Management Science

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

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

Emma von Essen Marieke Huysentruyt, Topi Miettinen

informs
http://pubsonline.informs.org/journal/mnsc

MANAGEMENT SCIENCE
Vol. 66, No. 12, December 2020, pp. 5861–5885
ISSN 0025-1909 (print), ISSN 1526-5501 (online)

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

Emma von Essen,[a,b] **Marieke Huysentruyt,**[c,d] **Topi Miettinen**[d,e,f]

[a] Department of Economics and Business Economics, Aarhus University, 8210 Aarhus V, Denmark; [b] The Swedish Institute for Social Research, Stockholm University, 10691 Stockholm, Sweden; [c] Strategy and Business Policy, HEC Paris, 78350 Jouy-en-Josas, France; [d] Stockholm Institute of Transition Economics, Stockholm School of Economics, 113 50 Stockholm, Sweden; [e] Department of Economics and Finance, Hanken School of Economics, FI-00100 Helsinki, Finland; [f] Helsinki Graduate School of Economics, FI-00100 Helsinki, Finland

Contact: emma.von.essen@sofi.su.se, https://orcid.org/0000-0002-1421-4582 (EvE); huysentruyt@hec.fr, https://orcid.org/0000-0003-3523-6323 (MH); topi.miettinen@hanken.fi, https://orcid.org/0000-0001-8133-3920 (TM)

Received: July 7, 2017
Revised: November 28, 2018; July 17, 2019
Accepted: August 20, 2019
Published Online in Articles in Advance: May 22, 2020

https://doi.org/10.1287/mnsc.2019.3494

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 player exploration on the optimality of the second player exploring as well. When agents have other-regarding preferences and imperfectly optimize, the encouragement effect is strongest. The explorative nature of the game raises the expected surplus compared with a payoff equivalent public goods game. We empirically confirm our main theoretical predictions using a novel experimental paradigm. Our findings are relevant for motivating and managing groups and teams innovating not only for private but also and especially so, for public goods.

History: Accepted by John List, behavioral economics.
Funding: The authors acknowledge financial support provided by the European Union through its Seventh Framework Program for research, technological development, and demonstration under FP7 Joint Technology Initiatives [Grant 613500] Social Entrepreneurship as a Force for more Inclusive and Innovative Societies (SEFORIS); Norges Forskningsråd [Grant 250506]; and the Hanken Foundation.
Supplemental Material: Data and the online appendices are available at https://doi.org/10.1287/mnsc.2019.3494.

Keywords: economics: behavior and behavioral decision making • economics: game theory and bargaining theory • economics: microeconomic behavior • industrial organization: firm objectives, organization and behavior • decision analysis: sequential

1. Introduction

Innovation plays a central role in the production of public goods, and positive payoff externalities play a central role in collaborative production of research and innovation. 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 healthcare, imminent environmental challenges, international terrorism, social fragmentation, and the chronic offending in low-income urban neighborhoods to name but a few. Furthermore, in many research-intensive environments, the outcome of exploration comprises a shared or public benefit accrued by all research team members. One promising trend in the provision of *innovations* with a public good benefit 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. Furthermore, 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 and externalities. With very few exceptions, previous work on the voluntary provision of public goods deals with situations that are static or involve limited exogenous uncertainty[1] 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, such as the importance of other-regarding preferences or reciprocal fairness norms in sequential public good provision, which generates a complementarity between the actions of early and late movers (Camerer 2003, Güth et al. 2007, Levati et al. 2007, Cappelen et al. 2015). By contrast, we specifically account for the uncertainty of a public good discovery and the complementarity driven by the informational and other-regarding preference channel, 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 on, 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 in finding 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 international standard. The key contribution of our model is to deliver sharpened insight into the strategic considerations that determine individual exploration decisions. Furthermore, 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. Using this structural model, we incorporate both informational and other-regarding complementarities, and exploiting an experimental design that allows disentangling of the two, we test to what extent the informational channel causally effects outcomes.

Our main results, both theoretical and experimental, are as follows. In our simple exploration game (EG), as in more complex strategic bandit models (Bolton and Harris 1999, Hörner and Skrzypacz 2016), the information externality induces an informational encouragement: a positive effect of first player exploration on the optimality of the second player 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, depend on the assumptions that we make about the agents' rationality. As 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 benefit, only exploration by the first player can trigger the second player to explore as well. In the subgame perfect equilibrium (SPE) of our baseline model, we thus expect (in theory) the first player's exploration to be nonmonotonic in the value of the public good. After we allow for imperfect optimization or other-regarding preferences,[2] which the extensive empirical and laboratory literature on strategic interactions and public goods provision, respectively, suggests 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 other's payoffs, because an informational effect and an other-regarding encouragement effect then coexist.[3] Together, our first set of theoretical predictions underscores the importance of information sharing as a nonmonetary channel that motivates exploration and highlights the critical role of social preferences for team exploration outcomes.

Next, to pin down the causal effect of the uncertainty of the outcome (that is, the informational encouragement effect) 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 provision game (PGG) that lacks the explorative nature. In this payoff equivalent game, the total value of the public good is the same as in the EG, and thereby, also the myopic incentives to contribute coincide in the two games. Yet, when it comes to dynamic incentives, the other-regarding encouragement appears both in the EG and the PGG, but the informational encouragement affects choices only in the EG. Comparing individual and aggregate expected outcomes in the EG with those in the PGG allows us to sharply identify the effect of the informational externality and thus, strengthen the internal validity of our main results. We establish theoretically that the expected surplus in the exploration game is weakly greater than in the payoff-equivalent PGG. Uncertainty in the public goods production process thus raises expected overall contributions.

Finally, our experimental results, based on the analysis of 13,760 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. Furthermore, we show that expected aggregate contributions in the EG consistently exceed expected aggregate contributions in the payoff equivalent PGG, and we find strong empirical support for the informational encouragement effect.

The stylized PGG and EG allow us to deliver sharpened insight into interactions within teams or groups with certain and uncertain public benefit outcomes, respectively, reliant on noncontractible, costly individual efforts with a limited set of independent alternatives. In some settings, there is a clear connection between EG and PGG, because team or group members can be organized to either follow relatively more certain but lower-return routines or share information regarding coordinated attempts to try out new approaches with an opportunity to realize a higher return. The setting can be a research-intensive environment, a cooperative, or a work team—say, applied econometricians at a middle-tier department, where departmental funding is contingent on the number of publications only, triggering the econometricians to use linear regression methods to publish high quantities in lower-tier journals, or where departmental funding is contingent on publication in top journals only, triggering the econometricians to interact, present to each other state-of-the-art identification methods that help each to find the best identification method for his or her respective applied research problem, and occasionally publish in top journals. Thus, the EG can be seen as a model of high-impact, breakthrough innovation (exploration), whereas the PGG can be usefully thought of as a model of generic incremental innovation or simply, exploitation.[4]

To the extent 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 is an important concern. Business leaders, for instance, who value opportunities for their employees to collaboratively work on projects that can transform team productivity or job satisfaction, let alone tackle societal challenges at large, are well advised to emphasize the inevitable uncertainty in these production processes (especially if the number of feasible solution is or can be made limited), to find ways to enhance the overall value of the public good created, and to promote information sharing and the dynamic nature of the discovery process. These strategies can induce more efficient outcomes.[5] 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 that wish to encourage their members to jointly search for approaches that enhance local social cohesion.

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 (nonrival and nonexcludable). This befits our focus on voluntary and joint search for the public good, and we extend 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.[6]

Second, our paper is related to the literature on moral hazard in teams (Holmström 1982), especially the theoretical analyses of sequential effort provision by Strausz (1999) and Winter (2006, 2009). The latter two works study the strategic incentives of team members when late movers observe the effort of early movers and efforts are complementary.[7] Winter (2009) shows how higher exogenous rewards can lead to lower efforts (the so-called incentive reversal effect), and our result regarding the nonmonotonicity 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 Plott and Smith 2008, part 6.1; Brown et al. 2011; and Klor et al. 2014, for instance). Relative to this literature, the key contribution of our paper is twofold. First, we develop a sequential, 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 as well as 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). They 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 coexist. They consider the so-called good news model, where exploration efforts are strategic substitutes. However, in a bad news model, the exploration efforts are typically strategic complements as pointed out by Hörner and Skrzypacz (2016) in a private goods setting.[8] This is more in line with our setup. Our model is a much simpler finite alternative model. In fact, in that regard, the exploration game is reminiscent of optimal search (Weitzman 1979) and recombinant innovation (Weitzman 1998), where old ideas can be reconfigured in new ways to make new ideas, much in the spirit of the way that agricultural scientists develop plant varieties by crosspollinating existing plant varieties.[9] Our theoretical and experimental setup also puts far fewer cognitive demands on laboratory participants than the canonical, multiarm bandit problems. Hence, this class of models is more amenable to laboratory testing and 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.[10] 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. In an independent experimental study of the bandit exploration model, Halac et al. (2016) and Deck and Kimbrough (2017) use similar approaches. 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).

Fourth, our paper ties into the literature on the voluntary provision of public goods, collective action, and prosocial behaviors originally studied by Olson (1965) in a self-regarding model and by Becker (1974) with altruistic preferences.[11] Our paper is most closely related to the work by McBride (2006) 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 nonmonotonicities. 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 (Palfrey and Rosenthal 1991, Anderson et al. 1998) to uncertainty inherent in the production process itself and as emphasized by Admati and Perry (1991) and Compte and Jehiel (2003, 2004), to the sequentiality and dynamic strategic interdependency of individual 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, and Section 5 concludes.

2. Theory

2.1. Basic Model of Exploration with Sequential Moves

Consider a simple two-stage, two-player exploration game with two partners (be it two employees, two coauthors, or two industry experts). There is a finite product space (of locations) and a unique public good (that is, treasure) in a single location within that space. The partners take turns to contribute to the exploration of the product space, and each can contribute by checking in one location whether the treasure is located there. Let K denote the number of locations. Ex ante, each location is likely to hold the treasure with probability $1/K$, and thus, without loss of generality, an action of player i can be denoted by a binary action $a_i \in \{0, 1\}$, where 0 indicates no contribution.

The valuation of the public good (that is, treasure size for i, α_i, with $i = 1, 2$) is ex ante known, nonexcludable, and obtained if and only if the public good 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 contribution costs. The cost of contributing, c_i, is borne by the relevant agent. Not contributing implies zero cost. Assume that $c_1 \geq c_2$, which is in line with the optimal incentive mechanism of Winter (2006).[12] The player in stage 2 can learn from the exploration of her partner in stage 1. The model assumes complete and perfect information (observable effort and outcomes), though no coordination device exists. We seek the subgame perfect equilibria of the game under different treasure size or public good value regimes.[13]

2.1.1. Subgame Perfect Equilibrium of the Exploration Game.

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

$$c_2 \leq \alpha/Y,$$

where $Y \leq K$ is the number of alternatives that have a positive probability of containing a treasure in the second stage. We call this player 2's myopic incentive to explore. Because the second stage is the last, the myopic incentive is also player 2's total incentive to explore. Likewise, player 1's myopic incentive to contribute is captured by $\alpha/K - c_1$. The myopic incentive to contribute is all that player 1 needs to consider if player 2's choice is not affected by that of player 1. There are two such cases.

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

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

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

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

capturing the simple intuition that, by contributing herself, she gets another contribution 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 1. *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 contributes when $\alpha < max\{\frac{c_1K}{1+\delta}, c_2(K-1)\}$,*
- *both players contribute when $max\{\frac{c_1K}{1+\delta}, c_2(K-1)\} < \alpha < Kc_2$,*
- *only player 2 contributes when $c_2K < \alpha < c_1K$, and*
- *both players contribute when $c_1K < \alpha$.*

This proposition reveals that player 1's equilibrium contribution decision is nonmonotonic in α owing to the encouragement effect, which is defined as the impact of player 1's contribution decision on player 2's contribution decision. Intuitively, when the value of the public good is very low, neither player finds it in his or her best interest to contribute. For somewhat higher values of the public good, neither player's myopic incentives to contribute are sufficient, but the dynamic encouragement effect triggers player 1 to explore and encourages player 2 to contribute 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 contribute. Player 1 knows this and free rides on player 2's contribution. Finally, after the value of the public good is so high that even player 1's myopic incentive dictates to contribute, then both players find it optimal to contribute. 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 cutoff contribution strategy in a two-player bad news Poisson bandit model cannot be a simple cutoff strategy; rather, it involves ranges of optimal encouragement and free riding, corresponding to bullets 2 and 3, 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 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 explore when $\alpha < 600$. Both contribute when $600 \leq \alpha < 800$. Only player 2 contributes when $800 \leq \alpha < 1{,}200$, and both players contribute when $\alpha \geq 1{,}200$. Notice that the expected total contributions are 1.75 units when both contribute, because player 2 contributes only when the treasure is not found in that case (that is, with probability 3/4).

Figure 1. (Color online) Theoretical Predictions in the Exploration Game, SPE, and Self-Interest

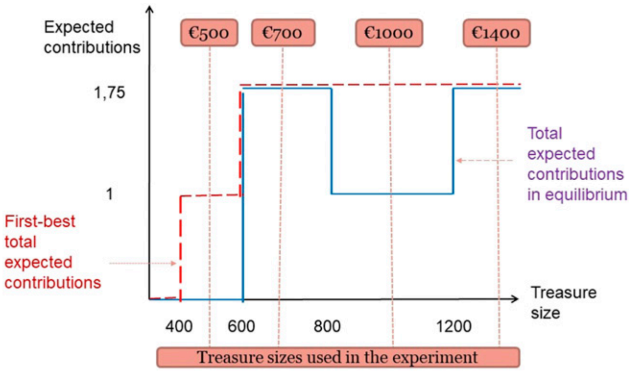

2.1.2. Subgame Perfect Equilibrium of the Voluntary Public Goods Game. The standard voluntary public goods game in its sequential two-player binary choice form is formally nested in our exploration game. To derive the standard voluntary public goods game from our exploration game, the value of the public 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 $\alpha_{PGG} = \alpha/K$. 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 with a marginal per capita return of α/K. 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 1's equilibrium contribution decision is monotonically increasing in α/K. Player 1 can no longer

Figure 2. (Color online) Theoretical Predictions in the Public Goods Game, SPE, and Self-Interest

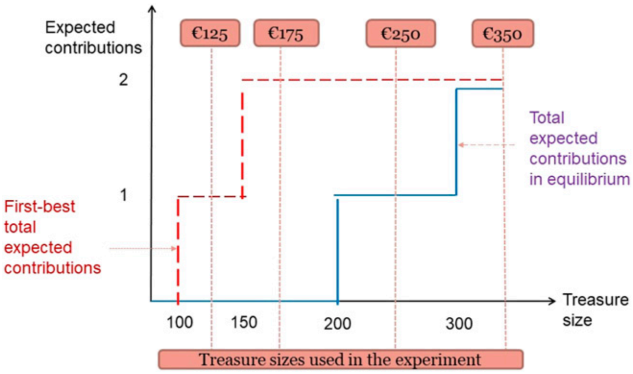

exert an influence on player 2's decision. The *information externality*, which is a distinct characteristic of the exploration game, disappears in the voluntary public goods game and with it, the encouragement effect. Thus, each player's myopic incentive dictates the optimal choices, and each player 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 or her best interest to contribute. For higher values of the public good, $c_2 < \alpha_{PGG} < c_1$, it is a dominant strategy for player 2 to contribute, and likewise, player 1 has a dominant strategy to free ride on player 2's contribution. Finally, after the value of the public good exceeds player 1's cost, then both players have a dominant strategy to contribute.[14]

As illustrated in Figure 2, the equilibrium level of contributions to the production of the public good is closer to first best in the *exploration game* than in the *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 EG (if the treasure is not found by player 1), but neither will contribute in the PGG.

From a welfare perspective, it is optimal that player i contributes if $c_i < (2\alpha)/K$, and this is true both in the EG and the standard PGG. 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.[15]

2.1.3. Comparison of the SPE Predictions. That the encouragement effect appears in the EG but not in the PGG yields two testable theoretical predictions.[16] These are the predictions that we have preregistered on the Open Science Framework platform at https://osf.io/ (name: exploration in partnership).

We summarize the main SPE hypotheses below.

1. *SPE contribution hypothesis.* Aggregate contributions across player types and treasure sizes will be weakly higher in the exploration game than in the public goods game. This prediction is primarily driven by first players contributing more in the exploration game compared with the public goods game when facing the second lowest treasure size.[17]

2. *SPE encouragement effect.* Player 2 will on average contribute more than player 1 in the public goods game, and for a limited range of treasure sizes, the wedge between player 1's and player 2's exploration efforts will be smaller in the exploration game than in the public goods game.

3. *SPE nonmonotonicity hypothesis.* In the exploration game, player 1's contributions will be nonmonotonic in treasure size. In the public goods game, player 1's contributions are monotonic in treasure size.

A key intuition behind all of these theoretical predictions is that there is an informational 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 contribute will trigger player 2 to not contribute as well. In reverse, there is also an implicit promise that his contribution will trigger player 2 to contribute as well (if the treasure is not found). Equivalently, the uncertainty in the production process of the public good invokes a complementarity between the two players' contribution decisions. As a result, aggregate equilibrium contributions or contributions to the public good are weakly higher in the exploration game than in the public goods game (Figure 3).

As a final remark, we illustrate that the self-interested, risk-neutral subgame-perfect equilibrium predicts that total surplus is weakly closer to first best in the EG than

Figure 3. Logical Connections Between SPE Hypotheses

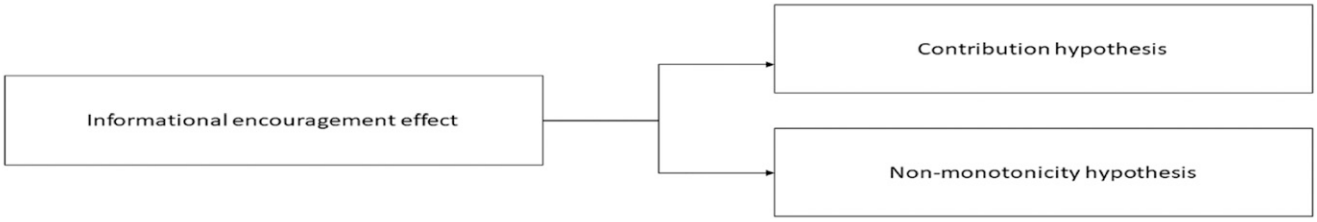

in the PGG for every treasure size when assuming $\alpha_{EG} = 4\alpha_{PGG} = 4\alpha$ (that is, that for risk-neutral players, the myopic incentive to contribute is the same in EG and PGG at each threshold treasure size) (see Section 2.1.2).

Consider the following levels of treasure size: VERY LOW ($\alpha_{PGG} < 100$ and $\alpha_{EG} < 400$), VERY LOW ($100 \le \alpha_{PGG} < 150$ and $400 \le \alpha_{EG} < 600$), LOW ($150 \le \alpha_{PGG} < 200$ and $600 \le \alpha_{EG} < 800$), HIGH ($200 \le \alpha_{PGG} < 300$ and $800 \le \alpha_{EG} < 1,200$), and VERY HIGH ($\alpha_{PGG} \ge 300$ and $\alpha_{EG} \ge 1,200$). Notice that these treasure size levels cover the entire space of potential treasure sizes.

Looking at the predicted contributions of player 1 and player 2 in Section 2, it is easy to calculate the expected total surplus in SPE. This latter is listed in the middle column of Table 1 for each treasure size class and for both the PGG and the EG. The first best contribution levels are even more straightforward to calculate: for a given number of treasures and options left, a player should contribute if two times the expected treasure size (each player receives the treasure value, and thus, the gross surplus equals 2α if the treasure is found) is greater than the private cost. Thus, in the PGG, player i should contribute iff $2\alpha_{PGG} > c_i$ (how we break the ties does not matter here). In the EG, player i should contribute if $2\alpha_{EG}/Y - c_i$, where Y is the number of alternatives that have a positive probability of containing a treasure. Thus, in the EG, player 1 should contribute if $\alpha_{EG}/2 > 300$. Player 2's first best contribution depends on whether player 1 contributed and whether a treasure was found. If player 1 contributed and found the treasure, then player 2 should not contribute. If player 1 contributed and did not find a treasure, player 2 should contribute iff in addition $2\alpha_{EG}/3 > 200$. If player 1 did not contribute, then player 2 should contribute iff $\alpha_{EG}/2 > 200$. The implied first best surplus is calculated in the rightmost column of Table 1. One can immediately see that the social surplus is weakly higher in the EG than in the PGG at all levels of treasure size. Thus, from the perspective of social surplus, the EG is predicted to reach a higher level of efficiency than the PGG.

Table 1. Expected Total Surplus

	Surplus at SPE	Surplus at First Best
VERY LOW		
PGG	0	0
EG	0	0
VERY LOW		
PGG	0	$2\alpha - 200$
EG	0	$2\alpha - 200$
LOW		
PGG	0	$4\alpha - 500$
EG	$4\alpha - 450$	$4\alpha - 450$
HIGH		
PGG	$2\alpha - 200$	$4\alpha - 500$
EG	$2\alpha - 200$	$4\alpha - 450$
VERY HIGH		
PGG	$4\alpha - 500$	$4\alpha - 500$
EG	$4\alpha - 450$	$4\alpha - 450$

2.2. Behavioral Model of Sequential Exploration

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[18] can be formally incorporated into our setup by deriving the logit quantal response equilibrium (QRE) (McKelvey and Palfrey 1998) instead of the SPE.[19] 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. Although 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 player i about the action profile a_j of the other player by $\hat{b}^i_j(a_j)$. In the quantal response equilibrium, player i chooses action a_i with probability

$$b_i(a_i) = \frac{exp\left((1/\mu)\left(\sum_{a_j}\hat{b}^i_j(a_j)v_i(a_i,a_j)\right)\right)}{\sum_a exp\left((1/\mu)\left(\sum_{a_j}\hat{b}^i_j(a_j)v_i(a,a_j)\right)\right)}. \qquad (4)$$

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

$$\frac{b_i(a_i')}{b_i(a_i'')} = \frac{exp\left((1/\mu)\left(\sum_{a_j} \hat{b}_j^i(a_j) v_i(a_i', a_j)\right)\right)}{exp\left((1/\mu)\left(\sum_{a_j} \hat{b}_j^i(a_j) v_i(a_i'', a_j)\right)\right)}, \quad (5)$$

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, $\hat{b}_i^j = b_i$ 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

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

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

In the sequel, we denote the probability of contributing or the contribution rate of player i by b_i. The probability of not contributing is thus $1 - b_i$. Moreover, we denote by $b_2(e)$ the contribution rate of player 2 in the contingency e.

The amount of contributions $b_1 + b_2$ (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 when aggregating over treasure sizes. The informational encouragement effect now appears for all treasure sizes (or values of the public good) in the exploration game, and yet, it 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 of player 2's actions (to contribute or not) occur with positive probability, and given the information externality, contribution by player 1 always increases the payoff of contribution to player 2. Furthermore, for sufficiently large μ, the informational encouragement effect is now increasing in treasure size. Intuitively, when the stakes are higher, player 2 has more to gain following contribution by player 1. In Online Appendix A, we present the formal analysis of the contribution behaviors in the QRE for the public goods and exploration games with selfish players.

2.2.2. Other-regarding Preferences. The behavioral and experimental economics literature on voluntary public goods provision in environments where there is no exogenous uncertainty provides a lot of evidence that other-regarding preferences must be invoked to understand the empirical contribution patterns. Although our focus is on understanding the implications of the informational encouragement effect, other-regarding preferences are likely to play a role. We thus want to understand the role in this context and in the end, be able to isolate the residual effect of the information channel when the other-regarding channel is controlled for. We, therefore, incorporate a more general model of preferences that embeds difference aversion and social welfare preferences as 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, to 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).[20] In our setting, this goal function can be written in the following form:

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

for $i = 1, 2$, where ρ and σ may be negative, 0, or positive and $\rho \geq \sigma$.[21] 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 to inequity or difference aversion. Irrespective of the specific distributional preferences that we consider, ρ is always understood to be greater than σ (Charness and Rabin 2002).[22]

2.2.3. Second Player Incentives and the Encouragement Effect. Let us consider first the public goods game. In the public goods game, every player's contribution yields a marginal per capita return of $\alpha_{PGG} = \alpha/K$, an ex ante fixed and certain value of public good. Suppose first that player 1 did not contribute. Then, player 2's payoff equals

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

if she contributes and 0 otherwise. Notice that the parameter σ indicates player 2's concern for player 1 if the payoff of player 2 is lower than that of player 1. Parameter σ appears here, because player 1 did not contribute ($a_1 = 0$), and therefore, player 2's payoff falls short of that of player 1 if the player 2 contributes.

In this case, the proportion of choice probabilities between contributing and not contributing equals

$$\frac{b_2(a_1 = 0)}{1 - (b_2(a_1 = 0))} = exp\Big(\big(1/\mu\big)\Big(\frac{\alpha}{K} - (1-\sigma)c_2\Big)\Big), \qquad (7)$$

and the log of the odds ratio between contributing and not contributing (7) is thus merely

$$\big(1/\mu\big)\Big(\frac{\alpha}{K} - (1-\sigma)c_2\Big). \qquad (8)$$

Suppose next that player 1 did contribute. Then, player 2's payoff equals

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

if she also contributes and

$$\alpha_{PGG} - \rho c_1 = \frac{\alpha}{K} - \rho c_1$$

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

$$\big(1/\mu\big)\Big(\frac{\alpha}{K} - \big(1-\rho\big)c_2\Big). \qquad (9)$$

In the public goods game, the only difference between expressions (8) and (9) is the behavioral other-regarding parameter terms in front of player 2's contribution cost. Because $\rho > \sigma$, we thus establish that an *other-regarding* player 2 is more likely to contribute if player 1 also contributed. By contrast, a *selfish* player 2's contribution decision (when $\sigma = \rho = 0$) remains unaffected by player 1's contribution choice in the standard public goods game (that is, 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 2 holds social preferences.

Consider next the exploration game. The analysis unfolds as in the public goods game: the other-regarding encouragement effect appears because of the second player's higher weight on the first player's welfare when the first player has contributed and thus, falls behind the second player in terms of payoff. Yet, unlike in the public goods game, first player contribution influences second player incentives also through an information channel. If the first player contributes and fails to find the treasure, then the second player has a higher chance of finding the treasure. Thus, the equation that describes the log odds of second player contribution to no contribution probabilities now reads

$$\big(1/\mu\big)\Big(\frac{\alpha}{K-1} - \big(1-\rho\big)c_2\Big) \qquad (10)$$

instead of (9). The second player log odds in the exploration game in the case that the first player does not contribute coincide with those in the public goods game: that is, (8).

Equations (8)–(10) allow us to decompose the total encouragement effect (10)–(8) into an informational effect (10)–(9) and an other-regarding effect (9)–(8). The greater these differences are, the greater the corresponding encouragement effects are. Clearly, encouragement should be stronger in the exploration game, because the information channel only appears there.

2.2.4. First Player and Aggregate Behavior. In the behavioral model, the other-regarding and informational encouragement effects influence second player behavior in a more continuous manner than in SPE theory, where agents are perfectly rational and self-interested. Independently of the treasure size, the second player will be more likely to contribute if the first player contributes (conditional on the treasure not being found) (see Online Appendix A for a formal analysis). The dynamic incentive effect will thus influence first player incentives to contribute at all treasure sizes, not just at the second lowest treasure size. One can show that, when parameters μ, ρ, and σ are sufficiently close to zero (and thus, players are almost perfectly rational and self-interested), the nonmonotonicity of the first player behavior still appears.

However, when μ is sufficiently high, the predicted first player behavior is monotone in treasure size (see the proof in Online Appendix A). The horizontal lines in Figures 4 and 5 in Section 3.5 depict the first player and second player contribution probabilities, respectively, for each treasure size in the two games for a representative agent model with parameter values $\mu = 18$, $\rho = 0$, and $\sigma = -1/6$. First mover behavior is monotone in treasure size in EG. The figures also show that contributions are, on aggregate, higher in the EG than in PGG. For the case $\mu = 18$, $\rho = 0$, and $\sigma = -1/6$, the expected total contributions equal 0.76 and 0.94 in the PGG and the EG, respectively. Moreover, one can derive a more general condition such that contribution rates in the EG are larger than in the PGG. Proposition A.4 in Online Appendix A shows that, quite generally, contributions in the EG are higher than in the PGG. It is natural to think that participants differ in characteristics relevant for μ, ρ, and σ, and thus, a model allowing for heterogeneity in this respect would be more realistic. Yet, as will be illustrated in

Figure 4. (Color online) First Player Contributions

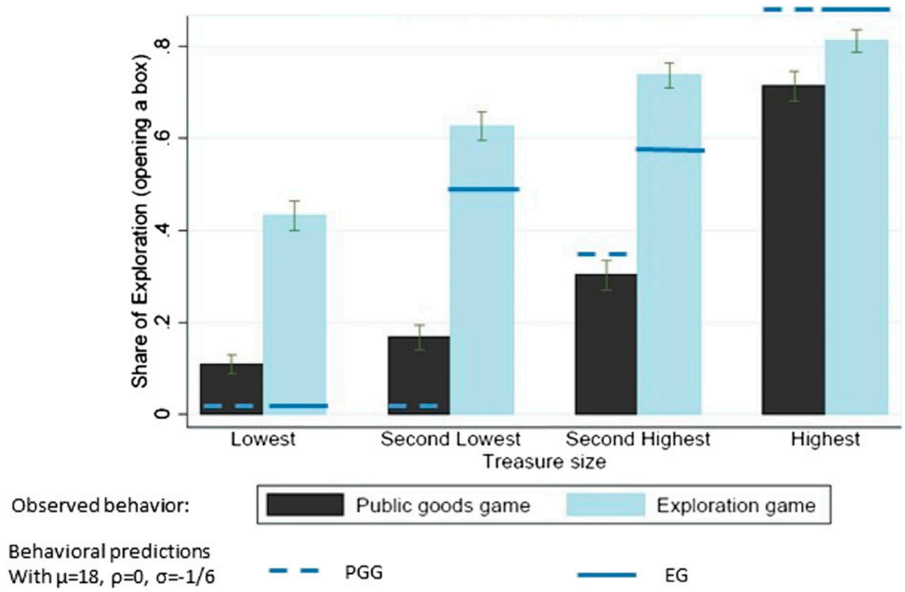

Section 4, we lack individual-level variation to empirically test such a model.

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

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

1. *Behavioral contribution hypothesis.* The aggregate contributions will be weakly higher in the exploration game than in the public goods game.[23]

2. *Behavioral encouragement effect.* In the exploration game, regardless of treasure size, the second player is more likely to contribute if the first player contributed but did not find a treasure compared with if the first player did not contribute, and this difference is higher than its counterpart in the public goods game.[24]

Figure 5. (Color online) Second Player Contributions

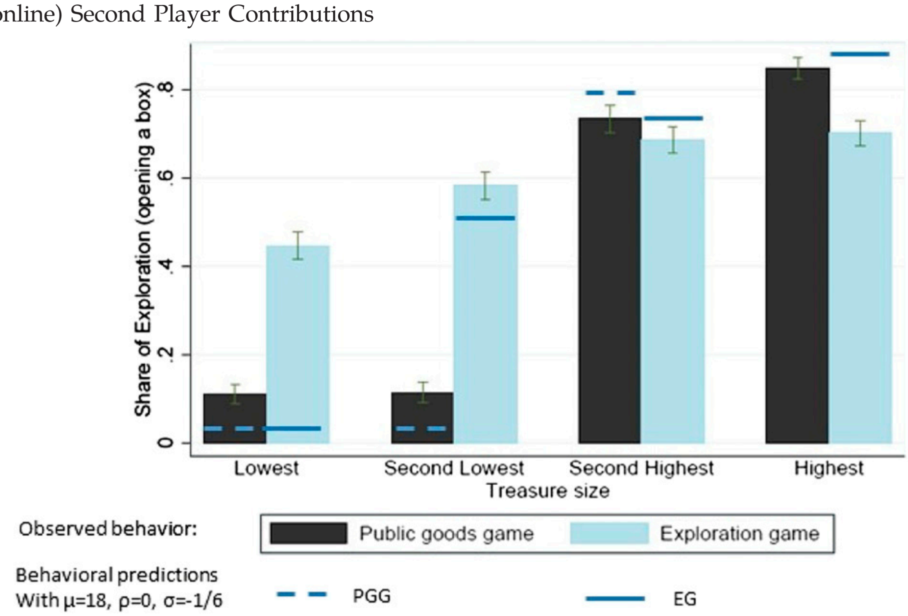

3. *Behavioral monotonicity hypothesis.* The contribution of the first player increases with the treasure size in the exploration game (because of stochasticity in second mover responses to first mover behavior).[25]

The behavioral analysis of Section 2.2 distinguishes the *informational* encouragement effect (compare Equations (10) and (9)) and other-regarding encouragement effects (compare Equations (9) and (8)). An other-regarding encouragement effect occurs in the public goods game provided that player 2 holds other-regarding preferences: player 2 reacts to the contribution by player 1 by increasing the probability of contributing. The informational behavioral effect occurs in the exploration game owing to the fact that first player contribution without finding the treasure increases the second player's chances of finding the treasure and thus, the probability of contributing. This informational effect, unlike in the SPE model, occurs for all treasure sizes in the exploration game. By subtracting the other-regarding behavioral encouragement effect from the informational behavioral encouragement effect, we can identify whether the informational effect plays a significant role or whether the encouragement is mainly driven by the other-regarding effect previously identified in the literature. The monotonicity effect merely notes that now, in the behavioral model, encouragement occurs for all treasure sizes; for sufficiently high level of randomness in second player behavior, the encouragement will increase in treasure size, and even first mover behavior will be monotone in treasure size. Notice also that, given first player behavior and the outcome of first player search, the probability of second player contribution is increasing in treasure size. Yet, of course, the more likely the first mover contribution in the EG is, the more likely it is that the treasure will be found, in which case the second player has fairly weak incentives to contribute. Thus, if the first mover probability of contribution increases drastically as the treasure size increases, then the unconditional probability of second player contribution may even fall.

Table 2 displays the testable theoretical predictions from both the SPE and behavioral models. We use the term B to denote the contribution by the respective players 1 and 2 facing the four different treasure sizes in the two games.

3. Study Design and Data

To test our hypotheses, we set up a laboratory experiment. We ran the experiment at the Cognition and Behavior Laboratory (Aarhus University, Denmark), the Aalto Choice Tank (Aalto University, Helsinki, Finland), and the Centre for Experimental Studies and Research (Bedriftøkonomisk Institutt Norwegian Business School, Oslo, Norway) during fall 2014 to spring 2016. In total, 436 subjects were recruited using identical recruitment procedures.[26] Each subject completed a 10-minute online survey at least five days before participating in the laboratory experiment. The laboratory session lasted on average 70 minutes. A 6.10 USD (United States dollar) participation fee and subsequent earnings, which averaged 7 USD, were paid in private at the end of the laboratory session.[27]

3.1. Online Survey

After signing up for 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), which is a six-item questionnaire where each question consists of a choice of one of nine possible allocations of money between oneself and another anonymous participant.

Table 2. Testable Theoretical Predictions from the SPE and Behavioral Model

Hypothesis	SPE	Behavioral
Contribution hypothesis	$B_{EG} > B_{PGG}$	$B_{EG} > B_{PGG}$
Encouragement effect		for all α
	$B_{2,EG,700}(a_1 = 1, Y = K-1) >$ $B_{2,EG,700}(a_1 = 0)$	$B_{2,EG,\alpha}(a_1 = 1, Y = K-1) >$ $B_{2,EG,\alpha}(a_1 = 0)$
Nonmonotonicity hypothesis	$B_{1,EG,700} > B_{1,EG,500}$	$B_{1,EG,700} > B_{1,EG,500}$
	$B_{1,EG,1,000} < B_{1,EG,700}$	$B_{1,EG,1,000} > B_{1,EG,700}$
	$B_{1,EG,1,400} > B_{1,EG,1,000}$	$B_{1,EG,400} > B_{1,EG,1,000}$

Notes. $B_{i,g,\alpha}$ refers to the contribution rate of player i in game g when the treasure size is α. $B_{2,g,\alpha}(e)$ specifies the (second player) contribution rate conditional on endogenous event e. EG refers to the exploration game, and PGG refers to the public goods game. $B_{EG} = \sum_{i=1}^{2}(1/4 \times B_{i,EG,500} + 1/4 \times B_{i,EG,700} + 1/4 \times B_{i,EG,1,000} + 1/4 \times B_{i,EG,1,400})$ and similarly, for B_{PGG}.

We used experimental currency units (ECUs) 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 0 and 60 in a lottery with a two-thirds probability of getting nothing and a one-third 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–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 participant won two and a half times the amount that he or 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 10, with 0 being risk averse and 10 being risk loving (Dohmen et al. 2011). To measure cognitive reasoning style, we used the cognitive reflection task (CRT) (Frederick 2005), which consists of three questions, without incentives. Finally, we asked the participants about their gender. The full questionnaire is in Online Appendix D.

3.2. Laboratory Experiment

The laboratory experiment was an internet-based game programmed for the purpose of this experiment.[28] The participants in each session were randomly assigned a game type-specific code on paper (both the exploration game and the public goods game were run in 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 assigned each player to 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. The participants played 32 rounds of either the sequential public goods or the sequential exploration game as either player 1 or player 2. They all encountered four levels of treasure sizes: that is, 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.

Before the laboratory experiment started, we made two random draws in public to establish the rewards tied to choices made in the online survey.[29] 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.[30] The participants then logged in to the game using the game type-specific codes and the anonymous unique identifiers that they had created at the onset of the online survey. The participants faced written instructions and control questions (see Online Appendix D). At any time of the experiment, the participants could revisit the instructions. On having correctly completed the control questions, the first game round could start.

At the beginning of the first game round, each participant received an endowment of 12,000 points. The participant's 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 was, however, unknown to the participants. In each session, the first and second players with the highest numbers of points in the public goods and exploration game, respectively, received a monetary prize of 13.68 USD. Each game round started with player 1 seeing four closed chests. In addition, the screen contained information on the participant's current number of points, the cost of contributing, the counterpart's cost of contributing, the size of the treasure (the treasure sizes were 500, 700, 1,000, and 1,400 points in the exploration game and 125, 175, 250, and 350 in the public good game), and the number of treasures left to explore. There was no information about other participants' current tally of points. The cost of contributing was kept constant throughout the session, and it always higher for player 1 than for player 2 (300 versus 200 points). See Figure B.1 in Online Appendix B for an image of the decision screen. Player 1 knew that player 2 would observe his or her choice before making his or 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 knew that only one 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 contribute equally across game types.

Player 1 then had to decide whether to pay to open a chest or not. Conditioned on player 1's choice, player 2 had the same choice to make. When the second player had made his or her choice, both players received feedback on the outcome of the game. Before each new round of the game, the participant was randomly rematched with another participant of the opposite player type within the same game type. When a participant had completed 32 rounds, the screen informed him or her of his or her total number of points. Finally, we announced the anonymous personal identifiers of the first and the second player winners publicly. A research assistant at each laboratory asked the participants for their anonymous personal identifier, found the individual specific amount of ECU that they had earned in both parts of the study, converted these into euros, added the show-up fee of 5 euros, and noted this on a separate piece of paper. Aarhus University then transferred the money to the participants' bank accounts. At Aalto University in Helsinki and BI Norwegian Business School in Oslo, the participants then received the earnings immediately in cash.

3.3. The Incentive Scheme

Players 1 and 2 were by design incentivized to work as a team, but neutral language was used. Each participant in a given player role was incentivized to compete against the other participants in that role in the group with the same game type but not against the participants in the opposing role that she was matched with. We expected such tournament incentives to (i) induce more self-interested behavior relative to a monetary compensation that is directly proportional to the tally of collected points, (ii) afford greater control over the self-interested encouragement threshold, and thus, (iii) produce a more favorable setting for the hypothesized nonmonotonicity effect to arise. To see this, consider a first mover who compares herself with the other first movers in the matching group and who has other-regarding preferences of the Charness and Rabin (2002, p. 851) form with the purely Rawlsian formulation (that is, $\delta = 1$). Now, if she earns the highest tally of points, she wins 13.68 USD, and her utility equals $(1 - \lambda)13.68 + \lambda 0$, where λ captures how much she cares about pursuing the social welfare versus her self-interest. However, if she is not the player with the highest tally of points, then she earns 0 euros, and her utility equals $(1 - \lambda)0 + \lambda 0 = 0$. Thus, the strength of the preference for winning depends on the unobserved parameter λ. Independently of the value of λ, the winning outcome gives a higher payoff than the losing outcome, and every losing outcome gives the same payoff independently of the identity of the winner. The tournament scheme implemented is thus designed to strengthen the incentive to behave as if self-interested and to downplay other-regarding motivation. It can easily be shown that this holds for all parameterizations of the outcome-based (consequentialist) versions of the Charness and Rabin (2002) model.

To ensure that the winner-takes-it-all part of the incentive scheme does not create the differences in contribution between the EG and the PGG, we ran an extra experiment in spring 2019. In particular, we wanted to test the contribution hypothesis in a design without the winner-takes-it-all part of the incentive scheme. In these additional sessions, the experimental design of the EG and PGG remained the same except that now participants were paid according to the total number of collected points. We introduced a conversion rate between the collected points and actual payoff that matched expectations in the main experiment so as be able to sharply identify the effect of the incentive scheme alone. In Norway, we used a conversion rate of 0.01, where 17,000 points translated into 170 NOK (Norwegian crown), and in Finland, we used a conversion rate of 0.001, where 1,700 points translated into 17 euros. The written instructions for this experiment are in Online Appendix E.

Before running the extra experiment, we calculated the optimal sample size needed to credibly detect an effect size for the contribution hypothesis. We thus assumed that the true effect size is the effect size of the contribution hypothesis in the main study.[31] Moreover, we assumed a power of 0.8 and a significance level (given that we aimed to test only one hypothesis) of 0.05. In this case, the optimal sample size is 36 participants in each of the two groups (EG and PGG). We ran two sessions in Oslo, Norway, and one session in Helsinki, Finland. A total of 54 participants completed the extra experiment. We then reestimated Table 4 using the data from the extra experiment (that is, with no group tournament incentives). The results are similar to those in our main study, which suggests that the group tournament incentive scheme is not the cause of the found differences in contribution between the two game types. See Online Appendix B, Table B.1.[32]

3.4. Ethics and Registration of Study

Because the data from the study are 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 Cognition and Behavior Laboratory Ethical Advisory Board.[33] 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.

3.5. Data

A total of 430 participants completed both the online survey and the laboratory session. Each participant

Table 3. Descriptive Statistics of Online Survey Variables Across Game Types

	n^a	Mean	Median	Standard deviation	Max	Min
Public goods game						
Gender (1 if woman, 0 otherwise)	190	0.55	1	0.50	1	0
Risky investment choice	183	33.05	30	19.14	60	0
Risk question	186	5.90	6	2.14	10	2
CRT score	186	1.94	2	1.08	3	0
Social value orientation	186	26.75	31	13.39	45	–9
Exploration game						
Gender (1 if woman, 0 otherwise)	239	0.55	1	0.50	1	0
Risky investment choice	230	33.83	30	19.21	60	0
Risk question	231	5.70	6	2.16	10	1
CRT score	231	1.94	2	1.06	3	0
Social value orientation	231	28.12	33	13.09	61	–16
Observations	430					

[a]Some participants did not answer all of the survey questions; the numbers of observations, therefore, vary.

completed 32 game rounds of play, implying a total of 13,760 observations overall. Table 3 displays a summary of the main variables from the online survey across the two game types separately. We confirm in Table B.2 in Online Appendix 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 (standard deviation: 1.10), and 60% of the sample answered correctly all three questions of the CRT.[34] The average SVO angle in our sample equaled 28° (standard deviation: 13.09). Following Murphy et al. (2011), the average participant should thus be classified as prosocial. Table B.2 in Online Appendix B shows the randomization check. None of the observable variables differ by game type.

Figures 4 and 5 show the raw results for first and second player contributions across the four treasure sizes; the darker bars depict average contributions in the public goods game, whereas the lighter bars depict those in the exploration game. We find that first player contribution increases with treasure size, and hence, it does not support the SPE nonmonotonicity hypothesis. Rather, observed effects between game types and other qualitative patterns in our data can be better explained with a QRE model that allows for other-regarding preferences. To illustrate this, we have superimposed theoretical predictions of first and second player behavior on the observed frequencies for the respective players using parameter values $\mu = 18$, $\rho = 0$, and $\sigma = -1/6$.[35]

3.6. Significance Level and Multiple Comparisons

The more null hypotheses that we test, the larger the probability of getting false rejections. When designing our experiment and calculating the sample size, we unfortunately did not take into account the multiple hypotheses testing. This section and the corrections were completed post hoc after reflecting on insightful comments from the editor and two referees. Before running the analyses but after the collection of the data, we preregistered three hypotheses and tests from SPE hypotheses. In post hoc correction for multiple comparison, we focused on part of these tests. The other tests we consider more exploratory, and we apply a conventional significance level of 0.05.

Simple adjustment for multiple corrections post hoc, such as Bonferroni, might increase the probability of type II error and reduce the power to detect an effect. We, therefore, used List et al. (2016) for a more sophisticated procedure of correction. This method does, however, not perfectly apply to our setting. We ran the correction code for the comparisons of the contributions between the groups that we randomized (type of game and player) using average contribution for each individual. Our results still hold for such a correction as shown in Table B.3 in Online Appendix B. As a complementary method, we also used a simple post-Bonferroni correction of the 39 regression coefficient tests that we present in Tables 4–7, implying an α of $0.05/39 = 0.0012$. Revisiting Tables 4–7, our

Table 4. ordinary Least Square Regression (OLS): Differences in Contributions Across Player Types

	(1)	(2)	(3)
Exploration game	0.244***	0.244***	0.155***
	(0.022)	(0.022)	(0.025)
First player		–0.030	–0.129***
		(0.022)	(0.031)
First player × Exploration game			0.178***
			(0.043)
Constant	0.387***	0.402***	0.452***
	(0.016)	(0.017)	(0.017)
Observations	13,760	13,760	13,760

Note. Robust standard errors are clustered on session and individual. $^{*}p < 0.05$; $^{**}p < 0.01$; $^{***}p < 0.001$.

Table 5. Second Player Contributions

	(1) Lowest	(2) Second lowest	(3) Second highest	(4) Highest
First player behavior	0.068	0.245***	0.135***	–0.016
	(0.035)	(0.038)	(0.038)	(0.031)
Exploration game	0.286***	0.343***	–0.076	–0.172***
	(0.026)	(0.030)	(0.039)	(0.039)
Encouragement (interaction)	0.233***	0.226***	0.188***	0.248***
	(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)
Observations	3,198	3,006	3,010	3,030

Note. OLS has robust standard errors clustered on session and individual.
*$p < 0.05$; **$p < 0.01$; ***$p < 0.001$.

results remain qualitatively the same. In the figures below, we now apply the corrected α to construct the error bars.

4. Results

To ease the readers' comprehension, we present our results in the order of our main testable hypotheses summarized in Table 3 and not in the order of preregistered plan versus post hoc analysis. This makes some of our results exploratory in nature, and they should be interpreted as such. Our results are consistent with the assumption that people imperfectly optimize and care not only about their own payoffs but also, about others' payoffs. Our analyses suggest the relevance of a behavioral model of sequential exploration for the public good.

4.1. Contribution Hypothesis

We start by analyzing how individual contribution behavior varies across game type (public goods versus exploration game). As an additional step, we look at contribution differences between player types. We estimate the following basic equation:

$$b_{\iota,g} = \gamma_{\iota,g} + \beta_1 G_{\iota,g} + \beta_2 T_{\iota,g} + \beta_3 G \times T_{\iota,g} + \varepsilon_{\iota,g}, \qquad (11)$$

where $b_{\iota,g}$ denotes whether individual ι contributes or not during game round g ($b_{\iota,g} = 1$ if player ι contributes and 0 otherwise), T denotes player type (taking the value 1 if the individual is a first player and 0 otherwise), G denotes the game type (equaling 1 if the game is the exploration game and 0 otherwise), and ε is the error term. Table 4 reports the regression results derived using a linear probability model with robust standard errors clustered by individual and session.[36] Column (1) of Table 4 reports the regressions results for game type, and column (2) of Table 4 reports for both game and player type. Aggregate contribution was about 24 percentage points larger ($p < 0.001$) in the exploration game compared with the public goods game. These results are consistent with the *contribution hypothesis* predicted by both the SPE and behavioral models. Overall, there is no significant difference between first and second players' exploration behavior.

Column (3) of Table 4 further includes the interaction term between game type and player type. Now, the first coefficient estimate reveals that, relative to the public goods game, a second player was 15 percentage points more likely to contribute in the exploration game ($p < 0.001$) than in the public goods

Table 6. Second Player Contributions—Individualistic Players

	(1) Lowest	(2) Second lowest	(3) Second 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

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

Table 7. Second Player Contributions—Prosocial Players

	(1) Lowest	(2) Second lowest	(3) Second 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	2,157	2,025	2,034	2,007

Note. OLS has robust standard errors clustered on individual.
$^{*}p < 0.05$; $^{**}p < 0.01$; $^{***}p < 0.001$.

game. The marginal effect of player type in the exploration game displays no such gap ($\beta = -0.049$, $p = 0.101$). The second coefficient shows that, in the public goods game, the first player was on average about 13 percentage points less likely to contribute than the second player. This finding reflects the fact that player 1's cost of contributing is higher than that of player 2. The first player thus has a larger myopic incentive to free ride. The interaction term suggests that the differences between player types is 18 percentage points larger in the public goods game compared with the exploration game. Taken together, these results confirm the contribution hypothesis, and the observed qualitative patterns indicate a possible encouragement effect.[37]

4.2. Encouragement Effect

To test the encouragement hypothesis, we begin by examining player 2's contribution behavior in the public goods game and exploration game separately. Figure 6 reveals that, in the PGG, the share of second players who contributed was always greater when the first player contributed than when she did not, except in the highest treasure size. This finding is consistent with the other-regarding encouragement effect.

Table 3 shows that the average values of the SVO angle measure equaled 27 and 28 for the respective subsamples: that is, subjects who played the public goods game and those who played the exploration game. The standard deviation was 13 in both groups. Had a subject consistently been self-regarding, the angle measure would yield a value of eight. Thus, our subjects seem predominantly other regarding. Notice that the SVO measure is a unitary measure. The other-regarding QRE predictions of our model, however, incorporate 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 Online Appendix C for a detailed description). The results show that, for the vast majority of our participants, $\rho \geq 0.5$ and $\sigma > 0$. In all, 96.5% of the participants had preferences $\rho > \sigma$.[38]

Figure 6. (Color online) Public Goods Game: Second Player Contribution Conditional on First Player Contribution

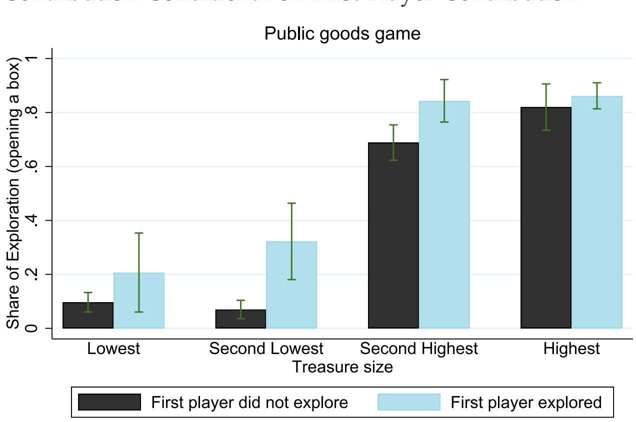

Note. Error bars based on Bonferroni-corrected p-value of 0.0012.

Figure 7. (Color online) Exploration Game: Second Player Contribution Conditional on First Player Contribution

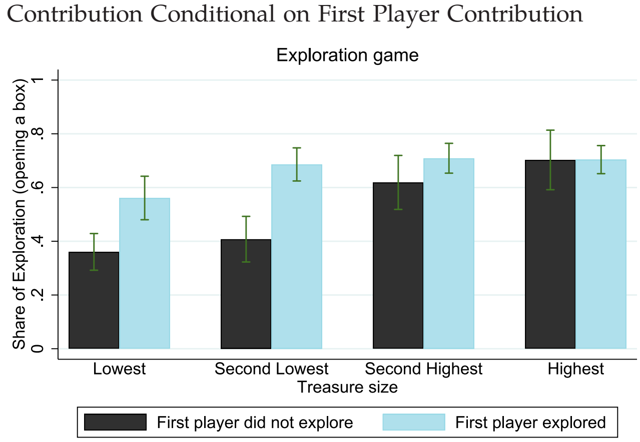

Note. Error bars based on Bonferroni-corrected p-value of 0.0012.

This suggests that the other-regarding encouragement effect should be controlled for. However, given that we lack variation in our individual measures to estimate individual-level differential effects of ρ and σ on contribution propensities, an other-regarding representative agent model is a good approximation.

We next turn to the analysis of player 2's decisions in the exploration game. Figure 7 shows that, when the first player did not contribute, the second player choice probabilities are consistent with what is optimal from a self-interest perspective in the sense that the observed contribution rate is below 50% when the contribution is suboptimal, and it is above 50% when the contribution is optimal. Notice that this was true also for the patterns in Figure 6 for the public good game. However, the choices are closer to the boundaries of the private optimum in the PGG, whereas they are closer to 50% in the EG. This may be because of the fact that, in the EG, one needs to calculate the gross expected benefit in the EG, whereas it is explicitly given in the PGG. The extra cognitive effort required in the EG may increase the level of noise in behavior and reduce the responsiveness to incentives.

Figure 7 also shows that those in the role of player 2 were more likely to contribute for all treasure sizes except the highest following a contribution by player 1 than following no contribution. This result is consistent with the informational encouragement effect, which is expected to occur in both the SPE and behavioral model of the exploration game: that is, the SPE informational encouragement effect (for the second lowest treasure size) and the QRE informational encouragement effect (for all treasure sizes), respectively. That the encouragement occurs for all treasure sizes is inconsistent with the informational encouragement in SPE and consistent with the informational encouragement in the behavioral model. However, given the prosocial character of the participants in our sample, these discrepancies in contribution decisions are possibly also driven by an other-regarding encouragement effect.

We then disentangle the informational encouragement effect from the other-regarding effect. To this end, we exploit the panel structure of our data and estimate an equation of a similar form as Equation (11). Now, $B_{g,\imath}$ corresponds to the contribution decision of second player $\imath$ during game round g. Let G again denote the game type (equaling one if the game is the exploration game and zero otherwise) and T denote the contribution decision taken by player 1. Variable T is defined differently depending on which game type we consider. In the EG, the variable takes on the value 1 if the first player contributed but did not find a treasure ($a_1 = 1, Y = 3$) and 0 if the first player did not contribute ($a_1 = 0$). In the PGG, the variable equals 1 if the first player contributed ($a_1 = 1, Y = 3$) and 0 if the first player did not contribute ($a_1 = 0$). In the exploration game, the difference in second player contribution rate when the first player did not find a treasure versus when the first player did not contribute reflects a combination of all encouragement effects discussed. In the public goods game, however, the difference can only capture a potential other-regarding encouragement effect. The estimated coefficient of the interaction term $G \times T$ in Equation (11) can now be interpreted as a measure of the informational encouragement effect:[39]

$$\begin{aligned} &\textit{Informational} \\ &= \overbrace{(B_{EG}(a_1 = 1, Y = 3) - B_{EG}(a_1 = 0)))}^{total} \\ &- \overbrace{(B_{PGG}(a_1 = 1, Y = 3) - B_{PGG}(a_1 = 0))}^{other-regarding}. \end{aligned} \tag{12}$$

Table 5 presents OLS regression results 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 contribute outweighs any behavioral considerations. The second coefficient in each regression captures the difference between second player contributions in the exploration game and the public goods game when the first player did not contribute. The positive coefficients at the two lowest treasure sizes and the negative ones at the two highest ones are likely to reflect the fact that it is simply easier to grasp the privately optimal behavior in the PGG where the gross benefit is explicitly given, than in the EG, where each participant must calculate the expected benefit based on the treasure size and the number of alternatives to explore. Therefore, behavior is closer to the private optimum in the PGG than in the EG (not to contribute for the two lowest and to contribute for the two highest when the first player did not contribute). The third coefficient estimates correspond to the value of the informational encouragement effect. We find that this effect is statistically significant and positive, lending support to the importance of the informational encouragement effect, which is not documented in the experimental literature before.[40]

4.3. Nonmonotonicity Hypothesis

The encouragement effect and the nonmonotonicity hypotheses are closely linked. We continue to estimate the nonmonotonicity by looking at player 1 behavior and player 2 behavior across the treasure sizes in the respective games.

Figures 4 and 5 present the average share of contributions of player 1 and 2, respectively, by treasure size. These raw averages indicate that individual contribution rate was monotonically increasing in treasure size. Tables B.6 and B.7 in Online Appendix B show the preregistered analysis of the contribution gap between pairwise treasure sizes across game types for the first and second players, respectively. Figures 4 and 5 as well as Tables B.6 and B.7 in Online Appendix B do not fully support the nonmonotonicity hypothesis. Despite there being empirical support for encouragement, greater rewards seem to invoke a higher contribution rate. These results cast doubts on the relevance of the SPE game-theoretic predictions in our setting. They lend support to a QRE model with other-regarding preferences for the exploration and public goods games.

In sum, our results empirically confirm the contribution hypothesis: aggregate contribution rate 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 contribute significantly more in the exploration game. Also, we find support for the encouragement hypothesis. We decompose the encouragement effect into two parts: the other-regarding encouragement effect, which is positive and significant, is already documented widely in the existing literature; the informational encouragement effect, which is also found to be positive and significant, is novel to the literature. However, we do not find indications of a nonmonotonic relationship between first player contribution rate and treasure size. This implies that we do not find an encouragement in the very narrow meaning of the definition (that is, according to the SPE model). This SPE model would predict a nonmonotonicity and incentive reversal owing to the range where the first player should contribute in order to encourage the second player to contribute for a lower treasure size and free ride when she knows that the second player's incentives to contribute are sufficient.

4.4. Additional Analyses: Social Value Orientation, Risk Aversion, and Cognitive Ability

The main results favor further exploring other-regarding preferences. We thus use the measures of preferences that we elicited in the online survey before the experiment. The covariates that 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 B.9 in Online Appendix B). To examine whether there are heterogeneous effects, Table 6 shows the encouragement effect for the subsample of individualistic participants, and Table 7 shows the encouragement effect for the prosocial part of the sample. In line with the theoretical prediction, the encouragement effect driven by other-regarding preferences (first coefficient) is much less pronounced among individualistic individuals compared with prosocial individuals. This lends support to a model with behavioral preferences.

In Table 8, we tried to understand whether the degree of prosociality mattered for contributions in the two games. 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 on the one hand and the deliberative and the intuitive thinkers on the other hand. Unfortunately, none of these individual characteristics help us further explain differences in contributions.

4.5. Robustness

To assess the robustness of our results, we conduct a number of additional tests. First, we estimate Equation (11), including dummy variables, for each of the 32 rounds (Table B.10 in Online Appendix B).

Second, in 13 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 Equation (11) again using only the sample where we randomly assigned participants to game type within the session. This subsample comprises 334 participants. The results are qualitatively similar (see Table B.11 in Online Appendix B).

Third, we tested whether the order in which we presented the treasure sizes affects contributions when shifting from one treasure size to another. On average, the first players facing an ascending order seem to contribute between 12% and 10% more than first players facing another order. However, this order effect does not change our results regarding the gap in contribution rates between the exploration and public goods game when comparing treasure sizes (that is, our main results remain). See Table B.12 in Online Appendix B.

Fourth, participants played eight rounds with each treasure size. To account for possible learning, we look at contributions in the last four rounds of play for each treasure size. This implies that we cut our sample in half. Table B.13 in Online Appendix B shows similar patterns as before, with greater contributions as treasure size increased. Holding treasure size constant, there is a greater contribution rate in the exploration game than in the public goods game. For second players, a large gap in contribution rates between the exploration game and public goods game prevailed for the smaller treasure sizes only.

Table 8. Contributions and Social Value Orientation (SVO) Angle

	(1) Lowest	(2) Second lowest	(3) Second highest	(4) Highest
Exploration game	0.318***	0.462***	0.026	–0.105**
	(0.036)	(0.033)	(0.041)	(0.033)
First player	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)
Player × SVO	–0.004	–0.001	–0.006	0.001
	(0.002)	(0.003)	(0.003)	(0.003)
Game type × Player × SVO	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	3,344	3,248	3,328	3,392

Note. OLS has robust standard errors clustered on individual.
$^*p < 0.05$; $^{**}p < 0.01$; $^{***}p < 0.001$.

Fifth, we take a final corollary result predicted by our behavioral 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 contribute even after the treasure was found by player 1. Figure 8 shows that there was in fact a small share of second players who contribute even when the treasure had been found. This share equaled about 20% of participants and was, as predicted, constant across treasure sizes.

4.6. Limitations

Researchers' ways of conducting quantitative studies may affect the results and interpretation of findings, such as the probability of false positives, leading to difficulties in interpreting research findings. This is an important discussion found across fields (Ioannidis 2005, Gelman and Carlin 2014, Maniadis et al. 2014).

Figure 8. (Color online) Second Player Contribution Conditional on First Player Finding

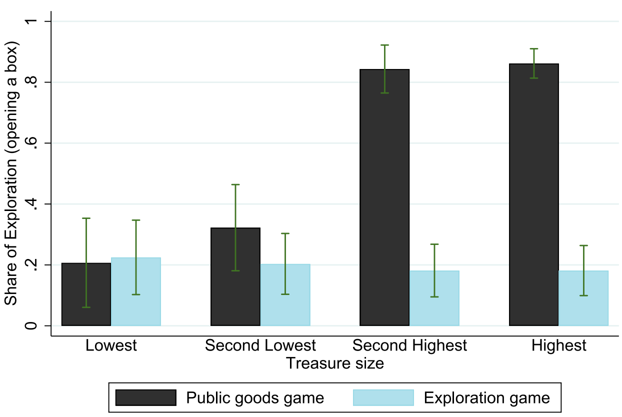

We preregistered hypotheses, variable coding, and some of the tests, which hopefully to some extent lessens our degree of freedom as researchers.

Maniadis et al. (2014) show that our prior beliefs about the hypothesis being true as well as the number of researchers currently exploring the question influence the probability of false positives. Using Maniadis et al. (2014), equation (2), we estimated the poststudy probability of a true relationship being reported. When comparing the average contribution between the exploration game and the public goods game, we have an effect size of –0.5, which is considered a medium effect (Cohen 1992). Using an α of 5% and our current sample size, we have a power of 100. Assuming that we are the only team exploring the research question and considering the following prior probabilities (10%, 50%, and 70%), the poststudy probabilities of a true relationship being reported as true are 69%, 95%, and 98%, respectively. These tests suggest that our results seem to be of relevance. However, as Maniadis et al. (2014) point out, future studies will decrease the probability of reporting false positives.

When interpreting our results, it should be noted that we had initially taken only the SPE predictions to the data. Ex post, because 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) and Goeree and Holt (2000, 2001) illustrate the power of this latter approach, and we reexpress the

Figure 9. (Color online) Second Player Contributions Conditional on First Player Not Contributing

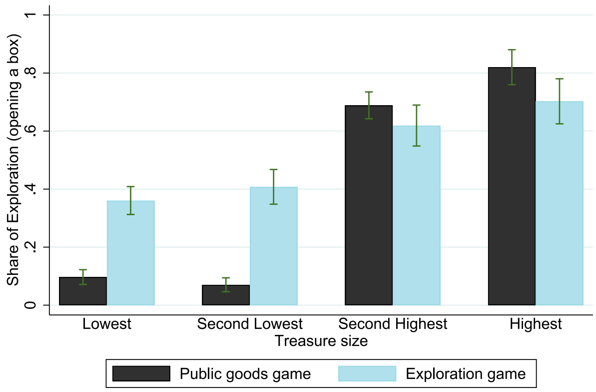

recommendation of Goeree and Holt (1999) and Camerer et al. (2004, footnote 5) that researchers in future-related theoretical and empirical work give more consideration to the QRE framework as an important theoretical benchmark.

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 also, too little variation in the implied social welfare utility parameters (Charness and Rabin 2002) ρ and σ (see Online Appendix C). There was an abundance of subjects with a tendency to share the earnings 50:50, but there were 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) (Figure 9).

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 us 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 that 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.

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).[41] We also study the robustness of these phenomena by extending them 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 contributions to exploration by player 1 motivate contributions by player 2. This encouragement effect, which we decomposed into an other-regarding part and an informational part, is at play for small and large public benefits to successful exploration and in theory, increases with the magnitude of the benefits. We provide evidence that not only establishes the other-regarding effect but also, establishes the entirely novel informational effect. Based on the informational effect, we theoretically derived that uncertainty in our game raises rather than decreases the aggregate level of exploration. Our experimental data robustly lend support to this contribution 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' contribution decisions, which brings equilibrium innovation closer to the social optimum. Future studies in less controlled field settings could potentially measure and test the social surplus directly rather than aggregate contributions. In practice, an organization's architecture (say, openness and interaction opportunities in a workspace), processes (say, whether interaction and exchange amongst peers are regularly organized), and culture (say, whether the organization strongly values openness to change versus conservation) as well as the rules and expectations set by external stakeholders (such as rules set by external funders) can strongly affect whether agents inside the organization are more likely to exploit versus explore within a known set of independent alternatives. The insights that we derive from our stylized model allow us to gain a better understanding of search in teams or groups, say by academics and scholarly output (see the example on econometricians and identification methods in Section 1) or farmers and biodiversity (say farmers in a co-op and their search for crop varieties that enhance biodiversity) to name a few examples.

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) and Winter (2006, 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 (that is, encouragement and discouragement) that need to be taken into 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 nonmonotonicity 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, assumed a very explicit complementarity between inputs, an increasing returns to scale technology, and asked whether sequential team production leads to incentive reversals (that is, nonmonotonicities). 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 nonmonotonicities suggests that the contextual differences influence behavior. Effects similar to ours can be observed in the experiment of Steiger and Zultan (2014), where experimental variation concerns the simultaneity versus sequentiality 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 nonmonotonicity aspect in the theoretical multiplayer 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 player can influence the second player probability of exploring (second player 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. Because of 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, one important implication of our results is that business leaders or governments that wish to harness 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 employees or citizens at large (for instance, through corporate culture or educational programs).[42] 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 difficult public goods problem lies, one can actually elicit greater voluntary contributions. Hence, when 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.

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, not 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.[43] 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 payoff to exploration prior to exploring opens a bridge between the literature of exploration and social learning (herding). However, if the locations contain public or private goods of variant values, the links to the search literature become obvious.[44]

The framework and methods proposed in this paper can also be used to study the effects of alternative knowledge production technologies or alternative incentive schemes on exploration behavior. For example, what if the knowledge production function is substitutional (that is, unsuccessful search lowers the probability of subsequent success), or what if there is a positive probability that none of the chests hold a treasure? Then, depending on the precise parameterization of these production functions, you could have encouragement or discouragement effects of not finding a treasure. Or what if the treasure is a lottery ticket either at one location (EG) or in all locations (PGG), which would allow us to control away the effect of risk aversion? Given different knowledge production technologies, what is the optimal mix of private and public benefits to encourage greater exploration for the good?

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

Acknowledgments

The authors thank Chloé Le Coq, Olivier Darmouni, Denis Gromb, Leif Helland, Tuomas Nurminen, Pauli Murto, Heiner Schumacher, and Ro'i Zultan as well as participants at seminars at Ben Gurion University; HEC Paris; Tilburg University; University of British Columbia; University of California, Santa Cruz; University of Turku; Helsinki Graduate School of Economics; Swedish Institute for Social Research; ESA European Meeting 2016; Economic Science Association World Meeting 2016 in Jerusalem; EARIE in Maastricht 2017; Nordic Conference of Behavioral and Experimental Economics in Gothenburg 2017; and European Economic Association Annual Conference 2018 in Cologne. They also thank Public Choice Research Center Lab at the University of Turku, Cognition and Behavior Laboratory at the University of Aarhus, Aalto Choice Tank at Aalto University, and the Centre for Experimental Studies and Research laboratory at the BI Norwegian Business School in Oslo for hospitality. They thank Olli Lappalainen, Knut-Eric Neset Joslin, Jason O'Donoghue, and Kristaps Dzonsons for excellent research assistance.

Endnotes

[1] Exceptions include Dickinson (1998), Levati and Morone (2013), Björk et al. (2016), and Vesely et al. (2017).

[2] We operationalize imperfect optimization by means of the 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.

[3] There are many empirical studies evidencing the other-regarding preference channel effect (Berg et al. 1995; Clark and Sefton 2001; Falk et al. 2003, 2008), and various theories have been put forward that rationalize such patterns. The other-regarding preference channel is predicted to be active both in the public goods game and in the exploration game. Yet, the informational encouragement effect appears in the exploration game only. This effect requires the information externality channel to operate.

[4] Compared with the returns from incremental innovation or exploitation, the returns to breakthrough innovation or exploration are bigger but systematically less certain (March 1991). See also Ederer and Manso (2013).

[5] 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). Our results suggest that this could be particularly effective if the employees are prosocially motivated and consequently, feel more engaged at work (Bolino and Grant 2016)

[6] In the case of a sequential search for a private good instead of a public good (keeping the information externality), the encouragement effect no longer occurs when players are self-interested. After we allow for other-regarding preferences, the encouragement effect kicks in but matters less than in the public goods case.

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

[8] See Hörner and Skrzypacz (2017, pp. 66–67).

[9] Recombinant innovation is pervasive in a wide number fields, such as genomics, agriculture, drug discovery, materials science, and particle physics. Weitzman (1998) discusses the example of researchers at Menlo Park searching for a material that can be carbonized and used as a filament to produce the "electric candle." The combinatorial-based knowledge production function is thus one type of knowledge production function that corresponds to our exploration game.

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

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

[12] 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 faces no implicit threat that his or her failure to innovate will trigger subsequent agents to shirk as well. Hence, player 2 should be provided with stronger incentives to exert effort than player 1.

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

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

[15] 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 contributions in equilibrium is monotonically increasing in α/K in both the public goods and exploration game.

[16] 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.

[17] If the sample is balanced across treasure sizes, the predicted average number of contributions in the EG equals $1/4 \times 0 + 1/4 \times 1.75 + 1/4 \times 1 + 1/4 \times 1.75 = 1.125$. In the PGG, the prediction equals $1/4 \times 0 + 1/4 \times 0 + 1/4 \times 1 + 1/4 \times 2 = 0.75$.

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

[19] 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 adapt their behavior accordingly.

Indeed, the game is played several times in each of the different public good value specifications: altogether, more than 30 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). However, 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 nonequilibrium models of strategic thinking and learning models typically rely on a higher number of parameters.

[20] Because the novelty and focus in our model and experiment are the information externality channel and because 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, Dufwenberg and Kirchsteiger 2004, Falk and Fischbacher 2006, Cox et al. 2007), and also, the social esteem model of Ellingsen and Johannesson (2008) would make very similar predictions as the social welfare utility model.

[21] Often, it is assumed that $0 \le \rho \le 1/2$ so that weight on the other is never greater than the weight on oneself. 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 ρ and the parameter for aversion for disadvantageous inequality $-\sigma$. The parameters in the model of Fehr and Schmidt (1999) are further constrained by $-\sigma \ge \rho \ge 0$.

[22] Initially, our preelicitation of other-regarding preferences using Murphy et al. (2011) aimed at making individual-specific predictions, but it turned out that there is too little variation in these across individuals (see Online Appendix C). According to the results of the elicitation, the average participant in our study is other regarding, thus justifying the preference model.

[23] There exist parameter values such that this happens, and Proposition A.4 in Online Appendix A shows that this holds quite generally.

[24] This latter is the theoretically predicted other-regarding encouragement effect (that is, the difference between Equations (8) and (9)).

[25] This holds for μ sufficiently high (see Online Appendix A).

[26] To recruit our subjects, we used Online Recruitment System for Economic Experiments (Greiner 2015) 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 rough sample size using List et al. (2011). We assumed a power of 80% and a significance level of 5%, and to have a minimum detectable effect size of 50%, we needed at least 64 observations in each group. At the time, we considered a sample of around 400 participants to be large enough for the normal distribution and a good approximation for the t distribution. In retrospect, we should have done this more carefully and used the effect sizes from previous literature when calculating the sample size. We, therefore, ran a postpower calculation in line with Gelman and Carlin (2014) for the contribution hypothesis. When we calculated the postpower analysis, we used the minimum detectable effect sizes from the results with binary outcomes in Klor et al. (2014) and Steiger and Zultan (2014) (that is, 20%) and our standard error from Table 4, column (1). The results indicate that we have do not have a power issue, and the postpower is high (0.99).

[27] 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.

[28] We thank the programmer Kristaps Dzonsons for his programming assistance.

[29] This did not generate uncontrolled variation, because we randomized groups within each session.

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

[31] Here, we used the effect size from the experiment run in Oslo.

[32] Because we did not fully reach the optimal sample size, we conducted a postpower calculation in line with Gelman and Carlin (2014) to ensure that the result is not biased by issues of low power. Here, we used the effect size from the contribution hypothesis in the main experiment and the standard error from the experiment without the group tournament incentive. The calculation shows that we do not have a power issue in this test (0.97).

[33] See http://bss.au.dk/en/cognition-and-behavior-lab/for-researchers/procedure-guide/.

[34] This percentage is higher compared with Frederick (2005), which can be because of, for example, learning and the fact that the CRT has become better known over time.

[35] 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. 2018 for instance). In this experiment, there are also qualitative patterns in the data that are consistent with the QRE model with other-regarding motivation but not with the QRE model with self-interested individuals or the SPE model with self-interested individuals. For example, we find that the first player contribution rate increases with treasure size and that, approximately 20% of the time, regardless of treasure size, the second player contributes even when the treasure has already been found by the first player in the exploration game. Thus, not only is the explanatory power in statistical terms higher, but also, the observed qualitative treatment effects and other qualitative patterns can be better explained with a QRE model that allows for other-regarding preferences.

[36] Angrist and Pischke 2009, chapters 3.3 and 3.4 support the choice of the linear specification rather than nonlinear alternatives when a saturated model with randomized treatments is used in a panel data setting.

[37] Given our exogenous experimental variation and a saturated model, the linear probability model is the correct specification (Angrist and Pischke 2009, sections 3.3 and 3.4). Yet, even 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 Tables B.4 and B.5 in Online Appendix B.

[38] Note that, because the slider elicitation method does not involve interactive choices, it does not allow estimation of the reciprocity parameter of the Charness and Rabin (2002) model. The second mover choices in the interactive games are likely to reflect the reciprocation preference, especially if the first mover does not contribute, which thus leads to an incentive of negative reciprocation.

[39] See Table 3 for definitions and explanations of $B_g(e)$.

[40] If we run a regular OLS regression, dropping the panel structure and cluster on session, the results stay the same. See Table B.8 in

Online Appendix B. Notice also that the number of observations in this table is lower than that in Table 8, for instance, because we have dropped the second player choices in the EG where the first player already found the treasure. Had we included those observations, we could not exploit the decomposition of the encouragement effect, but rather, the informational effect would be underestimated.

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

[42] See Andersson (2015) for an experiment suggesting causal effects of corporate values on prosocial organizational behavior and Kosse et al. (2018) for an example of such a program targeted at second grade children of low-socioeconomic status families.

[43] The novel experimental framework could, for instance, be used to study experimentation in a private goods setting as well.

[44] In fact, as opposed to the purely public good case presented in the paper, we also considered the case where rewards are purely private goods (see Online Appendix A.2). One can show that entirely privatizing the discovered good would have implications for the encouragement effects both through the other-regarding and through the informational channel. When $\sigma > 0$ and $\frac{\alpha}{K-1} - c_2 > 0$, then encouragement through both channels will be smaller, and thus, contributions will be reduced. The first mover's incentive to contribute is now lower, because there is reason to encourage the other only to the extent that the first mover is altruistic toward the second mover. In fact, if $\sigma \leq 0$, the first mover has no incentive to encourage the second mover, because she can only lose from the second mover's finding the treasure.

References

Admati AR, Perry M (1991) Joint projects without commitment. *Rev. Econom. Stud.* 58(2):259–276.

Aghion P, Tirole J (1997) Formal and real authority in organizations. *J. Political Econom.* 105(1):1–29.

Amabile TM, Hennessey BA, Grossman BS (1986) Social influences on creativity: The effects of contracted-for reward. *J. Personality Soc. Psych.* 50(1):14–23.

Anderson SP, Goeree JK, Holt CA (1998) A theoretical analysis of altruism and decision error in public goods games. *J. Public Econom.* 70(2):297–323.

Andersson O, Huysentruyt M, Miettinen T, Stephan U (2016) Person–organization fit and incentives: A causal test. *Management Sci*. 63(1):73–96.

Angrist JD, Pischke J-S (2009) *Mostly Harmless Econometrics: An Empiricist's Companion* (Princeton University Press, Princeton, NJ).

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

Becker GS (1974) A theory of social interactions. *J. Political Econom.* 82(6):1063–1093.

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

Björk L, Kocher M, Martinsson P, Nam Khanh P (2016) Cooperation under risk and ambiguity. Technical report, Gothenburg University, Sweden.

Bolino MC, Grant AM (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. *Acad. Management Ann.* 10(1):599–670.

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

Bonatti A, Hörner J (2011) Collaborating. *Amer. Econom. Rev.* 101(2): 632–663.

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

Brown M, Flinn CJ, Schotter A (2011) Real-time search in the laboratory and the market. *Amer. Econom. Rev.* 101(2):948–974.

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

Camerer C (2003) *Behavioral Game Theory: Experiments in Strategic Interaction* (Princeton University Press, Princeton, NJ).

Camerer CF, Ho T-H, Chong JK (2004) Behavioural game theory: Thinking, learning and teaching. Huck S, ed. *Advances in Understanding Strategic Behaviour* (Palgrave Macmillan, London), 120–180.

Cappelen AW, Reme B-A, Sørensen EØ, Tungodden B (2015) Leadership and incentives. *Management Sci.* 62(7):1944–1953.

Charness G, Rabin M (2002) Understanding social preferences with simple tests. *Quart. J. Econom.* 117(3):817–869.

Clark K, Sefton M (2001) The sequential prisoner's dilemma: Evidence on reciprocation. *Econom. J.* 111(468):51–68.

Cohen J (1992) A power primer. *Psych. Bull.* 112(1):155.

Compte O, Jehiel P (2003) Voluntary contributions to a joint project with asymmetric agents. *J. Econom. Theory* 112(2):334–342.

Compte O, Jehiel P (2004) Gradualism in bargaining and contribution games. *Rev. Econom. Stud.* 71(4):975–1000.

Cox JC, Friedman D, Gjerstad S (2007) A tractable model of reciprocity and fairness. *Games Econom. Behav.* 59(1):17–45.

Crawford VP, Costa-Gomes MA, Iriberri N (2013) Structural models of nonequilibrium strategic thinking: Theory, evidence, and applications. *J. Econom. Literature* 51(1):5–62.

Deck CA, Kimbrough EO (2019) Experimenting with contests for experimentation. *Southern Econom. J.* 84(2):391–406.

Dickinson DL (1998) The voluntary contributions mechanism with uncertain group payoffs. *J. Econom. Behav. Organ.* 35(4):517–533.

Dohmen T, Falk A, Huffman D, Sunde U, Schupp J, Wagner GG (2011) Individual risk attitudes: Measurement, determinants, and behavioral consequences. *J. Eur. Econom. Assoc.* 9(3):522–550.

Dufwenberg M, Kirchsteiger G (2004) A theory of sequential reciprocity. *Games Econom. Behav.* 47(2):268–298.

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

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

Erat S, Gneezy U (2016) Incentives for creativity. *Experiment. Econom.* 19(2):269–280.

Erev I, Haruvy E (2013) Learning and the economics of small decisions. Kagel JH, Roth AE, eds. *The Handbook of Experimental Economics*, vol. 2 (Princeton University Press, Princeton, NJ), 638–699.

Falk A, Fischbacher U (2006) A theory of reciprocity. *Games Econom. Behav.* 54(2):293–315.

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

Falk A, Fehr E, Fischbacher U (2008) Testing theories of fairness—intentions matter. *Games Econom. Behav.* 62(1):287–303.

Fehr E, Schmidt KM (1999) A theory of fairness, competition, and cooperation. *Quart. J. Econom.* 114(3):817–868.

Fershtman C, Rubinstein A (1997) A simple model of equilibrium in search procedures. *J. Econom. Theory* 72(2):432–441.

Frederick S (2005) Cognitive reflection and decision making. *J. Econom. Perspect.* 19(4):25–42.

Gelman A, Carlin J (2014) Beyond power calculations: Assessing types (sign) and type m (magnitude) errors. *Perspect. Psych. Sci.* 9(6):641–651.

Gilboa I, Postlewaite A, Samuelson L, Schmeidler D (2014) Economic models as analogies. *Econom. J.* 124(578):F513–F533.

Gneezy U, Potters J (1997) An experiment on risk taking and evaluation periods. *Quart. J. Econom.* 112(2):631–645.

Goeree JK, Holt CA (1999) Stochastic game theory: For playing games, not just for doing theory. *Proc. Natl. Acad. Sci. USA* 96(19): 10564–10567.

Goeree JK, Holt CA (2000) Asymmetric inequality aversion and noisy behavior in alternating-offer bargaining games. *Eur. Econom. Rev.* 44(4):1079–1089.

Goeree JK, Holt CA (2001) Ten little treasures of game theory and ten intuitive contradictions. *Amer. Econom. Rev.* 91(5):1402–1422.

Goeree JK, Holt CA, Palfrey TR (2010) Quantal response equilibria. Durlauf SN, Blume LE, eds. *Behavioural and Experimental Economics* (Palgrave Macmillan UK, London), 234–242.

Greiner B (2015) Subject pool recruitment procedures: Organizing experiments with ORSEE. *J. Econom. Sci. Assoc*. 1:114–125.

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

Güth W, Levati MV, Sutter M, Van Der Heijden E (2007) Leading by example with and without exclusion power in voluntary contribution experiments. *J. Public Econom.* 91(5):1023–1042.

Halac M, Kartik N, Liu Q (2016) Optimal contracts for experimentation. *Rev. Econom. Stud.* 83(3):1040–1091.

Harhoff D, Lakhani KR (2016) *Revolutionizing Innovation: Users, Communities, and Open Innovation* (MIT Press, Cambridge, MA).

Hollensbe E, Wookey C, Hickey L, George G, Nichols CV (2014) Organizations with purpose. *Acad. Management J.* 57(5):1227–1234.

Holmström B (1982) Moral hazard in teams. *Bell J. Econom.* 13(2): 324–340.

Hörner J, Skrzypacz A (2017) Learning, experimentation, and information design. *Adv. Econom. Econometrics* 1:63-98.

Ioannidis JP (2005) Why most published research findings are false. *PLoS Medicine* 2(8):e124.

Klor EF, Kube S, Winter E, Zultan R (2014) Can higher rewards lead to less effort? incentive reversal in teams. *J. Econom. Behav. Organ.* 97:72–83.

Kosse F, Deckers T, Pinger PR, Schildberg-Hörisch H, Armin F (2018) The formation of prosociality: Causal evidence on the role of social environment. Working Paper No. 7068, CESifo, Munich, Germany.

Levati MV, Morone A (2013) Voluntary contributions with risky and uncertain marginal returns: The importance of the parameter values. *J. Public Econom. Theory* 15(5):736–744.

Levati MV, Sutter M, Van der Heijden E (2007) Leading by example in a public goods experiment with heterogeneity and incomplete information. *J. Conflict Resolution* 51(5):793–818.

Levitt SD, List JA (2011) Was there really a Hawthorne effect at the hawthorne plant? an analysis of the original illumination experiments. *Amer. Econom. J. Appl. Econom.* 3(1):224–238.

Lindbeck A, Weibull JW (1988) Altruism and time consistency: The economics of fait accompli. *J. Political Econom.* 96(6):1165–1182.

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

List JA, Shaikh AM, Xu Y (2016) Multiple hypothesis testing in experimental economics. NBER Working Paper No. 21875, National Bureau of Economic Research, Cambridge, MA.

Mäki U (2005) Models are experiments, experiments are models. *J. Econom. Methodology* 12(2):303–315.

Maniadis Z, Tufano F, List JA (2014) One swallow doesn't make a summer: New evidence on anchoring effects. *Amer. Econom. Rev.* 104(1):277–290.

March JG (1991) Exploration and exploitation in organizational learning. *Organ. Sci.* 2(1):71–87.

McBride M (2006) Discrete public goods under threshold uncertainty. *J. Public Econom.* 90(6):1181–1199.

McCall JJ (1970) Economics of information and job search. *Quart. J. Econom.* 84(1):113–126.

McKelvey RD, Palfrey TR (1998) Quantal response equilibria for extensive form games. *Experiment. Econom.* 1(1):9–41.

Miettinen T, Kosfeld M, Fehr E, Weibull JW (2018) Revealed preferences in a sequential prisoners' dilemma: A horse-race between five utility functions. Technical report, CESifo, Munich, Germany.

Murphy RO, Ackermann KA, Handgraaf M (2011) Measuring social value orientation. *Judgment Decision Making* 6(8):771–781.

Olson M (1965) *The Logic of Collective Action*, vol. 124 (Harvard University Press, Cambridge, MA).

Osborn AF (1953) Applied imagination.

Palfrey TR, Rosenthal H (1991) Testing for effects of cheap talk in a public goods game with private information. *Games Econom. Behav.* 3(2):183–220.

Plott CR, Smith VL, eds. (2008) *Handbook of Experimental Economics Results*, vol. 1 (North-Holland, Amsterdam).

Samuelson L (2005) Economic theory and experimental economics. *J. Econom. Literature* 43(1):65–107.

Steiger E-M, Zultan R (2014) See no evil: Information chains and reciprocity. *J. Public Econom.* 109:1–12.

Stigler GJ (1961) The economics of information. *J. Political Econom.* 69(3):213–225.

Strausz R (1999) Efficiency in sequential partnerships. *J. Econom. Theory* 85(1):140–156.

Vesely S, Wengström E, et al. (2017) Risk and cooperation: Experimental evidence from stochastic public good games. Working paper, Lund University, Sweden.

Weitzman ML (1979) Optimal search for the best alternative. *Econometrica* 47(3):641–654.

Weitzman ML (1998) Recombinant growth. *Quart. J. Econom.* 113(2): 331–360.

Winter E (2006) Optimal incentives for sequential production processes. *RAND J. Econom.* 37(2):376–390.

Winter E (2009) Incentive reversal. *Amer. Econom. J. Microeconom.* 1(2): 133–147.