

Sequential Equilibrium in Monotone Games: A Theory-Based Analysis of Experimental Data*

Syngjoo Choi[†]
UCL

Douglas Gale[‡]
NYU

Shachar Kariv[§]
UC Berkeley

February 28, 2008

Abstract

A monotone game is an extensive-form game with complete information, simultaneous moves and an irreversibility structure on strategies. It captures a variety of situations in which players make partial commitments and allows us to characterize conditions under which equilibria result in socially desirable outcomes. However, since the game has many equilibrium outcomes, the theory lacks predictive power. To produce stronger predictions, one can restrict attention to the set of sequential equilibria, or Markov equilibria, or symmetric equilibria, or pure-strategy equilibria. This paper explores the relationship between equilibrium behavior in a class of monotone games, namely *voluntary*

*This research was supported by the Center for Experimental Social Sciences (C.E.S.S.) and the C. V. Starr Center for Applied Economics at New York University. We thank Tom Palfrey for detailed comments and suggestions. We also thank an associate editor and two anonymous referees for their comments. The paper has benefited from suggestions by the participants of seminars at several universities, the AEA 2007 annual meeting in Chicago, the ESA Asia-Pacific Regional Meeting at Osaka, and the Cowles Foundation Workshop on Coordination Games at Yale. For financial support, Gale acknowledges National Science Foundation for support under Grant No. SBR-0095109 and the C. V. Starr Center for Applied Economics at New York University, and Kariv thanks UC Berkeley for support under a COR Grant. Kariv is grateful for the hospitality of the School of Social Science in the Institute for Advances Studies.

[†]Department of Economics, University College London, Gower Street, London WC1E 6BT, UK (Email: syngjoo.choi@ucl.ac.uk, URL: <http://www.homepages.ucl.ac.uk/~uctpsc0>).

[‡]Department of Economics, New York University, 19 West 4th Street, #812, New York, NY, 10012 (E-mail: douglas.gale@nyu.edu, URL: <http://www.nyu.edu/econ/user/galed>).

[§]Department of Economics, University of California, Berkeley, 508-1 Evans Hall # 3880, Berkeley, CA 94720, USA (E-mail: kariv@berkeley.edu, URL: <http://socrates.berkeley.edu/~kariv/>).

contribution games, and the behavior of human subjects in an experimental setting. Several key features of the symmetric Markov perfect equilibrium (SMPE) are consistent with the data. To judge how well the SMPE fits the data, we estimate a model of Quantal Response Equilibrium (QRE) (McKelvey and Palfrey 1995, 1998) and find that the decision rules of the QRE model are qualitatively very similar to the empirical choice probabilities.

JEL Classification Numbers: *D82, D83, C92.*

Key Words: *experiment, monotone games, refinements, pure strategy, mixed strategy, Markov perfect, quantal response equilibrium.*

1 Introduction

A major concern of game theorists is to understand the conditions under which self-interested behavior gives rise to cooperative outcomes. To cite one famous example, the Folk Theorem for infinitely repeated games demonstrates that cooperation can be sustained in long run relationships. The Folk Theorem is only partly successful as a theory of cooperative behavior, because it guarantees the existence of a large class of equilibria, some of which are efficient and many of which are not. One response is to introduce more structure into the game. Aumann and Sorin (1989), Lagunoff and Matsui (1997), and Gale (1995) have all shown different ways of adding sufficient structure to guarantee efficient equilibrium outcomes in repeated coordination games, for example.

Here we follow the approach suggested by Gale (1995, 2001) in his study of *monotone games*. Formally, a monotone game is a repeated game in which players are constrained to choose strategies that are non-decreasing over time. This restriction converts a repeated game into a stochastic game. The state of the game at the beginning of the period is the profile of stage-game strategies from the previous period. Because strategies are non-decreasing, a player is making a commitment on his future actions when he chooses his current action. One can think of a monotone game as the dynamic implementation of a static game in which players gradually commit themselves to a final strategy.

In this paper, we explore the relationship between equilibrium behavior in a class of monotone games and the behavior of human subjects in an experimental setting. The class of games we choose to focus on are usually interpreted as *voluntary contribution games*. Suppose there is an indivisible public project, with cost K , and N players, each of whom has an endowment of E tokens. The players make irreversible contributions to the project at a

sequence of dates $t = 1, \dots, T$. At the end of T periods, the public project is carried out if and only if the sum of the contributions is large enough to meet the cost of the project. Each player assigns the value A to the project, so his utility if the project is completed is equal to A *plus* his endowment *minus* his contribution. If the project is not completed, his payoff equals his endowment *minus* his contribution.

The irreversibility of contributions plays a crucial role in the analysis of the monotone game. Although the game cannot be solved by backward induction (because of simultaneous moves), it is possible to use backward induction on the number of contributions to characterize the possible equilibrium outcomes. The central theoretical result states that, if the length of the game T is at least as great as the cost of the good K (and some side constraints are satisfied), then any sequential equilibrium of the game must involve provision of the good with positive probability (with probability one in a pure-strategy equilibrium). In this sense, the theory predicts positive provision in dynamic games satisfying $T \geq K$. By contrast, in static (one-shot) games, there always exists a Nash equilibrium with zero provision (as well as equilibria with positive and certain provision). So, in a static game, unlike the dynamic games, the theory cannot tell us whether to expect zero or positive provision.

Our experimental design is based on parametric examples of the game that serve as treatments in the laboratory experiments we ran. All the examples have three players ($N = 3$), so each example of a monotone game is defined by choosing different values of the horizon T , the value A , the cost K , and the endowment E . The experimental design has a baseline treatment, from which other treatments are derived by varying one parameter at a time. This design allows us to observe the sensitivity of the subjects' behavior to changes in each of the parameters.

Our analysis of the experimental data proceeds in three stages. In the first stage, we are mainly interested in describing the data. We consider the effect of variations in the parameters T , A , K and E on the provision of the public good (Section 5.1). We find significant effects for all of these parameters except for the endowment parameter E . Our central finding is that provision of the public good is high in *all* of the dynamic games. By contrast, in the static (one-shot) game, the provision rate falls close to zero for one parametric specification. The explanation of this observation is postponed until Section 7, after we have reviewed the evidence for an equilibrium interpretation of the data.

In the second stage of the analysis, we begin to explore the relationship between the theory and the experimental data (Section 5.2). Although there

are several nice results that can be proved for these games, there are important questions that theory alone cannot answer. The game has multiple sequential equilibria, so the theory cannot make unambiguous comparative static predictions, apart from the prediction that provision is positive when $T \geq K$. One way to reduce the set of equilibria and increase the predictive power of the theory is to impose restrictions such as symmetry and the Markov property on equilibrium strategies. This leads us to consider the set of symmetric, Markov-perfect equilibria (SMPE). We use aggregate provision data to test for the Markov property and find that in most cases the Markov property cannot be rejected. Even the SMPE is not necessarily unique, but the equilibrium set is drastically reduced and it is possible to make comparative static predictions in most cases. We find the predictions of the SMPE are, with few exceptions, qualitatively consistent with the empirical results from Section 5.1.

The preceding analysis detects both successes and failures of the SMPE in accounting for the features of the data. A more systematic comparison of the theory with the data requires a model that can account for the data in the sense of (a) reproducing important moments of the observed data and (b) explaining the data in terms of the theory. Observed behavior inevitably contains random mistakes, which can be interpreted, following Harsanyi and Selten, as the effect of a “trembling hand.” We cannot expect the theory to fit perfectly; but a version of the theory that allows for random errors may fit the data reasonably well.

For this purpose, we adopt the approach advocated by McKelvey and Palfrey (1995, 1998) and estimate a model of Quantal Response Equilibrium (QRE). The QRE can be thought of as a SMPE with “mistakes.” It incorporates symmetric, Markovian strategies and sequential rationality and adds an error structure to individual choice functions. The parameters of the QRE model measure the extent to which the choice probabilities are influenced by payoffs. A parameter close to zero implies random behavior; an unboundedly large parameter implies rational maximizing behavior. The parameter estimates are highly significant and positive, showing that the theory does predict the subjects’ behavior, and the decision rules generated by the QRE model are qualitatively similar to the empirical choice probabilities.

The rest of the paper is organized as follows. The next section provides a discussion of the closely related literature. Section 3 describes the theoretical model and presents some theoretical results that guided our experimental design. Section 4 describes the experimental procedures. Section 5 describes the experimental results. Section 6 provides the QRE analysis and Section

7 concludes. Proofs are gathered in Section 8.

2 Related literature

There is a small theoretical literature related to monotone games. Admati and Perry (1991) study a two-person game in which players make alternating contributions to the provision of an indivisible public good. They show that, under certain conditions, there is a unique (possibly inefficient) equilibrium. Marx and Matthews (2000) extend the model of Admati and Perry (1991) in a number of ways, including allowing for simultaneous moves. Gale (1995) demonstrates that, in the monotone game based on an infinitely repeated coordination game, all sequential equilibria are efficient in the limit as the length of the time period converges to zero. Gale (2001) introduces a general class of monotone games with positive spillovers, allowing for continuous and unbounded action spaces. He provides conditions under which cooperative outcomes can be supported as sequential equilibria in the absence of indivisibilities.

Duffy, Ochs and Vesterlund (2007), henceforth DOV, report the results of an experiment based on the model of Marx and Matthews (2000). The payoff from the public good has two components: one is proportional to the total contributions, up to the threshold amount needed for completion of the project; the other is a fixed benefit received only if the total contributions reach the threshold needed to complete the project. Their study is based on a 2×2 design in which they compare static versus dynamic games, with and without a positive completion benefit. DOV point out that, in the dynamic game with a completion benefit, there are pure-strategy, Nash equilibria in which the project is completed as well as pure-strategy, Nash equilibria in which no one contributes. In the other games, there is a unique zero-contribution, Nash equilibrium.¹

Experimentally, DOV observe higher contributions in dynamic games than in static games; however, they observe higher contributions even in dynamic games without a positive completion benefit.² They interpret this

¹DOV illustrate the completion equilibrium with what Marx and Matthews (2000) call the grim- g equilibrium. In a grim- g strategy profile, g is played in every period unless the aggregate contribution level is inconsistent with g , in which case no player contributes subsequently. Marx and Matthews (2000) show that g is a Nash equilibrium outcome if and only if the grim- g profile is a Nash equilibrium. However, a grim- g strategy is not sequentially rational after histories where it is feasible for a single player to complete the public project.

²To investigate alternative explanations for why increasing the number of periods in-

behavior, which is not predicted by the theory, as the result of irrationality or the “trembling hand”. We compare our findings (and interpretation) with DOV’s in the concluding section.

3 From theory to design

This section provides a bridge between the theory and the experimental design. First, we develop a number of theoretical results about the game. Then we define a series of simple parametric examples that serve as the treatments in the experimental design and describe their theoretical properties.

3.1 The model

The monotone game we study can be interpreted as a *voluntary contribution game*. The game is formally described using the following notation. There are N players, indexed $i = 1, \dots, N$. Each player is endowed with a finite number E of indivisible tokens. The game is divided into T periods, indexed $t = 1, \dots, T$. In each period, the players simultaneously choose how many tokens to contribute to the provision of the public good. The cost of the public good is K tokens and the public good is provided if and only if the total number of tokens contributed is at least K . If the public good is provided, each player receives A tokens in addition to the number of tokens retained from his endowment. The players have perfect information in the sense that, in each period, they know the full history of the game up to that period. Since contributions are irreversible, the player’s contribution is monotonically non-decreasing over time.

Each game is defined by five parameters (A, E, K, N, T) . These parameters are positive integers except for A , which is a non-negative real number. Each of the parameters influences the set of equilibria of the game in a distinct way. We assume that the aggregate endowment is at least as great as the cost of the public good ($NE \geq K$) so that provision of the public good is always feasible; and we assume that the aggregate value of the good is greater than the cost ($NA > K$) so that provision is always efficient. These assumptions are maintained throughout the rest of this section. We next develop a number of theoretical results about the game. Proofs are gathered in Section 8.

creases contributions, DOV study a variant of the dynamic game with no completion benefit, in which subjects are not informed about past aggregate contributions.

3.2 Coordination

The one-shot game ($T = 1$) provides a useful benchmark for our subsequent analysis of the dynamic games. To avoid trivial cases, we assume that $\min\{A, E\} < K$. This condition ensures that it is either infeasible ($E < K$) or individually irrational ($A < K$) for a single player to provide the public good.

Proposition 1 (one-shot) *Suppose that $T = 1$ and $\min\{A, E\} < K$. Then there exist*

- (i) *a Nash equilibrium in which the good is not provided;*
- (ii) *Nash equilibria in which the good is provided with probability one; and*
- (iii) *a mixed-strategy equilibrium in which the probability of provision is strictly between zero and one.*

In the one-shot game, non-provision of the public good is explained as a *coordination failure*. If a player thinks that no one else will contribute, it is not optimal for him to contribute. Conversely, if he thinks that his contribution is both necessary and sufficient for provision, then he will be happy to contribute. Hence, provision of the good can be supported, in spite of the free-rider problem, because each contributing player is pivotal (Bagnoli and Lipman, 1992).³

Turning to dynamic games ($T > 1$), two features are essential for ensuring cooperation. The first is the requirement that strategies be non-decreasing over time. The second is the assumption of sequential rationality.

³The importance of indivisibilities in voluntary contribution games with public goods was originally pointed out by Bagnoli and Lipman (1992) in the context of a static game. They show that the indivisibility of the public good makes each contributing player “pivotal” in the sense that, *at the margin*, his contribution is both necessary and sufficient for provision. The Bagnoli-Lipman, pivotal-player argument solves the free-rider problem in the sense that it guarantees the *existence* of efficient equilibria in games of any length; however, it is also consistent with the existence of inefficient and zero-provision equilibria. As in Bagnoli and Lipman (1992), the indivisibility of the public good is a crucial feature of the theoretical model used in this paper.

The model with continuous and unbounded strategy sets used in Gale (2001) requires an infinite horizon if cooperation is to be supported as a sequential equilibrium outcome. Obviously, in the laboratory an infinite horizon is problematic and this is why we have opted for a model with discrete contributions.

Andreoni (1998) examines a static threshold public good game. Propositions 2 and 3 in Andreoni (1998) follow from the Bagnoli-Lipman “pivotal” player argument, which applies to games of any length, and are similar to Proposition 1 in this paper. Adding taxation and government grants in Andreoni (1998) is equivalent to changing the cost of the public good (K) and the individual endowments (E) in our model.

In general, games that are infinitely repeated and/or involve simultaneous moves, cannot be solved by backward induction, but the monotonicity assumption makes it possible to use backward induction on the payoff-relevant state of the game, rather than the stage of the game. The irreversibility of contributions and the backward induction logic allow players to coordinate their actions and support cooperative outcomes. The sharpest result is obtained for the case of pure-strategy equilibria.

Proposition 2 (pure strategy) *Suppose that $A > E$ and $T \geq K$. Then, under the maintained assumptions $NE \geq K$ and $NA > K$, in any pure-strategy, sequential equilibrium of the game, the public good is provided with probability one.*

The logic of the proof can be illustrated by an example. Suppose that there are three players ($N = 3$) and each player has an endowment of one token ($E = 1$). There are two periods ($T = 2$), the cost of the public good is two ($K = 2$), and the value of the public good is greater than the individual endowment ($A > 1$). In any pure-strategy equilibrium, the probability of provision is either one or zero, so it is enough to show that the zero-provision equilibrium is not sequentially rational. Suppose, contrary to the claim in the proposition, that there exists a pure-strategy, sequential equilibrium with zero provision. Then every player's payoff is simply the value of his endowment $E = 1$. If one player contributes a token at period 1, one of the remaining players can earn at least $A > 1$ by contributing his endowment at period 2. Thus, the good must be provided at period 2 if one player contributes at period 1. Anticipating this response, it is clearly optimal for someone to contribute a token at period 1.

The assumption $T \geq K$ in Proposition 2 highlights the central role of backward induction in the analysis of the monotone game. It takes time to ensure cooperation and without enough time there may exist a no-provision equilibrium. The condition $A > E$ is needed to avoid equilibria in which players refuse to contribute because they are afraid of contributing too much. Proposition 5 below shows what happens when endowments are "too high."

Although pure-strategy equilibria give us a very clean result, they do not tell the whole story. To get a more robust result, we should take account of mixed strategies. Mixed strategies are interesting for at least two reasons. First, they are one way of rationalizing uncertain provision of the public good, something that is observed in the experimental data. Secondly, they expand the set of parameters for which there exists a zero-provision equilibrium. The preceding example, where $N = 3$, $A > E = 1$, and $T = K = 2$,

can be used to illustrate this possibility. If one player contributes a token in period 1, the continuation game at period 2 consists of two active players, only one of whom needs to contribute a token in order to provide the good. The continuation game possesses a symmetric, mixed-strategy equilibrium where each of the two active players contributes in period 2 with probability $0 < \lambda < 1$.

A necessary and sufficient condition for λ to be an equilibrium strategy is that each of the two uncommitted players be indifferent between contributing and not contributing. A simple calculation shows that indifference requires

$$\lambda = (A - 1)/A.$$

In this mixed-strategy equilibrium, the good is provided unless neither of the two players contributes, that is, the good is provided with probability $1 - (1 - \lambda)^2$. Then, if the player who contributes in period 1 anticipates his opponents will play the symmetric mixed-strategy equilibrium at period 2, it is rational for him to contribute if and only if

$$\left[1 - (1 - \lambda)^2\right] A \geq 1$$

or $A^2 - A \geq 1$. The critical value of A is thus $A^* \approx 1.618$. For any $A < A^*$ it is not optimal for a player to move first if he anticipates that the other two will play the mixed-strategy equilibrium, whereas in the pure-strategy case it is optimal to move first if $A > 1$. Thus, the use of mixed strategies in the continuation game can discourage an initial contribution and support an equilibrium with no provision. In the two-period game, $A > A^*$ is necessary and sufficient for provision of the good with positive probability. In fact, this condition is sufficient for positive provision in games of any length $T \geq K$. We summarize this discussion in the following proposition.

Proposition 3 (mixed strategy) *Suppose the parameters of the game are $N = 3$, $E = 1$, and $T \geq K = 2$. If $1 < A < A^*$, where $A^* \approx 1.618$, then there exists a mixed strategy equilibrium in which the good is provided with probability zero.*

A generalization of Proposition 3 is contained in Section 8 (Proposition 6).

3.3 Free-riding

Under-provision of the public good depends on two factors, coordination failure and the free-rider problem. As we have seen, coordination failure

occurs most readily in the one-shot game ($T = 1$), where players would all be better off if the good were provided, but are unable to escape from the bad equilibrium. In a dynamic game, the irreversibility of contributions and the backward induction logic allow players to coordinate their actions. The free-rider problem arises because each player would like someone else to contribute on his behalf. The severity of the free-rider problem depends on the extent to which it is possible to make unequal contributions in equilibrium. Games in which $K = NE$ provide a useful benchmark, because in these games each player must contribute his entire endowment in order for the public good to be provided. There is no possibility of taking a free ride on the contributions of other players.

To illustrate, suppose that $N = 3$, $A > E = 1$, and $T \geq K = 3$. In any sequential equilibrium of this game, each player can guarantee himself a payoff equal to $A > 1$ tokens. To see this, suppose that two players have already contributed by the end of period $T - 1$. Then the remaining player will contribute in period T and each of the players receives A tokens. Now suppose that one player has already contributed by the end of period $T - 2$. Then by the previous argument, either of the two other players can guarantee himself a payoff of A tokens by contributing in period $T - 1$. Clearly, any of the three players can guarantee himself A tokens by contributing in period $T - 2$. Our next result generalizes the preceding argument. Note that it does not rule out the use of mixed strategies, even along the equilibrium path.

Proposition 4 (no-free-riding) *Suppose that $K = NE$, $A > E$ and $T \geq K$. Then the good is provided with probability one in any sequential equilibrium of the game.*

Taking the condition $K = NE$ as a benchmark for the absence of free riding, the free-rider problem must be worse, in some sense, when the total endowment exceeds this level. In fact, if the total endowment is too large, non-provision is consistent with sequential equilibrium, as the following result shows. It provides a kind of converse to Proposition 2, which assumes an upper bound $E < A$ on the size of the individual endowment.

Proposition 5 (free-riding) *Suppose that $E > A$ and $T \geq K$. Then under the maintained assumptions, there exists a pure-strategy, sequential equilibrium of the game in which the public good is provided with probability zero.*

The essential ingredient in the construction of this equilibrium is the self-punishing strategy employed by Gale (2001). Whoever contributes first condemns himself to making a contribution so large that it outweighs his benefit from the public good. After the first contribution has been made, it is a sunk cost and the player cannot stop himself from making further individually rational contributions. The other players are, of course, only too happy to wait until he has punished himself for deviating. Anticipating the ultimate outcome, the player will never make that first, fatal contribution. To illustrate, note that if $E = K = 2$ and $1 < A < 2$, it is clearly not worthwhile for a single player to contribute two tokens since the value of the good is only $A < 2$. On the other hand, if a player is foolish enough to contribute one token, it is rational for him to contribute a second token later, since $A > 1$ and the first token is now a sunk cost. Hence, if a player contributes one token, he is condemned to contribute the second and will end up worse off than if he had never contributed at all.

3.4 Parametric examples

We next define a series of simple parametric examples that serve as the treatments in the experimental design. The four games we consider consist of a baseline game and three variants that are derived from the baseline by changing one parameter in each case. Throughout, we keep the number of players constant ($N = 3$) and consider two values of the public good, one *high* ($A = 3$) and one *low* ($A = 1.5$). The dynamic games come in two versions, a *long* version ($T = 5$) and a *short* version ($T = 2$). The set of sequential equilibria differs across treatments, particularly with respect to the existence of sequential equilibria with zero provision. In addition, varying the value of the good A and the length of the game T provides a test of the robustness of our results. The parametric examples are summarized in Table 2 below.

- The **baseline** treatment uses the parameters $A = 1.5, 3$; $E = 1$; $K = 2$; $T = 2, 5$. The game has a variety of sequential equilibria. The pure-strategy equilibria all involve provision of the good with probability one (Proposition 2), whereas the mixed-strategy equilibria allow for a positive probability that the good is not provided (Proposition 3). In the high-value treatments ($A = 3$), all sequential equilibria are characterized by a positive probability of the provision of the public good, whereas in the low-value treatments ($A = 1.5$) there exists a mixed-strategy sequential equilibrium in which the good is provided with probability zero (Proposition 3).

- ▶ The **high-cost** treatment is identical to the baseline treatment except that the cost of the public good has been increased to $K = 3$, so provision of the good requires every player to contribute ($K = NE$). In this game, there are two factors that affect the rate of provision of the public good. On the one hand, because every player must contribute in order to provide the public good, there is no possibility of taking a free ride here. On the other hand, the high cost of providing the good does allow for the existence of a zero-provision equilibrium when the horizon is short (Proposition 4 does not apply to this case because $K = 3 > T = 2$).
- ▶ The **high-endowment** treatment is the same as the baseline treatment except that the endowment is increased to $E = 2$. By increasing the individual endowments, we increase the potential asymmetry of contributions to the provision of the public good and hence the potential for free riding. According to Proposition 5, there exists a pure-strategy, zero-provision equilibrium of this game when the value is low ($A = 1.5$).
- ▶ The **one-shot** treatment is the same as the baseline treatment except that $T = 1$. Every game, static or dynamic, possesses a pure-strategy, sequential equilibrium in which the good is provided. The one-shot game differs from the dynamic games satisfying $T \geq K$, however, because the former always possesses a zero-provision equilibrium (Proposition 1) whereas in the latter, provision must be positive if $T \geq K$ and certain other conditions are satisfied.

Thus, if we observe positive provision in the dynamic games with $T \geq K$, but not in the static game ($T = 1$), we may take this as evidence of the effectiveness of backward induction reasoning.

3.5 Equilibrium refinements

Because each of the parametric examples has a large number of equilibria, the theory does not make strong predictions. For this reason, we need experimental data to assess first, whether the equilibria are useful in accounting for observed behavior and, secondly, whether some equilibria are more empirically relevant than others. Since we have too many equilibria to start with, we use standard equilibrium refinements to reduce the set of equilibria we have to deal with. One refinement is the restriction to symmetric strategies; the other is the restriction to Markov strategies. These

refinements are commonly used in the game theory literature and this is a valuable opportunity to test their empirical relevance.

A strategy is called *Markov* if it depends only on payoff-relevant variables. By limiting the variables on which behavior is conditioned, the Markov property reduces the set of sequential equilibria, sometimes substantially. When each player has an endowment of one token ($E = 1$), the payoff-relevant states of the game are denoted by (n, τ) , where n is the total number of contributions and τ is the number of periods remaining after the current period. When each player has an endowment of two or more tokens ($E \geq 2$), the payoff relevant states for subject i are denoted by (n, τ, n_i) , where n and τ have the usual meaning and n_i is the number of contributions to date by player i .

If we assume that strategies are symmetric and Markovian, we are led to consider the class of *symmetric Markov perfect equilibria* (SMPE), which take a relatively simple form. A general characterization of the SMPE is provided in Section 8 (Proposition 7). Table 1 reports the strategies corresponding to the SMPE in the baseline and high-endowment games.⁴

Note that a game with horizon $T < T'$ is isomorphic to a continuation game of the game with horizon T' , starting (with zero contributions) in period $T' - T$. So the SMPE of the short-horizon treatments ($T = 2$) are the same as the last two periods of the SMPE in the long-horizon treatments ($T = 5$), assuming there have been no contributions in the first three periods. Similarly, the SMPE of the static treatments correspond to the last period of the SMPE in the long-horizon treatments ($T = 5$), again assuming there have been no contributions in the first four periods.

Note also that, in order to provide the public good in the high-cost game, every player must contribute, but the timing of the contributions is immaterial. As a result, the game has a large number of SMPE.

[Table 1 here]

The main predictions of the SMPE presented in Table 1 can be summarized by four facts.

- First, there are no *pure-strategy* SMPE, although mixed strategies are only used off the equilibrium path when the value is low ($A = 1.5$).

⁴In the equilibria described in Table 1, we assume, as part of our definition of symmetric equilibrium, that players use symmetric Markov strategies both on and off the equilibrium path.

- ▶ Secondly, when the value of the good is high ($A = 3$) (top panels), there are *three* mixed-strategy SMPE in which the good is provided with positive probability. When the value is low (bottom panels), there is a *unique* SMPE, in which the good is provided with probability zero.
- ▶ Thirdly, within each game, the SMPE predict no provision of the public good in early periods when the value is high and no provision at all when the value is low. The use of mixed strategies in the continuation games at later periods discourages an initial contribution and, especially when the value of the public good is low enough, supports a unique SMPE with no provision.
- ▶ Fourthly, the SMPE predicts that, other parameters being equal, the contribution probability at each state is at least as high when the value is high ($A = 3$) as when the value is low ($A = 1.5$) because the probability is always zero in the low-value treatments.

Notice also that, for the baseline and high-endowment games, if the good is provided with probability one, the contributions must be asymmetrically distributed among the players in any pure-strategy sequential equilibrium. If the good is not provided, then mixed strategies are required off the equilibrium path in order to discourage a deviation. By contrast, when we increase the cost of the public good to $K = 3$ holding other parameters constant, provision of the good requires every player to contribute ($K = NE$). Thus, the high-cost game has pure-strategy SMPE in which every player contributes with probability one. On the other hand, the high cost of providing the good does allow for the existence of a pure-strategy SMPE with zero provision in the case where the value of the good is low ($A = 1.5$).

3.6 Summary

Table 2 summarizes the various games used in the experimental design and the most relevant equilibrium properties for each game. The right hand column (i) lists the propositions that characterize the most relevant properties for each treatment. Columns (ii)-(iv) list the parameters used in each game, and columns (v)-(vii) summarize the equilibrium properties. Recall that each game, static or dynamic, possesses both a pure-strategy sequential equilibrium with certain provision and, except the high-cost game when the horizon is long, a mixed-strategy sequential equilibrium with a positive probability of provision strictly less than one. Thus, columns (v) and (vi)

point out whether each game possesses a zero-provision, pure- and mixed-strategy, sequential equilibrium, respectively. As we have already seen, the use of mixed strategies in dynamic games, even if they are only used off the equilibrium path, can support zero-provision equilibrium for a strictly larger set of parameters than pure strategies can. Finally, the right hand column (vii) reports the probabilities that the good is provided in the SMPE.

[Table 2 here]

4 Experimental procedures

The experiment was run at the Experimental Economics Laboratory of the Center for Experimental Social Sciences (C.E.S.S.) at New York University. The subjects in this experiment were recruited from undergraduate classes at NYU. Throughout the experiment we ensured anonymity and effective isolation of subjects in order to minimize any interpersonal influences that could stimulate cooperation. After subjects read the instructions, the instructions were read aloud by an experimental administrator. Sample instructions are reproduced in an online appendix.

Each experimental session comprised 18 subjects (except for two sessions in which 15 subjects were used) and 15 independent decision-rounds (except for the one-shot treatment, which comprised 30 rounds). A single treatment (A, E, K, N, T) was used for each session and each treatment was used for one session. In each round, the subjects were randomly formed into six (respectively, five) three-person groups. So for each dynamic game we have observations on $6 \times 15 = 90$ (respectively, $5 \times 15 = 75$) different rounds and a total of $18 \times 90 = 1520$ (respectively $18 \times 75 = 1330$) individual decisions.

The three-person groups formed in each round were independent of the groups formed in any of the other rounds. At the beginning of the round, each subject had an endowment of E tokens. In the first period, each subject was asked to allocate his tokens to either an x -account or a y -account. Investing a token in the y -account was irreversible. After all subjects had made a decision, each subject observed the decisions of all the subjects in his group. In the second period, each subject was asked to allocate the token(s) remaining in his x -account between the two accounts. At the end of this period, each subject again observed the decisions of all the subjects in his group. This procedure was repeated until T decisions had been made.

When the first round ended, the computer informed subjects of their payoffs. Earnings in each round were determined as follows: if subjects contributed at least K tokens to their y -accounts, each subject received A

tokens plus the number of tokens remaining in his x -account. Otherwise, each subject received the number of tokens in his x -account only. This process was repeated until all decision rounds were completed.

Subjects were paid according to the number of tokens they earned in each round. Each token was worth 50 cents, so the total payment to a subject was equal to \$0.50 times the total number of tokens earned in 15 rounds, plus a \$5 show-up fee. Subjects received their payment privately as they left the experiment.

5 Experimental results

5.1 Provision

We begin with an analysis of the sensitivity of provision to changes in T and then consider the sensitivity of provision to changes in the other parameters, A , K and E . Table 3 presents the provision rates for each treatment.

[Table 3 here]

Unless otherwise noted, throughout this section, we use chi-square non-parametric tests (see Siegel and Castellan, 1988). These tests assume that the observations are independent random variables. This would be true, for example, if the subjects in a given session use identical mixed strategies. If there is heterogeneity among subjects, however, the outcomes of games in which the same subjects appear will not be independent and this will bias downward the standard errors, making it more likely that we will find a significant treatment effect. We have not found a simple way to account for dependence, so we have used the null of independence while recognizing that it may not be satisfied in this case.

5.1.1 The time effect (T)

The static or one-shot game differs from the baseline game only with respect to the number of periods T . The first thing we want to check is whether provision rates are higher in dynamic games. Referring to Table 3, when the value of the good is high ($A = 3$), the provision rates are 0.76 when $T = 1$, 0.74 when $T = 2$, and 0.81 when $T = 5$. A joint chi-square nonparametric test indicates that the difference in the provision rates is not statistically significant (p -value = 0.530). By contrast, when the value of the good is low ($A = 1.5$), the provision rate decreases sharply to 0.10 when $T = 1$, while remaining at 0.47 when $T = 2$ and 0.678 when $T = 5$. These provision rates

are significantly different at all conventional significance levels (p -value = 0.000). Hence, the effect of time length on public-good provision is significant when the value of the public good is low, but not when the value of the public good is high.

The data summarized in Table 3 also show differences in the provision rates corresponding to $T = 2$ and $T = 5$ for both the high-cost and high-endowment games. When the value of the good is high ($A = 3$), provision rates in the high-cost treatments are 0.88 and 0.99 when $T = 2$ and $T = 5$, respectively. Although the provision rates are both high, they are significantly different at the 1 percent significance level (p -value = 0.003). When the value of the good is low ($A = 1.5$), provision rates decrease to 0.30 and 0.64 for $T = 2$ and $T = 5$, respectively. The difference in provision rates is significant at the 1 percent level (p -value = 0.000). Finally, holding the value of the good A constant, the differences in the provision rates in the high-endowment treatments are significant at the 1 and 10 percent levels for $T = 2$ and $T = 5$, respectively (the respective p -values are 0.008 and 0.051). In summary, the data supports the following result.

Result 1 (time) *(i) Comparing the one-shot game with the baseline game, when the value of the good is low ($A = 1.5$), the provision rate is highest when $T = 5$ and lowest when $T = 1$. When the value of the good is high ($A = 3$), there are no significant differences between the provision rates in the static and baseline treatments.*

(ii) Within the high-cost and high-endowment treatments, the provision rate is significantly higher when $T = 5$ than when $T = 2$.

5.1.2 The value effect (A)

Since all games come in low-value ($A = 1.5$) and high-value ($A = 3$) versions, we can study the effect of the value of the good within each game. In the low-value version of the baseline game, the provision rates are 0.47 and 0.68 for $T = 2$ and $T = 5$, respectively. These provision rates are significantly lower than those in the high-value version of the baseline game at the 1 and 5 percent levels for $T = 2$ and $T = 5$, respectively (the p -values are 0.000 and 0.040). Similarly, the provision rates in the high-endowment treatments are 0.49 and 0.63 for $T = 2$ and $T = 5$, respectively, when the value is low. These rates are significantly lower than the corresponding rates when the value is high. The differences are significant at the 1 percent level (the respective p -values are 0.006 and 0.001). Finally, the provision rates in the low-value treatments are also significantly lower than those in the cor-

responding high-value treatments in both the one-shot and high-cost games (p -values = 0.000). Our next result summarizes these findings.

Result 2 (value) *Within each game, the provision rate is significantly higher when the value of the good is high ($A = 3$) than when it is low ($A = 1.5$), holding the game length T constant.*

5.1.3 The cost effect (K)

To explore the effects of the cost of the public good, we compare the provision rate in the baseline game ($K = 2$) with that in the high-cost game ($K = 3$), holding the game length T and the value of the good A constant. The provision rates in the high-value treatments are 0.88 and 0.99, when $T = 2$ and $T = 5$, respectively, confirming that the absence of the free rider problem ($K = NE$) leads to high provision. These provision rates are significantly higher than those in the corresponding baseline treatments at the 1 and 5 percent significance levels (the respective p -values are 0.000 and 0.022). When the value of the good is low ($A = 1.5$), there is no significant difference between the provision rates in the high-cost and baseline treatments when $T = 5$. Surprisingly, the provision rate is significantly lower in the high-cost treatment than in the low-cost treatment when $T = 2$ (p -value = 0.021). In summary, the data support the following result.

Result 3 (cost) *Comparing the high-cost game ($K = 3$) with the baseline game ($K = 2$) while holding the length T constant, the provision rates are higher in the high-cost treatment when the value of the good is high ($A = 3$). When the value of the good is low ($A = 1.5$), the provision rates are not significantly different when $T = 5$ and the provision rate is significantly lower in the high-cost treatment when $T = 2$.*

5.1.4 The endowment effect (E)

Finally, we compare the provision rate in the baseline game ($E = 1$) with that in the high-endowment game ($E = 2$), holding constant the game length T and the value of the good A . Referring to Table 3, when the value of the good is high ($A = 3$), provision rates in the high-endowment treatments with $T = 2$ and $T = 5$ are 0.69 and 0.86, respectively. When the value of the good is low ($A = 1.5$), the corresponding provision rates decrease to 0.49 and 0.63, respectively. However, none of the differences in the provision rates between the baseline and high-endowment treatments is statistically significant at any conventional significance level. We thus report the following result.

Result 4 (endowment) *The provision rates in the high-endowment game ($E = 2$) are not significantly different from those in the baseline game ($E = 1$), holding the value A and length T constant.*

5.2 Properties of equilibrium

In this section we discuss the evidence for the use of refinements. We first organize the data by calculating the relative frequencies of contributions for each of the payoff-relevant states in each dynamic treatment. The data are presented in Table 4. The payoff-relevant states of the game are represented by (n, τ) when $E = 1$ and by (n, τ, n_i) for each player i when $E = 2$, where n is the total number of contributions, τ is the number of periods remaining after the current period, and n_i is the total number of contributions to date by player i . The number in parentheses in each cell represents the number of subjects who have an endowment left for contribution. In a few cases subjects contributed a token when the public good had already been provided. These contributions appear to be simple mistakes of the “trembling hand” variety, rather than a systematic misunderstanding of the game or irrational behavior.

[Table 4 here]

5.2.1 Mixed and pure strategies

We are only interested in the aggregate contribution behavior and therefore focus on the average frequency of contributions in different states. From this point of view, if every subject uses a pure strategy, the average contribution rate will equal zero or one. By contrast, if every subject uses a mixed strategy, the contribution rate may lie between zero and one. As we have mentioned above, there exist mixed-strategy equilibria in which pure strategies are used exclusively on the equilibrium path. In any case, we are only looking for broad indications of pure- or mixed-strategy behavior, not trying to test for their use at the individual level, still less to test for equilibrium behavior. The most we can say from studying aggregate contribution rates is that rates close to one or zero are consistent with pure strategies and provision rates strictly between zero and one are consistent with mixed strategies.

In the high-cost game ($K = 3$), where most contributions are made in the first two periods, the overall pattern is consistent with the use of pure strategies. Referring to Table 4, when the horizon is short ($T = 2$) and the value of the good is high ($A = 3$), the relative frequencies of contributions

at states $(n, \tau) = (0, 0), (1, 0)$ are close to zero and one respectively. These cases approximate, on the one hand, a pure-strategy equilibrium with no provision and, on the other, a pure-strategy equilibrium with certain provision, at the corresponding states. Note that, given this behavior in the continuation games, mixing at the first period is clearly not a best response. By contrast, in other treatments the relative frequency of contributions in the last two periods is similar to the probability of contributions in (symmetric) mixed-strategy equilibria, as will be shown below. As pointed out in our theoretical discussion, the use of mixed strategies can discourage contributions in earlier periods and support equilibria with zero-provision or with a probability of provision strictly less than one. Even if we do not assume equilibrium play, the uncertainty of provision that is evident in the data may have the effect of discouraging subjects' willingness to contribute.

Referring back to Table 3, in all of the dynamic treatments, apart from the high-cost treatment, the provision rate is positive but significantly less than one. Such behavior is clearly not consistent with pure-strategy equilibrium. By contrast, in the high-cost game, when the horizon is long ($T = 5$) and the value of the good is high ($A = 3$), the probability the good is provided is virtually one and this is consistent with the use of pure-strategies. As noted above, even in the case where provision occurs with probability one, one cannot conclude that players are necessarily using pure strategies. There exist mixed-strategy equilibria, including some in which mixing occurs on the equilibrium path, in which the good is provided with probability one. The most that one can say is that this outcome is consistent with the use of pure strategies. The next result summarizes the findings.

Result 5 (mixed-strategies) *We find aggregate behavior consistent with*

- (i) non-provision, pure-strategy equilibria (e.g., the one-shot game with $A = 1.5$);*
- (ii) full-provision, pure-strategy equilibria (e.g., the high-cost game $A = 3$); and*
- (iii) mixed-strategy equilibria (e.g., the baseline game).*

5.2.2 Markov behavior

Our next question is whether subjects' behavior is consistent with Markov strategies. By definition, Markov strategies depend only on the payoff-relevant states and not on the complete history of play. In order to have a meaningful test for Markov behavior, it is necessary to identify states that have multiple histories leading to that state. This requirement eliminates

most states for the following reasons. First, in games with a short horizon ($T = 2$), all states have a single history, so we can only use data from the treatments with $T = 5$. Similarly, in games with a long horizon ($T = 5$), the states with $\tau = 4, 3$ all have a unique history. Secondly, all states with $n = 0$ have a single history, so we can only use states in which $n > 0$. Thirdly, when $n = 2$, it is never a best response to contribute and, with few exceptions, we observe no contributions in this case. This is an example of Markovian behavior, to be sure, but it is a rather trivial one and not worth subjecting to a formal test. So we are left with states in which $n = 1$ and $\tau = 0, 1, 2$. There are three such states, $(n, \tau) = (1, 2), (1, 1), (1, 0)$, where one token has already been contributed in the past. In the high-endowment game ($E = 2$), we only consider the case where the individual contribution variable $n_i = 0$ in each of the three states.

We focus on the baseline and high-endowment treatment. In the high-cost treatment ($K = 3$), the group outcomes are so highly clustered that it does not seem possible to have a meaningful test.

The procedure for testing the Markov property is as follows. There are $(4 - \tau)$ distinct histories leading to each of the payoff relevant states $(n, \tau) = (1, 2), (1, 1), (1, 0)$. In each of these histories, exactly one token has been contributed at one of the previous dates $t = 1, \dots, 4 - \tau$. Let $h(t, \tau)$ denote the history at period $T - \tau$ where one token was contributed at time period t and none at other dates. For example, when $(n, \tau) = (1, 2)$, there are two different histories reaching the current state: one where one token was contributed at $t = 1$, $h(1, 2)$, and the other where the contribution of one token was made at $t = 2$, $h(2, 2)$. The (joint) null hypothesis at each of the three states $(n, \tau) = (1, 2), (1, 1), (1, 0)$ is that the relative frequencies of contributions are the same for all histories reaching the state (n, τ) .

Table 5 below summarizes the relative frequencies of contributions from the different histories reaching each payoff-relevant state in the baseline and high-endowment treatments with $T = 5$. The number in parentheses in each cell represents the number of observations. At the last column of Table 5, p -values are reported under the null hypothesis. The numbers in parentheses in each cell of the last column represent the value of the chi-square test statistic and its degrees of freedom. In both high-endowment treatments, we cannot reject the Markov restriction in all three states at the 10 percent significance level. In each of the baseline treatments, we reject the null hypothesis only in one out of the three states at the 10 percent significance level. Thus, the data supports the following result.

Result 6 (Markov) *The overall outcomes in the test of the Markov prop-*

erty confirm that subjects' behavior supports the Markov specification.

[Table 5 here]

5.2.3 Symmetric Markov perfect equilibrium

If we add symmetry to the Markov property, we are led to consider the ability of the symmetric, Markov perfect, equilibrium (SMPE) to account for the broad features of the experimental data. In what follows, we again ignore the high-cost game ($K = NE$) where pure-strategy equilibria seem like a reasonable explanation for the data. For the other dynamic games, we compare the empirical patterns of contributions in Table 4 with the theoretical predictions of SMPE in Table 1. Perhaps surprisingly, we find the qualitative features of the theory and the experimental data very similar, although there are systematic deviations. Referring to Table 1 above, the relative frequencies of contributions in the first three periods are around 10 percent in both the baseline and high-endowment games. The strategic delay in early periods reflects the theory's prediction of zero contributions and is clearly robust to changes in the model parameters. Any deviation from the SMPE must increase contributions in the early periods and this probably increases the provision of the public good overall. The data also confirm that provision of the public good is higher than the SMPE predict. This is most clearly seen in the treatments where the value is low ($A = 1.5$).

Referring to Table 4, the SMPE predict that the contribution probabilities in a low-value treatment will be lower in the last two periods than those in the corresponding high-value treatment. It is interesting to note that the comparative-static prediction of the SMPE holds true: in the baseline treatments, the contribution probability at each state in the high-value treatment ($A = 3$) is no lower than that in the low-value treatment ($A = 1.5$). In summary, the dynamic patterns of contribution behavior are sensitive to changes in the parameters, especially the value of the public good. Such sensitivity of contribution behavior is often closely related to the predictions of SMPE. Nonetheless, the empirical contribution probabilities are often higher than those from the SMPE in the treatments where the value is low. The deviation in the subjects' behavior from SMPE in the last two periods is toward a lower probability of contribution. Hence, the analysis detects both successes and failures of the SMPE in accounting for different features of the data, as summarized in the next result.

Result 7 (SMPE) *The SMPE explains the qualitative patterns of contributions in the games. Nonetheless, there are systematic departures*

from the predictions of the SMPE. In early periods, they contribute more often than the SMPE predicts and, in later periods, less often than the SMPE predicts.

We can also use the SMPE to interpret the earlier empirical results about the sensitivity of provision rates to various parameter changes. Result 1 (time) contrasts the levels of provision in the static ($T = 1$) and baseline ($T = 2, 5$) games. When the value is high ($A = 3$) there is no significant difference in provision rates, whereas the difference is significant when the value is low ($A = 1.5$). Table 1 shows that when the value is low the SMPE predicts no provision, whereas it predicts positive provision when the value is high. Thus, at least when the value is high, the empirical behavior is consistent with the SMPE. When the value is low, by contrast, there is a clear deviation from the SMPE that will be examined in the context of the Quantal Response Equilibrium (QRE) in the next section.

The other results on provision rates are all consistent with the qualitative predictions of the SMPE. Result 2 (value) shows that higher values of the public good correspond to higher provision, *ceteris paribus*, and that is what the SMPE predicts as well. Result 3 (cost) shows that the provision rate is higher in the high-cost game ($K = 3$) than in the baseline game ($K = 2$). In the high-cost game, when the horizon is long ($T = 5$), all sequential equilibria result in provision of the good with probability one and this holds *a fortiori* for the SMPE. When the horizon is short ($T = 2$), there are SMPE with zero provision as well as with certain provision. Hence, the SMPE predicts the empirical results when the horizon is long and is at least consistent with them when the horizon is short. Result 4 (endowment) shows that provision rates are not significantly different between the high-endowment ($E = 2$) and baseline ($E = 1$) games. Table 1 indicates that the contribution probabilities are very similar for corresponding SMPE, albeit slightly lower when the endowment is high. It is fair to say that the SMPE predicts quantitatively similar provision rates in the baseline and high-endowment games.

6 Quantal Response Equilibrium (QRE)

The preceding analysis detects both successes and failures of the SMPE in accounting for different features of the data. A more systematic comparison of the theory with the data requires a model that can account for the data in the sense of (a) reproducing important moments of the observed data and (b) explaining the data in terms of the theory. Observed behavior inevitably

contains random mistakes, which can be interpreted, following Harsanyi and Selten, as the effect of a “trembling hand.” We cannot expect the theory to fit perfectly; but a version of the theory that allows for random errors may fit the data reasonably well.

The QRE model of McKelvey and Palfrey (1995, 1998) allows for the possibility of errors and yet retains a role for best-response behavior. At each state, a player’s contribution probability is assumed to be a function of the difference between the payoffs from contributing and not contributing. The predictions of the QRE model differ from those of the SMPE for two reasons: first, the QRE allows for the possibility of mistakes and, secondly, it assumes that players take into account the fact that others are also making mistakes.

6.1 Specification

To explain the subject’s propensity to choose an action different from the one predicted by the SMPE, we assume that each player’s payoff is perturbed by an idiosyncratic preference shock that has an extreme value distribution. For simplicity, suppose that each player has an endowment of one token ($E = 1$). Then, the contribution behavior of each uncommitted player at state (n, τ) follows a binomial logit distribution:

$$\lambda_{(n,\tau)} = \frac{1}{1 + \exp(-\beta_{(n,\tau)}\Delta_{(n,\tau)})}, \quad (1)$$

where $\lambda_{(n,\tau)}$ is the equilibrium probability of making a contribution, $\Delta_{(n,\tau)}$ is the difference between the expected payoffs from contributing and not contributing, and $\beta_{(n,\tau)}$ is a coefficient to be estimated. The choice of action becomes purely random as $\beta_{(n,\tau)}$ goes to zero, whereas the optimal action is chosen almost surely as $\beta_{(n,\tau)}$ goes to infinity. For positive values of $\beta_{(n,\tau)}$, the choice probability $\lambda_{(n,\tau)}$ is increasing in $\Delta_{(n,\tau)}$.

To calculate the payoff difference $\Delta_{(n,\tau)}$, we need to take account of the stochastic behavior of other players in the current and future periods. Specifically, it is assumed that players correctly predict the contribution probabilities in the continuation games as well as in the current state and use them to calculate the correct value of $\Delta_{(n,\tau)}$. In effect, we assume that players have rational expectations about the true error rates determined by the value of $\beta_{(n,\tau)}$ and use the estimated $\beta_{(n,\tau)}$ coefficients to approximate the true $\beta_{(n,\tau)}$. Furthermore, since the payoff differences $\Delta_{(n,\tau)}$ depend on the contribution probabilities $\lambda_{(n,\tau)}$ and *vice versa*, we need to solve for the

equilibrium values of payoff differences and contribution probabilities, taking as given the parameters (A, E, K, N, T) that define the game.

6.2 Solving for QRE: An example

We illustrate the method of solving for a QRE of the model using the baseline treatment with a short horizon ($T = 2$). The calculation of equilibrium proceeds by backward induction, beginning with the second and final period.

In the second period, the contribution probability is denoted by $\lambda_{(n,0)}$ where n is the number of contributions in the first period and $\tau = 0$ is the number of periods remaining. For $n = 1, 2$, the payoff from contributing is A and the payoff from not contributing is $1 + A(\lambda_{(n,0)})^{2-n}$. Thus, the difference between the expected payoffs from contributing and not contributing is given by

$$\Delta_{(n,0)} = A - 1 - A(\lambda_{(n,0)})^{2-n}$$

for $n = 1, 2$. Similarly, for $n = 0$, the payoff from contributing is $A \left[1 - (1 - \lambda_{(0,0)})^2 \right]$ and the payoff from not contributing is $1 + A(\lambda_{(0,0)})^2$. Then the payoff difference is

$$\Delta_{(0,0)} = A(1 - (1 - \lambda_{(0,0)})^2) - 1 - A(\lambda_{(0,0)})^2.$$

The logistic formula (1) can be solved numerically for specific values of the parameters A and $\beta_{(n,\tau)}$ to yield a *finite* number of values of the equilibrium contribution probabilities $\lambda_{(2,0)}$, $\lambda_{(1,0)}$ and $\lambda_{(0,0)}$.

Using these contribution probabilities $\lambda_{(2,0)}$, $\lambda_{(1,0)}$ and $\lambda_{(0,0)}$, we can calculate the continuation payoffs at the beginning of the second period. Let $\pi_{(n,0)}^0$ denote the continuation payoff for an *uncommitted* player when the total number of prior contributions is n . Then

$$\begin{aligned} \pi_{(0,0)}^0 &= \lambda_{(0,0)}A \left[1 - (1 - \lambda_{(0,0)})^2 \right] + (1 - \lambda_{(0,0)}) \left[1 + A(\lambda_{(0,0)})^2 \right], \\ \pi_{(1,0)}^0 &= \lambda_{(1,0)}A + (1 - \lambda_{(1,0)}) (1 + A\lambda_{(1,0)}), \\ \pi_{(2,0)}^0 &= \lambda_{(2,0)}A + (1 - \lambda_{(2,0)}) (1 + A). \end{aligned}$$

Similarly, let $\pi_{(n,0)}^1$ denote the continuation payoff for a *committed* player when the number of prior contributions is n . Then

$$\pi_{(1,0)}^1 = A \left[1 - (1 - \lambda_{(1,0)})^2 \right]$$

and $\pi_{(2,0)}^1 = \pi_{(3,0)}^1 = A$.

The equilibrium contribution probability in the first period is denoted by $\lambda_{(0,1)}$, where $n = 0$ is the number of prior contributions and $\tau = 1$ is the

number of periods remaining. Having calculated the continuation payoffs for every possible outcome in the second period, we can use equation (1) to solve for $\lambda_{(0,1)}$, where the difference between the expected payoffs from contributing and not contributing in the first period is given by

$$\begin{aligned} \Delta_{(0,1)} = & \left[(1 - \lambda_{(0,1)})^2 \pi_{(1,0)}^1 + 2\lambda_{(0,1)} (1 - \lambda_{(0,1)}) \pi_{(2,0)}^1 + (\lambda_{(0,1)})^2 \pi_{(3,0)}^1 \right] - \\ & - \left[(1 - \lambda_{(0,1)})^2 \pi_{(0,0)}^0 + 2\lambda_{(0,1)} (1 - \lambda_{(0,1)}) \pi_{(1,0)}^0 + (\lambda_{(0,1)})^2 \pi_{(2,0)}^0 \right]. \end{aligned}$$

Equation (1) can be solved numerically for each value of the parameters A and β and for each value of the second-period contribution probabilities $\lambda_{(2,0)}$, $\lambda_{(1,0)}$ and $\lambda_{(0,0)}$. The solutions for $\lambda_{(2,0)}$ and $\lambda_{(1,0)}$ are unique for each value of A and β , but there are multiple solutions for $\lambda_{(0,0)}$ when the value is high ($A = 3$). Furthermore, given a single set of parameters A and β and solutions $\lambda_{(2,0)}$, $\lambda_{(1,0)}$ and $\lambda_{(0,0)}$, the solution for $\lambda_{(0,1)}$ may not be unique. Thus, multiple QRE cannot be avoided in general.

6.3 Econometric results

We have estimated the QRE using data from the baseline game only. In the high-cost game ($K = 3$), the group outcomes are highly clustered and seem to be well explained by the pure-strategy equilibrium. In the high-endowment game ($E = 2$), there are more states (n, τ, n_i) and fewer observations for each state, so the estimates are bound to be less precise than in the baseline case. We estimated the QRE under several different restrictions on the parameters $\beta_{(n,\tau)}$: (i) a single parameter $\beta_{(n,\tau)} = \beta$ for all (n, τ) ; (ii) a single parameter for each τ , that is, we impose the restriction $\beta_{(n,\tau)} = \beta_{(n',\tau)}$ for all n, n' ; and (iii) a single parameter for each value of n , that is, we impose the restriction $\beta_{(n,\tau)} = \beta_{(n,\tau')}$ for all τ, τ' ; (iv) unrestricted $\beta_{(n,\tau)}$.

The experimental data consists of vectors of the form $\left\{ x_{(n,\tau)}^i \right\}_{i=1}^{m_{(n,\tau)}}$ for each state (n, τ) , where $n = 0, 1, 2$ and $\tau = 0, 1, \dots, T-1$, $x_{(n,\tau)}^i \in \{0, 1\}$ denotes a player's action, and $m_{(n,\tau)}$ denotes the total number of observations at state (n, τ) . Then the log-likelihood function for the entire data set is

$$\begin{aligned} & \mathcal{L} \left(\left\{ \beta_{(n,\tau)}; \left\{ x_{(n,\tau)}^i \right\}_{i=1}^{m_{(n,\tau)}} \right\}_{n=0,1,2 \text{ and } 0 \leq \tau \leq T-1} \right) \\ = & \sum_{\substack{n=0,1,2 \\ 0 \leq \tau \leq T-1}} \sum_{i=1}^{m_{(n,\tau)}} \left[x_{(n,\tau)}^i \log \lambda_{(n,\tau)} + (1 - x_{(n,\tau)}^i) \log (1 - \lambda_{(n,\tau)}) \right]. \end{aligned}$$

For each specification (i)-(iv), we seek to maximize the likelihood of the data with respect to the parameters $\beta_{(n,\tau)}$ subject to (1). The QRE restriction can be solved numerically for any value of β . We use the maximum likelihood (ML) method for estimating β (the program is available from the authors on request). For each value of the parameter β , we solve the model for the QRE contribution probabilities $\lambda_{(n,\tau)}$ for $n = 0, 1, 2$ and for $\tau = 0, 1$ when $T = 2$ (respectively, $\tau = 0, \dots, 4$ when $T = 5$). Then, using these probabilities, we calculate the likelihood of the empirical data we have observed. This procedure is repeated for each value of β as we search for the value that maximizes the likelihood. If there is more than one QRE solution at some states for a given β , we select the one that gives the highest likelihood.

The results are reported in Table 6. The parameter estimates of $\beta_{(n,\tau)}$ are positive and highly significant at each state (n, τ) , showing that the theory does help predict the subjects' behavior. Comparing the predictions of the most parsimonious single-parameter QRE specification with the empirical patterns of contributions in Table 4 and the theoretical predictions of SMPE in Table 1 shows that the QRE captures the subjects' tendency to make early contributions, something which the SMPE cannot reproduce. The estimated QRE also generates a pattern of contributions that is qualitatively similar to the SMPE and is very close in the last two periods, when most of the action occurs.

[Table 6 here]

Finally, Figure 1 compares the empirical contribution probabilities with the predicted logit contribution probabilities from the single-parameter model when the horizon is long ($T = 5$), split by high-value (Figure 1A) and low-value (Figure 1B) treatments. Each series is a function of the payoff differences estimated by the QRE. The predicted logit choice probabilities across treatments are graphed using the corresponding β estimates. The horizontal axis measures the difference between the expected payoffs from contributing and not contributing, $\Delta_{(n,\tau)}$, and the vertical axis measures the probability of making a contribution, $\lambda_{(n,\tau)}$. These graphical comparisons give a rough indication of goodness of fit and suggest that the logistic specification of the QRE model is consistent with the data.

[Figure 1 here]

7 Conclusion

We have undertaken an experimental investigation of a class of monotone games — voluntary contribution games — and focused on using several equilibrium refinements to interpret the data generated by the experiments. Our approach is to let theory drive design and our results suggest that the theory adequately accounts for several large-scale features of the data. Most importantly, although the multiplicity of equilibria means the theory lacks predictive power, several qualitative features of equilibrium match the data surprisingly well. Additionally, in sharp contrast to the experimental results in one-shot settings, there is a very high level of provision of the public good in the dynamic games. Furthermore, both pure- and mixed-strategy equilibria and several key features of the symmetric Markov perfect equilibrium are replicated in the data. Finally, to interpret the data more systematically, we estimate the model of QRE. The estimation results suggest that the QRE model does help predict the subjects' behavior. The QRE replicates the tendency of early contributions in games, which could not be captured by the SMPE, and at the same time predicts the choice behavior at later periods quite precisely.

DOV also find that increasing the number of periods appears to increase contributions, even in treatments where no player contributes in the unique Nash equilibrium outcome. These findings might suggest that the trembling hand plays a role in increasing the provision rates in our dynamic games. However, our finding that the pattern of contributions differs significantly across treatments, *while holding the time horizon constant*, implies that something other than pure trembling is needed to explain the high provision rates. For example, in the high-cost game ($K = NE$), most subjects contributed early, as if they anticipated the effect of the treatment on their opponents' willingness to contribute. Moreover, in the high-value treatment ($A = 3$) there is no horizon effect whereas in the low-value treatment ($A = 1.5$) there is a significant horizon effect. Neither of these features is explained by a tendency to “tremble.” In fact, the tendency in the baseline treatment when the horizon is long ($T = 5$) is to wait until the last two periods. This is inconsistent with an explanation based on the assumption that a certain number will contribute in each period, and emphasizes the role of backward induction in supporting cooperative outcomes.

The theory developed in the paper is complex and subtle and the experimental design we use produces a rich data set. We have tried to use the full spectrum of the games' theoretical properties in our empirical study of monotone games. At the same time, we emphasize that our objective is

not to test a single hypothesis or theoretical prediction. Nor do we consider alternative behavioral explanations or individual heterogeneity. We have examined a single data set from the point of view of the theory. More work is required, but we have made a start and shown that the equilibrium theory underlying the QRE takes us some way toward accounting for the aggregate behavior of experimental subjects.

8 Proofs

Let x_{it} denote the number of tokens contributed by player i up to the end of period t and let $x_t = (x_{1t}, \dots, x_{Nt})$ denote the profile of cumulative contributions at the end of period t . We refer to x_t as the state of the game in period t . Since a player's choice is irreversible, the state is monotonically non-decreasing over time, $x_t \geq x_{t-1}$, $\forall t$. By convention, $x_0 \equiv 0$. A player's payoff is a function of the terminal state x_T . The payoff function for player i is denoted by $U_i(x_T)$ and defined by

$$U_i(x_T) = \begin{cases} A + E - x_{iT} & \text{if } \sum_{i=1}^n x_{iT} \geq K; \\ E - x_{iT} & \text{if } \sum_{i=1}^n x_{iT} < K. \end{cases}$$

Proof of Proposition 1. If player i expects his opponents to contribute nothing, the best response is to contribute nothing. If he contributes $x_i > 0$, then either $x_i < K$ and his payoff is $E - x_i < E$ or $E \geq x_i \geq K > A$ and his payoff is $A + E - x_i \leq A + E - K < E$. In either case, he is better off contributing nothing. On the other hand, we can find at least one outcome x^* such that $\sum_{i=1}^N x_i^* = K$ and $A + E - x_i^* \geq E$ for every i . Note that a player will be worse off if he increases his contribution and a player who makes a positive contribution in equilibrium will be no better off if he reduces it because in that case the good is not provided. Thus, x^* is a Nash equilibrium. ■

Let k denote the number of tokens that need to be contributed in order for the public good to be provided and let τ denote the number of time periods remaining in the game, unless otherwise specified.

Proof of Proposition 2. If $k = 1$ and $\tau \geq 1$ the good will be provided in equilibrium as long as $A > 1$. If not, a single uncommitted player would provide the good and increase his payoff. Since only pure strategies are allowed, the good is provided with probability one in any sequential equilibrium of the game with $k = 1$ and $\tau \geq 1$.

Now consider an arbitrary $k < K$ and $\tau \geq k$ and assume that, for any sequential equilibrium of the game with k tokens needed and τ periods remaining, the good is provided for certain. Then consider the game with $k + 1$ tokens needed and $\tau + 1$ periods remaining. If a player contributes one token in the game, the good is provided for certain and his payoff will be at least A . Suppose, contrary to what we want to prove, that there exists a sequential equilibrium of this game in which the good is provided with probability zero. In such an equilibrium the player's payoff cannot be greater than E . Clearly, the player would be better off contributing the token and ensuring the good is provided. The desired result follows by induction. ■

Proof of Proposition 3. Suppose that $T = 2$. We construct the necessary mixed-strategy equilibrium. If $k = 0$, the best response for any uncommitted player is to contribute nothing. If $k = 1$, there is only one period left (one token has already been contributed in the first period) and there are two uncommitted players. We assume they play the unique symmetric mixed-strategy equilibrium of the one-shot game. If $k = 2$, we assume that no one contributes. To show that this strategy profile is a sequential equilibrium, it is sufficient to show that no player wants to deviate in the first period. But we have already shown that if $A < A^*$ and the player anticipates a mixed strategy equilibrium in the continuation game, it is not worth contributing in the first period. Thus, we have the required equilibrium.

The proof for $T > 2$ uses a variation of the argument in the proof for $T = 2$. Suppose, contrary to what we want to prove, there exists a sequential equilibrium in which provision is zero. Then the last two periods of the game are isomorphic to the game with $T = 2$. Suppose one player deviates by contributing a token in period $T - 1$. There is a unique continuation equilibrium at T in which both players randomize with the probability that makes the other indifferent. As we showed earlier, if $A > A^*$ the probability that the good is provided is sufficiently high that the first player is strictly better off by deviating. This contradicts the equilibrium conditions and proves the corollary. ■

Proof of Proposition 4. If $k = 1$ and $\tau \geq 1$ not contributing is strictly dominated by contributing a token since $A > E \geq 1$. Thus, the good must be provided with probability one in any sequential equilibrium of the game with $k = 1$ and $\tau \geq 1$. Note that we do allow the use of mixed strategies along the equilibrium path.

Now consider an arbitrary $k < K$ and $\tau \geq k$ and assume that, for

any sequential equilibrium of the game with k tokens needed and τ periods remaining, the good is provided for certain. Then consider the game with $k + 1$ tokens needed and $\tau + 1$ periods remaining. If an uncommitted player contributes one token in the game, the good is provided for certain and his payoff will be A . Suppose, contrary to what we want to prove, that there exists a sequential equilibrium of this game in which the good is provided with strictly less than probability one. In such an equilibrium, the player's payoff is less than the payoff A . Clearly, the player would be strictly better off contributing the token and ensuring the good is provided. The desired result follows by induction. ■

Proof of Proposition 5. We construct a pure-strategy equilibrium by considering a number of possible situations. If no one has contributed in the past, so that $k = K$, then no player contributes in the current or future periods. Now suppose that $0 < k < K$ and that exactly one player has contributed a positive amount M , say, where $M < [A + 1]$, the largest integer less than $A + 1$. Then that player is assumed to contribute $[A + 1] - M$ in the current period.

If $0 < k < K$ and $M \geq [A + 1]$ we can construct a pure strategy equilibrium along the lines of Proposition 2 to ensure that the good is provided with probability 1. The efficiency of provision implies that we can do this without requiring any player to contribute more than $[A + 1]$ tokens and we assume this property is satisfied in what follows. Finally, if $k = 0$ no player contributes in the current or future periods.

To show that this is a pure strategy equilibrium, we need to prove two facts. First, we need to show that it is optimal for the distinguished player to contribute $[A + 1] - M$ when called on to do so. By construction, this contribution is less than A and ensures that the good is provided in the pure-strategy continuation equilibrium, so the player is clearly better off making the contribution. Secondly, we need to show that no player wants to deviate by contributing when $k = K$. By construction, any player who deviates when $k = K$ will end up contributing $[A + 1] \geq A$, so the deviation cannot make him better off. This completes the proof. ■

Proposition 6 (provision) *For any positive integers E, K, N , and T satisfying $K < \min \{NE, T\}$ there exists a number $A^*(E, K, N, T)$ such that, for any $A \geq A^*(E, K, N, T)$, the probability that the public good is provided is positive in any sequential equilibrium of the game defined by (A, E, K, N, T) .*

Proof. If $k = 1$ and $\tau \geq 1$ the probability of provision in equilibrium must be positive as long as $A > 1$. If not, a single uncommitted player would provide the good and increase his payoff. Thus, the probability of provision must be positive in any sequential equilibrium of the game with $k = 1$ and $\tau \geq 1$.

Now consider an arbitrary $k < K$ and $\tau \geq k$ and assume that $A_k > 1$ and $\lambda_k > 0$ are such that, for any sequential equilibrium of the game with k tokens needed and τ periods remaining, the good is provided with probability at least λ_k if $A \geq A_k$. Then consider the game with $k + 1$ tokens needed and $\tau + 1$ periods remaining. If an uncommitted player contributes one token in the game, the good is provided with probability at least λ_k and his payoff will be at least A .

Suppose, contrary to what we want to prove, that there exists a sequential equilibrium of this game in which the good is provided with probability zero. In such an equilibrium the player's payoff cannot be greater than E . Then choose $A_{k+1} = \frac{2E}{\lambda_k}$. Then $\lambda_k A \geq 2E > E$ for any $A \geq A_{k+1}$. By forcing the provision of the good, the player can guarantee an equilibrium payoff greater than E , contradicting our assumption that the good is provided with probability zero. In fact, the probability λ that the public good is provided must satisfy

$$\lambda \geq \lambda_{k+1} \equiv \frac{\lambda_k A - E}{A} \geq \lambda_k - \frac{\lambda_k}{2} = \frac{\lambda_k}{2} > 0.$$

This proves by induction that there are numbers A_K and $\lambda_K > 0$ such that the probability of provision is at least λ_K in any game (A, E, K, N, T) satisfying $A \geq A_K \equiv A^*(E, K, N, T)$, as required. ■

Proposition 7 (SMPE) *A symmetric Markov perfect equilibrium (SMPE) of the voluntary contribution game is a choice probability $\lambda(m, \tau)$, satisfying the conditions:*

$$\sum_{k=0}^{N-1-n} B(k; N-n-1, \lambda(n, \tau)) [V(1, n+k+1, \tau-1) - V(0, n+k, \tau-1)] = 0,$$

if $1 > \lambda(m, \tau) > 0$;

$$\sum_{k=0}^{N-1-n} B(k; N-n-1, \lambda(n, \tau)) [V(1, n+k+1, \tau-1) - V(0, n+k, \tau-1)] \geq 0,$$

if $\lambda(m, \tau) = 1$; and

$$\sum_{k=0}^{N-1-n} B(k; N-n-1, \lambda(n, \tau)) [V(1, n+k+1, \tau-1) - V(0, n+k, \tau-1)] \leq 0,$$

if $\lambda(m, \tau) = 0$; where

$$V(a, n, 0) = U_i(a, n), \text{ for } a \in \{0, 1\}.$$

Proof. The SMPE strategy λ must satisfy the following equilibrium condition for every state (n, τ) :

$$\begin{aligned} & \sum_{k=0}^{N-1-m} B(k; N-n-1, \lambda(n, \tau)) \{ \lambda(n, \tau) V(1, n+k+1, \tau-1) \\ & \quad + (1 - \lambda(n, \tau)) V(0, n+k, \tau-1) \} \\ & \geq \sum_{k=0}^{N-1-n} B(k; N-n-1, \lambda(n, \tau)) \{ \lambda V(1, n+k+1, \tau-1) \\ & \quad + (1 - \lambda) V(0, n+k, \tau-1) \}, \end{aligned}$$

for any $\lambda \in [0, 1]$ and where $B(k; N-n-1, \lambda)$ represents the probability of k successes in $N-n-1$ independent Bernoulli trials with probability of success λ in each trial and $V(a, n, \tau)$ denotes the expected utility of a player who has chosen $a \in \{0, 1\}$ when the current state is (n, τ) .

In order to derive the equilibrium condition, we need to apply backward induction. First, it is easy to see that the value function at the terminal date is the same as the payoff function. Then, at the final date T and at the state x_T where $\sum_{i=1}^N x_{iT} = n < K$, player i who has not contributed yet should be indifferent between contributing and not contributing, given the fact that others who have not contributed use an equilibrium choice probability $\lambda(n, 0)$:

$$\sum_{k=0}^{n-1-n} B(k; N-n-1, \lambda(n, 0)) [U_i(1, n+k) - U_i(0, n+k)] = 0.$$

Since $U_i(1, n) - U_i(0, n) = A - 1$ when $n = K - 1$ and -1 when $n \neq K - 1$, we have

$$B(K-1-n; n-1-n, \lambda(n, 0)) = \frac{1}{A}.$$

And, analogously, at any date t and at the state (n, τ) where $\sum_{i=1}^N x_{it} = n < K$, in order to get an equilibrium choice probability $\lambda(n, \tau)$, we can apply the condition of the indifference between contributing and not contributing, which results in the equilibrium condition. ■

References

- [1] Admati A. and M. Perry (1991) "Joint Projects Without Commitment." *Rev. Econ. Stud.* **58**, pp. 259-276.
- [2] Andreoni J. (1998) "Toward a Theory of Charitable Fund-Raising." *J. of Pol. Econ.* **106**, pp. 1186-1213.
- [3] Aumann R. and S. Sorin (1989) "Cooperation and Bounded Recall." *Games and Econ. Behav.* **1**, pp. 5-39.
- [4] Bagnoli M. and B. Lipman (1992) "Private Provision of Public Goods Can Be Efficient," *Public Choice* **74**, pp. 59-78.
- [5] Duffy J., J. Ochs and L. Vesterlund (2007) "Giving Little by Little: Dynamic Voluntary Contribution Games." *J. of Pub. Econ.* **91**, pp. 1708-1730.
- [6] Gale D. (1995) "Dynamic Coordination Games." *Econ. Theory* **5**, pp. 1-18.
- [7] Gale D. (2001) "Monotone Games with Positive Spillovers." *Games and Econ. Behav.* **37**, pp. 295-320.
- [8] Lagunoff R. and A. Matsui (1997) "An 'Anti-Folk Theorem' for a Class of Asynchronously Repeated Games." *Econometrica* **65**, pp. 1467-1477.
- [9] McKelvey R. and T. Palfrey (1995) "Quantal Response Equilibria for Normal Form Games." *Games and Econ Behav.* **10**, pp. 6-38.
- [10] McKelvey R. and T. Palfrey (1998) "Quantal Response Equilibria for Extensive Form Games." *Exp. Econ.* **1**, pp. 9-41.
- [11] Marx L. and S. Matthews (2000) "Dynamic Voluntary Contribution to a Public Project." *Rev. Econ. Stud.* **67**, pp. 327-358.
- [12] Siegel and Castellan (1988). *Nonparametric Statistics for the Behavioral Science*. McGraw-Hill.

Table 1. The probability of contribution in SMPE across games

Baseline				High-endowment						
A=3				A=3						
τ/n	0			1	$\tau/(n, n_i)$	0			(1,0)	(1,1)
4	0.00				4	0.00				
3	0.00			0.00	3	0.00			0.00	0.00
2	0.00			0.00	2	0.00			0.00	0.00
1	0.56	0.55	0.00	0.00	1	0.50	0.48	0.00	0.00	0.00
0	0.00	0.21	0.79	0.67	0	0.00	0.21	0.79	0.42	0.42

A=1.5				A=1.5						
τ/n	0			1	$\tau/(n, n_i)$	0			(1,0)	(1,1)
4	0				4	0				
3	0			0	3	0			0	0
2	0			0	2	0			0	0
1	0			0	1	0			0	0
0	0			0.33	0	0			0.21	0.21

Table 2. Summary of the equilibrium properties

Game	Parameters			Equilibrium properties		
	E, K	T	A	Pure	Mixed	SMPE
Baseline (Proposition 2, 3)	1, 2	2	1.5	N	Y	0
			3	N	N	.62, .62, .89
		5	1.5	N	Y	0
			3	N	N	.62, .62, .89
High-cost (Proposition 5)	1, 3	2	1.5	Y	Y	0, 1
			3	Y	Y	1
		5	1.5	N	N	0, 1
			3	N	N	1
High-endowment (Proposition 6)	2, 2	2	1.5	Y	Y	0
			3	N	N	.63, .62, .89
		5	1.5	Y	Y	0
			3	N	N	.63, .62, .89
One-shot (Proposition 1)	1, 2	1	1.5	Y	N	0
			3	Y	N	0, .51, .99
(i)	(ii)	(iii)	(iv)	(v)	(vi)	(vii)

Table 3. Provision rate by treatment

Game	T	A	Provision	Std. Err.
Baseline ($E=1, K=2$)	2	1.5	0.467	0.053
	2	3	0.744	0.046
	5	1.5	0.678	0.050
	5	3	0.811	0.041
High-cost ($E=1, K=3$)	2	1.5	0.300	0.049
	2	3	0.878	0.035
	5	1.5	0.644	0.051
	5	3	0.989	0.011
High endowment ($E=2, K=2$)	2	1.5	0.489	0.053
	2	3	0.689	0.049
	5	1.5	0.633	0.051
	5	3	0.856	0.037
One-shot ($E=1, K=2$)	1	1.5	0.100	0.025
	1	3	0.761	0.032

Table 4. The relative frequencies of contribution at payoff-relevant states across treatments

Baseline
E=1, K=2

T=5, A=3			
τ/n	0	1	2
4	0.09 (270)		
3	0.08 (207)	0.11 (38)	0 (2)
2	0.11 (165)	0.07 (54)	0.25 (8)
1	0.37 (117)	0.07 (76)	0.10 (10)
0	0.36 (36)	0.60 (94)	0.08 (24)

T=5, A=1.5			
τ/n	0	1	2
4	0.09 (270)		
3	0.05 (207)	0.03 (36)	0 (3)
2	0.06 (177)	0.06 (54)	0.25 (4)
1	0.26 (144)	0.19 (70)	0.17 (6)
0	0.20 (57)	0.48 (88)	0.09 (23)

T=2, A=3			
τ/n	0	1	2
1	0.18 (270)		
0	0.62 (159)	0.54 (54)	0 (9)

T=2, A=1.5			
τ/n	0	1	2
1	0.18 (270)		
0	0.35 (150)	0.33 (64)	0 (7)

High-cost
E=1, K=3

T=5, A=3			
τ/n	0	1	2
4	0.76 (270)		
3	0.67 (6)	0.42 (24)	0.89 (36)
2	N/A	0.25 (8)	0.58 (12)
1	N/A	0.33 (6)	0.20 (5)
0	N/A	1 (2)	0.83 (6)

T=5, A=1.5			
τ/n	0	1	2
4	0.38 (270)		
3	0.15 (66)	0.37 (78)	0.79 (24)
2	0.10 (39)	0.23 (40)	0.68 (31)
1	0.74 (27)	0.94 (32)	0.29 (17)
0	0.00 (21)	0.10 (30)	0.33 (15)

T=2, A=3			
τ/n	0	1	2
1	0.68 (270)		
0	0.22 (9)	0.82 (38)	0.95 (39)

T=2, A=1.5			
τ/n	0	1	2
1	0.37 (270)		
0	0.93 (75)	0.35 (68)	0.75 (28)

High-endowment
E=2, K=2

T=5, A=3			
$\tau / (n, n_i)$	(0,0)	(1,0)	(1,1)
4	0.14 (270)		
3	0.03 (165)	0.02 (52)	0.12 (26)
2	0.07 (153)	0.04 (50)	0.08 (25)
1	0.3 (126)	0.08 (60)	0 (30)
0	0.53 (45)	0.46 (84)	0.26 (42)

T=5, A=1.5			
$\tau / (n, n_i)$	(0,0)	(1,0)	(1,1)
4	0.06 (270)		
3	0.05 (228)	0.09 (22)	0.00 (11)
2	0.13 (195)	0.05 (40)	0.15 (20)
1	0.21 (126)	0.07 (70)	0.00 (35)
0	0.04 (63)	0.39 (92)	0.07 (46)

T=2, A=3			
$\tau / (n, n_i)$	(0,0)	(1,0)	(1,1)
1	0.34 (270)		
0	0.44 (75)	0.34 (70)	0.11 (35)

T=2, A=1.5			
$\tau / (n, n_i)$	(0,0)	(1,0)	(1,1)
1	0.26 (270)		
0	0.13 (111)	0.38 (70)	0.00 (35)

Table 5. The relative frequencies of contributions from the different histories

Game	A	(n, τ)	$h(1)$	$h(2)$	$h(3)$	$h(4)$	p-value
Baseline ($E=1, K=2$)	1.5	(1,2)	0.03 (34)	0.10 (20)	–	–	0.63 (0.23, 1)
		(1,1)	0.06 (32)	0.25 (16)	0.32 (22)	–	0.05 (6.20, 2)
		(1,0)	0.54 (28)	0.25 (8)	0.30 (10)	0.52 (42)	0.30 (3.67, 3)
	3	(1,2)	0.00 (30)	0.17 (24)	–	–	0.07 (3.24, 1)
		(1,1)	0.00 (30)	0.06 (18)	0.14 (28)	–	0.21 (3.15, 2)
		(1,0)	0.47 (30)	0.75 (18)	0.60 (20)	0.64 (28)	0.27 (3.91, 3)
High endowment ($E=2, K=2$)	1.5	(1,2)	0.56 (18)	0.45 (22)	–	–	0.25 (1.33, 1)
		(1,1)	0.00 (10)	0.05 (20)	0.10 (40)	–	0.50 (1.40, 2)
		(1,0)	0.50 (10)	0.33 (18)	0.47 (32)	0.31 (32)	0.12 (5.91, 3)
	3	(1,2)	0.05 (44)	0.00 (6)	–	–	0.10 (2.70, 1)
		(1,1)	0.11 (38)	0.00 (6)	0.06 (16)	–	0.60 (1.02, 2)
		(1,0)	0.43 (30)	0.67 (6)	0.57 (14)	0.41 (34)	0.53 (2.23, 3)

Table 6A. QRE estimation results and the probability of contribution

A single parameter

$T=5, A=3$				$T=5, A=1.5$			
Log_lik = -472.52, $\beta=10.05$ (1.12)				Log_lik = -475.01, $\beta=12.34$ (0.95)			
τ/n	0	1	2	τ/n	0	1	2
4	0.11			4	0.08		
3	0.14	0.07	0.00	3	0.09	0.06	0.00
2	0.18	0.10	0.00	2	0.12	0.08	0.00
1	0.20	0.17	0.00	1	0.19	0.13	0.00
0	0.75	0.65	0.00	0	0.00	0.36	0.00

$T=2, A=3$				$T=2, A=1.5$			
Log_lik = -278.55, $\beta = 10.51$ (1.30)				Log_lik = -296.41, $\beta = 2.26$ (0.11)			
τ/n	0	1	2	τ/n	0	1	2
1	0.19			1	0.4		
0	0.76	0.65	0	0	0.3	0.42	0.09

Table 6B. QRE estimation results and the probability of contribution

A single parameter for each τ

$T=5, A=3$
Log_lik = -442.42

τ/n	0	1	2	β
4	0.09			20.08 (2.76)
3	0.09	0.05	0.00	23.94 (3.77)
2	0.11	0.14	0.00	16.30 (2.68)
1	0.31	0.13	0.00	461.88 (3.97E+25)
0	0.62	0.58	0.23	1.19 (0.33)

$T=5, A=1.5$
Log_lik = -406.19

τ/n	0	1	2	β
4	0.09			11.15 (1.33)
3	0.05	0.03	0.00	15.28 (1.98)
2	0.07	0.04	0.00	14.14 (1.73)
1	0.33	0.20	0.01	5.62 (3.14)
0	0.22	0.42	0.07	2.61 (0.27)

$T=2, A=3$
Log_lik = -271.76

τ/n	0	1	2	β
1	0.18			3.45E+10 (1.61E+12)
0	0.65	0.59	0.16	1.67 (0.41)

$T=2, A=1.5$
Log_lik = -267.39

τ/n	0	1	2	β
1	0.18			17.74 (7.05)
0	0.33	0.43	0.11	2.08 (0.25)

Table 6C. QRE estimation results and the probability of contribution

A single parameter for each n

$T=5, A=3$
Log_lik = -446.84

τ/n	0	1	2
4	0.11		
3	0.13	0.03	0.13
2	0.16	0.04	0.11
1	0.21	0.09	0.09
0	0.76	0.66	0.08
β	10.52 (0.64)	31.67 (16.76)	2.51 (1.11)

$T=5, A=1.5$
Log_lik = -457.43

τ/n	0	1	2
4	0.08		
3	0.09	0.04	0.16
2	0.10	0.06	0.13
1	0.19	0.10	0.10
0	0.00	0.35	0.09
β	11.64 (1.17)	20.09 (8.69)	2.36 (1.91)

$T=2, A=3$
Log_lik = -274.46

τ/n	0	1	2
1	0.18		
0	0.73	0.50	0.00
β	5.37 (0.20)	0.00 (0.00)	1546.5 (138.7)

$T=2, A=1.5$
Log_lik = -290.65

τ/n	0	1	2
1	0.38		
0	0.31	0.33	0.00
β	2.22 (0.13)	1.23E+11 (4.06E+11)	8581 (2089.3)

Table 6D. QRE estimation results and the probability of contribution

Unrestricted parameter

λ

β (Std. Err.)

T=5, A=3
Log_lik = -419

T=5, A=1.5
Log_lik = -395.72

τ/n	0	1	2
4	0.09 15.53 (4.38)		
3	0.08 13.06 (3.07)	0.11 18.25 (8.46E+06)	0.00 1752.8 (n/a)
2	0.11 10.20 (1.59)	0.07 40.68 (1106.5)	0.25 1.31 (28.97)
1	0.42 67188 (1.80E+49)	0.07 7.3E+05 (n/a)	0.10 2.37 (1.51E+14)
0	0.54 0.31 (0.35)	0.61 2.43 (13.74)	0.07 2.57 (11356)

τ/n	0	1	2
4	0.09 17.55 (18.40)		
3	0.05 28.54 (64.88)	0.03 14.86 (5.14E+13)	0.00 9218.5 (n/a)
2	0.06 29.21 (3.66E+05)	0.06 11.64 (2.98E+23)	0.25 1.44 (1.57E+06)
1	0.26 30.95 (1.23E+121)	0.19 5.65 (2.35)	0.17 1.76 (2.00E+10)
0	0.19 2.69 (0.18)	0.48 0.42 (9.05)	0.09 2.35 (7.62E+21)

T=2, A=3
Log_lik = -269.51

T=2, A=1.5
Log_lik = -265.19

τ/n	0	1	2
1	0.18 148.51 (3.07E+43)		
0	0.62 1.14 (0.29)	0.53 0.35 (0.89)	0.00 1.45E+22 (n/a)

τ/n	0	1	2
1	0.18 7.89 (4.15)		
0	0.35 1.98 (0.34)	0.33 4.19E+09 (3.79E+11)	0.00 132.52 (8.39E+4101)

Figure 1A. The predicted (QRE) and empirical contribution probabilities
(Baseline treatment with $T=5$ and $A=3$)




Figure 1B. The predicted (QRE) and empirical contribution probabilities
(Baseline treatment with $T=5$ and $A=1.5$)

