Exploration in Teams and the Encouragement Effect:

Theory and Experimental Evidence

3 May 2018, 11:15-12:30, bld. 72, room 465

Topi Miettinen (Hanken School of Economics)

Abstract: This paper analyzes a two-person, two-stage model of sequential exploration, where both information and payoff externalities exist, and tests the derived hypotheses in the laboratory. We theoretically show that even when agents are self-interested and perfectly rational, the information externality induces an encouragement effect: a positive effect of first-mover exploration on the optimality of the second-mover exploring as well. When agents have other-regarding preferences and imperfectly optimize, the encouragement effect is strongest. The explorative nature of the game raises overall efficiency compared to a payoff equivalent public goods game. We experimentally confirm our main theoretical predictions using a novel experimental paradigm. Our findings are relevant not only for motivating and managing innovation in fields like criminal justice, education, scientific research and health, but also for diverse applications such as prosocial behavior at work and collective action.

Exploration in Teams and the Encouragement Effect: Theory and Experimental Evidence

Emma von Essen[1], Marieke Huysentruyt[2, 4], and Topi Miettinen *[3, 4]

[1]Department of Economics and Business, Aarhus university

[2]Society and Organizations Center, HEC Paris

[3]Department of Economics, Hanken School of Economics, Helsinki

[4]SITE, Stockholm School of Economics

July 7, 2017

This paper analyzes a two-person, two-stage model of sequential exploration, where both information and payoff externalities exist, and tests the derived hypotheses in the laboratory. We theoretically show that even when agents are self-interested and perfectly rational, the information externality induces an encouragement effect: a positive effect of first-mover exploration on the optimality of the second-mover exploring as well. When agents have other-regarding preferences and imperfectly optimize, the encouragement effect is strongest. The explorative nature of the game raises overall efficiency compared to a payoff equivalent public goods game. We experimentally confirm our main theoretical predictions using a novel experimental paradigm. Our findings are relevant not only for motivating and managing innovation in fields like criminal justice, education, scientific research and health, but also for diverse applications such as prosocial behavior at work and collective action. (JEL: C91, D03, D83, O31)

*We thank Chloé Le Coq, Olivier Darmouni, Denis Gromb, Leif Helland, Tuomas Nurminen, Pauli Murto, Heiner Schumacher, Ro'i Zultan, and participants at seminars at HEC Paris, University of Turku, HECER in Helsinki, and at ESA World Meeting 2016 in Jerusalem. We gratefully acknowledge financial support provided by the European Union through its Seventh Framework Program for research, technological development, and demonstration under the grant agreement 613500 (SEFORIS project), Norwegian Research Council grant 250506, and Hanken Foundation. We also thank PCRClab at the University of Turku, COBELab at the University of Aarhus, Aalto Choice Tank at Aalto University, and the CESAR lab at the BI Norwegian Business School in Oslo for hospitality. We thank Olli Lappalainen, Knut-Eric Neset Joslin, Jason O'Donoghue and Kristaps Dzonsons for excellent research assistance.
Corresponding author: topi.miettinen@hanken.fi

1 Introduction

Innovation plays a central role in the production of public goods. Today, we observe a great need for new ideas that can help us meaningfully address the problems of poor or declining educational systems, unequal access to affordable health care, imminent environmental challenges, international terrorism, social fragmentation and the chronic offending in low-income, urban neighborhoods, to name but a few. One promising trend in the provision of *innovations* related to public goods is the rise of spontaneous, voluntary, often uncoordinated, yet joint search contributions by individuals, groups or organizations. Search is decentralized and distributed, but knowledge is freely revealed or shared. More specifically, agents privately incur costly exploration efforts in search of a solution, the benefits of which accrue to all and cannot be privately appropriated. Further, they share and update information about solutions that are potentially still feasible and about others that have been tested and abandoned during the search process. Thus, incentive design in such settings must deal with not only dynamic free-riding but also the inherent process of learning or information sharing.

With very few exceptions, previous work on the voluntary provision of public goods deals with situations that are static or involve limited exogenous uncertainty about the outcome of the contributions to public goods. This work has provided invaluable insights into the drivers of private contributions to a public good (Camerer, 2003; Cappelen et al., 2015; Güth et al., 2007; Levati et al., 2007). By contrast, we specifically account for the uncertainty of a public good discovery and learning over time, and we thus shift focus to the dynamics of searching for a public good.

In this paper, we provide a model of how teams of agents (more specifically pairs of individuals) sequentially and voluntarily explore for the public good in the following circumstances: when effort is privately costly but cannot be directly contracted upon, when the value of the discovery or public benefit is known (and shared), when a finite number of possible solutions exists, and when exploration is open (information is shared). For real life scenarios captured by the model we can consider two problem solvers who take turns to find a solution to a particular problem and can find it only by pursuing one of a given number of equally promising alternatives. They can be employees who are expected to improvise in order to make their peers feel more engaged at work, research teams in search of a solution to a theoretical modeling or empirical specification problem, or industry experts, say in telecommunications, in search of a new US or European-wide standard. The key contribution of our model is to deliver sharpened insight into the strategic considerations that determine individual exploration decisions. Further, our simple model lends itself not only to an analysis with self-interested and perfectly rational agents, but also to an analysis with agents who imperfectly optimize and/or hold other-regarding preferences. We put our main theoretical intuitions to the test in a computerized laboratory setting by, using a novel exploration game paradigm. The experiment, in turn, allows us to test the causal effect of uncertainty on the costly individual exploration decision and joint exploration outcomes.

Our main results, both theoretical and experimental, are as follows. In our simple exploration game, as in more complex strategic bandit models Bolton and Harris (1999) and Hörner and Skrzypacz (2016), the information externality induces, a standard encouragement: a positive effect of first-mover exploration on the optimality of the second-mover exploring as well. The novelty here is that we show that the expected occurrence and size of the encouragement effect not only depend on the value of the public good benefits, but also on the assumptions we make about the agents rationality. So long as we assume perfect rationality and self-interested preferences, which we do in the baseline model, the encouragement effect is only at play within a limited range of public benefit values. Intuitively, for a specific range of fairly low values of the public good benefits, only exploration by the first-mover can trigger the second-mover to explore as well. In the subgame perfect equilibrium of our baseline model, we thus expect (in theory) the first-mover's exploration to be non-monotonic in the value of the public good. Once we allow for imperfect optimization or other-regarding preferences, [1] which extensive empirical and laboratory literature on strategic interactions and public goods provision, respectively, suggest may matter in our setting, the existence of an encouragement effect readily extends to all possible values of the public good. Moreover, the encouragement effect in this behavioral model is largest when agents both imperfectly optimize and care about each others' payoffs.[2] Together, our first set of theoretical pre-

[1]We operationalize imperfect optimization by means of the quantal response equilibrium (QRE) (McKelvey and Palfrey, 1998), and other-regarding preferences by means of the social welfare utility model (Charness and Rabin, 2002). We present our rationale for these modelling choices in our theoretical section.

[2]There are many empirical studies evidencing the preference channel effect (Berg et al., 1995; Clark

dictions underscore the importance of information sharing as a non-monetary channel that motivates exploration.

Next, to pin down the causal effect of uncertainty of the outcome of the public goods on individual contributions, we theoretically compare equilibrium outcomes for the baseline as well as the behavioral model of joint exploration with those of a payoff equivalent, canonical voluntary public goods (henceforth, PG) provision game that lacks the explorative nature. In this payoff equivalent game the total value of the public good is the same as in the exploration game, and thereby also the myopic incentives to contribute coincide in the two games. Yet, whereas the dynamic encouragement effect appears in both the basic and behavioral version of the exploration game, it is absent in the standard PG provision game. And so we establish theoretically that aggregate exploration in the exploration game is weakly greater than aggregate contributions in the payoff equivalent PG provision game. Uncertainty in the public goods production process thus raises overall efficiency.

Finally, our experimental results, based on the analysis of 13,952 individual exploration decisions in a computerized laboratory environment, broadly confirm our main theoretical predictions. Observed behaviors in the laboratory are in fact best explained by the behavioral version of our exploration model, where we allow players to imperfectly optimize and hold other-regarding preferences. Further, we show that aggregate contributions in the exploration game consistently exceed aggregate contributions in the payoff equivalent PG provision game.

To the extent that our results are externally valid, one important implication of our results is that governments who wish to harness the decentralized voluntary search for the public good are well-advised to promote (i) information-sharing, for instance, by investing in improved technological infrastructure that can speed up the sharing of information and (ii) the development of social preferences amongst its citizens at large, for instance through educational programs. Interestingly though, the encouragement effect in our model leads even self-interested individuals to search for the public good. Another implication is that by emphasizing the uncertainty about where the solution to a

and Sefton, 2001; Falk et al., 2003, 2008) and various theories have been put forward that rationalize such patterns. The preference channel is predicted to be active both in the public goods game and in the exploration game. Yet, the imperfect optimization effect appears in the exploration game only. This effect requires the information externality channel to operate. We are not aware of previous experimental evidence identifying this effect, nor of theoretical arguments pointing it out.

difficult public goods problem lies, one can actually elicit greater voluntary contributions. And hence, as contributions to the public good can be framed as search contributions, this will raise and not lower, as one might have thought, overall contributions and bring aggregate contributions closer to the social optimum.

This paper is related to four strands of literature. First, our paper builds on a simple model of interactive search by Fershtman and Rubinstein (1997). We adapt this modeling framework to capture a situation where exploration is open (information is shared) and benefits in the event of discovery are public (non-rival and non-excludable). This befits our focus on voluntary and joint search for the public good and extends the model to allow for imperfect optimization and other-regarding preferences by relaxing standard assumptions in two relevant directions as suggested by the empirical and laboratory literature. There is a vast theoretical literature on search with seminal papers by Stigler (1961) and McCall (1970) who study fixed sample and sequential search, respectively, and analyze search that is carried out by individuals in isolation from each other. We study search where team members explore sequentially one after another, and the benefits of search are public accruing to the entire team.[3]

Second, our paper is related to the literature on moral hazard in teams (Holmstrom, 1982), especially the theoretical analyses of sequential effort provision by Strausz (1999), Winter (2006), and Winter (2009). The latter two study the strategic incentives of team members when late movers observe the effort of early movers and efforts are complementary[4]. Winter (2009) shows how higher exogenous rewards can lead to lower efforts (the so-called incentive reversal effect), and so our result regarding the non-monotonicity of exploration in the exploration game with subgame perfect equilibrium and with self-interest motivation can really be seen as a corollary of his finding. Many of these theoretical setups have also been studied experimentally (see Part 6.1 of Plott and Smith (2008), Klor et al. (2014), and Brown et al. (2011), for instance). Relative to this literature, the key contribution of our paper is twofold: First, we develop a sequential,

[3]In the case of sequential search for a *private* good instead of a public good (keeping the information externality), the encouragement effect no longer occurs in our baseline specification; nor does it arise when we search for the quantal response equilibrium. Once we allow for other-regarding preferences, the encouragement effect kicks in again for a specific, though much more limited range of private good values. In the case of *simultaneous* search for a public good, the encouragement effect remains at play for a limited range of public good parameter values in the SPE.

[4]Sequential moves also promote contributions when efforts are not complementary but asymmetries across parties typically erode the benefits or leadership in that case (Cappelen et al., 2015; Güth et al., 2007; Levati et al., 2007).

strategic model of search in teams where the returns of costly individual search efforts are uncertain. Second, our model is simple enough to lend itself not only to an analysis of perfect rationality but also to imperfect optimization and other-regarding preferences, and to an experimental study in the laboratory.

Third, our model can also be recast as a model of strategic experimentation. In this literature the paper most related to ours is Bonatti and Hörner (2011) who study a strategic bandit model where each of two team members must choose between costly exploration and a safe activity, and similarly where both informational and payoff externalities co-exist. They consider the so-called good news model, where exploration efforts are strategic substitutes. But, as pointed out by Hörner and Skrzypacz (2016), in a private goods setting exploration efforts are typically strategic complements in a bad news model.[5] This is more in line with our setup. Our model is a much simpler finite-alternative model. In fact, in that regard, our model is reminiscent of the project-selection setting of Aghion and Tirole (1997).[6] Our theoretical and experimental setup also puts far less cognitive demands on lab participants than the canonical, multi-arm bandit problems. Hence, this class of models is more amenable to laboratory testing and to the incorporation of imperfections and other-regarding preferences into the theoretical analysis of strategic exploration. By using such a model we promote the methodological ideals of Samuelson (2005) by exploiting the interplay between theory and experiments in order to advance human understanding of economic phenomena.[7] We further make it easier for experimental participants' to understand the setting by using intuitive and visually appealing video-instructions to explain the experimental design. The video-instruction itself constitutes a methodological contribution to the experimental literature.[8] In an independent experimental study of the Halac et al. (2016) bandit exploration model, Deck and Kimbrough (2017) use a similar approach. The novel exploration paradigm contributes to the experimental literature on innovation (see Ederer and Manso (2013), for instance, and Boudreau and Lakhani (2016) and Brüggemann and Bizer (2016) for two recent review articles).

Finally, our paper ties into the literature on the voluntary provision of public goods,

[5]See Hörner and Skrzypacz (2016), pp. 2-3.

[6]Their model, however, considers the optimal allocation of authority in order to promote incentives to acquire information. In their model, unlike ours, search efforts become strategic substitutes.

[7]See also Mäki (2005) and Gilboa et al. (2014).

[8]The video-instructions are available at https://dreambroker.com/channel/1ehcya5t/77qp05es

collective action, and prosocial behaviors originally studied by Olson (1965) in a self-regarding model and by Becker (1974) with altruistic preferences.[9] Our paper is most closely related to McBride (2006)'s work on the discrete version of the public goods game with symmetric uncertainty about the contribution threshold. In a self-regarding model, McBride (2006) finds, like us, that uncertainties in the public good provision environment may induce non-monotonicities. However, in his model, the encouragement effect does not arise. More generally, our paper shifts attention away from uncertainty about others degree of altruism, contribution costs, or valuations of the public good (Anderson et al., 1998; Palfrey and Rosenthal, 1991) to uncertainty inherent in the production process itself and, as emphasized by Admati and Perry (1991), Compte and Jehiel (2003), and Compte and Jehiel (2004), to the sequentiality and dynamic strategic interdependency of the contributions.

The rest of the paper is organized as follows: Section 2 presents the game-theoretic model of exploration for the public good; Section 3 explains the experimental procedure and data; Section 4 contains the experimental analysis; Section 5 concludes.

2 Theory

2.1 Basic Model of Exploration with Sequential Moves

Consider a simple two-stage, two-player exploration game in which partners (be it two employees, two co-authors, or two industry experts) take turns to explore. There is a finite product space (of locations) and a unique public good (i.e. treasure) in a single location within that space. Let K denote the number of locations to explore. Ex ante, each location is likely to hold the treasure with probability $1/K$. The valuation of the public good, i.e treasure size for i α_i, $i = 1, 2$, is ex ante known, non-excludable, and obtained if and only if the public goods is found (or a breakthrough is made). Without loss of generality, we can assume that $\alpha_1 = \alpha = \alpha_2$ and let the asymmetries between the players be reflected in the exploration costs. The cost of exploring c_i is borne by the relevant agent. Not exploring implies zero cost. Assume that $c_1 \geq c_2$, which is in line with Winter (2006)'s optimal incentive mechanism.[10] The player in stage two can

[9] Lindbeck and Weibull (1988) extend Becker's model to a dynamic setup

[10] According to Winter (2006), late movers should be given higher-powered incentives when there exist increasing returns to exploration. The basic intuition is that player 2 face no implicit threat that their

learn from the exploration of her partner in stage one. The model assumes complete and perfect information (observable effort and outcomes), though no coordination device exists. We seek for the subgame perfect equilibria of the game under different treasure size or public good value regimes.[11]

2.1.1 Subgame Perfect Equilibrium

Let us solve the subgame perfect equilibrium (SPE) of the model, by using backward induction. In stage two if the treasure has not been found, it is optimal for player 2 to explore iff

$$c_2 < \alpha/Y,$$

where $Y \leq K$ is the number of alternatives left to explore in the second stage. Likewise, player 1s myopic incentive to explore is captured by $\alpha/K - c_1$. This myopic incentive to explore is all that player 1 needs to consider if player 2s choice is not affected by that of player one. There are two such cases.

First, if $c_2 < \alpha/K < \alpha/(K-1)$, then player 2s exploration cost is so low that player two finds it optimal to explore whether player 1 explores an option or not (provided the public good is not found by player 1). Second, if $c_2 > \alpha/(K-1) > \alpha/K$, then player 2 cost of exploration is so high that player two finds it suboptimal to explore whether player one explores an option or not. Yet, player one, unlike player 2, needs also to consider dynamic effects of her choice on the exploration incentives of player 2, i.e. a potential *encouragement effect*. The encouragement effect is relevant and player 1s exploration may affect the incentives of player 2 if $\alpha/K < c_2 < \alpha/(K-1)$. Player 1 then prefers to explore if

$$\frac{\alpha}{K} - c_1 + \delta(\frac{K-1}{K})(\frac{\alpha}{K-1}) > 0, \tag{1}$$

where $0 < \delta \leq 1$ is player 1s discount factor. The above inequality can in equivalent terms be written simply as

$$c_1 < (1+\delta)\frac{\alpha}{K}, \tag{2}$$

failure to innovate will trigger subsequent agents to shirk as well. Hence, player 2 should be provided with stronger incentives to exert effort than layer 1.

[11] The public good could be interpreted as a project (as in Aghion and Tirole (1997)), as a mode of organization, a technological strandard for an industry, or a methodological breakthrough in academic collaboration (as suggested by Bonatti and Hörner (2011)), and so forth.

capturing the simple intuition that by investing in a unit of exploration herself, she gets another unit for free as an optimal reaction by player 2. We summarize these conditions and their behavioral implications for perfectly rational self-interested players in Proposition 1.

Proposition Let $c_1 > c_2$. Let $\alpha_1 = \alpha_2 = \alpha$. Conditional on the prize not having been found in the first stage, on the equilibrium path:

- Neither player explores when $\alpha < max\{\frac{c_1 K}{1+\delta}, c_2(K-1)\}$;
- Both players explore when $max\{\frac{c_1 K}{1+\delta}, c_2(K-1)\} < \alpha < Kc_2$;
- Only player 2 explores when $c_2 K < \alpha < c_1 K$;
- Both players explore when $c_1 K < \alpha$.

This proposition reveals that player one's equilibrium exploration decision is non-monotonic in α due to the *encouragement effect*, which is defined as the impact of player 1s exploration decision on player 2s exploration decision. Intuitively, when the value of the public good is very low, neither player finds it in his/her best interest to explore. For somewhat higher values of the public good, neither player's myopic incentives to explore are sufficient, but the dynamic encouragement effect triggers player 1 to explore and encourages player 2 to explore as well if player 1 does not find the public good. Then, for even higher values of the public good, it is a dominant strategy for player 2 to explore. Player 1 knows this and free-rides on player 2s contribution. Finally, once the value of the public good is so high that even player 1's myopic incentive dictates to explore, then both players find it optimal to explore. The basic intuition behind this result is precisely the same as the argument of Hörner and Skrzypacz (2016, pp. 2-3)for why the privately optimal best response to the opponent's simple cut-off exploration strategy in a two-player bad news Poisson bandit cannot be a simple cut-off strategy; rather it involves ranges of optimal encouragement and free-riding, corresponding to bullets two and three, respectively, in the above proposition. Thus our model allows us to test the basic encouragement intuition (Bolton and Harris (1999)) of the strategic experimentation models in a considerably simpler framework which is easily understood by the participants of our experiment.

Figure 3.1 illustrates the core theoretical insights by using a simple numerical example. Let $K = 4$, $c_2 = 200$, $c_1 = 300$ and $\delta = 1$. Then neither explores when $\alpha < 600$. Both explore when $600 \leq \alpha < 800$. Only player 2 explores when $800 \leq \alpha < 1200$ and both players explore when $\alpha \geq 1200$. Notice that the expected total contributions are 1.75 units when both explore, since player 2 explores only when the treasure is not found in that case, i.e. with probability 3/4.

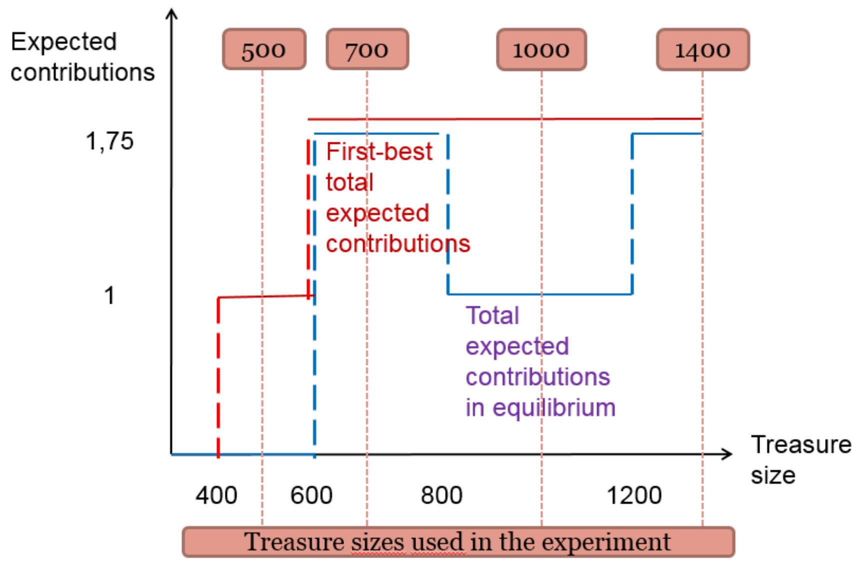

Figure 3.1: Theoretical predictions in the exploration game, SPE and self-interest

2.1.2 Subgame Perfect Equilibrium of the Voluntary Public Goods Game

The standard voluntary public goods game, in its sequential two player binary choice form, is formally nested in our exploration game. To derive the standard voluntary public goods game from our exploration game, the value of the public good is distributed evenly over the entire finite product space so that in each location the value of the public good is the same and equals $\alpha_C = \alpha/K$. Thus, in the standard public goods version of the game, $1/K$:th of the value of the unique treasure in the exploration game is produced for each individual contribution made in the standard game, and this value is produced with certainty for each contribution that costs c_1 for player 1 and c_2 for player 2. Indeed, in the canonical voluntary public goods game, every individual contribution generates a public good with certainty. In our special case where only a single contribution can be made by each player and choices are sequential, the standard voluntary public goods game is in fact a sequential prisoner's dilemma. The reaction functions and SPE for the sequential prisoner's dilemma game are very well known. Let us sketch the derivation here for purposes of comparison with the exploration game. Let α/K denote the public good produced for each individual contribution made. Then, player i will find it optimal to contribute to the production of a public good iff

$$c_i \leq \alpha/K. \tag{3}$$

Player 1s equilibrium exploration decision is monotonically increasing in α/K. Player 1 can no longer exert an influence on player 2s decision. The *information externality*, which is a distinct characteristic of the exploration game, disappears in the voluntary public goods game and with it the *encouragement effect*. Thus, each player's myopic incentive dictates the optimal choices and each player thus (generically) has a dominant strategy independent of the other player's choice.

Intuitively, when the value of the public good is very low, neither player finds it in his/her best interest to contribute. Then, given that $c_2 < c_1$, for higher values of the public good, it is a dominant strategy for player 2 to contribute and likewise player 1 has a dominant strategy to free-ride on player 2s contribution. Finally, once the value of the public good exceeds player 1s cost, then both players have a dominant strategy to contribute.[12]

As illustrated in Figure 3.2, the equilibrium level of contributions to the production of the public good is weakly higher in the *exploration game* than in the voluntary public goods game or *public goods game*. In particular, when $\alpha/K < c_2 < \alpha/(K-1)$ and $c_1 \leq \alpha(1+\delta)/K$ both players will contribute in the exploration game (if the treasure is not found by player 1), but neither will contribute in the public goods game.

[12]Strictly speaking, the game is a sequential prisoner's dilemma if and only if $\alpha/K < c_i < (2\alpha)/K$ for $i = 1, 2$.

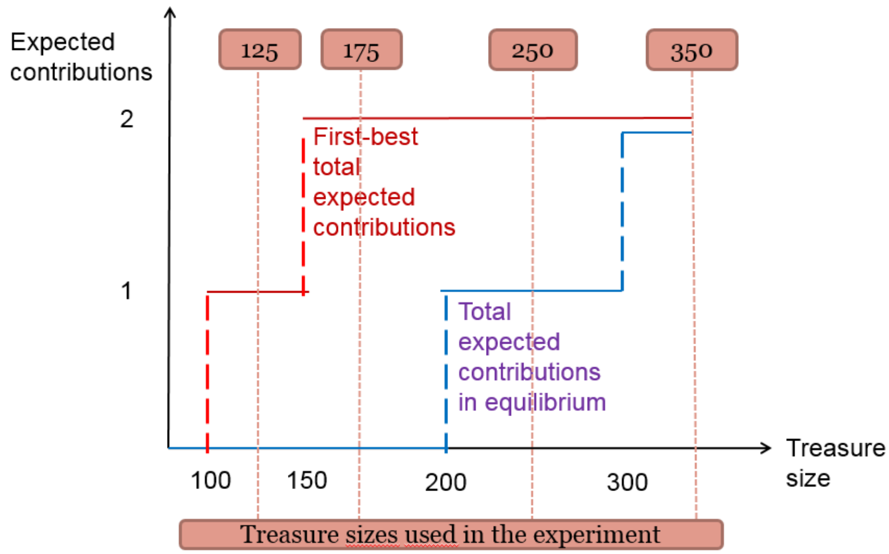

Figure 3.2: Theoretical predictions in the public goods game, SPE and self-interest

From a welfare perspective, it is optimal that player i contributes if $c_i < (2\alpha)/K$, and this is true both in the exploration game and the standard public goods provision game. Thus, the welfare properties of the two games coincide. Likewise, the myopic incentives in the two games are the same. The only difference is the presence/absence of the encouragement effect, and this property in regard to the efficiency of equilibrium play drives a wedge between the two games.[13]

2.1.3 SPE Predictions

A straightforward comparison between the SPE for the exploration game (henceforth, exploration game) and payoff equivalent voluntary public goods provision game (henceforth, public goods game) yields two testable theoretical predictions.[14] These are the predictions that we have pre-registered on the Open Science Framework platform at https://osf.io/ (Name: Exploration in partnership).

[13]Even if both the costs and benefits of contributing in the public goods game were symmetric, then still the exploration game would yield higher welfare when parameter values satisfy: $\alpha/K < c < min\{\alpha/(K-1), \alpha(1+\delta)/K\}$. In the fully symmetric parameters case though, the total amount of exploration in equilibrium is monotonically increasing in α/K in both the public goods and exploration game.

[14]Instead of considering the subgame perfect Nash equilibria, one can compare the sets of Nash equilibria in the two games. All Nash equilibria of the public goods game are also Nash equilibria of the exploration game, but the Nash equilibrium with encouragement in the exploration game is never a Nash equilibrium in the public goods game. Thus analogs of the listed hypotheses hold for the setwise comparison as well.

The hypotheses consider the subgame perfect equilibrium for the exploration game and public goods game concerning players who only care about their own payoff only.

1. *Efficiency Hypothesis:* Aggregate exploration (contributions to the public good) will be weakly higher in the exploration game than in the public goods game.

2. *Encouragement Effect and Non-monotonicity Hypothesis:* Player 2 will on average explore more than player 1 in the public goods game and, for a limited range of treasure sizes, the wedge between player 1s and player 2s exploration efforts will be smaller in the exploration game than in the public goods game. In the exploration game, player 1s exploration will be non-monotonic in treasure size.

A key intuition behind all these theoretical predictions is that there is an encouragement effect present in the relevant range of the value of the public benefit in the exploration game, contrary to the public goods game. Player 1 can face an implicit threat that his failure to explore will trigger player 2 to not explore as well. In reverse, there is also an implicit promise that his exploration will trigger player 2 to explore as well (if the treasure is not found). Equivalently, the uncertainty in the production process of the public goods invokes a complementarity between the two players' exploration decisions. As a result, aggregate equilibrium exploration or contributions to the public good are weakly higher in the exploration game than in the public goods game.

2.2 Behavioral Model of Sequential Exploration for the Public Good

Our basic model so far assumes that each player chooses the action with the highest payoff for sure (people always perfectly optimize) and only considers his or her own payoffs (people are selfish). In this section, we relax these two assumptions. We extend our model to allow people to imperfectly optimize and have social preferences.

2.2.1 Imperfect Optimization

While providing a useful benchmark for understanding choice behavior in our setting, the subgame perfect Nash equilibrium also presumes strong rationality assumptions about the capacity of implementing the optimal strategy with certainty. Real behaviors, however, are typically error-prone. In our setting, players might make errors and understand that others also make erroneous choices.

This notion of bounded rationality[15] can be formally incorporated into our set-up by deriving the logit quantal response equilibrium (QRE) McKelvey and Palfrey (1998) instead of the SPE.[16] In the logit-QRE, the choice probabilities reflect rationality in the sense that they are inversely related to the opportunity costs of the choices and the implied choice probabilities are correctly anticipated by the agents. While better strategies are more likely to be played than worse strategies, there is no guarantee that best response strategies and actions are played with certainty, and this fact is understood by all players. This relatively small departure from perfect rationality has been found to produce predictions that better fit data from laboratory experiments (Anderson et al. (1998), Goeree and Holt (2000), Goeree et al. (2010)).

In the logit quantal response model, the choice probabilities are proportional to the exponentials of the expected utilities, v_i, of the actions given the beliefs about the opponent's behavior. Let us denote the expectation of i about the action profile a_j of the other player by $\hat{\sigma}^i_j(a_j)$. In the quantal response equilibrium, player i chooses action a_i with probability

$$\sigma_i(a_i) = \frac{exp\left((1/\mu)\left(\sum_{a_j} \hat{\sigma}^i_j(a_j) v_i(a_i, a_j)\right)\right)}{\sum_a exp\left((1/\mu)\left(\sum_{a_j} \hat{\sigma}^i_j(a_j) v_i(a, a_j)\right)\right)} \tag{4}$$

. This formulation allows for considering both erratic decision making by self-interested agents (replace v_i with π_i, the pecuniary payoff of i) and other-regarding agents (use a more general value function v_i as we do in section 2.2.2 below).

Taking the ratio of choice probabilities of two different actions a'_i and a''_i (the odds ratio) yields

$$\frac{\sigma_i(a'_i)}{\sigma_i(a''_i)} = \frac{exp\left((1/\mu)\left(\sum_{a_j} \hat{\sigma}^i_j(a_j) v_i(a'_i, a_j)\right)\right)}{exp\left((1/\mu)\left(\sum_{a_j} \hat{\sigma}^i_j(a_j) v_i(a''_i, a_j)\right)\right)}, \tag{5}$$

[15]See Grüne-Yanoff (2007) for an encompassing discussion of the role of the concept of bounded rationality in economics and psychology.

[16]We motivate our decision to study the QRE as follows: In our experimental setting, participants have ample opportunity to learn about the population behavior and to adapt their behavior accordingly in our experiment. Indeed, the game is played several times in each of the different public good value specifications, altogether more than thirty times. The quantal response model, where players are assumed to have correct expectations about population behavior, thus strikes us as a more appropriate solution concept for the behavioral analysis than concepts analyzing inexperienced players (Crawford et al., 2013). On the other hand, models analyzing learning dynamics explicitly (Erev and Haruvy, 2013) seem unnecessarily complicated for our main focus. The QRE-model is a simpler one-parameter model whereas non-equilibrium models of strategic thinking and learning models typically rely on a higher number of parameters.

and thus the ratio of choice probabilities is proportional to the ratio of exponentials of expected utilities. Expectations and choice probabilities must coincide in equilibrium and thus $\hat{\sigma}_i^j = \sigma_i$ for $j \neq i$. The novel feature is noise, which increases in the noise parameter μ. As μ approaches infinity the choices are entirely random. As μ tends to zero (from above), the choice probabilities converge to a Nash equilibrium of the game. Thus with μ tending to zero and v_i replaced with π_i, we are back in the analysis of Section 2.1. The log of the odds ratio of choice probabilities in the QRE-model is merely

$$(1/\mu)\left(\sum_{a_j}\hat{\sigma}_j^i(a_j)v_i(a_i', a_j) - \sum_{a_j}\hat{\sigma}_j^i(a_j)v_i(a_i'', a_j)\right) \quad (6)$$

and therefore it perfectly linearly reflects the expected payoff difference between choosing the two actions given the expected behavior of others.

The amount of exploration (in probability mass terms) with imperfectly optimizing, selfish players in the exploration game is, as in SPE, weakly greater than in the public goods game. However, the encouragement effect now appears for *all* treasure sizes (or values of the public good) in the exploration game, yet continues to be absent in the public goods game. The intuition for this result is that now, for any given configuration of parameter values, *both* player 2s actions (to explore or refrain from exploring) occur with positive probability, and, given the information externality, exploration by player 1 always increases the payoff of exploration to player 2. Furthermore, the encouragement effect is now increasing in treasure size. Intuitively, when the stakes are higher, player 2 has more to gain following exploration by player 1. In Appendix A we present the formal analysis of the exploration behaviors in the QRE for the public goods and exploration games with selfish players.

2.2.2 Other-regarding Preferences

We now also incorporate a more general model of preferences that embeds difference aversion and social welfare preferences as identically parsimonious and tractable special cases. This more general model also nests purely selfish preferences as a limiting case.

We allow for people not only to be self-interested but also care about social efficiency and inequity by integrating the goal function in a generalized version of the social welfare model by Charness and Rabin (2002).[17] In our setting, this goal function can be written in the following form

$$v_i(a_i, a_j) = \begin{cases} (1-\rho) \cdot \pi_i(a_i, a_j) + \rho \cdot \pi_j(a_j, a_i) & \text{if } \pi_i(a_i, a_j) \geq \pi_j(a_j, a_i) \\ (1-\sigma) \cdot \pi_i(a_i, a_j) + \sigma \cdot \pi_j(a_j, a_i) & \text{otherwise} \end{cases},$$

for $i = 1, 2$, where ρ and σ may be negative, zero, or positive and $\rho \geq \sigma$.[18] The parameters ρ and σ allow for a range of different "distributional preferences" that rely solely on the outcomes and not on any notion of reciprocity. For instance, when $1 \geq \rho \geq \sigma > 0$, then these parameter values capture social welfare concerns; whereas when $1 > \rho > 0 > \sigma$, these parameter values correspond with inequity or difference aversion. Irrespective of the specific distributional preferences we consider, ρ is always understood to be greater than σ, and indeed this has also been widely corroborated empirically.

Lets consider first the *public goods game*. In the public goods game, every player's contribution yields $\alpha_C = \alpha/K$, an ex ante fixed and a certain value of public good. Suppose first that player 1 did not explore. Then player 2s payoff equals

$$\alpha_C - (1-\sigma)c_2 = \frac{\alpha}{K} - (1-\sigma)c_2$$

if she explores and 0 otherwise. Notice that the parameter σ indicates player 2s concern for player 1 if the payoff of player 2 is lower than that of player 1. Parameter σ appears here since player 1 did not explore, and therefore player 2s payoff falls short of that of player 1 if the player 2 explores. In this case, the proportion of choice probabilities

[17] Since the novelty and focus in our model and experiment is the information externality channel, and the well-documented other-regarding preference channel generates encouragement irrespective of the particular model specification, we decided to adopt a highly simplified consequentialist preference framework although it is known to abstract from some important nuances of human behavior (Falk et al., 2008). The inequity aversion model of Fehr and Schmidt (1999) the reciprocity models (Charness and Rabin, 2002; Cox et al., 2007; Dufwenberg and Kirchsteiger, 2004; Falk and Fischbacher, 2006) and also the social esteem model of Ellingsen and Johannesson (2008) would make very similar predictions as the social welfare utility model.

[18] Charness and Rabin (2002) provide very convincing evidence consistent with $\rho > \sigma$. For simplicity, we abstract from the reciprocity parameter of the original three-parameter model. Moreover, like Fehr and Schmidt (1999), we allow even negative values of σ and ρ. Notice indeed that the inequity aversion model of Fehr and Schmidt (1999) is a special case of this model with the parameter for aversion for advantageous inequality $\beta = \rho$ and the parameter for aversion for disadvantageous inequality $\alpha = -\sigma$. The parameters is the model of Fehr and Schmidt (1999) are further constrained by $\alpha \geq \beta \geq 0$ or in other terms $-\sigma \geq \rho \geq 0$.

between exploration and no exploration equals

$$\frac{\sigma_2(yes \mid no)}{\sigma_2(no \mid no)} = exp\left((1/\mu)\left(\frac{\alpha}{K} - (1-\sigma)c_2\right)\right) \tag{7}$$

, and the log of the odds ratio between exploration and no exploration (7) is thus merely

$$(1/\mu)\left(\frac{\alpha}{K} - (1-\sigma)c_2\right). \tag{8}$$

Suppose next that player 1 did explore. Then player 2s payoff equals

$$2\alpha_C - (1-\rho)c_2 - \rho c_1 = 2\frac{\alpha}{K} - (1-\rho)c_2 - \rho c_1$$

if she also explores and

$$\alpha_C - \rho c_1 = \frac{\alpha}{K} - \rho c_1$$

if she does not. Notice that the parameter ρ indicates player 2s concern for player 1 if the payoff of player 2 is higher than that of the player 1. Parameter ρ appears here since player 1 explored and has a higher exploration cost than player 2. Therefore player 2's payoff is higher than that of player 1 whether player 2 explores or not. The log of the odds ratio between exploration and free-riding is thus

$$(1/\mu)\left(\frac{\alpha}{K} - (1-\rho)c_2\right). \tag{9}$$

In the public goods game, the only difference between expressions (8) and (9) is the behavioral other-regarding parameter terms in front of player two's exploration cost. Since $\rho > \sigma$, we thus establish that an *other-regarding* player two is more likely to explore if player one also explored. By contrast, a *selfish* player two's exploration decision (when $\sigma = \rho = 0$) remains unaffected by player one's exploration choice in the standard public goods game, i.e. the public goods game. In sum, we find that an *other-regarding encouragement effect* now appears even in the public goods game, in opposition to the analysis of section 2.1, provided that player two holds social preferences.

Consider next the *exploration game*. As before, let K denote the number of alternatives or possible locations where the public good is located. In case player one did not explore, player two's payoff to exploration equals

$$\frac{(1-\sigma)\alpha+\sigma\alpha}{K}-(1-\sigma)c_2=\frac{\alpha}{K}-(1-\sigma)c_2,$$

and the payoff to no exploration equals 0. The proportion of choice probabilities between exploration and no exploration then equals

$$\frac{\sigma_2(yes \mid no)}{\sigma_2(no \mid no)}=exp\left((1/\mu)\left(\frac{\alpha}{K}-(1-\sigma)c_2\right)\right) \tag{10}$$

and the log of the odds ratio between exploration and no exploration is merely

$$(1/\mu)\left(\frac{\alpha}{K}-(1-\sigma)c_2\right).$$

In case player one explored but did not find the treasure, player two's payoff to no exploration equals $-c_1\rho$, and her payoff to exploration now equals $\frac{\alpha}{K-1}-(1-\rho)c_2-c_1\rho$. Taking the log of the odds ratio between exploration and no exploration yields

$$(1/\mu)\left(\frac{\alpha}{K-1}-(1-\rho)c_2\right).$$

Since $\frac{\alpha}{K-1}>\frac{\alpha}{K}$, there is an encouragement effect independently of the treasure size even if player two does not hold social preferences, i.e. even if $\sigma=\rho=0$. This is coined as *the quantal response encouragement effect.* Moreover, if $\rho>\sigma\geq 0$, there is a further encouragement effect called *the other-regarding encouragement effect.*

In case player 1 explored and found the treasure, then player 2s payoff to no exploration equals $\alpha-c_1\rho$ and payoff to exploration equals $\alpha-(1-\rho)c_2-c_1\rho$. Then the log of the odds ratio between exploration and no exploration equals

$$(1/\mu)\left(\alpha-(1-\rho)c_2-c_1\rho\right)-(1/\mu)\left(\alpha-c_1\rho\right)=(1/\mu)\left(-(1-\rho)c_2\right),$$

Thus, in our more general model of sequential behavior for the public good, player 2 is expected to explore with a positive probability even when player 1 already explored and found the treasure. Depending on the parameter values, this probability of mistakenly exploring lies between zero and 50%. If we let μ tend to zero, player two's probability of mistakenly exploring tends to zero. To sum up, in case player 1 does not explore or explores without finding the treasure, respectively, player 2s odds ratio between exploring and not exploring is greater than one iff

$$\frac{\alpha}{(1-\sigma)K} - c_2 > 0, \tag{11}$$

$$\frac{\alpha}{(1-\rho)(K-1)} - c_2 > 0. \tag{12}$$

If a decision maker has $\sigma = 0 = \rho$, then she has no intrinsic other-regarding preference. Even in the case of a purely selfish player 2, there is, as mentioned in the previous subsection, an encouragement effect in the exploration game. To see this, take the difference between (12) and (11) when imposing $\sigma = 0 = \rho$. This yields

$$\frac{\alpha}{(K-1)} - \frac{\alpha}{K} > 0,$$

and since $K - 1 < K$ player 2s incentive to explore in the exploration game is higher when player 1 explored and did not find the treasure than when the player 1 did not explore. By contrast in the public goods game, a self-interested player 2s incentive to explore is constant and equal to $\frac{\alpha}{K} - c_2$, regardless of whether player 1 explores. There is no interaction between the two players' incentives to explore in the public goods game if players are self-interested.

The expressions (11) and (12) illustrate that as long as $1 > \rho \geq \sigma > 0$ the behavioral other-regarding preferences strengthen player 2s incentives to explore following player 1s exploration decision. However, if $\sigma < 0$, as is the case when player 2 is averse to disadvantageous inequality, player 1s lack of initiative undermines player 2s incentive to explore.[19] Then player 2 is actually less likely to explore than a self-interested player 2. In all cases, the other-regarding preference can be understood as a multiplier factor magnifying the way the incentives of player 2 depend on player 1s exploration choice. From a player 1s perspective, other-regarding preferences also magnify his strategic incentives to explore.

It is easy to see that regardless of individual preferences or game type, player 2s probability of exploration is always increasing in α and decreasing in c_2. Moreover, if

[19]Introducing reciprocity preferences would further strengthen this other-regarding encouragement effect. For simplicity, and to focus on the effects through the informational channel, we decided to model the other regarding encouragement effect by means of a simple consequentialist model.

both (11) and (12) hold, then the probability of exploration is decreasing in μ. Finally, the greater is the difference between ρ the greater is the behavioral encouragement effect. and σ.

Notice that other-regarding preferences may also backfire. If σ is negative, player 2 is inequity averse, and equation (8) illustrates that the probability of exploration by player 2 is lower than that of exploration by a self-regarding player 2. The more averse player 2 is to disadvantageous inequality, the less likely it is that she will explore when player 1 did not explore. Similarly player 1 who is averse to disadvantageous inequality, and expects player 2 never to explore, would prefer not exploring either. Again, this preference for withdrawing exploration is stronger for a more inequality averse player 1.[20]

2.2.3 Behavioral Predictions

A straightforward comparison between the QRE for the exploration and public goods game, now allowing for a more general model of individual preferences, yields three distinct sets of testable theoretical predictions.

The hypotheses consider the logit quantal response equilibrium for the exploration game (uncertainty) and voluntary public goods game (certainty), with players who care about their own payoff and potentially also about the payoff of their counterpart:

1. *Efficiency Hypothesis.* All things equal, aggregate equilibrium exploration will be weakly higher in the exploration game than in the public goods game. Further, individual exploration will be increasing in treasure size and decreasing in the cost of exploration in both exploration and public goods games.

2. *Other-regarding Encouragement Effect in the Public goods game.* In the public goods game, an other-regarding player 2 will be more likely to explore than a selfish player 2 following exploration by player 1. In the public goods game, social

[20] There is however evidence that the relationship may not be that straightforward, but rather other-regarding players may be inclined to contribute as player 2 and as player 1 and that this latter may be partly driven by biased expectations about player 2s contributions (See Altmann et al. (2008), Blanco et al. (2014), Miettinen et al. (2017)). In our experiment, participants play the game for several rounds against varying opponents and thus beliefs about the population behavior should tend towards the actual behavior.

preferences imply a positive complementarity between the two players' exploration decisions, giving rise to an other-regarding encouragement effect.

3. *Standard, Quantal response, and Other-regarding Encouragement Effect in the exploration game.* Compared to the public goods game, player 2 is always more likely to explore in the exploration game following exploration by player 1, even if she is purely selfish, and this encouragement effect is stronger if choices are not implemented perfectly but rather with error (in the sense of the logit QRE). If player 2 also holds other-regarding preferences, then these preferences further increase the likelihood that player 2 explores following player 1 exploration. In the exploration game, other-regarding preferences augment the standard encouragement effect with a *quantal response* and an *other-regarding* encouragement effect.

Intuitively, greater expected rewards incentivize individual and joint exploration. Furthermore, formal analysis distinguishes the *standard* encouragement effect of section 2.1 from two *behavioral* encouragement effects. An *other-regarding* encouragement effect occurs in the public goods game provided player 2 holds other-regarding preferences: player 2 reacts to exploration by player 1. A purely standard encouragement effect occurs in the exploration game provided player 2 is selfish and the players play their optimal strategies with certainty. A *quantal response* encouragement effect occurs for all treasure sizes in the exploration game provided player 2 is selfish.

3 Experimental Design and Data

We ran the experiment at the Cobe Lab (Aarhus University, Denmark), the Aalto Choice Tank (Aalto University, Helsinki, Finland), and the Centre for Experimental Studies and Research (BI Norwegian Business School, Oslo, Norway). In total 436 subjects were recruited using identical recruitment procedures.[21] Each subject completed a 10-

[21] To recruit our subjects, we used ORSEE (?) in Norway and SONA in Denmark and Finland. The recruitment text included information about the duration, location and incentives for both parts of the study, the online survey and laboratory experiment. Before running the experiment we calculated a sample size. We assumed a power of 80% and a standard size of 5% and to have a minimum detecatble effect size of 20% we need approximately 200 observation in each treatment. The minimum detectable effect size was chosen through rough calculations of effect sizes from the results using binary outcomes in Klor et al. (2014) and Steiger et al. (2014). See List et al. (2011) for a discussion of sample size calculations of experiments.

We considered a sample of around 400 participants large enough for the normal distribution and a good approximation for the t-distribution.

minute online survey at least 5 days before participating in the laboratory experiment. The laboratory session lasted on average 70 minutes. A 6.10 USD participation fee and subsequent earnings, which averaged 7 USD, were paid in private at the end of the laboratory session.[22]

3.1 Online Survey

After signing up to the two-part study, participants could enter the online survey directly. At the outset, participants faced five questions so as to create an anonymous personal identifier. Later, participants used this identifier to sign into the laboratory experiment. This procedure allowed us to ensure the anonymity of the participants when merging their answers from the survey with their answers from the laboratory experiment. In the online survey, we measured Social Value Orientation (SVO), risk preferences, and cognitive reasoning style. Social value orientation was measured using the SVO Slider Measure (Murphy et al., 2011), a six-item questionnaire where each question consists of a choice of one out of nine possible allocations of money between oneself and another anonymous participant. We used Experimental Currency Units (ECU) as currency in the online survey, and after the study was completed they were converted to the local currency, with an exchange rate of 30 ECU = 1.14 USD. At the start of the laboratory experiment we randomly selected in public one of the six SVO questions to be subject to payment. To measure risk aversion, we relied on two measures. The first measure was the Gneezy and Potters (1997) investment task. The participants were given 60 ECU and could invest any amount between null and 60 in a lottery with a 2/3 probability of getting nothing and 1/3 probability of winning two and a half times the amount invested. At the laboratory experiment we also publicly announced whether the investment task was a success or a failure in the laboratory. We publicly showed the participants three cards with letters A to C that we placed in an empty urn and shuffled. We invited one of the participants to draw one of the three cards from the urn. If the A card was drawn, each participants won two and a half times the amount she had invested. We complemented this risk measure with a hypothetical question asking the participant to rate his or her general risk taking on a scale from 0 to ten, with 0 being risk averse and

[22] If a participant could only complete the online survey he or she was paid half the show up fee of 3.05 USD. In Norway, the average total earnings in the experiment were 42.18 USD. The higher rate was applied in order to meet the average earnings requirements of the local laboratory.

10 being risk loving (Dohmen et al., 2011). To measure cognitive reasoning style, we used the Cognitive Reflection Test (Frederick, 2005), which consists of three questions, without incentives. Finally, we asked the participants about their gender. The full questionnaire is in the online appendix.

3.2 Laboratory Experiment

The laboratory experiment was an internet software programmed for the purpose of this experiment by Kristaps Dzonsons.[23] By means of game type specific codes on paper that we handed out, within each session participants were randomly assigned to one of the two game types (both the uncertain and the certain game treatments were run parallel in each session to ensure control for day-of-the-week or hour-of-the-day and other session effects, see Levitt and List (2011)). The software also randomly allocated each player one of the two player types. On a few occasions, very few students signed up. Here, we randomly assigned participants to player types within sessions and randomized game type played in these sessions. We control for this in the analysis of Section 4.4. All participants played 32 rounds of either the public goods or the exploration game, and as either player 1 or player 2. They all encountered four levels of treasure sizes; i.e. eight rounds of each treasure size. Most participants faced the treasure sizes in ascending order. To study order effects we let a few randomly drawn sessions face another treasure size order.

At the beginning of the experiment each participant received an endowment of 12000 points. The participants tally of points was visible and updated automatically while playing, according to the outcome of each game round. In the instructions, we informed the participants to collect as many points as possible across the game rounds. The number of rounds were however unknown to the participants. The first- and the second-mover with the most number of points in the public goods and exploration game respectively received a monetary prize of 13.68 USD.[24]

[23]see http://www.kcons.eu/

[24]Player 1 and 2 were thus incentivized to work as a team and each participant in a given player role was incentivized to compete against each of the other participants in his or her player role in the treatment with the same game type, but not with any of the participants in the opposing role that she was matched with. We expected these tournament incentives to induce more self-interested behavior compared to having participant's monetary compensation directly proportional to the tally of collected points, and thus a greater control of the self-interested encouragement threshold and a most favorable setting for the hypothesized non-monotonicity effect to arise.

Before the laboratory experiment started, we made two random draws in public to establish the rewards tied to choices made during the online survey.[25] Next, we used streamed video instructions to facilitate the understanding of the laboratory game. In a simple way, the video described how the game rounds proceeded, how tournament incentives operated, and how we carried out the matching.[26] The participants then logged in to the game using the game type specific codes and the anonymous personal identifiers that they had created at the onset of the online survey. The participants faced written instructions and control questions (they are available in the online appendix). At any time of the experiment, the participants could revisit the instructions. Upon having correctly completed the control questions, the first game round could start.

Each game round started with player 1 seeing four closed chests. In addition, the screen contained information of the participant's current number of points, the cost of exploring, the counterpart's cost of exploring (a first-mover's cost of exploring was always 300 points and that of player 2 was always 200 points), the size of the treasure (the treasure sizes were 500, 700, 1000, and 1400 points in the exploration game and 125, 175, 250, and 350 in the standard public good game), and the number of treasures left to explore (but no information about any other participants' current tally of points). The cost of exploration was kept constant throughout the session, and always higher for player 1 than for player 2 (300 versus 200 points). See Figure B in appendix B for an image of the decision screen. Player 1 knew that player 2 will observe his or her choice before making his/her own choice. Participants playing the public goods game knew that there is a treasure of known size in each chest. In the exploration game participants know that only one out of four chests contains a treasure. Each size of the treasure in the exploration game was four times the corresponding size in the public goods game, thus keeping the expected total treasure value and the myopic incentive to explore equal across game types. Player one then had to decide whether to pay to open a chest or not. Conditioned on player ones choice, player two had the same choice to make. When the second-mover had made his /her choice, both players received feedback on the outcome of the game. Before each new round of the game the participant was randomly re-matched with another participant of the opposite player

[25]This did not generate uncontrolled variation since we randomized treatments within each session.

[26]The video-instructions were 14 minutes long; visit the following link to view the video: https://dreambroker.com/channel/1ehcya5t/77qp05es

type in the treatment with the same game type. When a participant had completed 32 rounds, the screen informed him/her of his/her total number of points. Finally, we announced the anonymous personal identifiers of the first- and the second-mover winners publicly. The online appendix includes screen shots from the experiment website.

A research assistant at each lab asked the participants for their anonymous personal identifier, found the individual specific amount of ECU they had earned in both parts of the study, converted these into EUROs, added the show up fee of 5 EURO, and noted this on a separate piece of paper. Aarhus University then transferred the money to the participants bank accounts. At Aalto University and BI Norwegian Business School in Oslo, the participants then received the earnings immediately in cash.

3.3 Ethics and Registration of Study

Since the data from the study is never connected to identifying information, the project was not considered for full ethical review according to current legislation in Denmark, Finland, and Norway. At Aarhus University, Denmark, the project underwent an informal ethical review process by the Cobe Lab Ethical Advisory Board.[27] In addition, before running the analyses but after the experimental data collection, we registered the SPE part of our study design and suggested analysis at the Open Science Foundation (Registration name: Exploration in partnership). In some cases, the analysis below deviates from the originally foreseen specification. We then report this and comment on this in the limitations section.

3.4 Data

A total of 436 participants completed both the online survey and the laboratory session. Each participant completed 32 game rounds of play, implying a total of 13952 observations overall. Table 3.1 displays a summary of the main variables from the online survey across the two game types separately. We confirm in Table B that none of the variables differ significantly by type of game. About half of our sample consisted of women and participants were on average neither risk averse nor risk-loving. The average CRT score was 1.87 (sd:1.10), and 60% of the sample correctly answered all three questions

[27] http://bss.au.dk/en/cognition-and-behavior-lab/for-researchers/ethics-and-data-protection/

of the Cognitive Reflection Test.[28] The average SVO angle in our sample equaled 28° (sd: 13.09). Following Murphy et al. (2011), the average participant should thus be classified as prosocial. Table B in appendix B shows the randomization check. None of the observable variables differ by treatment.

Table 3.1: Descriptive statistics of online survey variables across game types

	n*	mean	median	sd	max	min
Public goods game						
Gender (1 if woman, 0 otherwise)	195	0.55	1	0.50	1	0
Risky investment choice	190	33.33	30	19.38	60	0
Risk question	193	5.90	6	2.15	10	2
CRTscore	197	1.89	2	1.10	3	0
Social Value Orientation	193	27.01	31	13.37	45	-9
Exploration game						
Gender (1 if woman, 0 otherwise)	238	0.55	1	0.50	1	0
Risky investment choice	229	33.72	30	19.17	60	0
Risk question	230	5.71	6	2.16	10	1
CRTscore	239	1.87	2	1.11	3	0
Social Value Orientation	230	28.01	33	13.09	61	-16
Observations	436					

*Some participants did not answer all the survey questions, the number of observations therefore vary.

[28] This percentage is higher compared to Frederick (2005), this can be due for example to learning, the fact that the CRT has become better known over time.

4 Experimental Results

Overall, the theoretical predictions of the quantal response model with other-regarding motivation best explain observed behaviors.[29] We proceeded in two steps. We first tested the SPE predictions, which we had registered at OSF. As our results only partially supported those predictions, we next took our QRE predictions to the data. To ease the readers' comprehension, we present in this section our results by theme or main take-away, and not, in the order of pre-registered plan vs. post-hoc analysis. This makes some of our results exploratory in nature, and they should be interpreted as such. Our results are most consistent with the assumptions that people imperfectly optimize and care not only about their own payoffs but also about others' payoffs. We empirically confirm the relevance of a behavioral model of sequential exploration for the public good.

4.1 Efficiency and Non-monotonicity

We define $C_{g,\imath}$ as a measure of whether individual $\imath$ explores or not during game round g ($C_{g,\imath} = 1$ if player i explores). Then we estimate how individual exploration behavior varies across game type (public goods versus exploration game) and player type (player 1 versus player 2), We use the following equation, where T denotes player type (taking the value 1 if the individual is a first-mover and zero otherwise), G denotes the game type (equaling 1 if the game is the exploration game and zero otherwise), and ε is the error term:

$$C_{g,\imath} = \alpha_{g,\imath} + \beta_1 T_{g,\imath} + \beta_2 G_{g,\imath} + \beta_3 T * G_{g,\imath} + \varepsilon_i, . \tag{13}$$

In Table 4.1 this equation is estimated as a linear probability model (LPM) in ordinary least squares (OLS) with robust standard errors clustered by individual. Column (1) presents the results without the interaction term. We find that overall there was no significant difference between first- and second-movers' exploration behavior. Further,

[29] This model deviates from the self-interested subgame perfect equilibrium by introducing three new parameters (noise parameter and two other-regarding preference parameters) and these additional degrees of freedom increase explanatory power by construction (See Miettinen et al. (2017) for instance). Yet, in the present experiment there are qualitative patterns in the data that are consistent with the QRE model with other-regarding motivation but not with QRE with self-interest nor with SPE with self-interest. Thus not only is explanatory power in statistical terms higher, but also the observed qualitative treatment effects and other qualitative patterns (such as the first-mover exploration rate that increases with treasure size and the approximately 20% exploration rate, independent of the treasure size, by the second-mover when the treasure has already been found by the first mover in the exploration game).

relative to the public goods game, individuals in the exploration game were on average approximately 23% more likely to explore (p<0.001). This result is consistent with the *Efficiency Hypothesis*, predicted by the both the SPE and QRE of our models.

Table 4.1: OLS: Differences in exploration across player types

	(1) Exploration	(2) Exploration
First mover	-0.031	-0.124***
	(0.022)	(0.032)
Exploration game	0.239***	0.155***
	(0.023)	(0.025)
First mover X Exploration game		0.166***
		(0.046)
Constant	0.405***	0.452***
	(0.017)	(0.017)
Adjusted R^2		
Observations	13728	13728

Robust standard errors clustered on individual

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Column (2) of Table 4.1 includes the interaction term. The first coefficient shows that in the public goods game, player 1 were on average about 13% less likely to explore than player 2. This finding reflects the fact that play 1s cost of exploration is higher than that of player 2. The first-mover thus has a larger myopic incentive to free ride. The second coefficient reveals that relative to the public goods game, a second-mover was 15% more likely to explore in the exploration game (p<0.001) than in the public goods game. The interaction term, however, suggests that whereas the second movers explore more than the first-movers in the public goods game, there is no such gap in the exploration game. These patterns are in line with the encouragement effect.[30]

[30] Changing the specification to logit does not change our results. If we pool the data across game rounds and cluster on session the results stay the same. See Table B and A.3 in the appendix B.

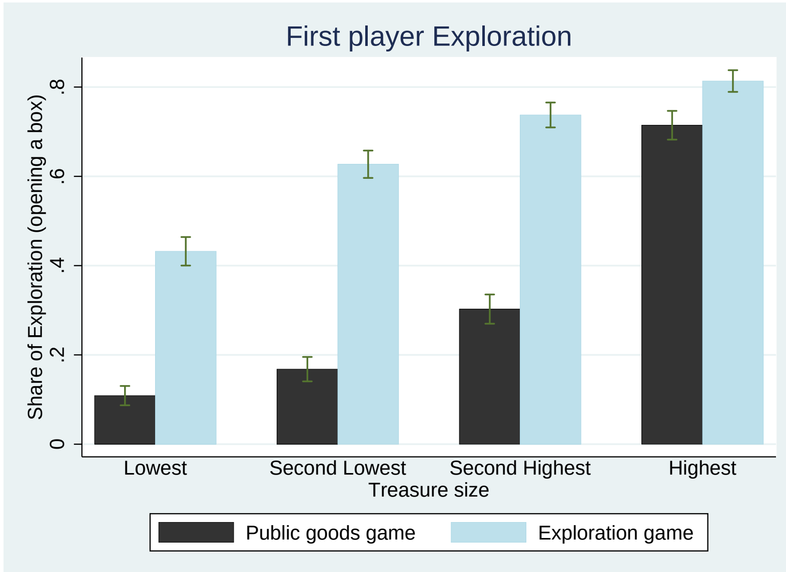

Figure 4.1: First player exploration

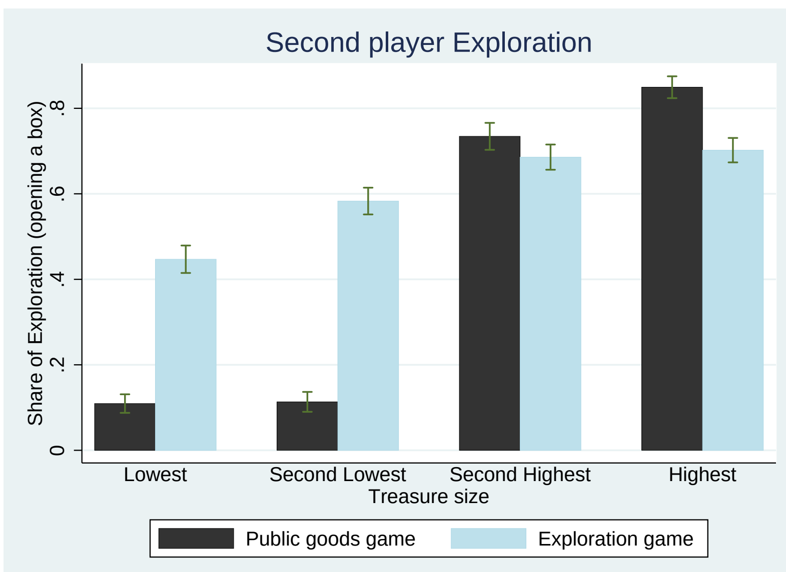

Figure 4.2: Second player exploration

Figures 4.1 and 4.2 present the average share of, respectively, exploration of player 1 and 2 by treasure size. These raw averages indicate that individual exploration was monotonically increasing in treasure size. Table A.4 and A.5 in the appendix B show the pre-registered analysis of the exploration gap between pairwise treasure sizes across game types for first and second player respectively. Figures 4.1 and 4.2 as well as Table A.4 and A.5 do not fully support the *Non-monotonicity Hypothesis*. Rather, greater rewards invoked more exploration. These results cast doubts on the relevance of the SPE game-theoretic predictions in our setting. They are, however, aligned with the theoretic predictions of the QRE model with other-regarding preferences for the exploration and public goods games.

In sum, our results empirically confirm the Efficiency Hypothesis. We find that aggregate exploration is significantly greater in the exploration than in the public goods game. Consistent with this Hypothesis, we also establish that relative to the public goods game, second players explore significantly more in the exploration game. However, we do not find indications of a non-monotonic relationship between exploration and treasure size. This implies that we do not find an encouragement effect in the very narrow meaning of the definition, i.e. according to the SPE model. This standard effect would predict a non-monotonicity and incentive reversal due to the range where the first-mover should explore in order to encourage the second mover to explore for a lower treasure size and free-ride when she knows that the second-mover's incentives to explore are sufficient. Next we probe further into understanding the encouragement effect.

4.2 Encouragement Effect

We first examine player 2s exploration behavior in the public goods game. Figure 4.3 reveals that the share of second players who explored was always greater if the first mover explored compared to not explored. This finding is consistent with the other-regarding encouragement effect, which we theoretically derived for the public goods game when individuals are other-regarding.

Indeed, our participants appear other-regarding given that 28 was the average value of the SVO angle in the relevant subsample, and thus the other-regarding encouragement effect for the representative agent is predicted to be either relevant or relevant for a vast majority of subjects. Notice, that our SVO measure is a unitary measure. The other-regarding QRE predictions, however, draw on a distinction in the regard for the other player when there is disadvantageous versus advantageous inequality. Therefore, we developed a protocol that delivers an estimate of ρ and σ for each individual, using the SVO six slider questions (See appendix C for a detailed description). Our findings suggest that for the vast majority of our participants, $\rho \geq 0.5$ and $\sigma > 0$. We lack variation in these parameters to be able to estimate differential effects of ρ and σ on exploration propensities.

These findings thus lends additional support to the predictions of our QRE model with other-regarding preferences.

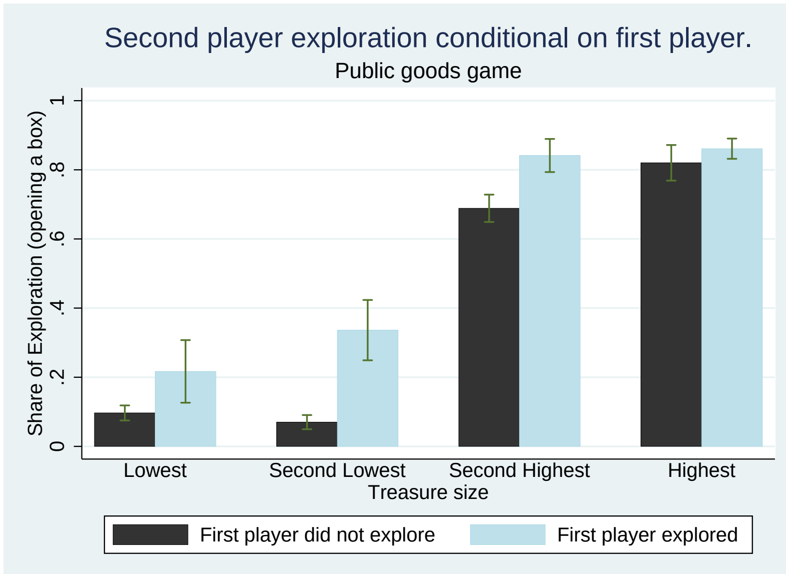

Figure 4.3: Public goods game: Second player exploration condition on first player exploration

We next turn to the analysis of player 2s exploration behavior in the exploration game. Figure 4.4 readily shows that those in role of player 2 were more likely to explore following exploration by player 1s than following no exploration. This result is consistent with the encouragement effect, expected to occur in both the SPE and QRE of the exploration game, i.e. the standard and the quantal response encouragement effect, respectively. However, given the prosocial character of the participants in our sample, these discrepancies in exploration decisions are possibly also driven by an other-regarding encouragement effect in favor of which we found evidence in the public goods game.

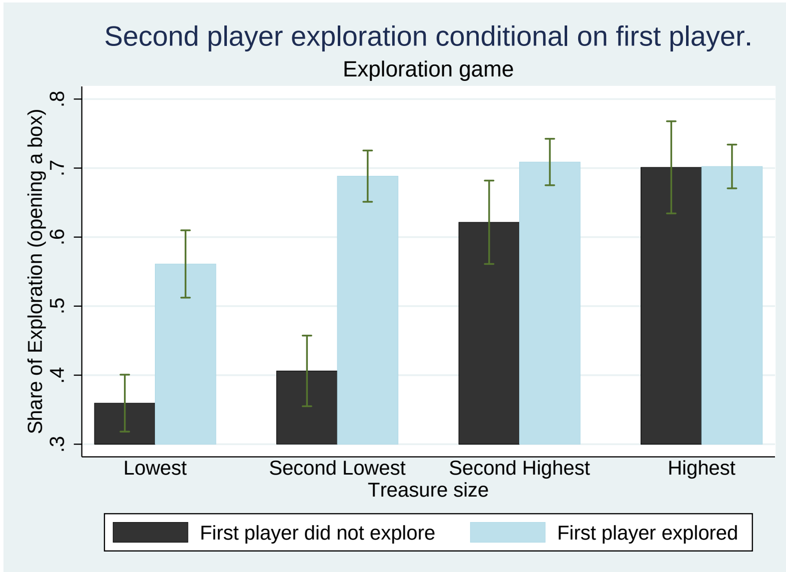

Figure 4.4: exploration game: Second player exploration conditional on first player exploration

We therefore disentangle the other-regarding encouragement effect from the standard and the quantal response effects. To this end, we again exploit the panel structure of our data, and estimate an equation of a similar form as equation (13). Now $C_{g,\imath}$ corresponds to the exploration decision of second-mover $\imath$ during game round g. Let G again denote the game type (equaling 1 if the game is the exploration game and zero otherwise) and T the exploration decision taken by player 1. Variable T is defined differently depending on which game type we consider. In the exploration game (eg) the variable takes on the value 1 if the first player explored but did not find a treasure (nf) and 0 if the first player did not explore (ne). In the public goods game (pg) the variable equals 1 if the first player explored (e) and 0 if the first player did not explore (ne). In the exploration game the difference in second-mover exploration when the first-mover did not find a treasure versus when the first-mover did not explore reflects a combination of all three encouragement effects. In the public goods game, the difference captures the other-regarding encouragement effect only. The estimated coefficient of the interaction term $G * T$ in equation (13) can now be interpreted as a measure of the sum of the standard and the quantal response (QR) effects:

$$Standard + QR = \overbrace{(C_{eg,\ nf} - C_{eg,\ ne}))}^{total} - \overbrace{(C_{pg,\ e} - C_{pg,\ ne})}^{other-regarding} \qquad (14)$$

Table 4.2 presents the results of the OLS regression estimates by treasure size. The first coefficient in each regression is an estimate of the other-regarding encouragement effect in the public goods game. This other-regarding encouragement effect is positive and significant, though it disappears for the highest treasure size. One plausible explanation for this is that for the highest treasure size, the self-interest motive to explore outweighs any behavioral considerations. The second coefficient in each regression captures the difference between second player exploration in the exploration game and the public goods game when the first player did not explore. The third coefficient estimate corresponds to the value of the standard and the quantal response encouragement effects. We find that this joint effect is statistically significant and positive.[31]

Table 4.2: Second player exploration.

	(1) Lowest	(2) 2nd lowest	(3) 2nd highest	(4) Highest
First player behavior	0.069*	0.243***	0.136***	-0.016
	(0.035)	(0.038)	(0.038)	(0.031)
Exploration game	0.285***	0.344***	-0.075	-0.173***
	(0.026)	(0.030)	(0.039)	(0.039)
Encouragement (interaction)	0.232***	0.229***	0.185***	0.249***
	(0.048)	(0.048)	(0.048)	(0.044)
Constant	0.103***	0.071***	0.695***	0.861***
	(0.012)	(0.011)	(0.026)	(0.025)
Adjusted R^2				
Observations	3190	3000	3003	3022

OLS has robust standard errors clustered on individual.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

4.3 Social Value Orientation, Risk Aversion and Cognitive Ability

The covariates we collected are orthogonal to the treatment status, and should not affect the results in the regressions. Including them in the regression does not change our results, see Table A.7 in appendix B. To examine whether there are heterogeneous effects, Table 4.3 shows the encouragement effect for the subsample of individualistic participants and Table 4.4 shows the encouragement effect for the prosocial part of the sample. Contrary to our initial expectations, we find that none of these individual

[31] If we pool the data across game rounds and cluster on session the results stay the same. See Table A.6 in the appendix B.

characteristics help us explain exploration.

Table 4.3: Second player exploration - individualistic players.

	(1) Lowest	(2) 2nd lowest	(3) 2nd highest	(4) Highest
First player behavior	0.015	0.125*	0.119*	-0.045
	(0.038)	(0.049)	(0.058)	(0.051)
Exploration game	0.311***	0.309***	0.014	-0.156*
	(0.046)	(0.055)	(0.068)	(0.065)
Encouragement (interaction)	0.247**	0.386***	0.120	0.297***
	(0.076)	(0.075)	(0.074)	(0.071)
Constant	0.092***	0.076***	0.701***	0.853***
	(0.017)	(0.019)	(0.047)	(0.044)
Adjusted R^2				
Observations	905	857	854	904

OLS has robust standard errors clustered on individual.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 4.4: Second player exploration - prosocial players.

	(1) Lowest	(2) 2nd lowest	(3) 2nd highest	(4) Highest
First player behavior	0.099*	0.280***	0.133**	-0.005
	(0.048)	(0.047)	(0.047)	(0.041)
Exploration game	0.291***	0.357***	-0.124*	-0.199***
	(0.033)	(0.037)	(0.049)	(0.050)
Encouragement (interaction)	0.239***	0.193**	0.223***	0.241***
	(0.063)	(0.059)	(0.060)	(0.057)
Constant	0.109***	0.072***	0.705***	0.872***
	(0.016)	(0.014)	(0.031)	(0.030)
Adjusted R^2				
Observations	2157	2025	2034	2007

OLS has robust standard errors clustered on individual.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 4.5: Exploration and SVO angle.

	(1) Lowest	(2) 2nd lowest	(3) 2nd highest	(4) Highest
Exploration game	0.318***	0.462***	0.026	-0.105**
	(0.036)	(0.033)	(0.041)	(0.033)
First mover	0.097	0.073	-0.172*	-0.109
	(0.067)	(0.064)	(0.083)	(0.070)
Social Value Orientation	0.003*	0.003*	0.001	0.000
	(0.001)	(0.001)	(0.002)	(0.001)
PlayerXSVO	-0.004	-0.001	-0.006	0.001
	(0.002)	(0.003)	(0.003)	(0.003)
GametypeXPlayerXSVO	0.001	-0.000	0.011***	0.005**
	(0.002)	(0.002)	(0.002)	(0.002)
Constant	0.028	0.034	0.655***	0.820***
	(0.040)	(0.035)	(0.063)	(0.046)
Adjusted R^2				
Observations	3344	3248	3328	3392

OLS has robust standard errors clustered on individual.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

We also performed a similar analysis for risk-aversion and cognitive ability but we failed to detect any difference between the risk-averse and the risk-neutral, one the one hand, and the deliberative and the intuitive thinkers, on the other hand.

4.4 Robustness and Further Analysis

To assess the robustness of our results, we conduct a number of additional tests. Firstly, we estimate equation 4.4 including dummy variables for each of the 32 rounds. Table A.8 in appendix B shows that results stay the same.

Secondly, in 13 out of the 36 sessions we had randomly assigned the game types between sessions instead of within sessions. To test that this does not affect our results we estimate the equation again using only the sample where we randomly assigned participants to game type within the session. This sub-sample comprises 334 participants. The results are qualitatively similar, see Table A.9 in the appendix B.

Thirdly, we tested whether the order in which we presented the treasure sizes affect exploration when shifting from one treasure size to another. On average, the first players facing an ascending order seem to explore between 12 % and 10 % more than first players facing another order. However, this order effect does not change our results regarding the gap of exploration between the exploration and public goods game when comparing treasure sizes, i.e. our main results remain. Please see Table A.10 in appendix B.

Fourthly, participants played eight rounds with each treasure size. To account for possible learning, we look at exploration in the last four rounds of play for each treasure size. This implies that we cut our sample in half. Table A.11 in appendix B show similar patterns as before with greater exploration as treasure size increased. Holding treasure size constant, there is greater exploration in the exploration game than in the public goods game. For second-players, a large gap in exploration between the exploration versus public goods game prevailed for the smaller treasure sizes only.

Finally, we take a final corollary result predicted by our QRE game-theoretic model to the data in an effort to further assess the relevance of this model in explaining observed behaviors. Theory predicts that with positive probability player 2 in the exploration game will explore even after the treasure was found by player 1. Figure 4.5 shows that indeed there was a small share of second-movers who explored even when the treasure had been found. This share equaled about 20% of participants and was, as predicted, constant across treasure sizes.

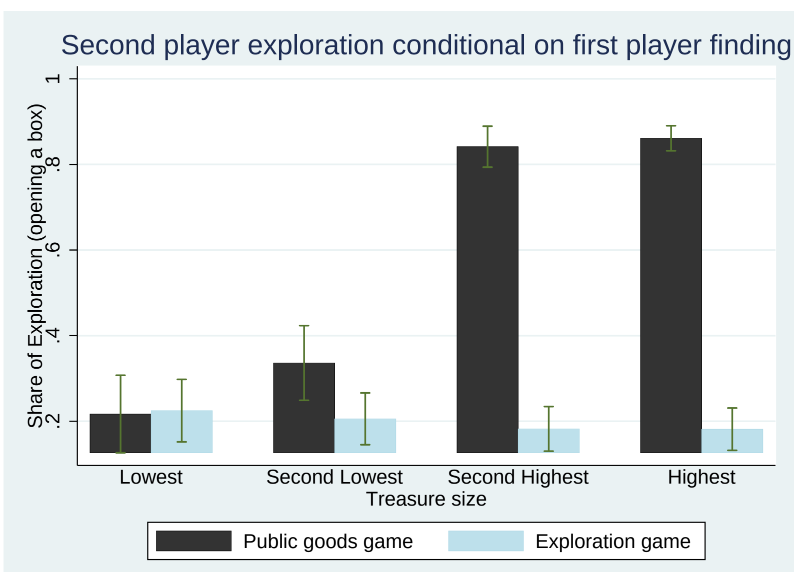

Figure 4.5: Second player exploration condition on first player finding

4.5 Limitations

When interpreting our results it should be noted that we had initially taken only to the data the SPE-predictions. Ex post, as those initial predictions were only partially validated empirically, we sought to extend our basic theoretic model and augmented its realism by allowing people to imperfectly optimize and hold other-regarding preferences, not only self-regarding preferences. As it turns out, the game-theoretic predictions that derive from this fuller version of our model best predict observed behaviors in the laboratory experiment. Indeed, Anderson et al. (1998), Goeree and Holt (2000), and Goeree and Holt (2001) illustrate the power of this latter approach and we re-express the recommendation of **?** and **?**, footnote 5 that researchers in future related theoretical and empirical work give more consideration to the QRE framework as an important theoretical benchmark.

Another limitation of our work is that our sample includes too little variation in the social value orientation (Murphy et al. (2011)) that was measured a week before the actual experiment, and thus too little variation in the implied social welfare utility parameters (Charness and Rabin (2002)), ρ and σ (see the appendix C). There was an abundance of subjects with a tendency to share the earnings fifty-fifty but few purely selfish or highly altruistic ones. This raises the importance of deriving the initial hypotheses within a framework with other-regarding preferences (and imperfect optimization); indeed, our representative subject has clearly significantly non-selfish preferences. The fact that there is too little variation limits our capacity to explore how heterogeneity along this dimension correlates with behavior in the exploration task and to test whether the variance is in line with the theoretical predictions.

Our exploration task is extremely simple, whereas the production function underlying innovation is admittedly anything but straightforward. With our design we are unable to separately identify the role that sensation seeking may have played in motivating exploration behaviors. That said, the simple and clear-cut model allows to decompose and carefully study the encouragement phenomenon. The experimental design served the purpose of providing clear answers for the particular hypotheses and research questions we were interested in. Our results are of course likely to be influenced by the particular context and design choices that we adopted, and further research is required to understand to what extent and when the results generalize. The setup also suggests further research where the encouragement effect and the level of uncertainty are varied in creativity research setups.

5 Conclusion

Using a novel experimental paradigm, we explored the factors that drive an individual's decision to interactively search for the public good - in particular, how willingness to search for the public good depends on exploration payoffs and uncertainty in the public goods' production process. Our focus is on the celebrated encouragement effect (first theoretically identified by Bolton and Harris (1999)) and the closely related incentive reversal effect (first pointed out by Winter (2009)).[32] We also study the robustness of these phenomena by extending it to a behavioral framework with imperfect optimization (McKelvey and Palfrey (1998)) and other-regarding preferences (Charness and Rabin (2002)).

We have shown that the behavioral patterns in the experimental data presented broadly conform to the theoretical predictions of our model of joint exploration under imperfect optimization and with other-regarding individuals, that exploration by player 1 motivates exploration by player 2. This encouragement effect is at play for small and large public benefits to successful exploration and in theory increase with the magnitude of the benefits. We theoretically derived that uncertainty in our game raises, rather than decreases, the aggregate level of exploration. Our experimental data robustly lends support to this efficiency hypothesis.

Our results underscore the role of uncertainty and learning in the provision of public goods. Learning or 'open innovation' induces a synergy between individuals' exploration decisions, which brings equilibrium innovation closer to the social optimum.

Our findings are also relevant to studying sequential team innovation, when individual effort cannot be observed by the principal and agents are rewarded based on joint output or success. As pointed out theoretically by Strausz (1999), Winter (2006), and Winter (2009), when (at least some) team members can observe other team members' effort, or the information structure can be at least partially designed, there are delicate incentive effects, i.e. encouragement and discouragement that need to be taken into

[32]See also Hörner and Skrzypacz (2016).

account when designing how the team operates.

The experiment of Klor et al. (2014) explicitly contrasted team production with simultaneous choices versus sequential choices. They found significant non-monotonicity effects in their sequential treatments. The difference between their design and ours is that they were not interested in team search per se but rather adopted a very explicit complementarity between inputs, an increasing returns to scale technology, and asked whether sequentially leads to incentive reversals, i.e. non-monotonicities. There was no exogenous uncertainty typical of any search process in their design. The key experimental variation in our study concerns precisely this certainty versus uncertainty (explorative nature) of returns to contribution to the public good. Yet, under subgame perfect Nash equilibrium and self-interest, the theoretical underpinnings are precisely the same. Thus, the fact that they observe a positive "incentive reversal" whereas we do not see much evidence of non-monotonicities suggests that the contextual differences influence behavior. Effects similar to ours can be observed in the experiment of Steiger et al. (2014) where experimental variation concerns the simultaneity vs. sequentially of choices, on the one hand, and the complementarity of effort, on the other hand.

The paper can also be seen as contributing to the understanding of the fundamental non-monotonicity aspect in the theoretical multi-player learning and experimentation literature in strategic two-arm bandit models (Hörner and Skrzypacz (2016), pp. 2-3), which lies at the heart of the encouragement effect theoretically discovered by Bolton and Harris (1999). In our setup, no exploration broadly corresponds to the safe arm and exploration to the risky arm. The first-mover can influence the second-mover probability of exploring (second-mover belief of high returns) by exploring. Our paper generally establishes the encouragement also empirically. Yet, the encouragement logic operates less perfectly and rationally than suggested by theory. Due to imperfect optimization, there is an encouragement effect not just around the belief threshold but rather independently of the parameter values.

Provided that our results are externally valid, the most direct out-of-sample implications of our results relate to situations in which motivating public goods provision are important concerns. Business leaders, for instance, who value opportunities for their employees to collaboratively work on projects that can transform society are well-advised to emphasize the inevitable uncertainty in these production processes, to find ways to enhance the overall value of the public good created, and invite contributions by social employees at a second (later) stage. These strategies can induce more efficient equilibrium outcomes.[33] These strategies can also be applied by public sector leaders. For instance, school principals who wish to encourage teachers to jointly search for approaches that effectively improve say parental engagement or community associations who wish to encourage their members to jointly search for approaches that enhance local social cohesion.

Interestingly, a rapidly rising share of experimentation for the public good actually occurs outside of mainstream organizations. More citizens than ever are voluntarily stepping up and jointly (openly) searching for novel ideas and solutions in a bid to make their societies more sustainable and more inclusive (Baldwin and Von Hippel (2011); Harhoff and Lakhani (2016)). Our research suggests that to enable these types of collective action it is recommended that citizens adequately appreciate in full the benefits of the public good. Also, by explicating the explorative character of these initiatives, citizens may well be more likely than less likely to contribute.

Let us finally discuss a few future related research paths that might prove particularly fruitful. The paper provides a complementary workhorse model to study some of the key questions instigated by the theoretical strategic experimentation literature in a simple setting. The novel experimental framework could, for instance, be used to study experimentation in a private goods setting as well. Generalizations to multiplayer teams, or endogenous ordering of exploration efforts seem straightforward. A setting where players have a common value for the good but where they receive private signals about the pay-off to exploration prior to exploring opens a bridge between the literature of exploration and social learning (herding). On the other hand, if the locations contain public or private goods of variant values, the links to the search literature become obvious.

It would also be of great interest to take steps away from the tightly controlled model-like laboratory settings toward more ecologically valid studies on creativity or innovativeness and to exogenously vary the uncertainty and the stakes and rewards related

[33]In their efforts to restore trust in businesses and straightforwardly build better businesses, many business leaders have sharpened their focus on purpose (Hollensbe et al. (2014)). They want their employees to have a prosocial impact and consequently feel more engaged at work (Bolino and Grant, 2016)

to the process of discovery. This class of studies encompasses both field experiments in collaboration with firms, non-profit organizations or public sector agencies, as well as more controlled studies in the laboratory using protocols established in creativity research (Osborn (1953), Amabile et al. (1986), Erat and Gneezy (2016)).

Bibliography

Admati, A. R. and M. Perry (1991). Joint projects without commitment. *The Review of Economic Studies 58*(2), 259–276.

Aghion, P. and J. Tirole (1997). Formal and real authority in organizations. *Journal of political economy 105*(1), 1–29.

Altmann, S., T. Dohmen, and M. Wibral (2008). Do the reciprocal trust less? *Economics Letters 99*(3), 454–457.

Amabile, T. M., B. A. Hennessey, and B. S. Grossman (1986). Social influences on creativity: The effects of contracted-for reward. *Journal of personality and social psychology 50*(1), 14.

Anderson, S. P., J. K. Goeree, and C. A. Holt (1998). A theoretical analysis of altruism and decision error in public goods games. *Journal of Public Economics 70*(2), 297–323.

Baldwin, C. and E. Von Hippel (2011). Modeling a paradigm shift: From producer innovation to user and open collaborative innovation. *Organization Science 22*(6), 1399–1417.

Becker, G. S. (1974). A theory of social interactions. *Journal of political economy 82*(6), 1063–1093.

Berg, J., J. Dickhaut, and K. McCabe (1995). Trust, reciprocity, and social history. *Games and economic behavior 10*(1), 122–142.

Blanco, M., D. Engelmann, A. K. Koch, and H.-T. Normann (2014). Preferences and beliefs in a sequential social dilemma: a within-subjects analysis. *Games and Economic Behavior 87*, 122–135.

Bolino, M. C. and A. M. Grant (2016). The bright side of being prosocial at work, and the dark side, too: a review and agenda for research on other-oriented motives, behavior, and impact in organizations. *Academy of Management Annals 10*(1), 599–670.

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

Bonatti, A. and J. Hörner (2011). Collaborating. *The American Economic Review 101*(2), 632–663.

Boudreau, K. J. and K. R. Lakhani (2016). Innovation experiments: Researching technical advance, knowledge production, and the design of supporting institutions. *Innovation Policy and the Economy 16*(1), 135–167.

Brown, M., C. J. Flinn, and A. Schotter (2011). Real-time search in the laboratory and the market. *The American Economic Review 101*(2), 948–974.

Brüggemann, J. and K. Bizer (2016). Laboratory experiments in innovation research: a methodological overview and a review of the current literature. *Journal of Innovation and Entrepreneurship 5*(1), 1–13.

Camerer, C. (2003). *Behavioral game theory: Experiments in strategic interaction.* Princeton University Press.

Cappelen, A. W., B.-A. Reme, E. Ø. Sørensen, and B. Tungodden (2015). Leadership and incentives. *Management Science 62*(7), 1944–1953.

Charness, G. and M. Rabin (2002). Understanding social preferences with simple tests. *The Quarterly Journal of Economics 117*(3), 817–869.

Clark, K. and M. Sefton (2001). The sequential prisoner's dilemma: evidence on reciprocation. *The economic journal 111*(468), 51–68.

Compte, O. and P. Jehiel (2003). Voluntary contributions to a joint project with asymmetric agents. *Journal of Economic Theory 112*(2), 334–342.

Compte, O. and P. Jehiel (2004). Gradualism in bargaining and contribution games. *The Review of Economic Studies 71*(4), 975–1000.

Cox, J. C., D. Friedman, and S. Gjerstad (2007). A tractable model of reciprocity and fairness. *Games and Economic Behavior 59*(1), 17–45.

Crawford, V. P., M. A. Costa-Gomes, and N. Iriberri (2013). Structural models of nonequilibrium strategic thinking: Theory, evidence, and applications. *Journal of Economic Literature 51*(1), 5–62.

Deck, C. A. and E. O. Kimbrough (2017). Experimenting with contests for experimentation. *Southern Economic Journal forthcoming*.

Dohmen, T., A. Falk, D. Huffman, U. Sunde, J. Schupp, and G. G. Wagner (2011). Individual risk attitudes: Measurement, determinants, and behavioral consequences. *Journal of the European Economic Association 9*(3), 522–550.

Dufwenberg, M. and G. Kirchsteiger (2004). A theory of sequential reciprocity. *Games and economic behavior 47*(2), 268–298.

Ederer, F. and G. Manso (2013). Is pay for performance detrimental to innovation? *Management Science 59*(7), 1496–1513.

Ellingsen, T. and M. Johannesson (2008). Pride and prejudice: The human side of incentive theory. *The American Economic Review 98*(3), 990–1008.

Erat, S. and U. Gneezy (2016). Incentives for creativity. *Experimental Economics 19*(2), 269–280.

Erev, I. and E. Haruvy (2013). Learning and the economics of small decisions. *The handbook of experimental economics 2*.

Falk, A., E. Fehr, and U. Fischbacher (2003). On the nature of fair behavior. *Economic Inquiry 41*(1), 20–26.

Falk, A., E. Fehr, and U. Fischbacher (2008). Testing theories of fairnessintentions matter. *Games and Economic Behavior 62*(1), 287–303.

Falk, A. and U. Fischbacher (2006). A theory of reciprocity. *Games and economic behavior 54*(2), 293–315.

Fehr, E. and K. M. Schmidt (1999). A theory of fairness, competition, and cooperation. *The quarterly journal of economics 114*(3), 817–868.

Fershtman, C. and A. Rubinstein (1997). A simple model of equilibrium in search procedures. *journal of economic theory 72*(2), 432–441.

Frederick, S. (2005). Cognitive reflection and decision making. *The Journal of Economic Perspectives 19*(4), 25–42.

Gilboa, I., A. Postlewaite, L. Samuelson, and D. Schmeidler (2014). Economic models as analogies. *The Economic Journal 124*(578), F513–F533.

Gneezy, U. and J. Potters (1997). An experiment on risk taking and evaluation periods. *The Quarterly Journal of Economics*, 631–645.

Goeree, J. K. and C. A. Holt (2000). Asymmetric inequality aversion and noisy behavior in alternating-offer bargaining games. *European Economic Review 44*(4), 1079–1089.

Goeree, J. K. and C. A. Holt (2001). Ten little treasures of game theory and ten intuitive contradictions. *American Economic Review*, 1402–1422.

Goeree, J. K., C. A. Holt, and T. R. Palfrey (2010). Quantal response equilibria. In *Behavioural and Experimental Economics*, pp. 234–242. Springer.

Grüne-Yanoff, T. (2007). Bounded rationality. *Philosophy Compass 2*(3), 534–563.

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*(5), 1023–1042.

Halac, M., N. Kartik, and Q. Liu (2016). Optimal contracts for experimentation. *The Review of Economic Studies 83*(3), 1040–1091.

Harhoff, D. and K. R. Lakhani (2016). *Revolutionizing innovation: Users, communities, and open innovation.* MIT Press.

Hollensbe, E., C. Wookey, L. Hickey, G. George, and C. V. Nichols (2014). Organizations with purpose. *Academy of Management Journal 57*(5), 1227–1234.

Holmstrom, B. (1982). Moral hazard in teams. *The Bell Journal of Economics*, 324–340.

Hörner, J. and A. Skrzypacz (2016). Learning, experimentation and information design. Technical report, Mimeo.

Klor, E. F., S. Kube, E. Winter, et al. (2014). Can higher rewards lead to less effort? incentive reversal in teams. *Journal of Economic Behavior & Organization 97*, 72–83.

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*(5), 793–818.

Levitt, S. D. and J. A. List (2011). Was there really a hawthorne effect at the hawthorne plant? an analysis of the original illumination experiments. *American Economic Journal: Applied Economics 3*(1), 224–238.

Lindbeck, A. and J. W. Weibull (1988). Altruism and time consistency: the economics of fait accompli. *Journal of Political Economy 96*(6), 1165–1182.

List, J., S. Sadoff, and M. Wagner (2011). So you want to run an experiment, now what? an introduction to optimal sample arrangements. *Experimental Economics 14*, 439–457.

Mäki, U. (2005). Models are experiments, experiments are models. *Journal of Economic Methodology 12*(2), 303–315.

McBride, M. (2006). Discrete public goods under threshold uncertainty. *Journal of Public Economics 90*(6), 1181–1199.

McCall, J. J. (1970). Economics of information and job search. *The Quarterly Journal of Economics*, 113–126.

McKelvey, R. D. and T. R. Palfrey (1998). Quantal response equilibria for extensive form games. *Experimental economics 1*(1), 9–41.

Miettinen, T., M. Kosfeld, E. Fehr, and J. W. Weibull (2017). Revealed preferences in a sequential prisoners' dilemma: A horse-race between five utility functions. Technical report, CESifo Working Paper.

Murphy, R. O., K. A. Ackermann, and M. Handgraaf (2011). Measuring social value orientation. *Judgment and Decision Making 6*(8), 771–781.

Olson, M. (1965). *The logic of collective action*, Volume 124. Harvard University Press.

Osborn, A. F. (1953). Applied imagination.

Palfrey, T. R. and H. Rosenthal (1991). Testing for effects of cheap talk in a public goods game with private information. *Games and economic behavior 3*(2), 183–220.

Plott, C. R. and V. L. Smith (2008). *Handbook of experimental economics results*, Volume 1. Elsevier.

Samuelson, L. (2005). Economic theory and experimental economics. *Journal of Economic Literature 43*(1), 65–107.

Steiger, E.-M. et al. (2014). See no evil: Information chains and reciprocity. *Journal of Public Economics 109*, 1–12.

Stigler, G. J. (1961). The economics of information. *Journal of political economy 69*(3), 213–225.

Strausz, R. (1999). Efficiency in sequential partnerships. *Journal of Economic Theory 85*(1), 140–156.

Winter, E. (2006). Optimal incentives for sequential production processes. *The RAND Journal of Economics 37*(2), 376–390.

Winter, E. (2009). Incentive reversal. *American Economic Journal: Microeconomics 1*(2), 133–147.

A Appendix: QRE of the public goods and exploration game

We now derive the quantal response equilibria for the exploration game and public goods game, reasoning by backward induction.

exploration game: We start by analysing player two's behavior. There are three possible cases to consider: (i) player one did not explore, (ii) player one did explore but did not find the treasure, (iii) player one did explore and found the treasure.

In case (i) [no exploration by player one], player two's payoff to no exploration is 0 and to exploration is

$$\frac{\alpha}{K} - c_2.$$

Thus the probability of exploring equals

$$\sigma_2(yes \mid no) = \frac{exp\left((1/\mu)\left(\frac{\alpha}{K} - c_2\right)\right)}{1 + exp\left((1/\mu)\left(\frac{\alpha}{K} - c_2\right)\right)} \tag{15}$$

and the proportion of choice probabilities between exploration and no exploration equals

$$\frac{\sigma_2(yes \mid no)}{\sigma_2(no \mid no)} = exp\left((1/\mu)\left(\frac{\alpha}{K} - c_2\right)\right) \tag{16}$$

and the log of the odds ratio is thus merely

$$\left((1/\mu)\left(\frac{\alpha}{K} - c_2\right)\right). \tag{17}$$

It is easy to see that the probability of exploring is increasing in α and decreasing in c_2. Moreover, if $\frac{\alpha}{K} - c_2 > 0$ then $\sigma_2(yes)$ is decreasing in μ. In case (ii) [failed exploration by player one], the payoff to no exploration still equals 0 but the payoff to exploration now equals $\frac{\alpha}{K-1} - c_2$ which is higher than $\frac{\alpha}{K} - c_2$. The implied probability of exploration now equals

$$\sigma_2(yes \mid yes, fail) = \frac{exp\left((1/\mu)\left(\frac{\alpha}{K-1} - c_2\right)\right)}{1 + exp\left((1/\mu)\left(\frac{\alpha}{K-1} - c_2\right)\right)} \tag{18}$$

which is larger than $\sigma_2(yes \mid no)$. Moreover

$$\frac{\sigma_2(yes \mid yes, fail)}{\sigma_2(no \mid yes, fail)} = exp\left((1/\mu)\left(\frac{\alpha}{K-1} - c_2\right)\right) \tag{19}$$

and the log of the odds ratio is thus merely

$$log\left(\frac{\sigma_2(yes \mid yes, fail)}{\sigma_2(no \mid yes, fail)}\right) = \left((1/\mu)\left(\frac{\alpha}{K-1} - c_2\right)\right). \tag{20}$$

The difference between expressions (.) and (.) reflect a positive encouragement effect

$$\frac{1}{\mu}(\frac{\alpha}{K(K-1)}. \tag{21}$$

This encouragement effect is increasing in the treasure size and decreasing in μ.

In case (iii) [successful exploration by player one], the payoff to no exploration equals α and the payoff to exploration equals $\alpha - c_2$. Thus the choice probability equals

$$\sigma_2(yes \mid yes, succeed) = \frac{exp\left((1/\mu)(\alpha - c_2)\right)}{exp\left((1/\mu)(\alpha)\right) + exp\left((1/\mu)(\alpha - c_2)\right)},$$

and the proportion of choice probabilities between exploration and no exploration equals

$$\frac{\sigma_2(yes \mid yes, succeed)}{\sigma_2(no \mid yes, succeed)} = \frac{exp\left((1/\mu)\left(\alpha - c_2\right)\right)}{exp\left((1/\mu)\left(\alpha\right)\right)},$$

and the odds ratio between those probabilities is thus

$$log\left(\frac{\sigma_2(yes \mid yes, succeed)}{\sigma_2(no \mid yes, succeed)}\right) = \left((1/\mu)\left(\alpha - c_2\right)\right) - \left((1/\mu)\left(\alpha\right)\right) = -\left(1/\mu\right)c_2,$$

Thus, the model predicts that even when player one did explore and found the treasure, player two explores with a positive probability. The probability of exploration is predicted to be below 50% and if we let μ tend to zero, the probability of mistakenly exploring tends to zero.

Let us next consider player one's incentives to explore. The payoff to no exploration equals

$$0 + \sigma_2(yes \mid no)[\alpha/K].$$

The payoff to exploration equals

$$\frac{\alpha}{K} - c_1 + \frac{K-1}{K}\sigma_2(yes \mid yes, fail)[\alpha/(K-1)] = \frac{\alpha}{K}(1 + \sigma_2(yes \mid yes, fail)) - c_1$$

Thus, the log of the odds ratio of exploration probabilities of the first mover equals

$$log\left(\frac{\sigma_1(yes)}{\sigma_1(no)}\right) = (1/\mu)\left(\frac{\alpha}{K}(1 + \sigma_2(yes \mid yes, fail) - \sigma_2(yes \mid no)) - c_1\right).$$

Since $\sigma_2(yes \mid yes, fail)) - \sigma_2(yes \mid no) > 0$, the player one will be more likely to explore even when

$$\frac{\alpha}{K} - c_1 < 0$$

and thus even when his myopic, self-interested preferences do not warrant exploration, or equivalently when he has no private incentives to explore in the public goods game . Moreover, in the present model where we allow for imperfectly rational expectations, player one can have dynamic incentives to explore even when $max\{\frac{c_1 K}{1+\delta}, c_2(K-1)\} < \alpha < Xc_2$ [the condition for the encouragement effect to be realised in the SPE] does not hold. When the encouragement effect outweighs a strictly negative private, shortsighted benefit-cost calculus, $\frac{\alpha}{K} - c_1 < 0$, then the likelihood that player one explores is increasing in treasure size.

Public goods Game: Again, reasoning by backward induction, we first consider player two's exploration decision. There now exist only two possible cases: (i) Player one does not explore or contribute; (ii) Player one contributes to the production of the public good.

In case (i) [no exploration by player one], player two's probability of exploration equals

$$\sigma_2(yes \mid no) = \frac{exp\left((1/\mu)\left(\frac{\alpha}{K} - c_2\right)\right)}{1 + exp\left((1/\mu)\left(\frac{\alpha}{K} - c_2\right)\right)}$$

In case (ii) [player one explore], player two's probality of exploration equals

$$\sigma_2(yes \mid yes) = \frac{exp\left((1/\mu)\left(2\alpha/K - c_2\right)\right)}{exp\left((1/\mu)\left(\alpha/K\right)\right) + exp\left((1/\mu)\left(2\alpha/K - c_2\right)\right)}$$

There is thus a small encouragement effect even in the public goods game (in the QRE with binary actions if payoffs to both actions are increased by a constant, then the choice probability of the higher payoff choice gets more probability mass due to the convexity of the exp-operator). This effect is yet of smaller magnitude than in the exploration game.

Proposition. QRE-Predictions. Second-mover.

- The Probability of exploration increases in treasure size.
- The probability of exploration is identical in the public goods game and in the exploration game if there is no explored alternative that did not contain a treasure. In reverse, the probability of exploration is lower in the public goods game than in the exploration game where there is an explored alternative that did not contain a treasure.
- The probability of exploration is positive even if there is no treasure left.

Proposition. QRE-Predictions. First-mover.

- Probability of exploration increases in treasure size.

Proposition. QRE-Predictions. Encouragement effect.

- Second-mover probability of exploration is higher if the first-mover explored alternative that did not contain a treasure than if the first-mover did not explore an alternative (compare to Scond-mover bullet point two above)
- This encouragement effect is increasing in treasure size.
- Probability of first-mover exploration increases more in treasure size than second-mover exploration.

B Appendix: Additional experimental results

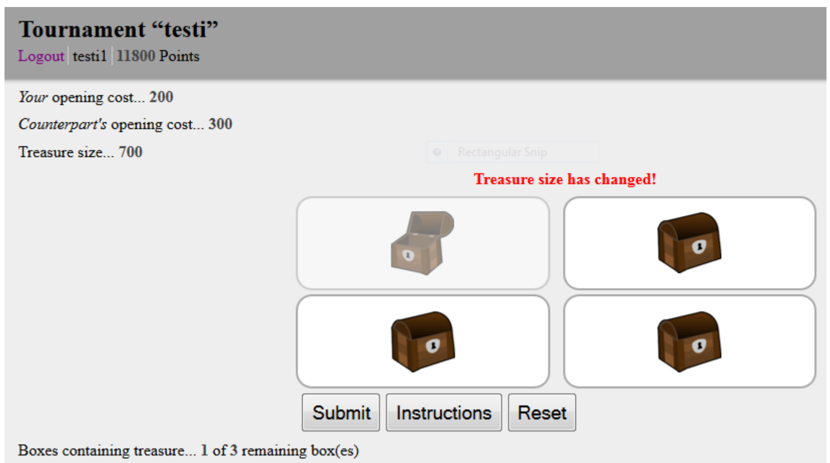

Table A.1: Differences between game types. Randomization check.

	(1) Female	(2) Risk question	(3) Cognitive reflection task	(4) Social Value Orientation
Exploration game	0.002	-0.154	-0.008	1.431
	(0.049)	(0.212)	(0.048)	(1.304)
Constant	0.550***	5.881***	0.605***	26.654***
	(0.036)	(0.158)	(0.036)	(0.972)
Adjusted R^2	-0.002	-0.001	-0.002	0.000
Observations	428	416	416	416

OLS regressions.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table A.2: Logit: Differences in exploration across player types.

	(1) Certainty game	(2) Uncertainty game
First mover	-0.033	-0.163***
	(0.029)	(0.041)
Exploration game	0.311***	0.178***
	(0.031)	(0.031)
First mover X Exploration game		0.244***
		(0.053)
Adjusted R^2		
Observations	13728	13728

Robust standard errors clustered on individual

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table A.3: OLS: Differences in exploration across player types

	(1) Exploration	(2) Exploration
First mover	-0.031	-0.133***
	(0.021)	(0.023)
Exploration game	0.246***	0.155***
	(0.027)	(0.025)
First mover X Uncertainty game		0.183***
		(0.031)
Constant	0.401***	0.452***
	(0.019)	(0.018)
Adjusted R^2	0.061	0.069
Observations	13728	13728

Robust standard errors clustered on session

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table A.4: First players. Pairwise OLS comparison of sizes of treasure and game types

	(1) Lowest/2nd lowest	(2) 2nd lowest/2nd highest	(3) 2nd highest/highest
Treasure size	0.051*	0.136***	0.416***
	(0.023)	(0.034)	(0.043)
Uncertainty game	0.295***	0.422***	0.420***
	(0.048)	(0.063)	(0.052)
Treasure size x Type of game	0.144***	-0.028	-0.339***
	(0.037)	(0.043)	(0.049)
Constant	0.125***	0.186***	0.308***
	(0.029)	(0.040)	(0.041)
Adjusted R^2			
Observations	3456	3448	3520

OLS has robust standard errors clustered on individual.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table A.5: Second players. Pairwise OLS comparison of sizes of treasure and game types

	(1) Lowest/2nd lowest	(2) 2nd lowest/2nd highest	(3) 2nd highest/highest
Treasure size	0.001	0.633***	0.114***
	(0.020)	(0.039)	(0.028)
Exploration game	0.334***	0.479***	-0.050
	(0.037)	(0.034)	(0.042)
Treasure size x Type of game	0.137***	-0.530***	-0.097**
	(0.030)	(0.046)	(0.035)
Constant	0.110***	0.103***	0.736***
	(0.020)	(0.022)	(0.036)
Adjusted R^2			
Observations	3408	3400	3472

OLS has robust standard errors clustered on individual.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

C Appendix: Deriving ρ and σ using the SVO sliders

What can be inferred about ρ and σ when looking at the choices in the six sliders tasks (see Figure A.1)? In Figure A.2, decision maker's own monetary compensation is measured along the horizontal axis, and the other's monetary compensation is measured along the vertical line. Each of the 6 tasks is presented as a line segment in Figure A.2. The marginal rate of substitution captures the individual rate at which the decision maker is indifferent between giving up a marginal amount against higher income for the other. Below the 45-degree line through the origin the decision maker's own payoff is higher than that of the other. Therefore the marginal rate of substitution between own earnings m against those of the other y equals

$$MRS(m, y) = -\frac{(1-\rho)}{\rho},$$

and above the 45-degree line where $y > m$,

$$MRS(m, y) = -\frac{(1-\sigma)}{\sigma}.$$

Thus when estimating σ from the SVO choice data, we should focus on tasks 3, 4, and 5 where at least a fraction of the corresponding line segment lies above the 45-degree

Table A.6: Second players. Pairwise pooled OLS comparison of sizes of treasure and game types

	(1) Lowest/2nd lowest	(2) 2nd lowest/2nd highest	(3) 2nd highest/highest
Treasure size	0.004	0.623***	0.114***
	(0.017)	(0.034)	(0.024)
Exploration game	0.338***	0.470***	-0.050
	(0.038)	(0.033)	(0.039)
Treasure size x Type of game	0.131***	-0.520***	-0.096**
	(0.032)	(0.046)	(0.030)
Constant	0.109***	0.112***	0.736***
	(0.018)	(0.020)	(0.033)
Adjusted R^2	0.193	0.216	0.019
Observations	3408	3400	3472

OLS has robust standard errors clustered on individual.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table A.7: Second players. Encouragement effect CG and UG.

	(1) Lowest	(2) 2nd lowest	(3) 2nd highest	(4) Highest
First player behavior	0.071*	0.241***	0.131***	-0.018
	(0.035)	(0.038)	(0.038)	(0.032)
Exploration game	0.290***	0.341***	-0.082*	-0.185***
	(0.027)	(0.030)	(0.040)	(0.039)
Encouragement (interaction)	0.232***	0.241***	0.184***	0.256***
	(0.049)	(0.049)	(0.048)	(0.045)
Risk question	-0.008	-0.001	-0.011	-0.008
	(0.006)	(0.005)	(0.006)	(0.005)
Cognitive reflection task	0.036	0.021	0.005	-0.003
	(0.024)	(0.021)	(0.025)	(0.021)
Social Value Orientation	0.001	0.001	-0.001	0.000
	(0.001)	(0.001)	(0.001)	(0.001)
Constant	0.092	0.030	0.777***	0.912***
	(0.051)	(0.042)	(0.053)	(0.042)
Adjusted R^2				
Observations	3093	2909	2911	2937

OLS has robust standard errors clustered on individual.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table A.8: First players: Pairwise comparison of size of treasure and game types.

	(1) Lowest/2nd lowest	(2) 2nd lowest/2nd highest	(3) 2nd highest/highest
Exploration game	0.279***	0.432***	0.422***
	(5.32)	(7.35)	(8.41)
Treasure size	0.105***	0.146***	0.462***
	(3.59)	(4.02)	(8.56)
Treasure size x Type of game	0.153***	-0.0342	-0.340***
	(4.08)	(-0.77)	(-6.81)
Observations	3392	3384	3456
Adjusted R^2			

OLS has robust standard errors clustered on individual.

32 dummies for round of the game are surpressed in the figure.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table A.9: Only within assignment: Pairwise comparison of size of treasure.

	(1) Lowest/2nd lowest	(2) 2nd lowest/2nd highest	(3) 2nd highest/highest
Exploration game	0.242***	0.408***	0.398***
	(0.062)	(0.067)	(0.058)
Treasure size	0.039	0.156***	0.405***
	(0.025)	(0.038)	(0.048)
Treasure size x Type of game	0.169***	-0.037	-0.334***
	(0.043)	(0.051)	(0.056)
Constant	0.142***	0.183***	0.324***
	(0.037)	(0.042)	(0.047)
Adjusted R^2			
Observations	2656	2656	2656

OLS has robust standard errors clustered on individual.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table A.10: Controlling for order: Pairwise comparison of size of treasure.

	(1) Lowest/2nd lowest	(2) 2nd lowest/2nd highest	(3) 2nd highest/highest
Exploration game	0.285***	0.433***	0.422***
	(0.050)	(0.057)	(0.050)
Treasure size	0.048*	0.143***	0.413***
	(0.023)	(0.035)	(0.044)
Treasure size x Type of game	0.148***	-0.032	-0.337***
	(0.037)	(0.044)	(0.050)
ascendingorder	0.124***	0.125**	0.104*
	(0.034)	(0.042)	(0.040)
Constant	0.052	0.098*	0.242***
	(0.034)	(0.043)	(0.043)
Adjusted R^2			
Observations	3392	3384	3456

OLS has robust standard errors clustered on individual.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table A.11: 4 last rounds: Pairwise comparison of size of treasure.

	(1) Lowest/2nd lowest	(2) 2nd lowest/2nd highest	(3) 2nd highest/highest
Exploration game	0.315***	0.449***	0.449***
	(0.048)	(0.049)	(0.052)
Treasure size	0.055*	0.136***	0.452***
	(0.026)	(0.039)	(0.048)
Treasure size x Type of game	0.123**	-0.006	-0.387***
	(0.044)	(0.051)	(0.055)
Constant	0.099***	0.145***	0.277***
	(0.027)	(0.032)	(0.042)
Adjusted R^2			
Observations	1696	1692	1728

OLS has robust standard errors clustered on individual.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

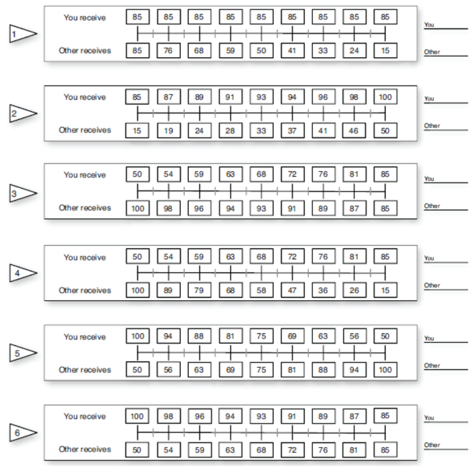

Figure A.1: SVO slider task.

line in Figure A.2. If a decision maker chooses an allocation strictly above the 45 degree line all these tasks, then the estimate satisfies $MRS \geq -3/7$, .i.e. $\sigma < 7/10$. If the decision maker chooses an allocation at the 45-degree line for slider task 3, but above the 45-degree line in tasks 4 and 5, then $\sigma \geq 1/2$. Otherwise $\sigma < 1/2$.

By a similar argument, all slider tasks apart from task 3 contribute to the estimation of ρ. If choices associated with sliders 4 and 5 at extreme south-east , then $\rho < 3/10$, otherwise $\rho \geq 3/10$. In this latter case, if moreover slider 5 lies at extreme south-east, then $1/2 > \rho \geq 3/10$. Otherwise $\rho \geq 1/2$. If slider 1 is in extreme south, then $\rho < 0$. If moreover slider 2 is in extreme south-west, then $\rho \leq -3/10$.

The connection between the exploration task and the slider tasks is as follows.

- The further to the northwest from the 45-degree line are sliders 3, 4, and 5, the more willing is player two to explore when the first-mover has not explored, with slider 3 to the extreme northwest indicating willingness to sacrifice and explore even for low treasure sizes and without first-mover exploration.

- The further away from extreme south and east are the sliders 4, 5, and 6 below the 45-degree line, the more willing is the player to explore in reaction to first-mover exploration. When sliders 1 and 2 are towards the south (and sliders 4, 5, and 6 are to the extreme south-east), a player prefers free-riding on the other's exploration effort.

Let us then consider first-mover incentives to explore. The encouragement effect is defined as a positive effect of first-mover exploration on second-mover probability of

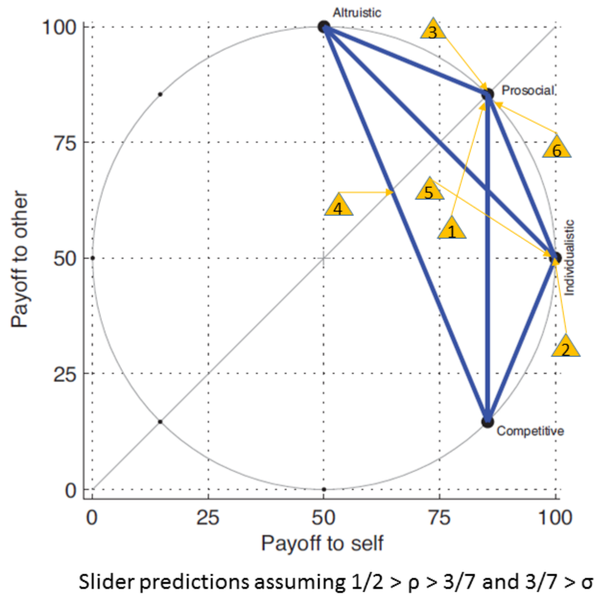

Figure A.2: SVO slider predictions.

exploration. Keep in mind that the analysis above concludes that for a self-interested second-mover, there is an encouragement effect in the exploration game but not in the public goods game. The encouragement effect appears in the public goods game if the second-mover is other-regarding, $1 > \rho \geq \sigma > 0$. Moreover, the behavioral other-regarding motivation always strengthens the incentives to explore. Thus from the first-mover perspective, the other-regarding motivation magnifies the strategic incentives to explore.

Let us first assume that both the first-mover and the second-mover are self-interested and they perfectly implement their optimal strategies and expect each other to do so (self-interest & subgame perfect equilibrium). Then there is never an encouragement effect in the public goods game but there is an encouragement effect in the exploration game if the conditions in Proposition 1 hold, that is when the treasure size is at the second-lowest level of 700 in our experiment.

Suppose then that the first mover and the second-mover are self-interested and they imperfectly implement their optimal strategies and rather choose according to the QRE so that the log-odds of the choice probabilities are as displayed in Section X. Then there is again no encouragement effect in the public goods game. Yet an encouragement effect appears in the exploration game for all treasure sizes, not just the second-lowest one.

Suppose then that the first-mover is self-interested but the second-mover is other-regarding and both implement their optimal strategies according to the QRE (and the first-mover knows that the second-mover is other-regarding). Then there is an encouragement effect both in the public goods game and in the exploration game for all treasure sizes. Yet, the encouragement effect is stronger in the exploration game.

If the first-mover is other-regarding and inequity averse, $\sigma < 0$, then the first-mover's intrinsic motivation generates a force that counteracts this indirect effect driven by the stronger encouragement effect. An inequity averse first-mover knows that she has a higher cost of exploring than the second-mover and the only way to reach equal payoffs or a position with advantageous inequality is by refraining from exploration. Yet, if the first-mover's σ parameter is positive, then the first-mover is efficiency concerned and more willing to contribute than a self-interested first-mover for all the treasure sizes (if there is a sufficient encouragement effect).

C.1 Estimating sigma and rho

To retrieve a crude estimate for σ we use slider task 3, 4, and 5. We first estimate a separate sigma for each slider task relevant ofr σ. Here the choice is coded as a share of other regarding behavior. Then we estimate an average of the estimated sigmas. For example, in task 3 there are 3 possible options lying strictly above the 45 degree line. If the individual chose the distribution in which the other gets the highest possible amount it is coded as 1, if the individual choses the second to highest amount for the other it is coded as 2/3 and if the individual chose the third to highest amount to the other we code it as 1/3. The rest of the options lying oin the 45 degree line or below are coded as 0. All options where the individual's outcome is large than the other's is relevant for rho, i.e., below the 45 degree line. We used the equivalent procedure to find an estimate for ρ as we did for estimating σ . Here, we use slider task 4, 5 and 6 to calculate the separate ρ.