



Learning dynamics for mechanism design: An experimental comparison of public goods mechanisms

Paul J. Healy*

Division of the Humanities and Social Sciences, California Institute of Technology, Pasadena, CA 91125, USA

Received 5 December 2003; final version received 17 March 2005

Available online 13 May 2005

Abstract

In a repeated-interaction public goods economy, incomplete information and dynamic behavior may affect the realized outcomes of mechanisms known to be efficient in a complete information one-shot game. An experimental test of five public goods mechanisms indicates that subjects with private information appear to best respond to recent observations. This provides predictions about which mechanisms will generate convergence to their efficient equilibrium allocations. These predictions match the experimental result that globally stable efficient mechanisms realize the highest efficiency in practice. The simplicity of the suggested best response model makes it useful in predicting stability of mechanisms not yet tested.

© 2005 Elsevier Inc. All rights reserved.

JEL classification: C72; C91; D83; H41

Keywords: Mechanism design; Experiments; Best response; Public goods; Dynamics

1. Introduction

Many mechanisms have been identified whose equilibria generate efficient allocations in economies with pure public goods. In general, mechanisms that require stronger equilibrium concepts are more restricted in their ability to select desirable outcomes. Theoretical results

* Fax: +1 626 405 9841.

E-mail address: pj@hss.caltech.edu.

are unclear about how these trade-offs should be resolved in practice. For example, consider an environment where agents have little to no information about each others' preferences and the level of a certain public good is to be re-evaluated at regular intervals. If a social planner were asked to choose a particular mechanism in this setting, which would she prefer? Are dominant strategy equilibria necessary in this environment? Will mechanisms with stable Nash equilibria converge quickly to an efficient outcome, even though preferences are private information?

In the current paper, five public goods mechanisms with various equilibrium properties are experimentally tested in a repeated game setting.¹ Specifically, the Voluntary Contribution, Proportional Tax, Groves–Ledyard, Walker, and Vickrey–Clarke–Groves mechanisms are all compared in an identical laboratory environment.² The goal of this research is to compare behavior across mechanisms and identify a simple learning dynamic that approximates actual behavior and correctly predicts when actions will converge to the efficient equilibria. Armed with this information, the social planner will then be able to select a mechanism whose desirable equilibrium properties should be realized in practice.

Previous experimental studies have concluded that learning dynamics play an important role in the repeated play of a mechanism. Two general observations suggest that behavior may be consistent with a learning model based on some form of best response play. First, convergence is observed only in game forms known to be supermodular, where best response play predicts convergence. Second, tests of dominant strategy mechanisms suggest that agents tend to play weakly dominated strategies that are best responses to previously observed strategy choices.³ Motivated by these observations, the current paper develops a simple model of best response play and finds that its predictions well approximate observed subject behavior.

The key six results of this study are as follows.

- (1) Subject behavior is well approximated by a model in which agents best respond to the average strategy choice over the last few periods. This model is shown to be significantly more accurate than the stage game equilibrium prediction.
- (2) Half of all decisions in the cVCG mechanism are at the demand-revealing dominant strategy point, while the remainder cluster around weakly dominated best response strategies that are payoff-equivalent to the dominant strategy prediction.
- (3) Behavior in the Groves–Ledyard and Voluntary Contribution mechanisms converges close to equilibrium.
- (4) Behavior in the Proportional Tax and Walker mechanisms does not systematically converge.
- (5) Because of the stability results, the cVCG mechanism is found to be the most efficient. The instability in the Walker mechanism often leads to payoffs below that of the initial endowment.

¹ This is, to the author's knowledge, the largest set of public goods mechanisms to be tested side-by-side to date.

² The cVCG mechanism refers to the Vickrey–Clarke–Groves mechanism in cases where the level of the public good is selected from a continuum. In contrast, the Pivot mechanism refers to the VCG mechanism when the public project choice is binary. The details of all five mechanisms appear in Section 5.

³ An overview of previous results is provided in the Section 2 of this paper, as well as in a survey by Chen [9].

- (6) Most strategy profiles observed to be stable or asymptotically stable are approximately equilibrium strategy profiles.

Note that the model presented in result 1 successfully predicts results 2–6. This indicates that the model is a reasonable and tractable tool for predicting subject behavior and convergence properties of public goods mechanisms.

A brief overview of the previous experimental literature is given in the next section. The learning model and its testable implications are then introduced in Section 3. Details of the experimental design are outlined in Section 4 and a complete description of each mechanism in use is given in Section 5. Results and data analysis appear in Section 6, and Section 7 offers concluding remarks.

2. Previous experiments

This section briefly summarizes previous experimental results on public goods mechanisms. One theme spanning these results is that behavior is, at least qualitatively, consistent with a model of best response play. This observation partially motivates the construction of the particular class of best response models in the following section.

The earliest studies of public goods provision have focused on the Voluntary Contribution mechanism. A wide variety of specifications and treatment variables have been examined, and this line of research continues to generate interesting results about preferences and behavior. A comprehensive summary of this literature is provided by Ledyard [23], who concludes that “in the initial stages of finitely repeated trials, subjects generally provide contributions halfway between the Pareto-efficient level and the free riding level,” and that “contributions decline with repetition.” For example, in an early paper by Isaac et al. [20], payoffs drop from 50% of the maximum in the first period to 9% by the fifth period. Strategies move toward the free-riding dominant strategy through repetition.

In the decades since the theoretical development of public goods mechanisms designed to solve the ‘free-rider’ problem, experimental tests have focused primarily on Nash mechanisms. Several studies, mostly due to Yan Chen, explore properties of the Groves–Ledyard mechanism. For example, Chen and Plott [11] study the effect of the punishment parameter and find that strategies converge rapidly to equilibrium for large parameter values; an observation consistent with known convergence results for best response dynamics in supermodular games. The authors conclude that a best response model that uses information from all previous periods is more accurate than one in which agents best respond to only the previous period.⁴ Chen and Tang [12] compare the Groves–Ledyard mechanism to the Walker mechanism and find that the Groves–Ledyard mechanism is significantly more efficient, apparently due to dynamically unstable behavior in the Walker mechanism. This instability is both observed in subject behavior and predicted by best response dynamic models.

The first well-controlled laboratory test of the dominant strategy Pivot mechanism (the VCG mechanism with a binary public choice) is run by Attiyeh et al. [1]. Subjects are given

⁴ This model, due to Carlson [6] and Auster [2], is a best response model (defined in Section 3) in which predictions equal the simple average (*not* the empirical frequency) of *all* previous periods.

positive or negative values for a proposed project and must submit a message indicating their demand. Although revealing one's demand is a (weak) dominant strategy, only 10 percent of observations are consistent with this prediction, with 13 of 20 subjects *never* revealing their true value.

Kawagoe and Mori [21] extend the Attiyeh et al. result by comparing the above treatment to one in which subjects are given a payoff table. The effect of having players choose from the table is significant, as demand revelation increases to 47%. Since the equilibrium is a weak dominant strategy in the sense that all agents have other strategies that are best responses to the equilibrium, the authors argue that subjects have difficulty discovering the undominated property of truth-telling because non-revelation strategies may also be best responses.

Most recently, Cason et al. [7] compare the Pivot and cVCG mechanisms with two players and show that deviations from truthful revelation, when they occur, tend to result in weakly dominated Nash equilibria. Specifically, assume the first subject in the Pivot mechanism announces her value truthfully. If the good is produced when the second player announces truthfully, it is also produced when he overstates his value. Thus, there exists a wide range of false announcements that can be equilibrium strategies for player two, even though they are weakly dominated by truth-telling. In the cVCG mechanism with only one preference parameter, on the other hand, agents have a strict dominant strategy to reveal truthfully. In the experiment, only half of the observed subject pairs play the dominant strategy equilibrium in the Pivot mechanism, while 81 percent reveal truthfully in the cVCG mechanism. Behaviorally, this explanation is consistent with an evolutive model where agents select payoff maximizing strategies rather than an educative model where agents solve for the equilibria of the game.

The results of these previous studies indicate that dynamically stable Nash equilibria and strict dominant strategies are good predictors of behavior, but unstable equilibria generate unstable behavior and weakly dominated best responses may draw players away from dominant strategy equilibria. These observations are consistent with a history-dependent best response model. The goal of the current paper is to refine this conjecture and identify a tractable model that, if not a perfect description of behavior, can at least predict the convergence properties of a mechanism in the repeated environment.

3. Setup and environment

The general environment in use is as follows. A set of agents $\mathcal{I} = \{1, \dots, n\}$ each have preferences for consumption of a private good $\mathbf{x} = (x_1, \dots, x_n)$ and a single public good y that can be represented by the differentiable function $u_i(y, x_i; \theta_i)$, where $\theta_i \in \Theta_i$ indicates the vector of utility parameters held by agent $i \in \mathcal{I}$. Each Θ_i is assumed to be convex and $\Theta = \times_{i=1}^n \Theta_i$. Unless stated otherwise, preferences are assumed throughout to be quasilinear, so that $u_i(x_i, y; \theta_i) = v_i(y; \theta_i) + x_i$ where $v_i(y; \theta_i)$ is strictly concave in y .

A public goods allocation is an $(n + 1)$ -tuple of the form (y, x_1, \dots, x_n) . No public good exists initially, although a linear technology can be used to build $y \geq 0$ units of the public good at a cost of κy units of the private good. Given an initial endowment of the private good

ω_i , consumption of the private good is given by $x_i = \omega_i - \tau_i$ for each i , where τ_i represents a transfer payment paid by agent i . Therefore, the public goods allocation is equivalently expressed as $(y, \tau_1, \dots, \tau_n)$. A vector of transfer payments is feasible if $\sum_i \tau_i \geq \kappa y$ and budget balanced if the constraint is met with equality.⁵

A mechanism is represented as a game form, indexed by g , in which agents choose a message $m_{g,i}$ from a strategy space $\mathcal{M}_{g,i}$ that is assumed here to be convex. The vector of all messages is denoted $\mathbf{m}_g \in \mathcal{M}_g = \times_{i=1}^n \mathcal{M}_{g,i}$. When there is no confusion, the g subscript will be dropped. For a given agent i and a vector of messages $\mathbf{m}_{-i} = (m_1, \dots, m_{i-1}, m_{i+1}, \dots, m_n)$, a best response message for agent i is that which maximizes i 's utility under the assumption that the other agents send messages \mathbf{m}_{-i} . The set of best responses to \mathbf{m}_{-i} in mechanism g is denoted $\mathcal{B}_{g,i}(\mathbf{m}_{-i}; \theta_i)$. Define $\mathcal{B}_g(\mathbf{m}, \theta) = \times_{i=1}^n \mathcal{B}_{g,i}(\mathbf{m}_{-i}; \theta_i)$ to be the set of message profiles that are best responses to the profile \mathbf{m} . Any fixed point of the best response profile mapping is a Nash equilibrium strategy profile of the game. Formally, any equilibrium strategy profile, denoted $\mathbf{m}^*(\theta) = (m_1^*(\theta), \dots, m_n^*(\theta))$, satisfies

$$\mathbf{m}^*(\theta) \in \mathcal{B}(\mathbf{m}^*(\theta), \theta).$$

Note that this solution concept requires each player's equilibrium strategy to be a function of the other players' types if m_i^* varies with θ_j for $j \neq i$. The equilibrium message does *not* depend on the types of other agents if the Nash equilibrium is in dominant strategies. The set of Nash equilibria for a given type profile and game is given by $\mathcal{E}_g(\theta)$.

The vector of received messages \mathbf{m} in mechanism g maps to a unique outcome of the form $(y_g(\mathbf{m}), \tau_g(\mathbf{m}))$, where $y_g : \mathcal{M}_g \rightarrow \mathbb{R}_+$ determines the level of the public good chosen and $\tau_g : \mathcal{M}_g \rightarrow \mathbb{R}^n$ determines the vector of transfer payments of the private good to be paid by each agent.⁶ The strategy space (inputs) and outcome function (outputs) completely characterize the mechanism. All mechanisms considered here are feasible and some are budget balanced.

The objective of the mechanism designer is to implement a social choice correspondence $\mathcal{F} : \Theta \rightarrow \mathbb{R}_+ \times \mathbb{R}^n$ with certain desirable properties. Let $(y^{\mathcal{F}}(\theta), \tau^{\mathcal{F}}(\theta)) \in \mathcal{F}(\theta)$ represent a particular public goods allocation satisfying the properties for preference parameters θ . In the public goods environment, the appropriate social choice correspondence for a utilitarian planner is the mapping $\mathcal{P}(\theta)$ that picks the set of Pareto optimal allocations $(y^{\mathcal{P}}(\theta), \tau^{\mathcal{P}}(\theta))$ satisfying

$$y^{\mathcal{P}}(\theta) \in \arg \max_{y \in \mathbb{R}_+} \left[\sum_{i=1}^n v_i(y; \theta_i) - c(y) \right]$$

and such that $\tau^{\mathcal{P}}(\theta)$ is budget balanced. A mechanism g implements \mathcal{F} if, for every $\theta \in \Theta$, the outcome function selects allocations in $\mathcal{F}(\theta)$ at every equilibrium message $\mathbf{m}^*(\theta)$. If g implements $\mathcal{P}(\theta)$, then it is said to be efficient. If $y_g(\mathbf{m}^*(\theta)) = y^{\mathcal{P}}(\theta)$ for some $(y^{\mathcal{P}}(\theta), \tau^{\mathcal{P}}(\theta)) \in \mathcal{P}(\theta)$ and $\tau_g(\mathbf{m}^*(\theta))$ is feasible but not budget balanced, then the

⁵ Individual budget constraints are not imposed in the following analysis, so that τ_i may be larger than ω_i .

⁶ Set-valued mechanisms may be defined, but implementation in this context is assumed to require the selection of a unique outcome.

mechanism is only outcome efficient. The surplus transfer payments in this case are assumed to be wasted and yield no value to any agent in the economy.

3.1. A best response model of behavior

A history-based best response learning model assumes that each agent i forms predictions about the strategies others will use in period t based on the observed strategies of the other players in periods 1 through $t - 1$, denoted \mathbf{m}_{-i}^1 through \mathbf{m}_{-i}^{t-1} . Agent i 's prediction about the strategy to be played by agent j in period t is represented by the function $\psi_j^i(m_j^1, \dots, m_j^{t-1}) \in \mathcal{M}_j$, which maps each possible history of agent j into a unique pure strategy for every $j \neq i$.⁷ Agent i is then assumed to select a best response to his predictions. Letting $\psi : \mathcal{M}^{t-1} \rightarrow \mathcal{M}$ represent the vector of predictions generated from the history of play up to period t , the strategy profile occurring in period t will be an element of $\mathcal{B}(\psi(\mathbf{m}^1, \dots, \mathbf{m}^{t-1}), \theta)$. If ψ_j^i is undefined for some $j \neq i$, then let the best response model predict any strategy in \mathcal{M}_j . In short, best response learning models assume that players are utility maximizers, but that their predictions are myopic and may be inaccurate.

When \mathcal{M} is a convex set in \mathbb{R}^n , a k -period average best response dynamic assumes that

$$\psi_j^i\left(\left\{m_j^s\right\}_{s=1}^{t-1}\right) = \bar{m}_j^{t,k} = \frac{1}{k} \sum_{s=t-k}^{t-1} m_j^s \quad (1)$$

for all $i, j \in \mathcal{I}$ when $t > k$, and $\psi_j^i = \emptyset$ when $t \leq k$. Let $\bar{\mathbf{m}}^{t,k} = (\bar{m}_j^{t,k})_{j=1}^n$. In this model, agents best respond to the prediction that the average message of the previous k periods will be played in the current period. Note that $\psi_j^i \in \mathcal{M}_j$ by the convexity of \mathcal{M}_j .

Behaviorally, the k -period average model implies that agents best respond to an estimate of the current *trend* in the messages of other agents. Here, the estimate of trend is given by a simple moving average filter. Other filters may be used to determine the current trend in \mathbf{m}_{-i} , such as exponential smoothing or time-weighted moving averages. Although these various trend models will produce slightly different results, the implication of such models is that agents form unique, pure-strategy predictions about the decisions of others using previous observations that are not highly sensitive to period-by-period fluctuations in the history of play.

One simple fact immediate from the definition of the k -period dynamic is that strictly dominated strategies will not be observed. This provides the first testable proposition.

Proposition 1. *In the k -period best response dynamic, no strictly dominated strategy is observed in any period $t > k$.*

Given that this dynamic is suggested as a model capable of predicting convergence in public goods mechanisms, it is of interest to study the limiting behavior of this process.

⁷ Most dynamic models (such as fictitious play) are based on mixed strategy predictions over a finite strategy space. The best response models suggested here generate pure strategy predictions over a continuous strategy space.

The following propositions and corollaries establish the relationship between the k -period average best response model and Nash equilibrium. Note that several of these theoretical results will be verified empirically in Section 6.

Proposition 2. *If a strategy is observed in $k + 1$ consecutive periods of the k -period average best response dynamic then it is a Nash equilibrium.*

Proposition 2 immediately implies the following important corollary.

Corollary 3. *All rest points of the k -period best response dynamic are Nash equilibria.*

The following proposition shows that convergence implies that the limit point is a Nash equilibrium.

Proposition 4. *Given some $\theta \in \Theta$, let $\{\mathbf{m}^t\}_{t=1}^\infty$ be a sequence of strategy profiles consistent with the k -period average best response dynamic that converges to a profile $\mathbf{q} \in \mathcal{M}$. If the best response correspondence $\mathcal{B}(\cdot, \theta)$ is upper hemi-continuous at \mathbf{q} and non-empty on \mathcal{M} , then \mathbf{q} is a Nash equilibrium strategy profile at θ .*

This follows from that fact that $\{\bar{\mathbf{m}}^{t,k}\}_{t=1}^\infty$ must also converge to \mathbf{q} , so $\{\mathcal{B}(\bar{\mathbf{m}}^{t,k}, \theta)\}_{t=1}^\infty$ converges to a set containing \mathbf{q} .

Corollary 5. *Given some $\theta \in \Theta$, let $\{\mathbf{m}^t\}_{t=1}^\infty$ be consistent with the k -period average best response dynamic that converges to a profile $\mathbf{q} \in \mathcal{M}$. If y_g , τ_g , and each $u_i(\cdot; \theta_i)$ are continuous and single-valued and \mathcal{M} is compact, then \mathbf{q} is a Nash equilibrium strategy profile at θ .*

This corollary is a simple application of the Theorem of the Maximum, which guarantees that the best response correspondence is upper hemi-continuous and non-empty under the given conditions. The notion of asymptotic stability requires that the dynamic path from all initial points in some neighborhood of \mathbf{q} converge to \mathbf{q} . By Proposition 4, this is clearly sufficient for \mathbf{q} to be a Nash equilibrium.

Corollary 6. *If, for some $\theta \in \Theta$, $\mathbf{q} \in \mathcal{M}$ is asymptotically stable according to the k -period average best response dynamic and \mathcal{B} is upper hemi-continuous and non-empty, then \mathbf{q} is a Nash equilibrium strategy profile at θ .*

One might conjecture that the dynamic is more stable (in a global sense) under larger values of k . However, simple games can be constructed in which cycles can occur under a particular value of k , but globally stability obtains for the $k - 1$ and $k + 1$ dynamics.⁸ It is natural to then ask what properties of a game are sufficient for global stability to obtain for

⁸ Simply pick a 2-player game with $\mathcal{M}_i = [0, 1]$ and where the best response function for each player equals 0 on some very small neighborhood around $(k - 1) / k$ and equals 1 everywhere else. Start the dynamic with each player playing 0 for the first k periods. Cycles then emerge that jump in and out of the unique equilibrium *ad infinitum*. See also Bear [4].

all values of k . In the class of supermodular games (see Topkis [33]) the monotonicity of the best response correspondence guarantees global stability of the 1-period best response dynamic. The following proposition demonstrates that this result extends to the k -period dynamic.

Proposition 7. *In a supermodular game, if, for some $\theta \in \Theta$, $\{\mathbf{m}^t\}_1^\infty$ is consistent with a k -period average best response dynamic, then $\liminf \mathbf{m}^t \geq \underline{\mathcal{E}}(\theta)$ and $\limsup \mathbf{m}^t \leq \bar{\mathcal{E}}(\theta)$, where $\underline{\mathcal{E}}(\theta)$ and $\bar{\mathcal{E}}(\theta)$ are the smallest and largest pure strategy Nash equilibrium profiles at θ .*

Corollary 8. *If a supermodular game has a unique pure strategy Nash equilibrium, then the k -period average best response dynamic is globally asymptotically stable.*

The proof of Proposition 7 appears in the appendix.⁹ This result is of particular significance because it is consistent with the claim of Chen [8] and Chen and Gazzale [10] that supermodularity is sufficient for convergence in a variety of environments tested in the laboratory.

There are also two sets of ‘dominant diagonal’ conditions under which the 1-period dynamic is known to be globally stable. In both cases, this stability result can be shown to extend to all k -period dynamic models. The first is due to Gabay and Moulin [16], and is summarized in the following proposition.¹⁰

Proposition 9. *Assume that $\mathcal{M} = [0, +\infty)^n$, and for some $\theta \in \Theta$, each u_i is twice continuously differentiable and, for every $\mathbf{m} \in \mathcal{M}$, u_i satisfies pseudo-concavity $\left(\frac{\partial u_i}{\partial m_i}(\mathbf{m}, \theta_i) \cdot (m_i - m'_i) \geq 0 \Rightarrow u_i(\mathbf{m}, \theta_i) \geq u_i(m'_i, \mathbf{m}_{-i}, \theta_i)\right)$, coercivity $\left(\lim_{m_i \rightarrow +\infty} \left|\frac{\partial u_i}{\partial m_i}(\mathbf{m}, \theta_i)\right| = +\infty\right)$, and strict diagonal dominance $\left(\left|\frac{\partial^2 u_i}{\partial m_i^2}(\mathbf{m}, \theta_i)\right| > \sum_{j \neq i} \left|\frac{\partial^2 u_i}{\partial m_i \partial m_j}(\mathbf{m}, \theta_i)\right|\right)$. Then there exists a unique Nash equilibrium $\mathbf{m}^*(\theta)$ and every sequence $\{\mathbf{m}^t\}_1^\infty$ consistent with a k -period best response dynamic converges to $\mathbf{m}^*(\theta)$.*

The second, more direct condition for stability of the 1-period model requires that the best response correspondence be a single-valued, linear function of the form $\mathcal{B}(\mathbf{m}, \theta) = \mathbf{A}(\theta) \mathbf{m} + \mathbf{h}(\theta)$, where $\mathbf{h}(\theta) \in \mathbb{R}^n$ and $\mathbf{A}(\theta) = [a_{ij}(\theta)]_{i,j=1}^n$ is a real matrix for which there exists a strictly positive n -vector \mathbf{d} such that $d_i > \sum_{j=1}^n d_j |a_{ij}(\theta)|$ for $i = 1, \dots, n$. When $\mathbf{A}(\theta)$ is non-negative, this positive dominant diagonal condition is both necessary and sufficient for global stability (see Murata [26, Chapter 3] for details). Using the distributed lag methods of Bear [3,4], this same condition is easily shown to be necessary

⁹ As an alternative method of proof, it can be established that any sequence $\{\mathbf{m}^t\}_1^\infty$ consistent with the k -period dynamic must satisfy the *adaptive dynamics* conditions of Milgrom and Roberts [24]. Proposition 7 then follows from Theorem 8 of that paper.

¹⁰ A similar result was previously established by Rosen [28] using the *diagonal strict concavity* condition on utilities.

and sufficient for global convergence of any best response dynamic whose predictions are weighted averages of past observations, including the k -period dynamics.

Note that the global stability properties of the k -period dynamic are similar to properties of other well-known learning models. Consider, for example, the fictitious play dynamic, which is best suited for games with small, finite strategy spaces. As with the k -period model, fictitious play has stable Nash equilibrium points (Brown [5]) and is globally stable in supermodular games (Milgrom and Roberts [25, Theorem 8]) but is also capable of off-equilibrium limit cycles (Shapley [30]).

4. Experimental design

The five public goods mechanisms under consideration were tested in a laboratory environment using human subjects. All experiments were run at the California Institute of Technology during the 2002–2003 academic year using undergraduates recruited via E-mail. Most subjects had participated in economics experiments, though none had experience with the particular game forms in the current study. Four sessions were run with each mechanism for a total of 20 sessions.¹¹ Each session consisted of five subjects interacting through computer terminals. Subjects only participated in one session, in which they played a single mechanism 50 times against the same four cohorts. Each iteration of the mechanism is referred to as a period. Multiple sessions were run simultaneously, so that more than five subjects would be in the lab at the same time. Each subject knew she was grouped with four others, but could not discern which individuals were in her group.¹² Instructions were given to the subjects at the beginning of the experiment and read aloud by the experimenter. Participants were then given time to ask any clarifying questions.¹³

Before the experiment, subjects were given their preference parameters and initial endowments privately on a slip of paper. An incomplete information environment was used because the interest of this study is to identify mechanisms whose efficiency properties are robust to the assumptions of complete information. If the learning process of agents converges rapidly to a mechanism's efficient equilibrium, then a social planner need not worry about the informational assumptions of the equilibrium concept.

After receiving their private information, subjects logged into the game from their computer terminal using an Internet browser program. The software interface includes two useful tools for the subjects to use at any time. First, a history window is available that displays the results of all past periods. Subjects can see all previous outcomes, including the message they sent, taxes paid, public good levels, and profits. The entire vector of messages submitted by the other agents in previous periods is not shown; only the relevant variables used in calculating the tax, value and payoff functions are provided. Subjects can open this

¹¹ Four additional sessions with the cVCG mechanism were run, but had to be discarded due to a software failure. These data are very similar to the reported sessions and feature a slightly higher frequency of demand revelation.

¹² In one cVCG session, only one group of subjects was in the laboratory at the same time. The subjects were well separated to prevent communication or other out-of-experiment effects.

¹³ Instructions and data are available at <http://kakutani.caltech.edu/pj>. Subjects were not deceived in any way during this experiment.

Table 1
Preference parameters $\theta_i = (a_i, b_i)$ and ω_i used in all sessions

Player	1	2	3	4	5
a_i	1	8	2	6	4
b_i	34	116	40	68	44
ω_i	200	140	260	250	290

window at any time and are also shown the same information at the conclusion of each period. The second tool, called the ‘What-If Scenario Analyzer,’ allows subjects to enter hypothetical messages into a calculator-like program to view what levels of $y(\mathbf{m})$, $\tau_i(\mathbf{m})$ and profit would result. Each subject is shown only her own hypothetical tax, value, and profits so that subjects cannot deduce the preference parameters of other subjects. Instead of a practice period, subjects were given 5 min to experiment with the ‘What-If Scenario Analyzer’ and ask questions.

The benefit of the ‘What-If Scenario Analyzer’ is that it enables subjects to perform searches over the strategy space before selecting their strategies. In this sense, it is similar to giving subjects a complete payoff table, although the current tool provides more feedback than payoffs alone. The interest of this study is to understand the learning dynamics involved as subjects resolve the uncertainty about the strategies of others. With inexperienced subjects, this process may be confounded with subjects figuring out how strategies map to outcomes. By experimenting with the ‘What-If’ tool before the experiment, the subjects become well informed about the game they are playing, but may be unsure what strategies others will be using. This reduces the potential confound in the observed dynamics.

Once the experiment begins, subjects enter their message (twice for confirmation) into the computer each period.¹⁴ The feedback at the end of the period is identical to the history information described above. Total earnings are kept at the bottom of the screen at all times, along with the current period number and the total number of periods in the game. Subjects’ earnings were tallied in ‘francs’ and converted to dollars at the end of the experiment. Conversion rates from francs to actual dollar earnings varied by mechanism between 650 to 800 francs per dollar so that typical subjects would earn 40–50 cents for each of their 50 decisions, plus a \$5 show-up fee. Sessions typically lasted from 90 min to 2 h.

Each subject in a session was assigned a unique player type, differing only by their utility parameters and endowments. The same five player types were used in every session. Quasilinear preferences are induced with concave quadratic values for the public good given by $v_i(y; \theta_i) = -a_i y^2 + b_i y$ and an initial endowment of the private good, ω_i . The vector $\theta_i = (a_i, b_i)$ and the endowment ω_i are positive for all $i \in \mathcal{I}$. The quasilinear, quadratic structure of the preferences is common knowledge across all subjects, although the vectors of individual coefficients $\theta_i = (a_i, b_i)$ and endowment ω_i are private information. The chosen player type profile $\theta = (\theta_1, \dots, \theta_5)$ and endowments are identical across all periods, sessions and mechanisms. These values are given in Table 1.

The marginal cost of the public good is chosen to be constant at $\kappa = 100$ in every session. As will be shown in the next section, these parameter values have been chosen to provide

¹⁴ Subjects were restricted to enter messages with no more than two decimals.

distinct predictions between various mechanisms. Given the quasilinear preferences, the Pareto optimal level of the public good is uniquely solved by $y^P(\theta) = 4.8095$. From an experimental design standpoint, a non-focal value for the Pareto optimum is preferred so that public good levels observed at or near Pareto optimal levels cannot alternatively be explained by subjects choosing integer strategies, for example.

5. The mechanisms

The following section describes each of the five mechanisms in detail. The outcome functions and strategy spaces are presented in Table 2. A reader familiar with the details of public goods mechanisms may skip the discussion of the Voluntary Contribution, Proportional Tax, Groves–Ledyard, and Walker mechanisms, although the cVCG mechanism in use here has interesting properties that are critical in understanding the results presented below.

5.1. Voluntary Contribution mechanism

In this simple mechanism, each player i announces m_i , the number of units of the public good to be added to the total. The sum of the contributions represents the realized level of the public good and the tax paid by agent i is the cost of her contribution to the public good. In this mechanism, each agent has a ‘Robinson Crusoe’ ideal point, denoted \tilde{y}_i , representing the amount of public good she would contribute in the absence of contributions by others.

Table 2
The five mechanisms tested

Mechanism	Outcome functions	Strategy space
Voluntary Contribution	$y(\mathbf{m}) = \sum_i m_i$ $\tau_i(\mathbf{m}) = \kappa m_i$	$\mathcal{M}_i = [0, 6]$
Proportional Tax	$y(\mathbf{m}) = \sum_i m_i$ $\tau_i(\mathbf{m}) = \kappa y(\mathbf{m})/n$	$\mathcal{M}_i = [0, 6]$
Groves–Ledyard	$y(\mathbf{m}) = \sum_i m_i$ $\tau_i(\mathbf{m}) = \frac{\kappa y(\mathbf{m})}{n} + \frac{\gamma}{2} \left(\frac{n-1}{n} (m_i - \mu_i)^2 - \sigma_i^2 \right)$ $\mu_i = \frac{1}{n-1} \sum_{j \neq i} m_j \quad \sigma_i^2 = \frac{1}{n-2} \sum_{j \neq i} (m_j - \mu_i)^2$	$\mathcal{M}_i = [-4, 6]$
Walker	$y(\mathbf{m}) = \sum_i m_i$ $\tau_i(\mathbf{m}) = \left(\frac{\kappa}{n} + m_{(i-1) \bmod n} - m_{(i+1) \bmod n} \right) y(\mathbf{m})$	$\mathcal{M}_i = [-10, 15]$
cVCG	$y(\hat{\theta}) = \arg \max_{y \geq 0} (\sum_i v_i(y; \hat{\theta}_i) - \kappa y)$ $\tau_i(\hat{\theta}) = \frac{\kappa y(\hat{\theta})}{n} - \sum_{j \neq i} \left(v_j(y(\hat{\theta}); \hat{\theta}_j) - \frac{\kappa y(\hat{\theta})}{n} \right)$ $+ \max_{z \geq 0} \sum_{j \neq i} \left(v_j(z; \hat{\theta}_j) - \frac{\kappa z}{n} \right)$	$\mathcal{M}_i = \Theta_i$ $= \mathbb{R}_+^2$

The best response function for each agent is then given by

$$\mathcal{B}_i(\mathbf{m}_{-i}; \theta_i) = \tilde{y}_i - \sum_{j \neq i} m_j. \quad (2)$$

In the quadratic environment, $\tilde{y}_i = (b_i - \kappa) / 2a_i$. Using the parameters of the experiment, the vector of Robinson Crusoe ideal points is

$$\tilde{\mathbf{y}} = \left(-33, 1, -15, -2\frac{2}{3}, -7 \right).$$

Since player 2 is the only for which $\tilde{y}_i > 0$, he is the only player who does not have a dominant strategy of contributing zero. The unique Nash equilibrium of this game is therefore $\mathbf{m}^*(\theta) = (0, 1, 0, 0, 0)$, which results in a suboptimally low level of the public good. Under the k -period dynamic, this equilibrium must obtain by period $2k + 1$. Note that if the message space were unbounded, then no equilibrium would exist and the k -period model would diverge.

5.2. Proportional Tax mechanism

The Proportional Tax mechanism is an alternative to the Voluntary Contribution mechanism in which each agent must pay an equal share of the total cost. Under this scheme, agents' Robinson Crusoe ideal points are given by $\tilde{y}_i = (b_i - \kappa/n) / 2a_i$, which is necessarily larger than under the Voluntary Contribution mechanism. Specifically,

$$\tilde{\mathbf{y}} = (7, 6, 5, 4, 3).$$

This mechanism has the same message space and best response function as the Voluntary Contribution mechanism (Eq. (2)) but no agents have a dominant strategy in this mechanism since $\tilde{\mathbf{y}} \gg \mathbf{0}$. The unique pure strategy Nash equilibrium is the corner strategy profile $\mathbf{m}^*(\theta) = (6, 0, 0, 0, 0)$, which results in a suboptimally *large* level of the public good.¹⁵ Note that although no agent has a dominant strategy, players 3, 4, and 5 are only willing to contribute when the total contributions of all others is well below their ideal points.

The Proportional Tax mechanism is also of interest because it provides the foundation of both the Groves–Ledyard and the Walker mechanisms. One can show that if the Proportional Tax mechanism had an interior Nash equilibrium, it would select an optimal level of the public good. The problem is that, generically, interior equilibria do not exist. The Groves–Ledyard and Walker mechanisms are essentially variants of the Proportional Tax mechanism with an additional ‘penalty’ term in the transfer function chosen to guarantee the existence of an interior equilibrium point.

5.3. Groves–Ledyard mechanism

The mechanism of Groves and Ledyard [18] was the first constructed whose Nash equilibria yield fully efficient outcomes. The mechanism requires all agents pay an equal share

¹⁵ If $\mathcal{M} = \mathbb{R}$, then no equilibrium would exist. If the upper bound of the message space were chosen to be greater than 7, then the equilibrium would be $(7, 0, 0, 0, 0)$.

of the cost plus a penalty term based on deviations from the average of the others' contributions and on the variance of those contributions. Each agent's unique best response message in the quadratic, quasilinear environment is given by

$$\mathcal{B}_i(\mathbf{m}_{-i}; \theta_i) = \frac{b_i - \kappa/n}{2a_i + \gamma \frac{n-1}{n}} + \left(\frac{\gamma/n - 2a_i}{2a_i + \gamma \frac{n-1}{n}} \right) \sum_{j \neq i} m_j.$$

In the experimental environment, this mechanism has a unique pure strategy equilibrium of

$$\mathbf{m}^*(\theta) = (1.0057, 1.1524, 0.9695, 0.8648, 0.8171),$$

which results in $y(\mathbf{m}^*(\theta)) = 4.8095 = y^{\mathcal{P}}(\theta)$.¹⁶ The message space $\mathcal{M}_i = [-4, 6]$, which is identical to that used by Chen and Plott [11] and Chen and Tang [12], is sufficiently wide so that the equilibrium is not near a corner of the strategy space.

One nice property of the Groves–Ledyard mechanism is that it can possess the global stability properties of Propositions 7 and 9. In the quadratic environment, the equilibrium is supermodular if $\gamma > 2n \max_i (a_i)$ and satisfies the positive dominant diagonal condition if $\gamma > [(n-2)/(n-1)]n \max_i (a_i)$. Using the experiment parameters, these conditions are $\gamma > 80$ and $\gamma > 30$, respectively. Since $\gamma = 100$ is used in the experiment, both sufficient conditions for stability are satisfied. Therefore, all k -period best response dynamics are globally stable in this setting.

5.4. Walker mechanism

One important theoretical drawback of the Groves–Ledyard mechanism is that it may select some efficient allocations that do not Pareto dominate the initial endowment. Partially in response to this issue, Walker [36] developed a ‘paired difference’ mechanism that implements Lindahl allocations in Nash equilibrium. Since Lindahl allocations are guaranteed to Pareto dominate the initial endowment and tax each agent based on their marginal willingness to pay, Walker’s mechanism appears to provide the most desirable solution to the free-rider problem.

In the quadratic environment, the unique best response function is given by

$$\mathcal{B}_i(\mathbf{m}_{-i}; \theta_i) = \frac{b_i - (\kappa/n + m_{(i-i) \bmod n} - m_{(i+1) \bmod n})}{2a_i} - \sum_{j \neq i} m_j.$$

Solving for the equilibrium with the given parameters,

$$\mathbf{m}^*(\theta) = (12.276, -1.438, -6.771, -2.200, 2.943),$$

which gives the Lindahl allocation of $y = 4.8095$ and

$$\tau = (117.26, 187.8, 99.855, 49.469, 26.567)$$

¹⁶ Note that the equilibrium strategy profile has less variance for larger values of γ . Here, γ is chosen to be large for stability reasons.

To accommodate the disperse equilibrium messages, the message space is expanded to $\mathcal{M}_i = [-10, 15]$ for each $i \in \mathcal{I}$.

Although this mechanism implements Lindahl allocations in Nash equilibrium, its equilibria are known have instability problems. In the quadratic environment, the 1-period best response dynamic can be represented by the system of difference equations $\mathbf{m}^t = \mathbf{A}\mathbf{m}^{t-1} + \mathbf{h}$, where the row sums of \mathbf{A} all equal $-(n-1)$. This matrix is irreducible and non-positive, so by Harriff et al. [19, Corollary 1], its dominant eigenvalue must then equal $-(n-1)$, which is greater than 1 in absolute value. Thus, Cournot best response is unstable under any parameter choice θ .¹⁷ Chen and Tang [12] also argue that the cost of small deviations from equilibrium in the Walker mechanism are lower than in the Groves–Ledyard mechanism, making the former less robust to experimentation.

5.5. Continuous VCG (cVCG) mechanism

The cVCG mechanism represents a particular selection from the class of dominant strategy incentive compatible mechanisms developed independently by Vickrey [35], Clarke [13], and Groves [17]. In these direct mechanisms (where $\mathcal{M}_i = \Theta_i$), truth-telling weakly dominates all other strategies. However, given any \mathbf{m}_{-i} , there exist messages $\mathbf{m}_i \neq \theta_i$ such that $\mathbf{m}_i \in \mathcal{B}_i(\mathbf{m}_{-i}; \theta_i)$. As will be demonstrated, with two preference parameters, all points on a particular line in the strategy space that intersects the ‘truth-telling’ strategy are best responses to \mathbf{m}_{-i} . Consequently, a best response learning model predicts that agents select messages from the best response line that are not necessarily the truth-telling equilibrium. Given that the dominant strategy equilibrium is a zero-dimensional set, the best response set is one-dimensional, and the strategy space is two-dimensional, it is easy to distinguish between equilibrium, best response, and random (or, unexplained) strategy choices.

In the cVCG mechanism, $\mathcal{M}_i = \Theta_i$, which equals \mathbb{R}_+^2 in the experiment, and any message \mathbf{m}_i can be equivalently expressed as an announced parameter value $\hat{\theta}_i = (\hat{a}_i, \hat{b}_i)$. Agents are free to misrepresent preferences by announcing $\hat{\theta}_i \neq \theta_i$. The outcome function takes the vector of announced parameter values $\hat{\theta}$ and solves for the Pareto optimal level of the public good on the assumption that $\hat{\theta}$ is the true vector of preference parameters.

Each agent’s tax is comprised of three parts: an equal share of the cost of production, a reward equal to the net utility of all other agents assuming their preference announcements are truth-telling, and a penalty equal to the maximum possible net utility of all others under their given preference announcement. The third term necessarily dominates the second, so the sum of transfers is always weakly greater than the cost of the project.¹⁸

This mechanism is constructed so that each agent i prefers an announcement $\hat{\theta}_i$ that yields a Pareto optimal level of the public good under the assumption that $\hat{\theta}_{-i} = \theta_{-i}$. To

¹⁷ If \mathbf{A} were non-negative, instability of the one-period model would be sufficient (and necessary) for instability in the k -period model. This result does not extend to the case of non-positive matrices. However, for the experiment parameters, the k -period model is unstable for all values of k in the range under consideration, and it is conjectured that no value of k will guarantee stability in this environment. See Harriff et al. [19, p. 359] for the relevant theorems and counter-examples.

¹⁸ This mechanism is known to be budget-balanced in the quadratic environment when all agents have the same slope parameter. This is not true in the current environment.

see this directly, note that the first-order condition for utility maximization with quadratic preferences is that

$$\frac{dy}{d\hat{\theta}_i} \left(\left[b_i + \sum_{j \neq i} \hat{b}_j - \kappa \right] - y(\hat{\theta}) \left[2 \left(a_i + \sum_{j \neq i} \hat{a}_j \right) \right] \right) = 0.$$

Since $dy/d\hat{\theta}_i \neq 0$ for all $\hat{\theta}_i \in \Theta_i$ and all $i \in \mathcal{I}$, utility maximization is achieved by setting the term in parentheses to zero through manipulation of $y(\hat{\theta})$. The necessary and sufficient condition for maximization is therefore

$$y(\hat{\theta}_i, \hat{\theta}_{-i}) = \frac{b_i + \sum_{j \neq i} \hat{b}_j - \kappa}{2 \left(a_i + \sum_{j \neq i} \hat{a}_j \right)} = y(\theta_i, \hat{\theta}_{-i}). \quad (3)$$

Thus, any announcement by player i that results in the same level of the public good as would have obtained under truth-telling is necessarily a best response.

Since $\hat{\theta}_i = \theta_i$ satisfies (3) for all $\hat{\theta}_{-i}$, truth-telling is a dominant strategy. Given a particular value of $\hat{\theta}_{-i}$, however, there exists a range of $\hat{\theta}_i \neq \theta_i$ that satisfy condition (3). This set of values is given by the best response correspondence

$$\mathcal{B}_i(\hat{\theta}_{-i}; \theta_i) = \left\{ (\hat{a}_i, \hat{b}_i) \in \Theta_i : (\hat{b}_i - b_i) = 2(\hat{a}_i - a_i) y(\theta_i, \hat{\theta}_{-i}) \right\}. \quad (4)$$

Clearly, the manifold of best responses to $\hat{\theta}_{-i}$ is a line through Θ_i that must intersect θ_i . The slope of this line depends on $\hat{\theta}_{-i}$, so the best response line rotates about θ_i as $\hat{\theta}_{-i}$ varies.¹⁹ If an agent holds a prediction that places non-zero probability on multiple values of $\hat{\theta}_{-i}$, then the dominant strategy point becomes the unique best response.

It is important to reiterate the fact that the set of Nash equilibria of this mechanism extends beyond the dominant strategy equilibrium. As a simple example, if $n - 1$ agents submit $\hat{\theta}_i = \theta_i$ while the n th agent announces $\hat{\theta}_n \in \mathcal{B}_n(\theta_{-n}; \theta_n) \setminus \{\theta_n\}$, then a Nash equilibrium with a weakly dominated strategy has obtained. Since $y(\hat{\theta}_n, \theta_{-n}) = y(\theta) = y^P(\theta)$, this equilibrium is also outcome efficient.

Although this mechanism is known to be inefficient due to its lack of budget balance, the size of the predicted inefficiency varies with the parameter choices. The results obtained in the laboratory may be sensitive to the choice of preference parameters. In the current experiment, equilibrium efficiency is over 99%. The discussion in Section 6.5 will highlight the significance of this (or any) fixed parameter choice in analyzing the results.

¹⁹ Note that $\hat{\theta}_i$ affects agent i 's utility only through the value of $y(\hat{\theta}_i, \hat{\theta}_{-i})$, so indifference curves in agent i 's strategy space correspond to level curves of the $y(\cdot, \hat{\theta}_{-i})$ function. The set $\mathcal{B}_i(\hat{\theta}_{-i}; \theta_i)$ is therefore the level set of i 's most preferred quantity of the public good given $\hat{\theta}_{-i}$.

6. Results

6.1. Calibrating the parameter k

Using the observed data, best response model predictions for each period $t > k$ are generated for $k \in \{1, \dots, 10\}$ and compared to the observed message. To focus further analysis of the best response models, the value of k that minimizes the mean absolute deviation between the best response prediction and the data is selected from $k \in \{1, \dots, 10\}$. Define k^* to be the parameter that minimizes

$$\sum_{g,i=1}^5 \sum_{s=1}^4 \frac{1}{51 - t_{\min}} \sum_{t=t_{\min}}^{50} \left(\inf_{\hat{m}_i \in \mathcal{B}_{g,i}(\bar{\mathbf{m}}_{g,s,-i}^{t,k}; \theta_i)} \|m_{g,s,i}^t - \hat{m}_i\| \right),$$

where g represents each of the five mechanisms under consideration, s indexes the 4 identical sessions of each mechanism, and $\|\cdot\|$ is the standard Euclidean norm.²⁰ Since the first k periods of each model are used to seed the initial beliefs, they must be excluded from analysis. Consequently, t_{\min} must be strictly larger than k . In four of the five mechanisms, $\mathcal{B}_{g,i}(\bar{\mathbf{m}}_{g,s,-i}^{t,k}; \theta_i)$ is unique and $\mathcal{M}_{g,i} \subseteq \mathbb{R}^1$, so the term in parentheses reduces to a simple absolute difference. In the cVCG mechanism, this term represents the orthogonal distance from the observed message to the appropriate best response line.

Table 3 reports the average deviation for various values of k and t_{\min} . Note that for every value of k considered, the average score decreases in t_{\min} , indicating that the models are less accurate in early periods than in later periods. Therefore, comparisons between models should only be made for fixed values of t_{\min} .

Given that messages are serially dependent and the nature of this dependence is unknown, no appropriate notion of significance is applicable to this analysis.²¹ The objective of this subsection is to make further analyses tractable by selecting a single value k^* to represent the class of k -period best response models. Therefore, statistical significance of the difference in quality of fit between best response models is unimportant in this context; choosing the minimum-deviation model is sufficient.

Result 1. *Among the k -period best response models with $k \in \{1, \dots, 10\}$, the 5-period model is estimated to be the most accurate.*

²⁰ A ‘scoring rule’ such as the quadratic scoring rule characterized by Selten [29] would be more appropriate if the behavioral models generated probabilistic predictions of play. With deterministic behavioral models and a continuum strategy space, the scoring rule simply counts the number of observations that *exactly* match the prediction. In the Walker mechanism, for example, 47,025 individual predictions are generated across the 10 models and only one of them is exactly correct. Since the strategy space is endowed with a distance metric, a notion of error based on that metric is used rather than a measure of error in the space of (degenerate) probability distributions. This is similar in spirit to the notion of error in econometric models, where mean squared deviation is most often used because it is tractable and can be interpreted as an estimate of variance. These considerations do not apply here, and the mean absolute deviation measure is more robust to outliers.

²¹ Serial dependence is clear from inspection of correlograms. Several models of serial dependence were estimated, including various time trend regressions, GARCH models, and a variety of stochastic differential equations. None of these procedures fit the data well or generated an uncorrelated error structure.

Table 3

Calculated average quadratic score for various k -period best response models

k	First time period used in calculating average minimum deviation (t_{\min})									
	2	3	4	5	6	7	8	9	10	11
1	1.407	1.394	1.284	1.151	1.104	1.088	1.072	1.054	1.054	1.049
2	—	1.240	1.135	0.991	0.967	0.949	0.932	0.922	0.913	0.910
3	—	—	1.097	0.963	0.940	0.925	0.904	0.888	0.883	0.875
4	—	—	—	0.952	0.932	0.915	0.898	0.877	0.866	0.861
5	—	—	—	—	0.924	0.911	0.895	0.876	0.860	0.853
6	—	—	—	—	—	0.911	0.897	0.881	0.868	0.854
7	—	—	—	—	—	—	0.899	0.884	0.873	0.863
8	—	—	—	—	—	—	—	0.884	0.874	0.864
9	—	—	—	—	—	—	—	—	0.879	0.870
10	—	—	—	—	—	—	—	—	—	0.875

Bold-faced entries represent, for each value of t_{\min} , the smallest average quadratic score among the 10 models tested. Note that the measure cannot be calculated for $k \geq t_{\min}$ since k periods are used to “seed” the model.

Support. The result follows immediately from inspection of the average deviation measures in Table 3. The measures are strictly decreasing in k for all $t_{\min} \leq 5$ (for which the $k = 5$ model cannot be calculated). For every $t_{\min} \geq 6$, $k = 5$ minimizes the average score, with one minor exception.²²

Other types of best response models were also considered. For example, a more general k -period model that includes a discount factor so that more recent observations receive greater weighting slightly outperforms the undiscounted $k = 5$ model, but the increase in accuracy is marginal considering the added parameter. Empirical analysis of the best response models will henceforth be limited to the undiscounted 5-period model.²³

6.2. Best response in non-VCG mechanisms

Due to the substantial difference between the structure of the first four mechanisms and that of the cVCG mechanism, results pertaining to the latter will be considered separately.

Given that each of these public goods mechanisms was developed under the assumption that agents play Nash equilibrium strategies, the static Nash equilibrium serves as a key benchmark against which the best response models may be tested, even though the experimental environment is one of incomplete information. This is particularly true for the Groves–Ledyard and Walker mechanisms, where the Nash equilibrium is the only point at which efficient outcomes obtain. If a dynamic best response model is found to provide significant improvement in predictive power over Nash equilibrium, then mechanism design theory is improved by insisting that mechanisms converge quickly under this dynamic.

²² Table 3 was also generated using a squared deviation metric. In this case, $k = 8$ yields slightly smaller error measures than $k = 5$ when $t_{\min} = 8$ or 10, but $k = 5$ is more often the minimizer.

²³ Complete analysis was performed on all models in $k \in \{1, \dots, 10\}$, and results for $k \geq 2$ are similar to the case of $k = 5$. As is apparent from Table 3, the $k = 1$ model is notably less accurate than the others because the smoothing achieved by the $k \geq 2$ models provides a better fit.

6.2.1. Comparison of best response and equilibrium models

The goal of this section is to determine whether the error of the best response model is significantly smaller than the error of the Nash equilibrium prediction. Standard parametric tests are inappropriate for this data because a one-sample runs test for randomness indicates that neither time series of errors is randomly drawn from a zero-median distribution, and tests for correlation indicate that the errors are serially dependent.²⁴ This nonstationarity implies that statistics aggregated across time may be easily misinterpreted.²⁵ For example, the average prediction error across all periods does not estimate the expected error in any one period; analysis of the average must be considered specific to the length of the experiment.²⁶ For these reasons, empirical analysis is performed on each player type in each period individually, with data aggregated only across the four sessions of each mechanism. The results of these period-by-period tests cannot be aggregated across time.

The prediction error of each model, averaged across the four sessions, is presented in Figs. 1–4. Around each data point is a 95% confidence interval generated by the bias-corrected and accelerated bootstrapping method.²⁷ These graphs begin to illustrate the superiority of the best response model over the Nash equilibrium model. While the two predictions are often very similar, there are certain player types for whom the equilibrium model systematically under- or over-predicts the observed strategies. The best response model appears both more accurate and more precise than the equilibrium model whenever differences between the two are observed.

The statistical analysis is aimed at testing the null hypothesis

$$H_0 : \mathbb{E} [|m_i^t - m_i^*(\theta)|] \leq \mathbb{E} [|m_i^t - \mathcal{B}_i(\bar{\mathbf{m}}_{-i}^{t,k}; \theta_i)|] \quad (5)$$

for each player type i and period $t > k$, where the expectation is taken across the four sessions of each mechanism. A non-parametric permutation test for a difference in means between the two model errors is performed in each period for each player type in each mechanism. Each test was based on an simulated distribution of 2000 draws, more than enough to minimize the variation in estimated p -values due to random sampling.

The power of the permutation test depends on the difference between the predictions of the two models. If the two models have very similar predictions, the outcome of the test will not yield strong posteriors about the truth of the alternative hypothesis. In the following analysis, tests will only be run when there is enough power to conclude that

²⁴ The runs test indicates that the best response model errors for 16 of the 20 total player types are not evenly scattered about zero at a significance level of 5%. Each of the 4 player types with model errors apparently randomly drawn from a zero-median distribution were from different mechanisms, indicating that the assumption of mean-zero random errors for all player types in any one mechanism is likely invalid. The errors of the equilibrium model were not evenly scattered about zero for 19 of 20 player types at the 5% significance level. Tests of first-order correlation indicate that the errors are serially correlated for all 20 player types in both models.

²⁵ The dependence also implies that neither model fully captures the true dynamics of subject behavior in repeated games.

²⁶ This is a point occasionally forgotten in past analyses of time series data in experiments, leading to results that likely depend on the somewhat arbitrary choice of experiment length.

²⁷ Two thousand draws are used in each period for each player type in each mechanism, which is more than enough to eliminate any bootstrap sampling error. See Efron and Tibshirani [15] for details on the bootstrapping method and related statistical tests.

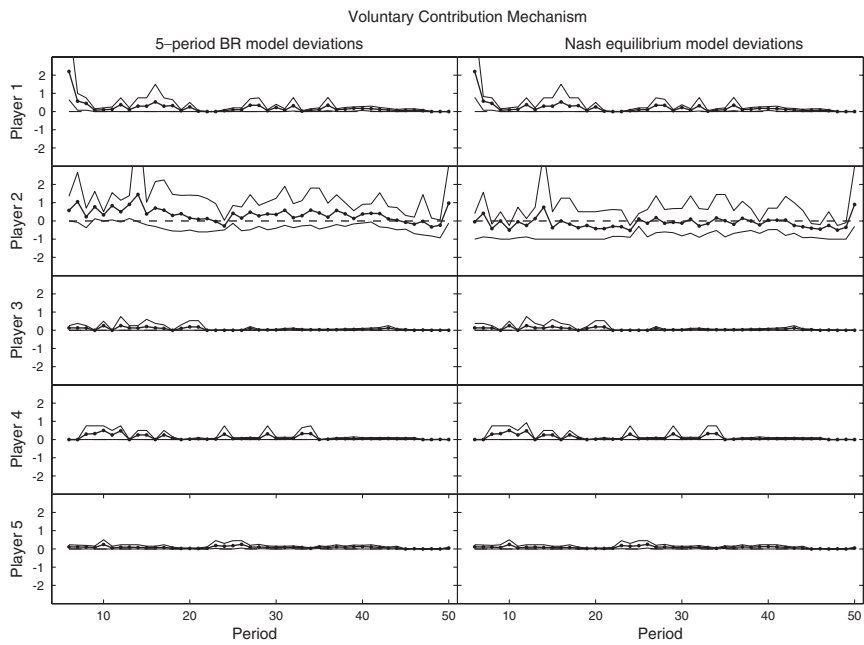


Fig. 1. Average model error for the 5-period best response model and Nash equilibrium models in the Voluntary Contribution mechanism.

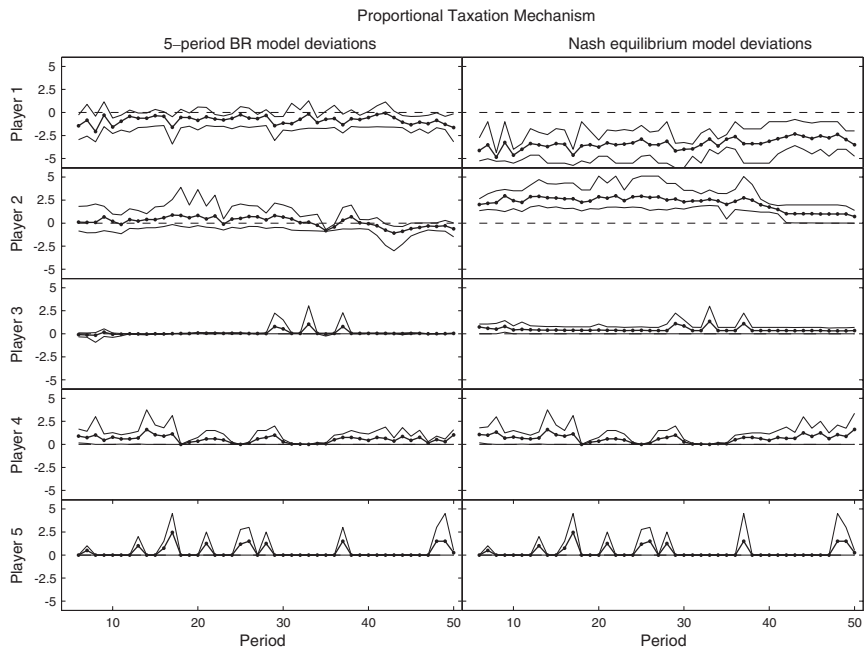


Fig. 2. Average model error for the 5-period best response model and Nash equilibrium models in the Proportional Taxation mechanism.

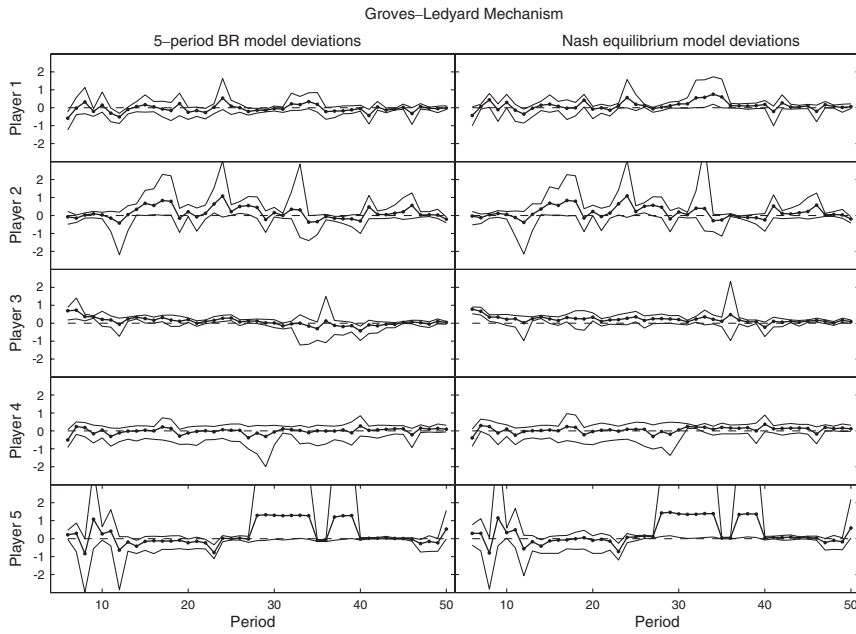


Fig. 3. Average model error for the 5-period best response model and Nash equilibrium models in the Groves–Ledyard mechanism.




Fig. 4. Average model error for the 5-period best response model and Nash equilibrium models in the Walker mechanism.




Fig. 5. Simulated power of the permutation test given in equation for various differences in model predictions as a ratio of the data's standard deviation.

$\mathbb{P}[H_A \text{ True} | \text{Reject } H_0] \geq 90\%$. Assuming diffuse priors on the truth of H_A and a standard significance level of 5%, this is equivalent to requiring a test power of at least 45%. Without this level of power, it is likely that a rejection of H_0 is due to random sampling rather than an actual difference in model errors. If one additionally required that $\mathbb{P}[H_A \text{ True} | \text{Do Not Reject } H_0] \leq 10\%$, a power of 89.4% would be needed.

In order to identify what difference between model predictions is necessary to guarantee a test power of 45%, a simulation of the permutation test is performed for various differences between model predictions. Specifically, four independent, normally distributed random messages (w_1, \dots, w_4) are generated with mean μ_a and variance σ_w^2 . These represent four observations of a particular player type in a particular period. This is repeated 100 times and the permutation test is performed in each repetition on the hypotheses $\tilde{H}_0 : \mathbb{E}[|w - \mu_b|] \leq \mathbb{E}[|w - \mu_a|]$, where μ_a and μ_b represent the predictions of two different models, the first of which is correct in the sense that it predicts the true mean of the data. An estimate of the power of the permutation test is given by the percentage of simulated tests that correctly reject \tilde{H}_0 . The simulation is repeated for various values of $(\mu_a - \mu_b)$ and σ_w^2 and the estimated power of the test is plotted as a function of $(\mu_a - \mu_b)/\sigma_w$ in Fig. 5.²⁸

²⁸ It should be noted that if the mean of the data were $\mu_w \neq \mu_a$ and $\mu_w > \mu_a > \mu_b$, then the test would have more power. If $\mu_a > \mu_w \geq (\mu_a + \mu_b)/2$, the test would have less power.

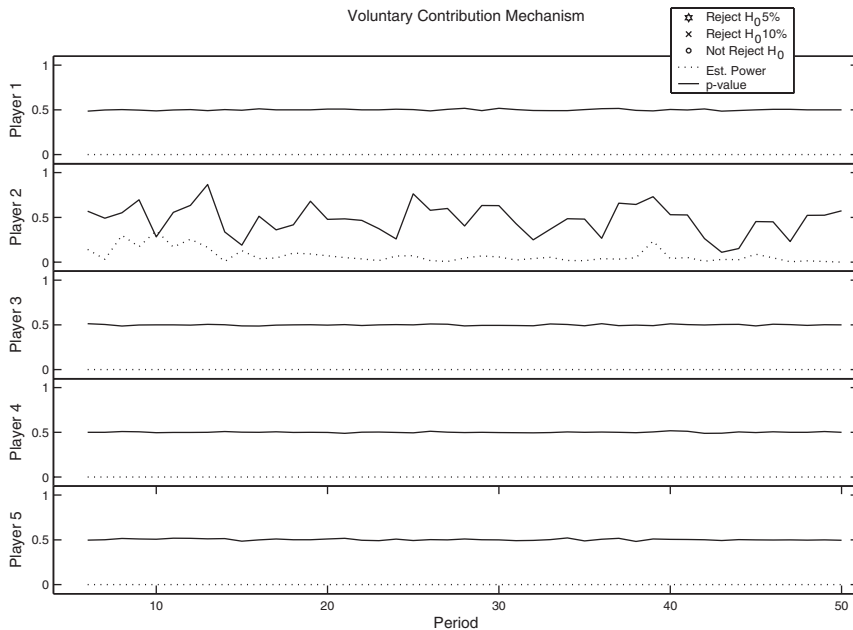


Fig. 6. p -values and estimated power for the permutation tests in the Voluntary Contribution mechanism. Stars, Xs, and Os represent test results for those tests with a power of at least 0.45. Star represents rejection of H_0 at the 5% level, X represents rejection at the 10% level, and O represents no rejection of H_0 . Note that this mechanism has no test with power $\geq 45\%$.

From this graph, it is clear that the distance between the two predictions should be at least 1.75 standard deviations of the data in order to keep the probability that a rejection of H_0 is incorrect under 10%.

Figs. 6–9 display the p -value of the permutation test for each player type in each period, along with the estimated power of each test (from Fig. 5) and, for those tests with power greater than 45%, whether or not the test rejects the null hypothesis at the 5% and 10% significance levels.

Result 2. *The 5-period best response model is overall a more accurate model than the Nash equilibrium model for the non-VCG mechanisms.*

Support. In the Voluntary Contribution mechanism (Fig. 6), players 1, 3, 4, and 5 have a strict dominant strategy, so the power of the test is zero for these player types. Although the p -values for player 2 never indicate a rejection of the null hypothesis, the power of the test is always below 40%, rendering its conclusions ambiguous.

In the Proportional Tax mechanism (Fig. 7), players 3, 4, and 5 have little incentive to contribute given the contributions of others, so that the free-riding equilibrium strategy is most often a best response. For players 1 and 2, best response is occasionally far from equilibrium, providing enough power for the permutation tests to be conclusive. For player

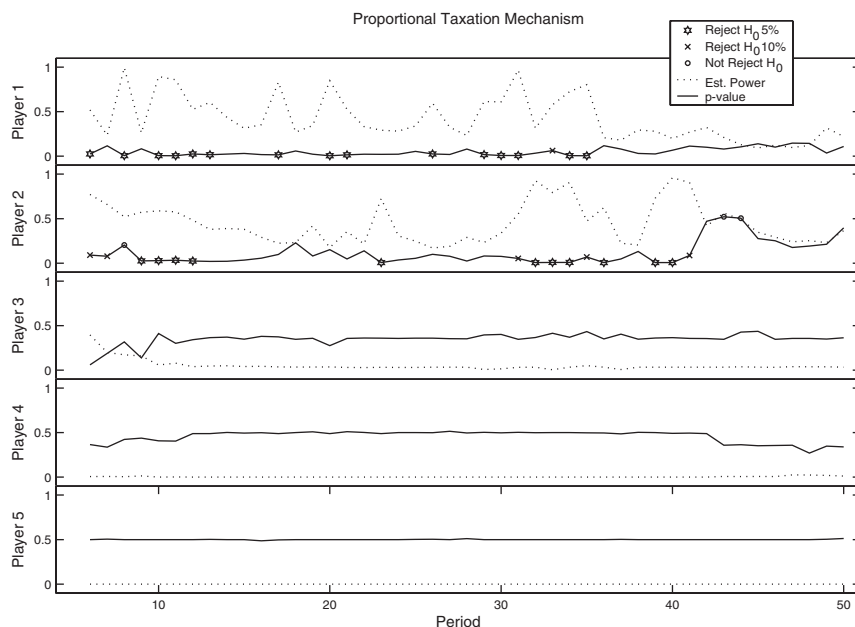


Fig. 7. p -values and estimated power for the permutation tests in the Proportional Tax mechanism. Stars, Xs, and Os represent test results for those tests with a power of at least 0.45. Star represents rejection of H_0 at the 5% level, X represents rejection at the 10% level, and O represents no rejection of H_0 .

1, all 16 tests with sufficient power reject the null hypothesis at the 10% level, and 15 of 16 reject at the 5% level. Player 2's results are similar, although the data revert toward equilibrium in the final periods (see Fig. 2 as well).

The rapid convergence of the Groves–Ledyard mechanism (Fig. 8) to equilibrium, which is accompanied by the convergence of best response predictions to equilibrium, reduces the power of the test in most periods. The 10 tests with sufficient power (out of 175 possible) do favor the Nash equilibrium model, though the test power being well below 89.4% in all 10 tests prevents conclusive rejection of H_A .

The Walker mechanism (Fig. 9) provides the most testing power due to the lack of convergence of the data. Of the 41 tests due to player 3, only 2 fail to reject the null hypothesis at the 10% level, while 37 tests reject H_0 at the 5% level. Of the 37 tests due to player 5, 8 fail to reject H_0 at the 10% level, while 12 tests reject H_0 at the 10% level and 17 tests reject at the 5% level. The p -values of all 37 tests are below 0.20. These rejections are scattered evenly throughout the session, indicating no particular pattern over time. Of the other 3 player types, players 1 and 2 have low p -values on average, with five rejections of H_0 and only 1 failure to reject. Player 4 shows mixed support overall, and in the few tests with sufficient power, shows fairly strong support for equilibrium behavior. This can also be seen in Fig. 4.

The significance of this result lies in its implications for implementation in a repeated interaction setting, where the assumption that agents play the stage game equilibrium is less

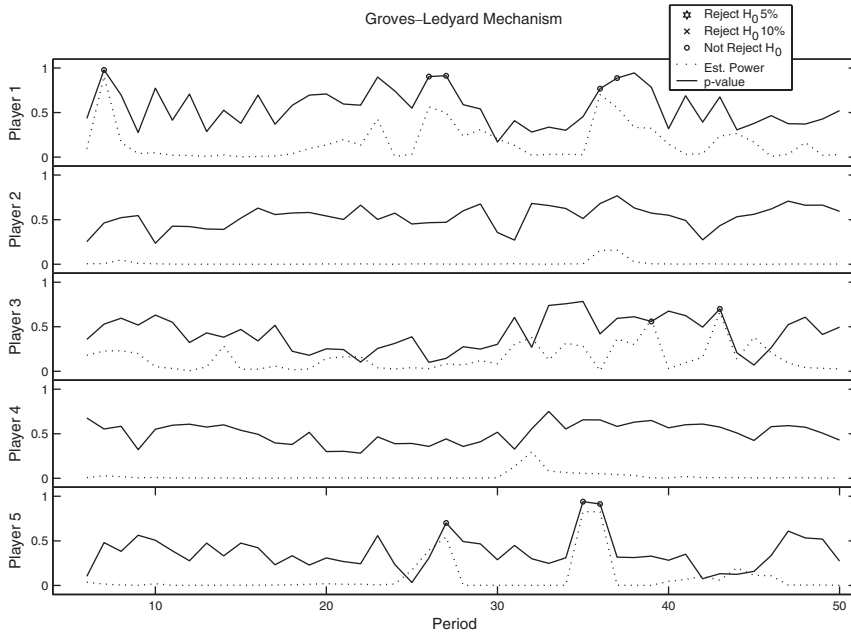


Fig. 8. p -values and estimated power for the permutation tests in the Groves–Ledyard mechanism. Stars, Xs, and Os represent test results for those tests with a power of at least 0.45. Star represents rejection of H_0 at the 5% level, X represents rejection at the 10% level, and O represents no rejection of H_0 .

accurate than a simple best response behavioral assumption. Mechanisms constructed under the assumption of equilibrium behavior may fail to implement the desired outcome due to instability in the behavioral process. Although the best response model does not provide a complete description of human behavior, a mechanism designer who assumes this simple dynamic will be able to more accurately implement the desired outcomes than a designer who assumes static equilibrium behavior.

6.3. Best response in the cVCG mechanism

In this direct mechanism, agents announce both \hat{a}_i and \hat{b}_i . The decisions of subjects can be grouped into three categories: ‘full’ revelation ($\hat{a}_i = a_i$ and $\hat{b}_i = b_i$, which is the dominant strategy), ‘partial’ revelation ($\hat{a}_i = a_i$ or $\hat{b}_i = b_i$), and no revelation. The equilibrium model predicts that all messages will be of the first type. The best response model predicts that all three types of messages are possible, but messages that are not fully revealing must lie along the best response surface. The following subsections look at (a) what percentage of messages are fully revealing, and (b) whether non-revealing messages are scattered randomly or are centered around the line of best responses.

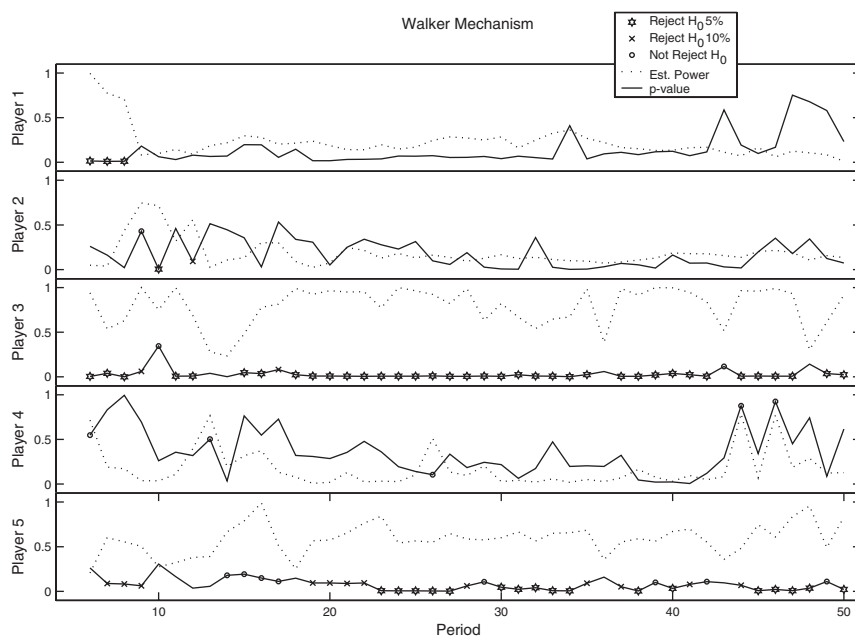


Fig. 9. p -values and estimated power for the permutation tests in the Walker mechanism. Stars, Xs, and Os represent test results for those tests with a power of at least 0.45. Star represents rejection of H_0 at the 5% level, X represents rejection at the 10% level, and O represents no rejection of H_0 .

6.3.1. Frequency of revelation

Previous experimental tests of dominant strategy mechanisms with a weak dominant strategy indicate that around half of all subjects play their dominant strategy. Table 4 indicates that this result holds true in the current study. Rates of both full and partial revelation in the cVCG mechanism are reported. Note a message that varies from truth-telling by any amount is encoded as non-revealing.

Result 3. *Truthful revelation in the cVCG mechanism is observed in the majority of decisions. This frequency increases in the final periods.*

Support. Refer to Table 4. On average, 54% of all observed messages are full revelation strategies, with the frequency increasing to 59% in the final 10 periods. Average partial revelation rises from 64% to 72% over the last 10 periods. Half of the 20 subjects fully reveal in at least nine of the last 10 periods, and 12 at least partially reveal. Three subjects never choose full revelation, and an additional three subjects reveal fully only twice. Every subject reveals partially at least twice over the course of the experiment.

Analysis of individual data also reveals that in 98% of the cases where a subject only partially reveals, it is the \hat{b}_i term that is misrepresented. This is likely due to the fact that altering the linear term has a more transparent effect on payoffs than the quadratic term.

Table 4

Frequency of revelation of both parameters ('Full' revelation) and of only one parameter ('Partial' revelation) by each subject in the cVCG mechanism for all 50 periods and for the last 10 periods

Session	Periods	Player 1		Player 2		Player 3		Player 4		Player 5		Average	
		Full	Partial	Full	Partial	Full	Partial	Full	Partial	Full	Partial	Full	Partial
Session 1	All 50	0.82	0.94	0.76	0.78	0.00	0.04	0.88	0.90	0.62	0.76	0.62	0.68
	Last 10	1.00	1.00	1.00	1.00	0.00	0.00	1.00	1.00	1.00	1.00	0.80	0.80
Session 2	All 50	0.92	1.00	0.28	0.30	1.00	1.00	0.00	0.04	0.88	0.94	0.62	0.66
	Last 10	1.00	1.00	0.60	0.60	1.00	1.00	0.00	0.00	0.70	0.70	0.66	0.66
Session 3	All 50	0.40	0.80	0.04	0.06	1.00	1.00	0.48	0.48	0.70	0.72	0.52	0.61
	Last 10	0.00	1.00	0.00	0.00	1.00	1.00	0.50	0.50	0.90	0.90	0.48	0.68
Session 4	All 50	0.00	0.68	0.94	0.98	0.92	1.00	0.04	0.08	0.04	0.34	0.39	0.62
	Last 10	0.00	1.00	1.00	1.00	1.00	1.00	0.00	0.10	0.00	0.60	0.40	0.74
Average	All 50	0.54	0.86	0.51	0.53	0.73	0.76	0.35	0.38	0.56	0.69	0.54	0.64
	Last 10	0.50	1.00	0.65	0.65	0.75	0.75	0.38	0.40	0.65	0.80	0.59	0.72

6.3.2. Misrevelation and weakly dominated best responses

Recall from Section 5.5 that, given $\hat{\theta}_{-i}$, agent i has a line of best responses through θ_i that are payoff equivalent to truth-telling. The slope of this line depends on $\hat{\theta}_{-i}$, so the best response line is sensitive to a player's prediction about the strategies of the others. Since the k -period best response model provides a point prediction of $\hat{\theta}_{-i}$, it is easily testable in this framework. In particular, the model cannot be rejected if misrevelation messages are centered around the particular best response line suggested by the k -period average prediction.

A convenient method for analyzing the data is to convert each two-dimensional message $(\hat{a}_i^t, \hat{b}_i^t)$ into polar coordinates $(\hat{\phi}_i^t, \hat{r}_i^t)$ with the origin at the truthful revelation point (a_i, b_i) . Here, $\hat{\phi}_i^t$ represents the angle from (a_i, b_i) to $(\hat{a}_i^t, \hat{b}_i^t)$ and \hat{r}_i^t represents the distance between these points. Fully revealing observations are not included in this analysis since they are consistent with both best response and equilibrium play.

For each period $t > k$ and player i , the k -period best response model identifies a particular angle $\phi_i^B(\hat{\theta}_{-i}^k; \theta_i)$ such that any announcement $\hat{\theta}_i^t$ with $\hat{\phi}_i^t = \phi_i^B$ is a best response to the average of the previous k values of $\hat{\theta}_{-i}$. Fig. 10 graphs $(\hat{\phi}_i^t - \phi_i^B, \hat{r}_i^t)$ in polar coordinates, where ϕ_i^B is the angle generated from the 5-period model. In this figure, the origin ($\hat{r}_i^t = 0$) represents a full revelation announcement and the horizontal axis ($\hat{\phi}_i^t - \phi_i^B$) represents a message consistent with the 5-period best response model. If subjects play the dominant strategy equilibrium, the data should scatter evenly around the origin. If subjects follow the 5-period best response model, the data should scatter around the horizontal axis.




Fig. 10. Polar-coordinate representation of the cVCG data with the origin at the truth-telling point and the horizontal axis corresponding to the 5-period best response surface. Two different scalings of the same graph are presented.

Unfortunately, the removal of all full-revelation observations reduces the average sample size to less than two observations per period—not enough to perform a statistical test. Qualitatively, the evidence strongly supports the 5-period best response model. Half of all non-equilibrium observations are within 1.3° of the best response line and 81% are within 10° of the prediction. This analysis includes all partial revelation observations for which $\hat{\phi}_i^t$ is necessarily a multiple of $\pi/2$. After removing these observations, just over half of the remaining data are within 0.83° of the best response prediction, and 79% are within 10° .

Fig. 11 shows the time-series representations of \hat{r}_i^t and $\hat{\phi}_i^t - \phi_i^B(\bar{\theta}_{-i}^{t,k}; \theta_i)$ for each player type in the cVCG mechanism. The 95% confidence intervals are again formed by the bias-corrected and accelerated bootstrapping method with 2000 draws. The average distance from truth-telling is frequently large, highly variable, and does not converge toward zero for three of the five player types. The graphs of $\hat{\phi}_i^t - \phi_i^B(\bar{\theta}_{-i}^{t,k}; \theta_i)$ across time show that the off-equilibrium data are centered at or near the best response prediction, with more stability in later periods. Again, small sample sizes prevent clean statistical analyses.

The tendency for the angular deviation to be slightly positive by about 6° , which is visible in both figures, arises from the partial revelation observations. Around 87% of the best response lines are between 83° and 85° while 20% of all off-equilibrium observations are partial revelation strategies located at 90° .

6.4. Testing theoretical properties of the best response model

In Section 3, various theoretical properties of the k -period average best response model are derived. Each of these may be tested empirically to confirm that the important implications of this behavioral assumption are observed in the laboratory.

In the cVCG mechanism, the best response line for each player is characterized by Eq. (4). Given $\hat{\theta}_{-i}$, each of these is a single equation of two variables, so the set of Nash equilibria of the game is the set of solutions to a system of n equations with $2n$ variables. Clearly, there




Fig. 11. Time series of average distance from equilibrium and angular deviation from the best response line for each player type in the cVCG mechanism with 95% confidence intervals.

exists a large number of such solutions. An observed strategy profile $\hat{\theta}$ is an ε -equilibrium if, for each $i \in \mathcal{I}$, the loss in utility between announcing θ_i and $\hat{\theta}_i$ is less than ε . This does not necessarily indicate that $\hat{\theta}_i$ is close to θ_i since $\hat{\theta}_i$ may be in a neighborhood of $\mathcal{B}_i(\hat{\theta}_{-i}; \theta_i)$ that is far from θ_i .

Result 4. *Weakly dominated ε -Nash equilibria are observed, while the dominant strategy equilibrium is not.*

Support. Setting $\varepsilon = 1$, 30.5% of observed strategy profiles in the cVCG mechanism are weakly dominated ε -Nash equilibria.²⁹ At $\varepsilon = 5$, 67% of the profiles are ε -equilibria. Across the last 12 periods, ε -equilibria are observed 93.8% of the time for $\varepsilon = 5$. In the first session, subjects play a particular ε -equilibrium (for $\varepsilon = 1/2$) in each of the final 19 periods. In *none* of the 200 repetitions of the cVCG mechanism is the truth-telling dominant strategy equilibrium observed.

Beyond providing further support for a best response model of behavior, this result has greater implications; it suggests that elimination of weakly dominated strategies leads to

²⁹ Agents typically earn over 300 francs per periods, so $\varepsilon = 1$ represents a deviation from optimality of less than 0.33%.

the elimination of certain Nash equilibria that are observed in the laboratory. This equilibrium selection technique is consequently inappropriate as a realistic equilibrium selection algorithm. The following result indicates that elimination of *strictly* dominated strategies is consistent with observed behavior.

Result 5 (Proposition 1). *Messages quickly converge to and do not significantly deviate from strictly dominant strategies after period 5.*

Support. Players 1, 3, 4, and 5 in the Voluntary Contribution mechanism are the only players with strictly dominant strategies. From Fig. 1, it is clear that these players quickly converge to their equilibrium strategies. The bootstrapped confidence intervals must lie in the strategy space. Of the 180 confidence interval lower bounds after period 5, only 15 are different from zero, and only 1 is greater than 0.1. In the last 9 periods, the *upper bound* of the intervals are all no greater than 0.25, and no greater than 0.1 in the last 4 periods.

The following results indicate that convergence to or repetition of a message profile is often indicative of a Nash equilibrium, as predicted by the best response model of behavior.

Result 6 (Proposition 2). *If a strategy profile is observed in 6 consecutive periods, then it is most likely a Nash equilibrium.*

Support. In the non-cVCG mechanisms, there are 754 messages m_i^t such that $m_i^t = m_i^{t-1} = \dots = m_i^{t-5}$. Of those, 74.8% are Nash equilibrium messages. 80.1% of such messages are within 1 unit of Nash equilibrium. In the cVCG mechanism, 45% of the 375 such messages are ε -equilibria with $\varepsilon = 1$. Setting $\varepsilon = 5$ increases the frequency to 82.1%.

Result 7 (Proposition 4). *If a sequence of strategy profiles converges to a point q , then q is most likely a Nash equilibrium strategy profile.*

Support. Of the 20 groups across the 5 mechanisms, only one played the same strategy profile in all of the last 10 periods, indicating convergence to a particular strategy profile; the first session of the cVCG mechanism converged to an ε -equilibrium (for $\varepsilon \geq 1/2$) in all of the final 19 periods. One group in the Proportional Tax mechanism played a particular non-equilibrium strategy in 15 of the final 25 periods while another group played the Nash equilibrium profile in 7 of the final 10 periods.

Finally, Propositions 7 and 9 are confirmed empirically by the convergence of the data to equilibrium in the Groves–Ledyard mechanism, which is supermodular *and* satisfies the dominant diagonal condition for the given parameters.

Overall, the above results indicate that the dynamic properties of observed behavior are generally in line with the theoretical properties of the k -period best response dynamic. The k -period dynamic is apparently a reasonably accurate yet tractable model for predicting repeated game behavior and convergence in these settings.




Fig. 12. The average level of public good and realized efficiency for each mechanism in each period along with 95% confidence intervals. PO represents the Pareto Optimal level of the public good (4.8095) and IR represents the efficiency of the initial endowments (71.18%).

6.5. Efficiency and public good levels

The ability to compare data across a fairly large number of mechanisms leads to the natural question of which mechanisms generate the most efficient outcomes. In fact, this study provides a unique opportunity to do so since no other experiment to date has tested as many processes side-by-side. It should be understood, however, any experimental result may be very sensitive to changes in parameters. For example, it may be the case that if the efficiency of the dominant strategy equilibrium of the cVCG mechanism were lower, then subjects would play it less often, further reducing realized efficiencies. With this note of caution, the average public good level and realized efficiency are presented in Fig. 12 along with 95% bootstrap confidence intervals.

Result 8. *For the given parameters, the average public good levels are closest to the Pareto optimal level in the cVCG mechanism, followed by the Groves–Ledyard mechanism. Overall efficiency in these two mechanism is higher than the others, with the Walker mechanism often resulting in efficiency below that of the initial endowment.*

Support. The average public good level is not significantly different from the Pareto optimum ($y^P = 4.8095$) in 43 of the 50 periods for the cVCG mechanism and 35 of 50

periods in both the Groves–Ledyard and Walker mechanisms. However, the realized public good level in the Walker mechanism is highly variable. The average public good level in the cVCG mechanism is not significantly different from the Pareto optimum in 22 of the final 25 periods, whereas the same is true in only 15 of the final 25 periods for the Groves–Ledyard mechanism. As predicted, the Voluntary Contribution mechanism significantly under-provides and the Proportional Tax mechanism over-provides the public good.

The cVCG mechanism is the most efficient, followed by the Groves–Ledyard mechanism. The average efficiency of the Walker mechanism is not significantly greater than the efficiency of the initial endowments (71.18%) in 38 of the 50 periods, and is significantly lower in 4 periods. This trend does not disappear across time.

Recall from Section 5.5 that there exists a large set of Nash equilibria of this mechanism that are necessarily outcome efficient, so the high level of efficiency achieved by the cVCG mechanism is due to best response behavior and not preference revelation. Also, the fact that the Walker mechanism often realizes efficiencies at or below the initial endowment is particularly surprising given that it is the only outcome-efficient mechanism tested whose equilibrium guarantees outcomes that Pareto dominate the endowment.

6.6. *Open questions*

The fundamental difficulty of testing the efficiency of mechanisms in the laboratory lies in their sensitivity to parameter choice. Past research has focused on changing punishment parameters of the mechanisms, but it is unknown how behavior differs when preferences are varied within a mechanism. In particular, the role of the incentive-efficiency trade-off in guiding behavior is not known.

The current set of experiments makes use of the ‘What-If Scenario Analyzer’ tool that enables subjects to calculate hypothetical payoffs. This tool is provided as an alternative to the payoff tables often provided in experiments. The effect of this tool on dynamic behavior is not well understood. The version of the software in use for this study does not store data about what hypothetical scenarios subjects considered; only actual decisions are tracked. The possibility of studying hypothetical explorations is an exciting extension to this research that may provide additional understanding about the learning dynamics in use.

On the theoretical front, much work remains with respect to dynamics in public goods mechanisms. The literature is far behind that of market dynamics for private goods, where stability has been extensively analyzed for several decades. Kim [22] shows that mechanisms Nash implementing Lindahl allocations must be locally unstable for some environment. However, restricting attention to quasilinear environments, a dynamically stable game form is introduced that implements the Lindahl correspondence. Similarly, de Tranquayle [14] and Vega-Redondo [34] introduce mechanisms that converge to the Lindahl allocation under Cournot best response behavior. Some mechanisms such as that studied by Smith [31] use convergence of tatonnement-like dynamics as a stopping rule in each period of the mechanism. However, there have been only limited attempts to seriously consider dynamic issues in the theoretical mechanism design literature.

7. Conclusion

Motivated by the observation that many of the results of previous experimental studies are consistent with a simple best response dynamic model, this paper experimentally compares five different public goods mechanisms in order to test this conjecture. In particular, dynamically stable and unstable Nash mechanisms are compared along with a weak dominant strategy mechanism whose best response properties provide an opportunity to distinguish between best response and equilibrium behavior. This latter mechanism, though tested in simpler forms in previous experiments, has never been tested in the laboratory with more than one preference parameter.

The results of these experiments support the best response behavioral conjecture, particularly as an alternative to the static equilibrium hypothesis. Strategies converge to Nash equilibria that are asymptotic attractors in a best response dynamic and diverge from equilibria that are not. In the weak dominant strategy mechanism, behavior tends to track a rotating best response line through the strategy space, implying that subjects who do not understand the undominated properties of truthful revelation instead seek a best response strategy, resulting in convergence to weakly dominated Nash equilibria. This result implies that elimination of weakly dominated strategies is an inappropriate tool for positive game theoretic analysis.

Although outcomes and efficiency are sensitive to parameters, the continuous VCG mechanism performs well in both categories, as does the Groves–Ledyard mechanism. The Walker mechanism, due to its instability, generates efficiencies often below that of the initial endowment.

The implications for mechanism design are straightforward. Most theorists have ignored dynamic stability in designing mechanisms. The significant contribution of this paper is that it bridges the behavioral hypotheses that have existed separately in dominant strategy and Nash equilibrium mechanism experiments. The finding that a 5-period average best response dynamic is a reasonably accurate behavioral model in all of these settings implies not only that dynamic behavior should be considered in theoretical research, but it also provides some guidance as to which behavioral models are appropriate. In particular, Nash implementation mechanisms should satisfy dynamic stability and dominant strategy mechanisms should satisfy the strict dominance property if either is to be considered for real-world use in a repeated interaction setting.

Acknowledgments

The author wishes to thank John Ledyard for financial support and helpful advice. This research has benefited from the comments of an anonymous referee, as well as Federico Echenique, Ken Binmore, Tim Cason, and Yoshi Saijo.

Appendix

Proof of Proposition 7. Let $\underline{\mathbf{m}} = \inf \mathcal{M}$ and $\overline{\mathbf{m}} = \sup \mathcal{M}$. Define $\underline{\mathcal{B}}(\mathbf{m}, \theta) = \inf \mathcal{B}(\mathbf{m}, \theta)$ and $\overline{\mathcal{B}}(\mathbf{m}, \theta) = \sup \mathcal{B}(\mathbf{m}, \theta)$ as the infimal and supremal best responses to a given message profile \mathbf{m} . Since the game is supermodular, $\underline{\mathbf{m}}$ and $\overline{\mathbf{m}}$ are finite elements of \mathcal{M} , and $\underline{\mathcal{B}}$ and $\overline{\mathcal{B}}$

are elements of \mathcal{B} for all \mathbf{m} and θ . Furthermore, $\underline{\mathcal{B}}$ and $\overline{\mathcal{B}}$ are non-decreasing functions of \mathbf{m} and there exists a smallest Nash equilibrium $\underline{\mathcal{E}}(\theta)$ and a largest Nash equilibrium $\overline{\mathcal{E}}(\theta)$.³⁰

Consider the sequence $\{\underline{\mathbf{m}}^t\}_1^\infty$ where $\underline{\mathbf{m}}^t = \underline{\mathbf{m}}$ for $t = 1, \dots, k$ and $\underline{\mathbf{m}}^t = \underline{\mathcal{B}}(\underline{\mathbf{m}}^{t-k}, \theta)$ for $t > k$. In order to establish by induction that the sequence is monotone increasing, assume that $\underline{\mathbf{m}}^t \geq \underline{\mathbf{m}}^s$ for some $t > k$ and all $s < t$. This is certainly true for $t = k + 1$. Since $\underline{\mathbf{m}}^t \geq \underline{\mathbf{m}}^{t-k}$, then $\underline{\mathbf{m}}^{t+1} \geq \underline{\mathbf{m}}^t$. By monotonicity of $\underline{\mathcal{B}}$, it must be that $\underline{\mathbf{m}}^{t+1} \geq \underline{\mathbf{m}}^t$. This implies that $\underline{\mathbf{m}}^{t+1} \geq \underline{\mathbf{m}}^s$ for all $s < t$. By induction, it is established that $\{\underline{\mathbf{m}}^t\}_1^\infty$ is a monotone increasing sequence. Since $\overline{\mathbf{m}}$ is finite, $\{\underline{\mathbf{m}}^t\}_1^\infty$ must converge to some point $\underline{\mathbf{m}}^* \in \mathcal{M}$. By Proposition 4, $\underline{\mathbf{m}}^*$ is a Nash equilibrium profile.

Assume that for some $t > k$, $\underline{\mathbf{m}}^s \leq \underline{\mathcal{E}}(\theta)$ for all $s < t$, which is true for $t = k + 1$. Then $\underline{\mathbf{m}}^t \leq \underline{\mathcal{E}}(\theta)$ and, by monotonicity of $\underline{\mathcal{B}}$, $\underline{\mathbf{m}}^t \leq \underline{\mathcal{B}}(\underline{\mathcal{E}}(\theta), \theta) \leq \underline{\mathcal{E}}(\theta)$. This implies that $\underline{\mathbf{m}}^s \leq \underline{\mathcal{E}}(\theta)$ for all $s \leq t + 1$. Therefore, the sequence $\{\underline{\mathbf{m}}^t\}_1^\infty$ is bounded above by $\underline{\mathcal{E}}(\theta)$.

Since $\{\underline{\mathbf{m}}^t\}_1^\infty$ converges to some equilibrium point, it must be that $\lim \underline{\mathbf{m}}^t = \underline{\mathcal{E}}(\theta)$. Similar induction arguments establish that the sequence $\{\overline{\mathbf{m}}^t\}_1^\infty$ of k -period average best responses starting from $\overline{\mathbf{m}}$ must converge to $\overline{\mathcal{E}}(\theta)$.

Now consider any arbitrary sequence $\{\mathbf{m}^t\}_1^\infty$. If $\underline{\mathbf{m}}^s \leq \mathbf{m}^s \leq \overline{\mathbf{m}}^s$ for all s less than some $t > k$, then by monotonicity of $\underline{\mathcal{B}}$ and $\overline{\mathcal{B}}$, it must be that $\underline{\mathbf{m}}^t \leq \mathbf{m}^t \leq \overline{\mathbf{m}}^t$. Since this hypothesis is true for $t = k + 1$, induction implies that $\underline{\mathbf{m}}^t \leq \mathbf{m}^t \leq \overline{\mathbf{m}}^t$ for all t . These bounds establish the result in the limit. \square

Proof of Proposition 9. Recall that $\mathcal{M}_i = [0, \infty)$ and, for each θ in some $\Theta_0 \subset \Theta$,

$$\lim_{m_i \rightarrow +\infty} \left| \frac{\partial u_i}{\partial m_i}(\mathbf{m}) \right| = +\infty$$

and $\left[\partial^2 u_i / \partial m_i \partial m_j \right]_{i,j=1}^n$ satisfies diagonal dominance on a set $\Theta_0 \subseteq \Theta$.

Fix $\theta \in \Theta_0$. Gabay and Moulin [16, Theorem 4.1] show that there must exist a unique Nash equilibrium $\mathbf{m}^*(\theta)$ and that diagonal dominance implies $\mathcal{B}(\mathbf{m}, \theta)$ is single-valued and strictly non-expansive in the sup-norm, so that for all $\mathbf{m}, \mathbf{m}' \in \mathcal{M}$,

$$\|\mathcal{B}(\mathbf{m}, \theta) - \mathcal{B}(\mathbf{m}', \theta)\|_\infty < \|\mathbf{m} - \mathbf{m}'\|_\infty, \quad (\text{A.1})$$

where $\|\mathbf{m}\|_\infty = \sup_i |m_i|$. If $\{\mathbf{m}^t\}_1^\infty$ is consistent with the k -period dynamic, then by (A.1),

$$\begin{aligned} \|\mathbf{m}^t - \mathbf{m}^*(\theta)\|_\infty &< \|\overline{\mathbf{m}}^t - \mathbf{m}^*(\theta)\|_\infty \\ &= \left\| \frac{1}{k} \sum_{s=1}^k (\mathbf{m}^{t-s} - \mathbf{m}^*(\theta)) \right\|_\infty \\ &\leq \frac{1}{k} \sum_{s=1}^k \|\mathbf{m}^{t-s} - \mathbf{m}^*(\theta)\|_\infty \\ &\leq \sup_{1 \leq s \leq k} \|\mathbf{m}^{t-s} - \mathbf{m}^*(\theta)\|_\infty \end{aligned}$$

³⁰ See Zhou [37] for details.

for every $t > k$. If

$$\lim_{t \rightarrow \infty} \sup_{1 \leq s \leq k} \|\mathbf{m}^{t-s} - \mathbf{m}^*(\theta)\|_{\infty} = 0, \quad (\text{A.2})$$

then convergence of $\{\mathbf{m}^t\}_1^{\infty}$ to $\mathbf{m}^*(\theta)$ is established.

Take any $\mathbf{q} \in \mathcal{M}^k \subseteq \mathbb{R}^{nk}$. For any such \mathbf{q} , there exists a unique sequence $\{\mathbf{m}^t\}_1^{\infty}$ consistent with the k -period dynamic such that $(\mathbf{m}^1, \dots, \mathbf{m}^k) = \mathbf{q}$. Define $G(\mathbf{q}, \theta) = (\mathbf{m}^{k+1}, \dots, \mathbf{m}^{2k})$ to be the next k terms of the k -period dynamic, all of which can be uniquely determined given \mathbf{q} and θ . Iterated application of G generates a sequence $\{\mathbf{q}^r\}_1^{\infty}$ where $\mathbf{q}^1 = \mathbf{q}$ and $\mathbf{q}^{r+1} = G(\mathbf{q}^r, \theta)$. Define $\mathbf{q}^*(\theta) = (\mathbf{m}^*(\theta), \dots, \mathbf{m}^*(\theta))$ and note that $\mathbf{q}^*(\theta)$ is a fixed point of $G(\cdot, \theta)$. Condition (A.2) can now be rewritten as $\lim_{r \rightarrow \infty} \|\mathbf{q}^r - \mathbf{q}^*(\theta)\|_{\infty} = 0$.

The following demonstrates that G is strictly non-expansive. Pick any points $\mathbf{q} = (\mathbf{m}^1, \dots, \mathbf{m}^k)$ and $\hat{\mathbf{q}} = (\hat{\mathbf{m}}^1, \dots, \hat{\mathbf{m}}^k)$. If $G(\mathbf{q}, \theta) = (\mathbf{m}^{k+1}, \dots, \mathbf{m}^{2k})$ and $G(\hat{\mathbf{q}}, \theta) = (\hat{\mathbf{m}}^{k+1}, \dots, \hat{\mathbf{m}}^{2k})$, then by (A.1),

$$\begin{aligned} \|\mathbf{m}^{k+1} - \hat{\mathbf{m}}^{k+1}\|_{\infty} &< \left\| \frac{1}{k} \sum_{s=1}^k (\mathbf{m}^s - \hat{\mathbf{m}}^s) \right\|_{\infty} \\ &\leq \frac{1}{k} \sum_{s=1}^{\infty} \|\mathbf{m}^s - \hat{\mathbf{m}}^s\|_{\infty} \\ &\leq \sup_{1 \leq s \leq k} \|\mathbf{m}^s - \hat{\mathbf{m}}^s\|_{\infty}. \end{aligned}$$

Similarly,

$$\|\mathbf{m}^{k+2} - \hat{\mathbf{m}}^{k+2}\|_{\infty} < \sup_{2 \leq s \leq k+1} \|\mathbf{m}^s - \hat{\mathbf{m}}^s\|_{\infty}.$$

By replacing $\|\mathbf{m}^{k+1} - \hat{\mathbf{m}}^{k+1}\|_{\infty}$ in the argument of the supremum with $\sup_{1 \leq s \leq k} \|\mathbf{m}^s - \hat{\mathbf{m}}^s\|_{\infty}$, the inequality becomes

$$\begin{aligned} \|\mathbf{m}^{k+2} - \hat{\mathbf{m}}^{k+2}\|_{\infty} &< \max \left\{ \sup_{2 \leq s \leq k} \|\mathbf{m}^s - \hat{\mathbf{m}}^s\|_{\infty}, \sup_{1 \leq s \leq k} \|\mathbf{m}^s - \hat{\mathbf{m}}^s\|_{\infty} \right\} \\ &= \sup_{1 \leq s \leq k} \|\mathbf{m}^s - \hat{\mathbf{m}}^s\|_{\infty}. \end{aligned}$$

Applying this reasoning to $\|\mathbf{m}^{k+t} - \hat{\mathbf{m}}^{k+t}\|_{\infty}$ for all $t = 2, \dots, k$ gives

$$\|\mathbf{m}^{k+t} - \hat{\mathbf{m}}^{k+t}\|_{\infty} < \sup_{1 \leq s \leq k} \|\mathbf{m}^s - \hat{\mathbf{m}}^s\|_{\infty} \quad \forall t = 1, \dots, k$$

or

$$\sup_{1 \leq s \leq k} \|\mathbf{m}^{k+s} - \hat{\mathbf{m}}^{k+s}\|_{\infty} < \sup_{1 \leq s \leq k} \|\mathbf{m}^s - \hat{\mathbf{m}}^s\|_{\infty}.$$

This is equivalent to

$$\|G(\mathbf{q}, \theta) - G(\hat{\mathbf{q}}, \theta)\|_{\infty} < \|\mathbf{q} - \hat{\mathbf{q}}\|_{\infty},$$

so $G(\cdot, \theta)$ is strictly non-expansive for each $\theta \in \Theta_0$ and for all $\mathbf{q}, \hat{\mathbf{q}} \in \mathcal{M}^k$. By an application of Edelstein's Theorem (see Ortega & Rheinboldt [27, p. 404]), the sequence $\{\mathbf{q}^r\}_1^\infty$ must converge to $\mathbf{q}^*(\theta)$.³¹ This implies that (A.2) holds, completing the proof. \square

References

- [1] G. Attiyeh, R. Franciosi, R.M. Isaac, Experiments with the pivot process for providing public goods, *Public Choice* 102 (2000) 95–114.
- [2] R.D. Auster, The invariably stable cobweb model, *Rev. Econ. Stud.* 38 (1971) 514–529.
- [3] D.V.T. Bear, The matrix multiplier and distributed lags, *Econometrica* 31 (1963) 514–529.
- [4] D.V.T. Bear, Distributed lags and economic theory, *Rev. Econ. Stud.* 33 (1966) 235–243.
- [5] G.W. Brown, Iterative solution of games by fictitious play, in: T.C. Koopmans (Ed.), *Activity Analysis of Production and Allocation*, Wiley, New York, 1951.
- [6] J.A. Carlson, The stability of an experimental market with supply-response lag, *The Southern Econ. J.* 33 (1967) 305–321.
- [7] T. Cason, T. Saijo, T. Sjöström, T. Yamato, Secure implementation experiments: Do strategy-proof mechanisms really work? California Institute of Technology, Department of Humanities and Social Sciences, Working Paper, 2003.
- [8] Y. Chen, Dynamic stability of nash-efficient public goods mechanisms: Reconciling theory and experiments, in: R. Zwick, A. Rapoport (Eds.), *Experimental Business Research*, vol. II, Kluwer Academic Publishers, Norwell, MA, 2004, forthcoming.
- [9] Y. Chen, Incentive compatible mechanisms for pure public goods: A survey of experimental research, in: C.R. Plott, V. Smith (Eds.), *The Handbook of Experimental Economics Results*, Elsevier Press, Amsterdam, 2005.
- [10] Y. Chen, R. Gazzale, When does learning in games generate convergence to nash equilibria? The role of supermodularity in an experimental setting, *Amer. Econ. Rev.* 94 (2004) 1505–1535.
- [11] Y. Chen, C.R. Plott, The Groves–Ledyard mechanism: An experimental study of institutional design, *J. Public Econ.* 59 (1996) 335–364.
- [12] Y. Chen, F.-F. Tang, Learning and incentive-compatible mechanisms for public goods provision: An experimental study, *J. Polit. Econ.* 106 (1998) 633–662.
- [13] E. Clarke, Multipart pricing of public goods, *Public Choice* 2 (1971) 19–33.
- [14] P. de Trenchay, Stable implementation of Lindahl allocations, *Econ. Lett.* 29 (1989) 291–294.
- [15] B. Efron, R. Tibshirani, *An Introduction to the Bootstrap*, Chapman & Hall, New York, 1993.
- [16] D. Gabay, H. Moulin, On the uniqueness and stability of nash-equilibria in noncooperative games, in: A. Bensoussan, P. Kleindorfer, C.S. Tapiero (Eds.), *Applied Stochastic Control in Econometrics and Management Science*, North-Holland, Amsterdam, 1980.
- [17] T. Groves, Incentives in teams, *Econometrica* 41 (1973) 617–633.
- [18] T. Groves, J.O. Ledyard, Optimal allocation of public goods: A solution to the ‘free-rider’ problem, *Econometrica* 45 (1977) 783–809.
- [19] R. Harrieff, D.V.T. Bear, J. Conlisk, Stability and speed of adjustment under retiming of lags, *Econometrica* 48.
- [20] R.M. Isaac, K.F. McCue, C.R. Plott, Public goods provision in an experimental environment, *J. Public Econ.* 26 (1985) 51–74.
- [21] T. Kawagoe, T. Mori, Can the pivotal mechanism induce truth-telling? An experimental study, *Public Choice* 108 (2001) 331–354.
- [22] T. Kim, Stability problems in the implementation of Lindahl allocations, Ph.D. Thesis, University of Minnesota, 1987.

³¹ Edelstein's Theorem requires compactness. In this case, the set $\{\mathbf{m} \in M : \|\mathbf{m} - \mathbf{m}^*(\theta)\|_\infty \leq \|\mathbf{m}^1 - \mathbf{m}^*(\theta)\|_\infty\}$ is compact for each $\{\mathbf{m}^t\}_1^\infty$ and, by (A.1), all subsequent elements of the sequence lie in this compact set. By restricting attention to this set, Edelstein's Theorem applies.

- [23] J.O. Ledyard, Public goods: A survey of experimental research, in: J. Kagel, A. Roth (Eds.), *Handbook of Experimental Economics*, Princeton University Press, Princeton, NJ, 1995.
- [24] P. Milgrom, J. Roberts, Rationalizability learning and equilibrium in games with strategic complementarities, *Econometrica* 58 (1990) 1255–1277.
- [25] P. Milgrom, J. Roberts, Adaptive and sophisticated learning in normal form games, *Games Econ. Behav.* 3 (1991) 82–100.
- [26] Y. Murata, *Mathematics for Stability and Optimization of Economic Systems*, Academic Press, New York, 1977.
- [27] J.M. Ortega, W.C. Rheinboldt, *Iterative Solution of Nonlinear Equations in Several Variables*, Academic Press, New York, 1970.
- [28] J.B. Rosen, Existence and uniqueness of equilibrium points for concave n -person games, *Econometrica* 33 (1965) 520–533.
- [29] R. Selten, Axiomatic characterization of the quadratic scoring rule, *Exper. Econ.* 1 (1998) 43–61.
- [30] L.S. Shapley, Some topics in two-person games, in: M. Dresher, L.S. Shapley, A.W. Tucker (Eds.), *Advances in Game Theory*, Princeton University Press, Princeton, 1964.
- [31] V. Smith, Incentive compatible experimental processes for the provision of public goods, in: V. Smith (Ed.), *Research in Experimental Economics*, JAI Press, Greenwich, CT, 1979.
- [33] D.M. Topkis, *Supermodularity and Complementarity*, Princeton University Press, Princeton, NJ, 1998.
- [34] F. Vega-Redondo, Implementation of Lindahl equilibrium: An integration of static and dynamic approaches, *Math. Soc. Sci.* 18 (1989) 211–228.
- [35] W. Vickrey, Counterspeculation, auctions and competitive sealed tenders, *J. Finance* 16 (1961) 8–37.
- [36] M. Walker, A simple incentive compatible scheme for attaining Lindahl allocations, *Econometrica* 49 (1981) 65–71.
- [37] L. Zhou, The set of nash equilibria of a supermodular game is a complete lattice, *Games Econ. Behav.* 7 (1994) 295–300.