making a contribution to a specific good is that the remaining budget in the private account is reduced. There is no possibility that any contributed tokens are forfeited, unless the threshold is reached and the corresponding bonus is paid out. Hence, an early small contribution as a signal that players should coordinate on a particular good, is relatively cheap.

Subjects have all necessary information available in the user interface. At each moment in the game, each player is informed about T (and hence $T - t$), the number of players who have contributed to each good, their total contributions to each good C^t_g, and the player's own contribution to each good c^t_{jg} as well as his remaining tokens $e_j - \sum_g \left(c^t_{jg} \right)$, In addition, the game instructions are accessible from the user interface throughout the experiment.

In the one-shot version of a single linear threshold public goods game with no rebate and full refund, there exist two sets of (pure-strategy) Nash equilibria. The first set contains inefficient equilibria. Such equilibria satisfy the *No Deviation Constraint* (Croson and Marks, 2000). That is, the threshold is not met and cannot be met unless players contribute more than the bonus. Given full refund, any equilibrium with non-zero contributions in this set is payoff-equivalent to the zero-contributions equilibrium, which is always included in this set. For any player, contributing less would not affect payoffs while contributing more would either not affect payoffs (if such additional contribution does not fill the gap to the threshold), or would lead to a lower payoff (if such additional contribution would fill the gap, given that the gap is larger than the size of the bonus).

The third set contains efficient equilibria. In such equilibria, the threshold is exactly met with the sum of contributions being divided over the players with no player contributing more than the bonus from the good (cf. Bagnoli and Lipman, 1989; Croson and Marks, 2000). For any player, contributing more would be wasteful by no rebate, while contributing less involves not reaching the threshold and thereby foregoing the bonus.

Moving from one to multiple public goods in the one-shot setting, contributions to the goods can be treated as separate games so that additional equilibria occur, each of which is a combination of the above equilibria for single goods. The number of equilibria is only constrained by the players' endowments and this constraint will be more strict as the number of goods increases. Efficient equilibria correspond to those where, given endowment constraints, the sum of bonuses is maximized. In our experiment, parameters for group size, endowments, and thresholds are selected such that each group can maximally reach two thresholds (as explicitly stated in the instructions),[6] which allows us to narrow down efficient equilibria to those where the two goods with the highest bonus reach their threshold.

Finally, moving from one-shot to a real-time game, we know that multiple-period frameworks generally allow for any individually rational outcome to be supported in a subgame-perfect Nash equilibrium. Given no rebate and full refund, the set of individually rational outcomes corresponds – in terms of contributions and payoffs – to the static Nash equilibria described above.

[6] All projects are efficient with thresholds being less than five times the bonus, while the total endowments of a group suffice to fund two goods, but never three.

3.2 Hypotheses

Below we motivate our treatments and formulate hypotheses based on expected game behavior. A general observation is that, given our game specification and parameterization, equilibria with one or two thresholds reached payoff-dominate those with, respectively, zero or one. This dominance could lead to a relatively high share of successful projects, boosted by our game features of no rebate and full refund. Coordination problems and cheap-riding, however, are likely to compromise this outcome, especially in the treatments with six projects. In treatments SIA72 and SEED72 we expect that coordination problems are resolved by signaling focal projects. As a result, we expect cheap-riding to be the dominant obstacle to reach efficient outcomes in these treatments.

In the below hypotheses we will compare outcomes between different treatments based on two criteria, measured at the group level. The weaker criterion is the average level of contributions to specific projects. The stronger criterion is the average number of projects for which the threshold is reached. We will refer to both criteria together as treatments being more 'successful' (i.e. more thresholds reached and higher contributions) or less.

First, we present our hypothesis on overall game behavior, which combines three hypotheses as formulated by Bagnoli et al. (1992):

Hypothesis 1. (a) The groups will reach the threshold of exactly two goods; (b) There will not be any contributions in excess of the goods' thresholds; (c) The subjects' contributions will be individually rational, i.e. no subject contributes more to a good than the level of its bonus.

Since our game is an on-line field experiment over the course of four days, we have much less control than in the lab. Accordingly, we expect much more noise in our data compared to e.g. Bagnoli et al. (1992).

Next, we present our hypotheses on specific differences between treatments.

Hypothesis 2. Project success is lower under SIX than under BEN

Hypothesis 3. Seeded projects under SEED72 and SIXSEED78 are more successful than their unseeded counterparts in BEN and SIX.

Hypothesis 4. The dominating good (i.e. 72_{18}) is more successful under SIA72 than under BEN.

Regarding Hypothesis 2, the SIX treatment features an additional four goods, each of which is dominated by one of the two BEN goods. Increasing the number of available goods is expected to obstruct coordination (cf. Corazzini et al., 2015), even though two goods dominate the other four.

Regarding Hypothesis 3, the SEED72 treatment features good 92_{18}, that receives 20 seeded tokens (and a similar increase in its threshold compared to the BEN good 72_{18}). Seeding a project is expected to signal that contributions should be made to the seeded good (cf. List and Lucking-Reiley, 2002; Rondeau and List, 2008). Increased focus on good 72_{18}, combined with full refund, is expected to simplify coordination on the seeded good, compared to the BEN good 72_{18}). A similar expectation holds for the seeded good 98_{17} under SIXSEED78 compared to 78_{17} SIX.

Regarding Hypothesis 4, the SIA72 treatment features a 'seemingly irrelevant alternative' 72_{17} that is dominated by one of the other goods, 72_{18}, for which it serves as a decoy. Adding this decoy is expected to lead to an 'attraction effect', a signal that contributions should be made to the dominating good. This effect is described and tested in a large body of research in psychological and marketing science (cf. Ariely and Wallsten, 1995). Recently, (Frederick et al., 2014) found that in less abstract choice situations, the attraction effect may disappear or even lead to a 'repulsion effect'. Since the choice situation in our experiment is featured by a high level of abstraction, we hypothesize that the attraction effect will dominate.

4 Results

We start the description of our results by an analysis of overall game behavior. We provide statistics related to Hypothesis 1 and we present additional contribution characteristics that are important in order to correctly interpret the game's results at the group level. Next, we assess Hypotheses 2–4 by comparing treatment outcomes in terms of their success.

Despite our efforts to assure full groups of active contributors (as described in Section 2.3), only 406/450 subjects participated actively with at least one login. Since at least 5 subjects are required to reach both goods' thresholds, for the group-level analysis, we drop any group that has less than 5 active subjects. This results in 8 dropped groups: 2 for SEED72, 3 for SIA, and 3 for SIX. Group-level results and test statistics are provided for the remaining independent observations at the group level with $5 \times 15 - 8 = 67$ groups, between 12 and 15 per treatment. For the individual-level analysis, consistent with the conditions for payment, we drop any subjects that logged in less than 4 times. This results in 116 dropped subjects, leaving $406 - 116 = 290$ subjects for our analysis of individual game behavior, in which we control for group size and cluster standard errors at the group level.

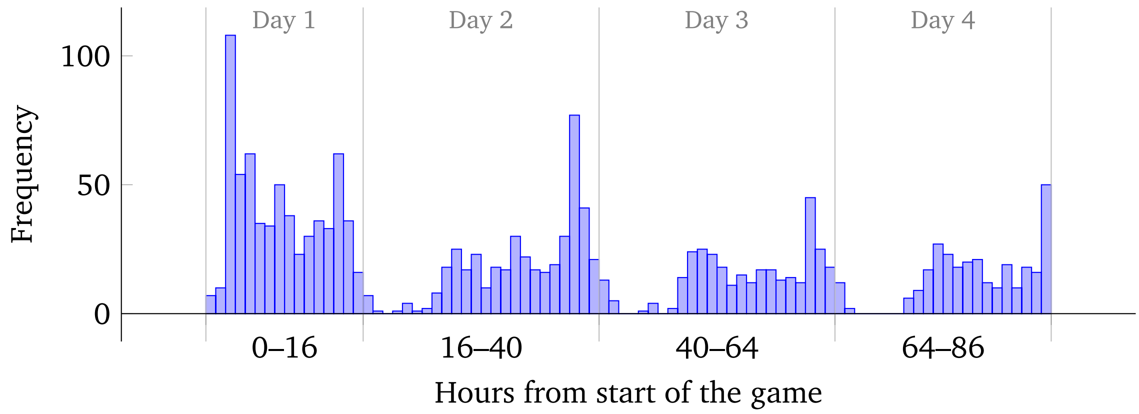

Figure 1: Investment frequency across all treatments in one-hour bins. Hour 0 corresponds to THU 08:00 and hour 86 corresponds to SUN 22:00.

4.1 Overall game behavior

In this section we provide an overview of overall game behavior which allows us to assess Hypothesis 1. Before doing so, we first illustrate the timing and levels of subjects' contributions in Figures 1 and 2. The figures reveal several tendencies that make us confident that subjects took the game seriously, despite the low stakes and despite the length of the game.

Figure 1 shows the frequency of investments made across all treatments in one-hour bins. The figure shows spikes at the start of the game, and each night between 21:00–22:00. The figure also show a modest decrease in investment frequency over the course of the four days, which is partly explained by subjects running out of tokens: at the end of the game, 63% of subjects had invested all of their tokens. Given this constraint, there is no strong decline in interest in playing the game as the game progresses, which we interpret as subjects taking the game seriously.

Figure 2 shows the frequency of (non-aggregated) investment levels across all treatments. The figure shows spikes at focal numbers 5 and 10. The small spike at 34 indicates subjects that went all-in in one project, an apparently irrational investment decision as discussed below. The dominance of small contributions, less than 5 tokens, illustrates that subjects were either signaling their interest in a specific project, trying to cheap-ride, or both. We interpret this as another signal that the game was taken seriously by most subjects.

We now turn to the assessment of Hypothesis 1, for which we find only weak support, in line with our expectation that our data would be relatively noisy due to the lab-in-the-field setting. We start with Part (a) of Hypothesis 1 which states that groups will reach the

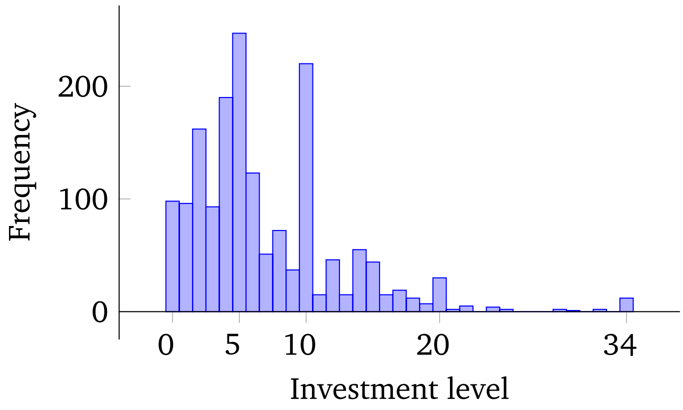

Figure 2: (Non-aggregated) investment level frequency across all treatments.

threshold of exactly two goods. We find very limited support for this part. Only 20 of 67 groups, or 30%, manage to coordinate on this efficient equilibrium. Another 37/67 (55%), reaches the threshold of exactly one good while 10/67 (15%) reaches no threshold. Recall that outcomes with 0 or 1 successful project may constitute (inefficient) equilibria, as long as the *No Deviation Constraint* is satisfied. The difference in success rate with e.g. Bagnoli et al. (1992) who find a success rate of 48% may be due to the field-setting or the lack of repetition (and hence scope for learning) in our experimental design.

Part (b) of Hypothesis 1 states that there will not be any contributions in excess of the goods' thresholds. We find moderate support for this part, illustrated by Figure 3. This figure shows total investments per project, normalized with the threshold set to zero. Hence, the horizontal axis displays the number of over-invested tokens per project.

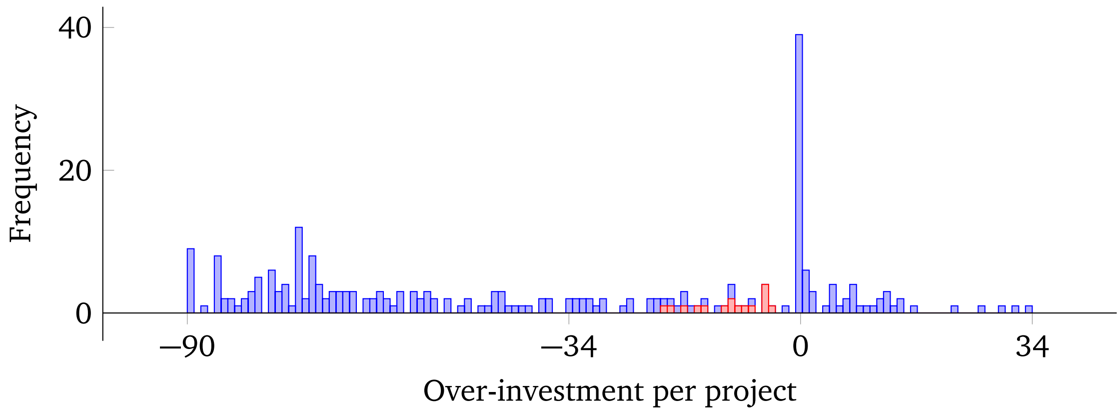

Figure 3: Project over-investment frequency in tokens across all treatments; red bars mark the 16 projects that violate the *No Deviation Constraint*.

The equilibrium behavior that was discussed in Section 3 is visible in Figure 3. The 39 projects that make up the spike at 0 were successful without wasting any tokens, i.e.

efficient equilibrium behavior. There are 177 projects to the left of 0. These projects did not reach the threshold. Of these projects, 161 satisfy the *No Deviation Constraint*, i.e. inefficient equilibrium behavior. The 16 remaining projects, illustrated by red bars in Figure 3, violate the *No Deviation Constraint*. These are projects where a single subject could have made the project successful (considering his remaining tokens, the remaining gap to the threshold, and the bonus upon success). Finally, there are 38 projects with a strictly positive over-investment, i.e. to the right of 0. These are projects where one or more subjects displayed seemingly irrational behavior by contributing more than what was required to reach the threshold.

There are several possible explanations for over-investments. One is that we cannot rule out coordination failures in exactly reaching the threshold, for instance when two contributions are made shortly after one another such that feedback on others' contributions was not yet visible on the user interface, or possibly overlooked. Other alternative explanations include the possibility that subjects derive utility other than the monetary bonus from making a good successful, calculation errors, and better-safe-than-sorry contributions (e.g. 'should we reach the threshold or surpass it?', or 'perhaps I do not receive the bonus if I did not contribute any token?'). All in all, 200 of 254 projects, or 79%, have outcomes that are predicted by equilibrium behavior. Related to Part (b) of Hypothesis 1, we find that 38 of 254 projects, or 15%, feature contributions in excess of the projects' thresholds.

Part (c) of Hypothesis 1 states that no subject contributes more to a good than the level of its bonus. We find weak support for this part. Only 68% of subjects displays individually rational behavior in terms of contributions. However, as with Part (b) of Hypothesis 1, it is well possible that part of the apparent irrational behavior by the other 32% is driven by alternative explanations. Such explanations receive suggestive support from the contribution behavior of the 94 subjects that contributed more to a good than the level of its bonus. Of these, 60% over-contributed just 4 tokens or less while 14% over-contributed exactly 10 tokens. We interpret these over-contributions as possibly arising from calculation errors or better-safe-than-sorry contributions.

The results on individually rational contribution behavior are partly illustrated by Figure 4, which shows payoff in tokens per participant. Payoffs lower than 34, subjects' initial endowments, indicate that an apparently irrational investment was made. Small spikes at 18 and 24 include subjects that went all-in in one successful BEN good, and received the bonus of, respectively, 18 or 24 tokens. The spike at 34 includes subjects that did not invest or whose investments were not successful and got refunded so they ended up with their initial endowment of 34 tokens. The spike at 42 includes subjects that went all-in in two successful BEN goods, and received two bonuses that sum up to $18 + 24 = 42$ tokens.

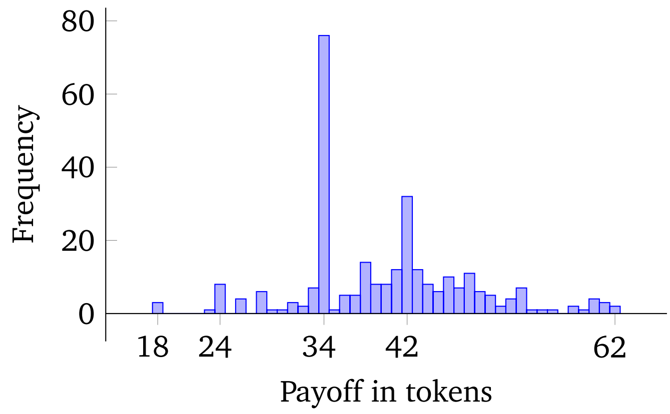

Figure 4: Payoff frequency in tokens across all treatments.

All payoffs lower than 34 tokens indicate that subjects were making apparently irrational investments. All payoffs higher than, approximately, $34-(72+84)/6+(18+24)=50$ (and 45 for groups of 5) include subjects that were either lucky with other group members contributing relatively more or they were successfully cheap-riding.

Combining the results for the three parts of Hypothesis 1, we find weak support only. We do find evidence, however, that subjects took the game seriously, despite the low stakes and despite the length of the game. This allows us to continue with the assessment of treatment results.

4.2 Treatment results

In this section we compare treatment outcomes in terms of their success in order to assess Hypotheses 2–4. We first provide graphical evidence of treatment impacts in Figure 5. The top chart of Figure 5 shows the average number of successful projects per group across all treatments. At first sight there appear to be some stark differences between treatments, not so much in the total number of successful projects, but rather in the distribution of success over the different projects. Under BEN, subjects appear to have a slight preference for the more efficient 84_{24} compared to 72_{18}. This preference is stronger under SIA72 and SIXSEED78, is weakened under SIX, and reversed under SEED72. Under SIA72, 72_{17} draws successful projects from 72_{18}. SIX and SIXSEED78 display quite a bit of successful projects other than the two BEN projects.

The bottom chart of Figure 5 shows the related average total contributions per group across treatments. Similar to the total number of successful projects, total contributions (i.e. the aggregate of contributions to all projects) do not appear to differ much across treatments. Most of the variation that occurred in the top chart is also visible in the bottom

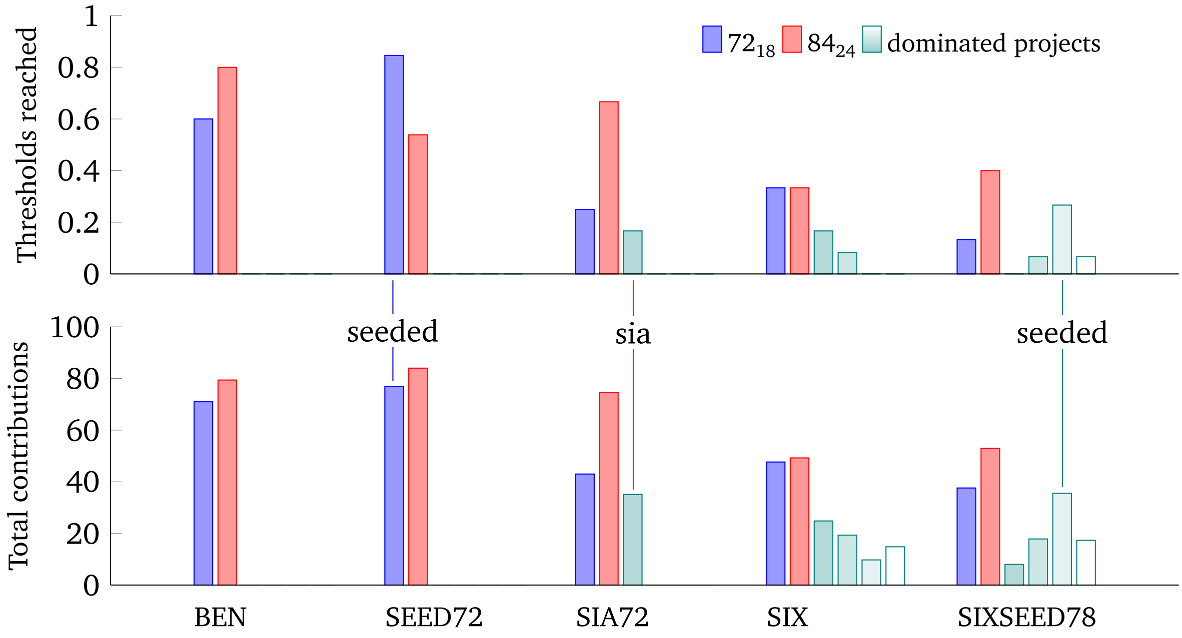

Figure 5: Thresholds reached per group (top chart) and total investments in projects per group (bottom chart) across treatments, see Table 1.

chart. Nevertheless, there are subtle differences between the two charts, driven by whether or not changes in total contributions materialize into changes in the number of thresholds reached. One notable difference is the relatively high number of successful seeded projects for SEED72 in the top chart relative to its received contributions in the bottom chart. Another difference are the non-zero contributions in the bottom chart for some dominated projects under SIX and SIXSEED78 that do not materialize into successful projects in the top chart.

We proceed with a statistical analysis of treatment differences.

We start with Hypothesis 2, stating that project success is lower under SIX than under BEN. We assess this hypothesis using the regressions presented in Table 2. In this table we provide an analysis of treatment effects on total contributions and thresholds reached. The dependent variable in linear regression Models (1) and (3) is total contributions, measured as the sum of contributions to all available projects, at the group level. The dependent variable in ordered probit Models (2) and (4) is thresholds reached, measured as the sum of thresholds reached of all available projects, at the group level. Recall from Section 3 that our parameter selection is such that each group can maximally reach two thresholds. Hence, the value of the dependent variable in these models is any integer between 0 and 2. Both dependent variables are regressed on treatment dummies and the number of subjects per group (recall that groups consist of 5 or 6 subjects for all treatments). Models (1) and (2) consider all available projects while Models (3) and (4) focus on the two BEN projects only.

For Models (2)–(4), sign and significance of the coefficients for the SIX dummy indicate strong support for Hypothesis 2, both in terms of total contributions and in terms of thresholds reached. Model (1) – regressing total contributions to all available projects at the group level – is the exception, displaying a positive sign. Despite these higher contributions, however, Model (2) shows that thresholds reached are lower, which we interpret as a sign of coordination problems. In addition, the number of subjects per group is positively related to project success. We find that having a sixth group member makes a substantial difference in total contributions per group, which translates into a 20% higher probability of reaching an additional threshold.[7]

The parametric results of Table 2 are confirmed by Mann-Whitney rank-sum tests. Based on all available projects, total contributions are higher under SIX compared to BEN ($z = -1.59$, $p3 = 0.11$), while the number of thresholds reached is lower under SIX compared to BEN ($z = 1.64$, $p = 0.10$). Focusing on the two BEN projects, total contributions are lower under SIX compared to BEN ($z = 4.05$, $p = 0.00$), while the number of thresholds reached is also lower under SIX compared to BEN ($z = 2.83$, $p = 0.00$).

Going beyond Hypotesis 2, Table 2 shows that identical results occur for the other two treatments with more than 2 projects, i.e. SIA72 and SIXSEED78. This result confirms the severity of coordination problems that occur when multiple projects are available (Corazzini et al., 2015, cf.).

Next, we assess Hypothesis 3, stating that seeded projects under SEED72 and SIXSEED78 are more successful than their unseeded counterparts under, respectively, BEN and SIX. We assess this hypothesis using the regressions presented in Table 3. In this table we continue on Table 2 by analysing treatment effects on total contributions and thresholds reached. The difference is that we now focus on the target project of each treatment. Recall from Table 1 that the target project for SEED72 and SIA is project 72_{18} (or 92_{18} when seeded) and the target project for SIXSEED78 is project 78_{17} (or 98_{17} when seeded). Total contributions and thresholds reached are regressed on the relevant treatment dummies and the number of subjects per group. We exclude a model that regresses tresholds reached for the target project under SIXSEED78 since this treatment is a perfect predictor of thresholds reached, as illustrated by the top panel of Figure 5.

In addition to this perfect prediction, sign and significance of the coefficient for the SIXSEED78 dummy in Model (3) indicates strong support for Hypothesis 3. The coefficients

[7]This last effect cannot be observed directly from Table 2; it is computed using average marginal effects for Models (2) and (4).

Table 2: Project success as measured by total contributions and thresholds reached.

	All projects		BEN projects only	
	(1) Contributions	(2) Thresholds	(3) Contributions	(4) Thresholds
main				
SEED72	3.180	−0.318	6.733	−0.247
	(7.694)	(0.466)	(8.696)	(0.474)
SIA72	0.643	−0.745	−33.69***	−1.189**
	(7.772)	(0.469)	(8.785)	(0.492)
SIX	13.73*	−1.085**	−54.27***	−1.791***
	(7.772)	(0.473)	(8.785)	(0.530)
SIXSEED78	20.92***	−0.917**	−58.82***	−1.995***
	(7.332)	(0.445)	(8.287)	(0.513)
# Subjects/group	30.80***	1.113***	15.74***	0.929***
	(5.077)	(0.321)	(5.738)	(0.326)
Constant	−20.03		63.30*	
	(28.57)		(32.29)	
cut1				
Constant		4.336**		3.045*
		(1.723)		(1.757)
cut2				
Constant		6.280***		5.354***
		(1.813)		(1.847)
Observations	67	67	67	67

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Coefficient estimates from linear regression models in (1) and (3) and ordered probit models in (2) and (4) (and standard errors in parentheses). All models are at the group level. Contributions in (1) and (3) equal the sum of contributions to all available projects (1) or BEN projects only (3). Similarly, thresholds in (2) and (4) equal the sum of thresholds reached. Treatment names represent dummy variables, with BEN as the baseline treatment.

for the SEED72 dummy in Models (1) and (2), however, indicate no effect of seeding under this treatment. These results are robust to alternative specifications of the dependent variable. Tables 6 and 7 in the Appendix confirm the results of Table 3 using difference and ratio of success of the target project relative to the other available projects as the dependent variable. The only difference with Table 3 is that measuring success in relative terms reveals a positive impact of seeding under SEED72, not in terms of contributions, but only in terms of thresholds reached, an effect that is also observable in Figure 5. The

Table 3: Treatment impacts on project success of the target project.

	Project 72_{18}		Project 78_{17}
	(1) Contributions	(2) Thresholds	(3) Contributions
main			
SEED72	5.937	0.681	
	(6.751)	(0.547)	
SIA72	-27.98^{***}	-0.957^{*}	
	(6.762)	(0.518)	
SIXSEED78			26.38^{**}
			(10.57)
# Subjects/group	−0.386	0.316	5.115
	(5.828)	(0.456)	(10.51)
Constant	73.14^{**}	−1.491	−18.81
	(32.56)	(2.537)	(59.20)
Observations	40	40	27

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Coefficient estimates from linear regression models in (1) and (3) and probit model in (2) (and standard errors in parentheses). All models are at the group level. The target project in (1) and (2) is project 72_{18} (or 92_{18} when seeded). The target project in (3) is project 78_{17} (or 98_{17} when seeded). Contributions in (1) and (3) equal the sum of contributions to the target project. Thresholds in (2) equals the sum of thresholds reached of the target project. Treatment names represent dummy variables, with BEN as the baseline treatment for (1) and (2) and SIX as the baseline treatment for (3).

results of Table 3 are also confirmed by non-parametric Mann-Whitney rank-sum tests. For the seeded good under SEED72, we find that total contributions to 92_{18} under SEED72 do not differ from those to 72_{18} under BEN ($z = -1.33$, $p = 0.18$) and the same holds for the number of thresholds reached ($z = -1.41$, $p = 0.16$). For the seeded good under SIXSEED78, we find that total contributions to 92_{18} under SEED72 are higher than those to 78_{17} under SIX ($z = -2.75$, $p = 0.01$) and the same holds for the number of thresholds reached ($z = -1.90$, $p = 0.06$). Combined, we find moderate support for a seeding effect; weaker under SEED72 and stronger under SIXSEED78.

Seeding appears to work with six available projects but not with two. A straightforward explanation is that the multitude of available projects and their similarity in terms of characteristics (see Table 1), may trigger boundedly rational behavior. If so, subjects may pay less or no attention to specific feautures of the problem as in e.g. Gabaix (2014). Under these conditions, choosing a seeded project is an easy default (Carroll et al., 2009). Such default behavior is not induced when the required attention to make the decision is low, such as when only two projects have to be compared. In that case, subjects will make a fully-informed comparison of potential payoffs. In line with results found by Corazzini et al. (2015), in doing so they ignore the possible coordination benefits offered by focusing on the seeded project. Subjects in our sample have a slight preference for the more efficient project 84_{24} (see the BEN results in Figure 5), which was the reason for seeding the less efficient project 72_{18} under SEED72. Since the seeded project is still the least efficient, subjects may ignore the seeded project. On top of this, recall that our selection of parameters for group size, endowments, and thresholds allows each group to maximally reach two thresholds. When faced with two available projects subjects that expect to reach two thresholds may therefore consider the seeding as irrelevant.

Finally, we assess Hypothesis 4, stating that the dominating good 72_{18} is more successful under SIA72 than under BEN. Again, we can use the regressions presented in Table 3. Sign and significance of the coefficient for the SIA72 dummy in models (1) and (2) of Table 3 indicate a negative effect of adding a seemingly irrelevant project on project success of project 72_{18}. Similar to the impact of seeding, this result is robust to alternative specifications of the dependent variable. Tables 6 and 7 in the Appendix confirm the results of Table 3. The only difference with Table 3 is that the negative impact of SIA on thresholds reached is not statistically significant when success is measured in relative terms. The results of Table 3 are confirmed by non-parametric Mann-Whitney rank-sum tests, which provide no support for our hypothesis either. The number of thresholds reached for 72_{18} under SIA72 is lower than under BEN ($z = 1.79$, $p = 0.07$), and the same holds for total contributions ($z = 2.85$, $p = 0.00$). All in all we find a strong effect of adding the seemingly irrelevant project, but the effect goes in the unexpected direction. Combined, we find no support for Hypothesis 4.

The unexpected effect of adding a seemingly irrelevant project reveals a repulsion- rather than an attraction effect. Importantly, the presence of repulsed subjects does not cause an increase in success of project 84_{24}, but rather causes some success for the irrelevant alternative itself. In addition to possible increased coordination problems with three rather than two projects, a reasonable explanation for the success of the irrelevant alternative is that the 1-token difference in bonus between the target project and its decoy is perhaps too small to rule out contributions to the irrelevant project, An explanation for the observed repulsion effect is not clear. Given the state of the literature on this effect, and especially the lack of experimental work on the repulsion effect, we can only speculate that there is some type of 'black sheep' effect at play, where adding the irrelevant alternative makes the competitor option more unique (Frederick et al., 2014).

Summarizing our results, we find that (a) adding projects strongly decreases overall project success, even when the added projects are dominated; (b) providing a project with seed money may strongly increase project success (but not unequivocally); and (c) adding a seemingly irrelevant project that serves as a decoy has a strong effect, but not in the expected direction: a repulsion effect occurs where an attraction effect was expected.

5 A conservation experiment

In order to verify our seeding and repulsion results in a setting closer to the field, as well as closer to possible applications, we conducted a follow-up experiment. In cooperation with a Dutch NGO, Natuurmonumenten, we ran 3 more sessions of our experiment that differed in one important aspect from our earlier experiment. This difference is that we replaced the abstract projects from our earlier experiment by actual nature conservation projects. Upon reaching a project threshold, the bonus of the project is not paid out in tokens to the group members but rather the tokens are transferred to money and paid out to the NGO. This change brings about two implications. One is that we move from a monetary bonus to an in-kind bonus, the size of which is described to the subjects as representing a 'nature development score'. The second is that we move from a local public good (i.e. only group members benefit) to a global public good (i.e. society as a whole benefits from any use and non-use values derived from the conserved nature area).

In order to keep differences in instructions and user interface as small as possible, the conservation projects were described in general terms, without revealing their name, location, or characteristics, while all parameter values for projects' threshold and bonus are kept constant between the two experiments.

Three treatments were repeated in this conservation experiment: BEN, SEED72, and SIA. The two treatments with six projects were not repeated since corresponding nature conservation projects could not be credibly included, given the fixed parameters for thresholds and nature conservation scores. Sampling and procedures were identical to the earlier experiment. The experiment took place in three sessions, each lasting from Thursday 08:00 to Sunday 22:00 in the period of May-June 2017. The conservation experiment had a lower share of active subjects; 224/270 subjects participated actively with at least one login. As a result, there were more groups with less than 5 subjects. These were dropped: 4 for BEN, 4 for SEED72, and 2 for SIA. Only 110 subjects are included in our analysis of individual game behavior.

Overall game behavior was roughly similar to the earlier experiment.[8] There are two

[8] See Figures 7–9 in the Appendix.

differences though. One is a higher frequency of smaller contributions, i.e. 5 tokens or less. The second difference is the lower share of successful projects. This last difference is illustrated in Figure 6 that provides graphical evidence of treatment impacts in the conservation experiment. Similar to Figure 5, the top chart of Figure 6 shows the average number of successful projects per group across all treatments, while the bottom chart shows the related average total contributions per group across treatments. Visual comparison to Figure 5 reveals similarities in total contributions and treatment impacts, but not in thresholds reached. These appear to be much lower in the conservation experiment despite the small differences in total contributions.

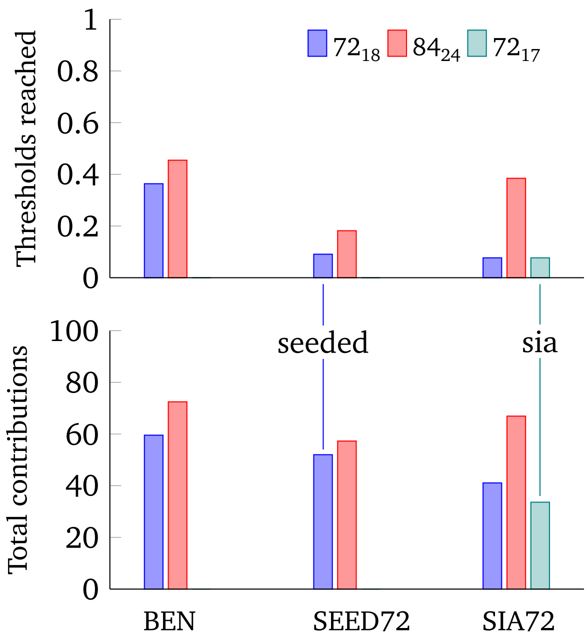

Figure 6: Conservation experiment: Thresholds reached per group (top chart) and total investments in projects per group (bottom chart) across treatments. Compare Figure 5.

This apparant difference in contributions and thresholds between treatments is confirmed by the regressions presented in Table 4. In this table, we repeat Models (1) and (2) of Table 2 in order to analyse experiment and treatment effects on total contributions and thresholds reached, pooling data from both experiments. Sign and significance of the coefficients for the conservation experiment dummy indicate a strong experiment effect which causes lower total contributions and thresholds reached in the conservation experiment. Comparing the two experiments, total contributions in the conservation experiment are slightly lower (on average -20%) while thresholds reached are much lower (on average -53%).

We proceed with a statistical analysis of treatment differences, repeating Models (1)

Table 4: Pooled experiments: Project success as measured by total contributions and thresholds reached. Compare Table 2.

	All projects	
	(1) Contributions	(2) Thresholds
main		
Conservation experiment	-17.00^{***}	-1.035^{***}
	(5.833)	(0.302)
SEED72	-8.398	-0.827^{**}
	(6.793)	(0.366)
SIA72	4.189	-0.749^{**}
	(6.718)	(0.354)
# Subjects/group	30.28^{***}	1.585^{***}
	(5.833)	(0.325)
Constant	-14.38	
	(32.83)	
cut1		
Constant		6.949^{***}
		(1.718)
cut2		
Constant		8.684^{***}
		(1.820)
Observations	75	75

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$
Coefficient estimates from linear regression model in (1) and ordered probit model in (2) (and standard errors in parentheses). All models are at the group level. Contributions in (1) equal the sum of contributions to all available projects.. Similarly, thresholds in (2) equal the sum of thresholds reached. Treatment names represent dummy variables, with BEN as the baseline treatment and the earlier experiment as the baseline experiment.

and (2) of Table 3 for the conservation experiment. In Table 5 we analyse treatment effects on total contributions and thresholds reached, focusing on project 72_{18} (or 92_{18} when seeded), the target project for SEED72 and SIA. We find qualitatively similar results of project success in the conservation experiment. The coefficients for the SEED72 dummy in Models (1) and (2) indicate no effect of seeding under this treatment, confirming our result for Hypothesis 3 under the earlier experiment. Sign and significance of the coefficient for the SIA72 dummy in models (1) and (2) indicate a negative effect of adding a seemingly irrelevant project on project success of project 72_{18}.

Table 5: Conservation experiment: Treatment impacts on project success of the target project. Compare Table 3.

	Project 72_{18}	
	(1) Contributions	(2) Thresholds
main		
SEED72	−7.545	−1.342
	(8.491)	(0.817)
SIA72	−20.30**	−1.633**
	(8.199)	(0.828)
# Subjects/group	16.39**	1.514**
	(7.301)	(0.698)
Constant	−26.89	−8.365**
	(38.96)	(3.716)
Observations	35	35

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Coefficient estimates from linear regression model in (1) and probit model in (2) (and standard errors in parentheses). All models are at the group level. The target project in (1) and (2) is project 72_{18} (or 92_{18} when seeded). Contributions in (1) equal the sum of contributions to the target project. Thresholds in (2) equals the sum of thresholds reached of the target project. Treatment names represent dummy variables, with BEN as the baseline treatment.

Similar to the earlier experiment, these results are robust to alternative specifications of the dependent variable. Tables 8 and 9 in the Appendix confirm the results of Table 3. The only difference with Table 5, also found in the earlier experiment, is that the negative impact of SIA on thresholds reached is not statistically significant when success is measured in relative terms. Finally, we also find qualitatively similar results if we pool the data from both experiments. Table 10 in the Appendix confirms the results of Tables 3 and 5.

All in all, the conservation experiment confirms the results of our earlier experiment, with the notable difference that total contributions are slightly lower and thresholds reached are significantly lower. We attribute this difference to the changed nature of the bonus under the conservation experiment. The bonus is now in-kind rather than monetary and global rather than local, decreasing individual incentives to reach the threshold.

6 Conclusion

We assess the impact of crowd-funding design on the success of crowd-funded public goods using a lab-in-the-field experiment. Specifically, we analyze whether mechanisms that signal focal projects can be used to increase project success. Summarizing our results on mechanisms in terms of our predicted results, we find strong support for Hypothesis 2, mixed support for Hypothesis 3, and no support for Hypothesis 4.[9] In words, we find that (a) adding projects strongly decreases project success of target projects, even when the added projects are dominated; (b) providing a project with seed money may strongly increase project success (but not unequivocally); and (c) adding a seemingly irrelevant project that serves as a decoy has a strong effect, but not in the expected direction: a repulsion effect occurs where an attraction effect was expected.

All in all, we find that investment behavior on crowd-funding websites can be directed in order to mitigate coordination problems. Our lab-in-the-field results partly confirm earlier lab findings, they point to the scope for design alternatives that improve coordination, and give further support for crowd-funding as a promising mechanism for the private provision of public goods.

[9]Hypothesis 1 does not analyze crowd-funding design specifically, and receives weak support.

References

Agrawal, J., C. Catalini, and A. Goldfarb (2014). Some simple economics of crowdfunding. In J. Lerner and S. Stern (Eds.), *Innovation Policy and the Economy*. Chicago: University of Chicago Press.

Alekseev, A., G. Charness, and U. Gneezy (2017). Experimental methods: When and why contextual instructions are important. *Journal of Economic Behavior and Organization 134*, 48–59.

Ariely, D. and T. S. Wallsten (1995). Seeking subjective dominance in multidimensional space: An explanation of the asymmetric dominance effect. *Organizational Behavior and Human Decision Processes 63*(3), 223–232.

Bagnoli, M., S. Ben-David, and M. McKee (1992). Voluntary provision of public goods: The multiple unit case. *Journal of Public Economics 47*(1), 85–106.

Bagnoli, M. and B. Lipman (1989). Provision of public goods: Fully implementing the core through private contributions. *Review of Economic Studies 56*(4), 583–601.

Belleflamme, P., N. Omrani, and M. Peitz (2015). The economics of crowdfunding platforms. *Information Economics and Policy 33*, 11–28.

Bøg, M., H. Harmgart, S. Huck, and A. Jeffers (2012). Fundraising on the internet. *Kyklos 65*(1), 18–30.

Cadsby, C. and E. Maynes (1999). Voluntary provision of threshold public goods with continuous contributions: Experimental evidence. *Journal of Public Economics 71*(1), 53–73.

Carlsson, F., D. Daruvala, and O. Johansson-Stenman (2005). Are people inequality-averse, or just risk-averse? *Economica 72*(287), 375–396.

Carroll, G., J. Choi, D. Laibson, B. Madrian, and A. Metrick (2009). Optimal defaults and active decisions. *Quarterly Journal of Economics 124*(4), 1639–1674.

Corazzini, L., C. Cotton, and P. Valbonesi (2015). Donor coordination in project funding: Evidence from a threshold public goods experiment. *Journal of Public Economics 128*, 16–29.

Cornes, R. and T. Sandler (1996). *The Theory of Externalities, Public Goods, and Club Goods, 2nd Edition*. Cambridge: Cambridge University Press.

Crosetto, P. and T. Regner (2014). Crowdfunding: Determinants of success and funding dynamics. Jena Economic Research Papers 2014-035, Friedrich-Schiller-University Jena.

Croson, R. and M. Marks (2000). Step returns in threshold public goods: A meta-and experimental analysis. *Experimental Economics 2*(3), 239–259.

Davies, R. (2015). Three provocations for civic crowdfunding. *Information, Communication and Society 18*(3), 342–355.

Devaraj, S. and P. Patel (2016). Influence of number of backers, goal amount, and project duration on meeting funding goals of crowdfunding projects. *Economics Bulletin 36*(2), 1242–1249.

Dorsey, R. (1992). The voluntary contributions mechanism with real time revisions. *Public Choice 73*(3), 261–282.

Frederick, S., L. Lee, and E. Baskin (2014). The limits of attraction. *Journal of Marketing Research 51*(4), 487–507.

Gabaix, X. (2014). A sparsity-based model of bounded rationality. *Quarterly Journal of Economics 129*(4), 1661–1710.

Goren, H., R. Kurzban, and A. Rapoport (2003). Social loafing vs. social enhancement: Public goods provisioning in real-time with irrevocable commitments. *Organizational Behavior and Human Decision Processes 90*(2), 277–290.

Holt, C. and S. Laury (2002). Risk aversion and incentive effects. *American Economic Review 92*(5), 1644–1655.

Isaac, R., D. Schmidtz, and J. Walker (1989). The assurance problem in a laboratory market. *Public Choice 62*(3), 217–236.

Kurzban, R., K. McCabe, V. Smith, and B. Wilson (2001). Incremental commitment and reciprocity in a real-time public goods game. *Personality and Social Psychology Bulletin 27*(12), 1662–1673.

List, J. and D. Lucking-Reiley (2002). The effects of seed money and refunds on charitable giving: Experimental evidence from a university capital campaign. *Journal of Political Economy 110*(1), 215–233.

Marks, M. and R. Croson (1998). Alternative rebate rules in the provision of a threshold public good: An experimental investigation. *Journal of Public Economics 67*(2), 195–220.

Mehta, J., C. Starmer, and R. Sugden (1994). The nature of salience: An experimental investigation of pure coordination games. *American Economic Review 84*(3), 658–673.

Mollick, E. (2014). The dynamics of crowdfunding: An exploratory study. *Journal of Business Venturing 29*(1), 1–16.

Palfrey, T. and H. Rosenthal (1984). Participation and the provision of discrete public goods: A strategic analysis. *Journal of Public Economics 24*(2), 171–193.

Pitschner, S. and S. Pitschner-Finn (2014). Non-profit differentials in crowd-based financing: Evidence from 50,000 campaigns. *Economics Letters 123*(3), 391–394.

Rondeau, D. and J. List (2008). Matching and challenge gifts to charity: Evidence from laboratory and natural field experiments. *Experimental Economics 11*(3), 253–267.

Schelling, T. (1960). *The Strategy of Conflict*. Harvard University Press.

Solomon, J., W. Ma, and R. Wash (2015). Don't wait!: How timing affects coordination of crowdfunding donations. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*, pp. 547–556. ACM.

Spencer, M., S. Swallow, J. Shogren, and J. List (2009). Rebate rules in threshold public good provision. *Journal of Public Economics 93*(5), 798–806.

Stiver, A., L. Barroca, S. Minocha, M. Richards, and D. Roberts (2015). Civic crowdfunding research: Challenges, opportunities, and future agenda. *New Media and Society 17*(2), 249–271.

Strausz, R. (2017). A theory of crowdfunding - A mechanism design approach with demand uncertainty and moral hazard. Forthcoming in *American Economic Review*.

van de Kragt, A., J. Orbell, and R. Dawes (1983). The minimal contributing set as a solution to public goods problems. *American Political Science Review 77*(1), 112–122.

Wash, R. and J. Solomon (2014). Coordinating donors on crowdfunding websites. In *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing*, pp. 38–48. ACM.

Appendix: Supplementary regressions

We check robustness of the results reported in Table 3 using two additional sets of regressions. In Tables 6 and 7 we provide an analysis of treatment effects on the difference and ratio of success between the target project and all other available projects. Recall that the target project for SEED72 and SIA, used in models (1) and (2) of both tables, is project 72_{18} (or 92_{18} when seeded) and the target project for SIXSEED78, used in Model (3) of both tables, is project 78_{17} (or 98_{17} when seeded). Contribution-difference in Table 6 is measured as contributions to the target project minus the sum of contributions to all other available projects, at the group level. Threshold-difference in Table 6 is measured as thresholds reached of the target project minus the sum of thresholds reached of all other available projects, at the group level. Given that each group can maximally reach two thresholds, the value of this dependent variable is any integer between -2 and 1. Contribution-ratio and threshold-ratio in Table 7 are measured in a similar way, but taking ratios rather than differences.

The results from both tables point in the same direction, confirming the results presented in the main text. Sign and significance of the coefficients for the treatment dummy in Models (2) and (3) indicate strong support for Hypothesis 3, while sign and significance of the coefficients for the treatment dummy in Model (1) indicates strong support for Hypothesis 4.

We check robustness of the results reported in Table 5 using three additional sets of regressions for the conservation experiment. In Tables 8 and 9 we provide an analysis of treatment effects on the difference and ratio of success between the target project and all other available projects. These tables and their interpretation are otherwise similar to Tables 6 and 7. Table 10 is similar to Table 5, but using pooled data from both experiments. The results from all three tables point in the same direction, confirming the results presented in the main text.

Table 6: Treatment impacts on project success of the target project: differences. Compare Table 3.

	Project 72_{18}		Project 78_{17}
	(1) Contribution-difference	(2) Threshold-difference	(3) Contribution-difference
main			
SEED72	5.649	1.031**	
	(13.90)	(0.458)	
SIA72	−57.25***	−0.526	
	(13.92)	(0.431)	
SIXSEED78			43.54**
			(20.63)
# Subjects/group	−18.66	−0.816**	−37.96*
	(12.00)	(0.390)	(20.52)
Constant	94.87		65.78
	(67.04)		(115.6)
cut1			
Constant		−6.370***	
		(2.257)	
cut2			
Constant		−4.967**	
		(2.192)	
cut3			
Constant		−3.516	
		(2.149)	
Observations	40	40	27

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Coefficient estimates from linear regression models in (1) and (3) and ordered probit model in (2) (and standard errors in parentheses). All models are at the group level. The target project in (1) and (2) is project 72_{18} (or 92_{18} when seeded). The target project in (3) is project 78_{17} (or 98_{17} when seeded). Contribution-difference in (1) and (3) equals contributions to the target project minus the sum of contributions to all other available projects. Threshold-difference in (2) equals thresholds reached of the target project minus the sum of thresholds reached of all other available projects. Treatment names represent dummy variables, with BEN as the baseline treatment for (1) and (2) and SIX as the baseline treatment for (3).

Table 7: Treatment impacts on project success of the target project: ratios. Compare Table 3.

	Project 72_{18}		Project 78_{17}
	(1) Contribution-ratio	(2) Threshold-ratio	(3) Contribution-ratio
SEED72	0.0125	0.300^{**}	
	(0.0443)	(0.139)	
SIA72	-0.187^{***}	−0.152	
	(0.0444)	(0.152)	
SIXSEED78			0.155^{**}
			(0.0613)
# Subjects/group	−0.0627	−0.196	−0.00105
	(0.0382)	(0.128)	(0.0610)
Constant	0.826^{***}	1.484^{**}	0.0595
	(0.214)	(0.718)	(0.344)
Observations	40	36	27

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Coefficient estimates from linear regression models (and standard errors in parentheses). All models are at the group level. The target project in (1) and (2) is project 72_{18} (or 92_{18} when seeded). The target project in (3) is project 78_{17} (or 98_{17} when seeded). Contribution-ratio in (1) and (3) equals the ratio of contributions to the target project over the sum of contributions to all available projects per treatment. Threshold-ratio in (2) equals the ratio of thresholds reached of the target project over the sum of thresholds reached of all available projects per treatment. Treatment names represent dummy variables, with BEN as the baseline treatment for (1) and (2) and SIX as the baseline treatment for (3).

Table 8: Conservation experiment: Treatment impacts on project success of the target project: differences. Compare Table 5.

	Project 72_{18}	
	(1) Contribution-difference	(2) Threshold-difference
main		
SEED72	7.636	0.0262
	(13.25)	(0.527)
SIA72	−45.65***	−0.611
	(12.80)	(0.509)
# Subjects/group	−8.084	−0.424
	(11.40)	(0.439)
Constant	29.71	
	(60.81)	
cut1		
Constant		−3.169
		(2.361)
cut2		
Constant		−0.851
		(2.353)
Observations	35	35

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Coefficient estimates from linear regression model in (1) and ordered probit model in (2) (and standard errors in parentheses). All models are at the group level. The target project in (1) and (2) is project 72_{18} (or 92_{18} when seeded). Contribution-difference in (1) equals contributions to the target project minus the sum of contributions to all other available projects. Threshold-difference in (2) equals thresholds reached of the target project minus the sum of thresholds reached of all other available projects. Treatment names represent dummy variables, with BEN as the baseline treatment.

Table 9: Conservation experiment: Treatment impacts on project success of the target project: ratios. Compare Table 5.

	Project 72_{18}	
	(1) Contribution-ratio	(2) Threshold-ratio
SEED72	0.0351	−0.0980
	(0.0625)	(0.302)
SIA72	-0.164^{**}	−0.255
	(0.0604)	(0.231)
# Subjects/group	−0.0186	0.0784
	(0.0538)	(0.231)
Constant	0.548^{*}	−0.0392
	(0.287)	(1.304)
Observations	35	15

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Coefficient estimates from linear regression models (and standard errors in parentheses). All models are at the group level. The target project in (1) and (2) is project 72_{18} (or 92_{18} when seeded). Contribution-ratio in (1) equals the ratio of contributions to the target project over the sum of contributions to all available projects per treatment. Threshold-ratio in (2) equals the ratio of thresholds reached of the target project over the sum of thresholds reached of all available projects per treatment. Treatment names represent dummy variables, with BEN as the baseline treatment.

Table 10: Pooled experiments: Treatment impacts on project success of the target project. Compare Tables 3 and 5.

	Project 72_{18}	
	(1)	(2)
	Contributions	Thresholds
main		
SEED72	−2.060	−0.128
	(5.580)	(0.375)
SIA72	-24.81^{***}	-1.168^{***}
	(5.501)	(0.417)
# Subjects/group	11.50^{**}	0.991^{***}
	(4.556)	(0.328)
Constant	3.763	-5.370^{***}
	(25.00)	(1.791)
Observations	75	75

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Coefficient estimates from linear regression model in (1) and probit model in (2) (and standard errors in parentheses). All models are at the group level. The target project in (1) and (2) is project 72_{18} (or 92_{18} when seeded). Contributions in (1) equal the sum of contributions to the target project. Thresholds in (2) equals the sum of thresholds reached of the target project. Treatment names represent dummy variables, with BEN as the baseline treatment.

Appendix: Conservation experiment figures

The below figures repeat Figures 1–3 using data from the conservation experiment.

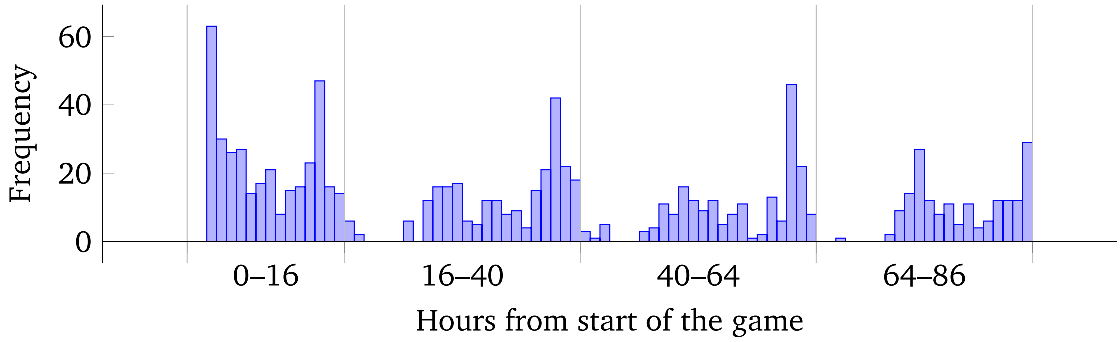

Figure 7: Conservation experiment: Investment frequency across all treatments in one-hour bins. Hour 0 corresponds to THU 08:00 and hour 86 corresponds to SUN 22:00. Compare Figure 1.

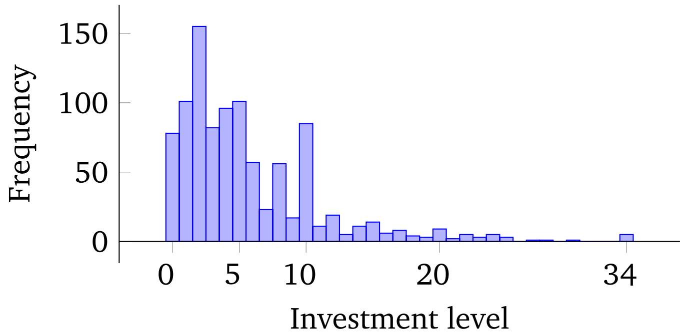

Figure 8: Conservation experiment: (Non-aggregated) investment level frequency across all treatments. Compare Figure 2.

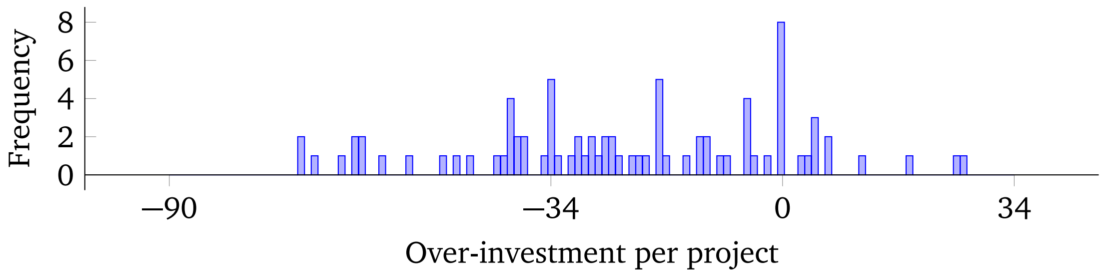

Figure 9: Conservation experiment: Project over-investment frequency in tokens across all treatments. Compare Figure 3.

Appendix: User Interface

Figure 10 shows a screenshot of the user interface at the start of the experiment under BEN, with information as described in Section 2.3. Via a tab subjects can switch between this screen and the instructions, provided below for treatment BEN (translated back from Dutch).

Figure 10: Screenshot of the user interface.

Appendix: Instructions

Introduction

Welcome and thank you for joining. Your participation is completely anonymous. Any results will solely be used for research purposes and will not be shared with others.

You will participate in a game. In the game you will be shown a number of projects in which you and the other participants in the game can invest. You and the other participants receive a number of tokens to invest in these projects. By investing, the group of participants, or part of the group, can make sure that a project is successful. Successful projects generate a payoff to all members of the group, including those that did not invest. Bottom-line is that it is your decision what to do with the tokens that you receive. It is not possible to make mistakes.

At the end of the game, on Sunday-night 22:00, the tokens that you earned in the game will be converted to money and this money will be credited to your bank-account by TNS-NIPO. The number of tokens that you earn depend on the decisions that you and the other participants make in the game. Hence, you are making real investment decisions. After the game ends, you will be sent a survey to fill out.

The game is further introduced below. It is important to read this information carefully. Next you can start the game via the "Game" button in the top right corner of your screen.

Projects in the game

All participants have the option to invest in the projects up to the end of the game. There will be no additional projects and no projects will be canceled. The description of each project is based on two characteristics: threshold and payoff. The thresholds and payoffs are different for each project and they remain unchanged throughout the game. Both characteristics are further explained below.

Threshold: The threshold is the number of tokens required to make the project succeed. The threshold can be reached by investments from you and the other participants in the game. Two important rules are the following:

1. When the total level of investments is equal to or higher than the threshold, the project is successful. The invested tokens will be used to carry out the project.
2. When the total level of investments is equal to or higher than the threshold, the project is not successful. Any investments made will be returned to the players.

Payoff: The payoff is the number of tokens that is given to each participant if the project is successful. If a project is successful, EVERYONE receives the payoff, including those that did NOT invest in this project. We do this since many projects, including projects financed by crowd-funding, involve benefits to people that did not invest in the project themselves.

The main rules of the game

1. You are member of a group of 6 participants. Each one of you receives, once-only, 34 tokens at the start of the game. This implies that the group as a whole owns 204 tokens at the start of the game.

2. The value of each token is 20 Euro-cents. Your 34 tokens at the start of the game are worth 6.80 Euro. By investing tokens in projects you can earn more tokens.
3. You may invest but you are not required to do so. During the game, no projects will be added or dropped. Also, the characteristics of the project remain unchanged throughout the game.
4. You may invest at any moment during the course of the game.
5. You may invest multiple times in the same project.
6. You may invest in more than one project. The group as a whole owns a sufficient number of tokens to successfully fund two projects.
7. Once you have made an investment, this action cannot be canceled. When the project is not successful, your investment will be returned to you. When the project is successful, each player receives the corresponding payoff.
8. The game ends on Sunday night 22:00. At that time the successful projects will be identified, your final number of tokens will be calculated, and converted into money units (20 Euro cents per token).
9. Your final number of tokens equals the sum of: (start tokens) - (invested tokens) + (returned tokens for unsuccessful projects) + (payoff tokens for successful projects).

Example

The characteristics of project A are: threshold 60 and payoff 15. You do not invest in project A. Because other participants do invest in project A, the threshold is reached and the project is successful. Each participant receives a payoff of 15 tokens.

The characteristics of project B are: threshold 82 and payoff 20. During the game you invest a total of 10 tokens (6 tokens on Saturday and 4 tokens on Sunday) in project B. Other participants also invested in project B, but the threshold is not reached and the project is not successful. Your invested 10 tokens are returned to you.

Your final number of tokens at the end of the game is calculated as follows:

Start tokens	34
− invested tokens: 10 tokens in project B	− 10
+ returned tokens for unsuccessful projects: 10 tokens for project B	+ 10
+ payoff tokens for successful projects: 15 tokens for project A	+ 15
= final number of tokens at the end of the game	= 49

In this example you would have earned 15 more tokens compared to the number of start tokens (the 34 tokens at the start of the game), and your bank account would have been credited by 9.80 Euro (=49*20 Euro cents).

Requirements

There are three requirements that you should meet:

1. For four days, you should log in at least once each day on the game website.
2. You should decide at least once per day whether you want to make an investment, and if so, in which project or projects, and make this decision on the game website.
3. After the game ends you will be sent a survey that you should fill out completely.

What to do when you log in?

Each time that you log in we ask you to take the following actions:

1. Have a close look at the projects. The thresholds and payoffs remain unchanged throughout the game. The total investments and the number of investors may of course change, depending on your investments and those of the other participants. This information is updated continuously.
2. Determine whether you want to make an (additional) investment during this login, and if so, how many tokens. Take into account that the sum of your investments cannot exceed 34 tokens.
3. Make your decision by choosing your investment from the drop-down list of the relevant project.

Frequently asked questions

1. **Can I lose money in this game?** ANSWER: No, this is not possible. It is possible to lose part of your 34 tokens that you receive at the start of the game. This is possible only when you invest more tokens in a project than the level of the payoff and the project happens to be successful (if it were not successful your invested tokens would be returned to you).
2. **What happens when investments exceed the threshold?** ANSWER: Nothing. Suppose a project has a threshold of 100 tokens and the total investments in this project equal 105 tokens. Every participant receives this project's payoff but the 5 tokens 'over-investment' are not returned.
3. **What is the best moment to invest and in which project?** ANSWER: This depends on your personal preferences. You can use all information provided on the game website to decide when and in which project to invest. This information includes the threshold and payoff, but also past investments by you and the other participants in the game.
4. **Can I invest more than 34 tokens?** ANSWER: No, you cannot invest more than the 34 tokens you receive at the start of the game. It is also not possible to purchase more tokens.
5. **Can I see which other players have already invested?** ANSWER: The game website shows per project the total level of investments made as well as the number of players that has already invested. The website does not show the identity of these players, when they invested, or how many was invested by each player.

For other questions, please contact < Contact Information >.