

**Learning and Incentive-Compatible Mechanisms for Public Goods
Provision: An Experimental Study**



Yan Chen; Fang-Fang Tang

The Journal of Political Economy, Vol. 106, No. 3. (Jun., 1998), pp. 633-662.

Stable URL:

<http://links.jstor.org/sici?sici=0022-3808%28199806%29106%3A3%3C633%3ALAIMFP%3E2.0.CO%3B2-%23>

The Journal of Political Economy is currently published by The University of Chicago Press.

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/about/terms.html>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/journals/ucpress.html>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is an independent not-for-profit organization dedicated to creating and preserving a digital archive of scholarly journals. For more information regarding JSTOR, please contact support@jstor.org.

Learning and Incentive-Compatible Mechanisms for Public Goods Provision: An Experimental Study

Yan Chen

University of Michigan

Fang-Fang Tang

University of Bonn and National University of Singapore

This is the first systematic experimental study of the comparative performance of two incentive-compatible mechanisms for public goods provision: the basic quadratic mechanism by Groves and Ledyard and the paired-difference mechanism by Walker. Our experiments demonstrate that the performance of the basic quadratic mechanism under a high punishment parameter is far better than that of the same mechanism under a low punishment parameter, which, in turn, is better than that of the paired-difference mechanism. We estimate three individual behavioral models: an

We would like to thank John Ledyard, Charles Plott, Reinhard Selten, Mark Walker, and Luis Corchon for helpful discussions and comments; Robin Hanson for sharing the design and data of his pilot experiment on the paired-difference mechanism; Mark Olson for his help in organizing the Amsterdam experiments; Sergey Tsyplakov for making the payoff charts; Bruce Bell for computer programming for the experiments; and Klaus Abbink for his assistance in the initial data sorting. Seminar participants at Bielefeld, Bonn, Michigan, Wisconsin, Saarland, and the 1996 and 1997 North American summer meetings of the Econometric Society provided helpful comments. An anonymous referee and the editor provided many valuable suggestions. The research support provided by Rackham grant 386491 awarded to Chen, financial support to Chen and Tang from the Deutsche Forschungsgemeinschaft, SFB303 at the University of Bonn, and the support of the Laboratory for Experimental Economics and Political Science at Caltech and the Center for Research in Experimental Economics and Political Decision-Making at the University of Amsterdam are gratefully acknowledged. Chen also thanks the Department of Economics at the University of Bonn for its hospitality. Any remaining errors are our own.

[*Journal of Political Economy*, 1998, vol. 106, no. 3]

© 1998 by The University of Chicago. All rights reserved. 0022-3808/98/0603-0003\$02.50

exponentialized relative payoff sum model outperforms the generalized fictitious play model. We also provide a sufficient condition for convergence under the basic quadratic mechanism.

I. Introduction

How to design decentralized institutions to facilitate cooperation in an environment with public goods has been a challenging problem for economists for a long time. Natural processes, such as the voluntary contribution mechanisms, have been shown both theoretically and experimentally to be unable to solve the “free-rider” problem (Ledyard 1995). Therefore, since the 1970s, economists have been seeking decentralized mechanisms that are nonmanipulable and achieve Pareto-optimal allocation of resources with public goods.

By now it is well known that it is impossible to design a mechanism for making collective allocation decisions that is informationally decentralized, nonmanipulable, and Pareto-optimal (Green and Laffont 1977; Roberts 1979; Walker 1980). There are many mechanisms that preserve Pareto optimality at the cost of nonmanipulability, some of which preserve “some degree” of nonmanipulability. In particular, some mechanisms that have the property that Nash equilibria¹ are Pareto-optimal have been discovered. They can be found in the work of Groves and Ledyard (1977), Hurwicz (1979), and Walker (1981).

All these “next-best” mechanisms have very similar static properties, which leads one to consider properties other than optimality of Nash equilibria in an effort to distinguish among them. One important additional dimension of performance is the dynamics induced by these mechanisms in a laboratory. Any actual implementation is necessarily a dynamic process, starting somewhere off the equilibrium path. The fundamental question concerning implementation of a specific mechanism is whether the dynamic processes will actually converge to one of the equilibria promised by theory. If the dynamic processes do not converge, then the nice properties in equilibrium cannot be achieved. Therefore, it is crucial to study the dynamic properties of a mechanism and to extract the properties of static mechanisms that induce convergence. This motivates the research reported in this paper.

We select two Nash-efficient mechanisms to implement in a labo-

¹ Other implementation concepts include perfect Nash equilibrium (Bagnoli and Lipman 1989), undominated Nash equilibrium (Jackson and Moulin 1992), etc.

ratory: the basic quadratic mechanism by Groves and Ledyard (1977) and the paired-difference mechanism by Walker (1981). While the basic quadratic mechanism has been studied in laboratories, the paired-difference mechanism has not been systematically studied in laboratories. A comparison of these two families of mechanisms has not been performed either. This comparison allows us to abstract the properties that induce convergence when a mechanism is implemented among boundedly rational agents, that is, to answer the question, What properties of a static mechanism can induce the subjects to learn to play an equilibrium strategy?

To study the dynamic learning processes induced by these mechanisms, we use two major families of learning models: an exponentialized relative payoff sum model and a generalized fictitious play model. A static equilibrium model is also analyzed as a benchmark.

The paper is organized as follows. In Section II we review the theoretical properties of the basic quadratic and paired-difference mechanisms. Section III goes over the experimental design. Section IV summarizes the group-level results. Section V introduces the learning models and uses them to analyze the data. Section VI discusses two additional aspects of the mechanisms that induce good dynamics and provides a sufficient condition for convergence under the basic quadratic mechanism. Section VII reviews previous implementation of these mechanisms and compares the findings in this paper with those in Chen and Plott (1996). Section VIII concludes the paper.

II. The Mechanisms—Static and Dynamic Properties

Two families of mechanisms are studied in the same environment: the basic quadratic mechanism and the paired-difference mechanism. These two mechanisms have very similar static properties. Both are Nash-efficient and balanced with the same dimension of message space. The paired-difference mechanism is also individually rational in equilibrium. These properties are introduced in turn.

A. The Basic Quadratic Mechanism: Static Properties

The basic quadratic mechanism is the first mechanism in a general equilibrium model in which through a government allocation-taxation scheme the behavioral equilibria (Nash) are Pareto-optimal. And it balances the budget both on and off the equilibrium path.

The basic quadratic mechanism specifies each individual's tax share by

$$T_i^{\text{BQ}}(x_i | \mu_{-i}, \sigma_{-i}^2) = \frac{X}{I} \cdot b + \frac{\gamma}{2} \left[\frac{I-1}{I} (x_i - \mu_{-i})^2 - \sigma_{-i}^2 \right],$$

where $\gamma > 0$ is the punishment parameter; I is the number of people in the economy; x_i is individual i 's message, indicating her proposed addition to the total amount of the public good provided; and $X = \sum_i x_i$ is the total amount of the public good. Define $S_{-i} = \sum_{j \neq i} x_j$ as the sum of proposed increments by all other members of the group except i , $\mu_{-i} = S_{-i}/(I-1)$ as the mean of others' messages, and $\sigma_{-i}^2 = \sum_{h \neq i} (x_h - \mu_{-i})^2 / (I-2)$ as the squared standard error of the mean of others' messages. Production of the public good exhibits constant returns to scale, and b denotes the per unit cost of the public good.

Therefore, an individual's tax share is composed of three parts: the per capita cost of production, $(X \cdot b)/I$; a positive multiple, $\gamma/2$, of the difference between her own message and the mean of others' messages, $[(I-1)/I] \times (x_i - \mu_{-i})^2$; and the squared standard error of the mean of others' messages, σ_{-i}^2 . While the first two parts guarantee that Nash equilibria of the mechanism are Pareto-optimal, the last part ensures that the budget is balanced both on and off the equilibrium path. Note that the free parameter, γ , determines the magnitude of punishment when an individual deviates from the mean of others' messages. Although it does not affect any of the static theoretical properties of the mechanism, as we shall see from the experimental evidence, varying γ can induce very different dynamics.

The basic quadratic mechanism has two drawbacks: It does not satisfy the individual rationality constraint (i.e., an individual can be worse off as a result of participating in the process), and in a general environment, multiple equilibria can exist (Bergstrom, Simon, and Titus 1983). The way we deal with the first problem is to give every subject an initial endowment. For the second problem, a quasi-linear environment is used in which there exists a unique Nash equilibrium. The equilibrium selection problem in a general environment is left for future research.

B. The Paired-Difference Mechanism: Static Properties

The paired-difference mechanism implements Lindahl allocations as Nash equilibrium outcomes. Therefore, besides all the nice properties of the basic quadratic mechanism, it is also individually ratio-

nal in equilibrium; that is, no individual will be worse off as a result of participating in the mechanism.

The paired-difference mechanism specifies each individual's tax share by

$$T_i^{\text{PD}}(x_i | S_{-i}, d_{-i}) = \left(\frac{b}{I} + d_{-i} \right) (x_i + S_{-i}) \equiv \left(\frac{b}{I} + x_{i-1} - x_{i+1} \right) X,$$

where the level of individual i 's tax, T_i^{PD} , depends on her proposed addition, x_i ; the sum of proposed additions of other participants, S_{-i} ; and the difference between the amounts proposed by her two neighbors, d_{-i} . The amount $(b/I) + x_{i-1} - x_{i+1}$ is i 's Lindahl price.

Therefore, an individual's tax share is composed of two parts: the per capita cost of production, $(X \cdot b)/I$, and an amount determined by the messages of her two neighbors, $(x_{i-1} - x_{i+1})X$.

So far the two families of mechanisms have very similar static properties. The paired-difference mechanism has one more advantage over the basic quadratic mechanism in that it is individually rational in equilibrium. An interesting question is whether they will induce similar dynamic paths and properties.

C. *Dynamic Properties of the Mechanisms*

As in Hurwicz (1972), the Nash equilibrium of a game form can be viewed as a stationary point of some decentralized iterative adjustment process. In such a process, players may have incomplete information but continually revise their actions until a point is reached at which unilateral deviation no longer pays. In most situations of economic interests, when the provision of public goods is involved, individual agents do not know the characteristics of others. Therefore, we implement the mechanisms as finitely repeated games of incomplete information.

Presumably one could conduct a traditional equilibrium analysis of the repeated games by solving for the sequential equilibrium of each game and check whether players followed the sequential equilibrium. However, the informational and rationality requirements needed to reach such an equilibrium are extreme: not only are the rationality of the players and the payoffs common knowledge, but the beliefs they hold about each other's behavior also need to be commonly known, which does not seem plausible in this experimental setting. Furthermore, in our experimental design, players know only their own payoffs, so the sequential equilibrium analysis is not applicable.

When summarizing some of the lessons emerging from the large

accumulation of experimental findings on the behavior of subjects in games, Smith (1990) suggests that (1) in a one-shot game, behavior is not well predicted by equilibrium; (2) in a repeated, *complete* information setting, players tend to "cooperate" and thus repeated game effects emerge, though they are hard to predict; and (3) in a repeated setting with *incomplete* information, where players have knowledge only of their own payoffs, the best predictors of long-run behavior are the equilibria of the one-shot game with complete information.

Since the environment in point 3 is exactly the environment for this set of experiments and is also the environment that most learning processes postulate, throughout the paper we adopt the working hypothesis prevalent in the learning literature that subjects focus on the stage game strategies (see, e.g., Crawford 1995). We explore the hypothesis that stage game equilibrium is reached via a process of gradual adjustment by boundedly rational players who encounter each other in a repeated setting. The key question then is whether a particular dynamic process will converge to an equilibrium. Typical examples are Cournot best-response dynamics and the fictitious play learning process. Precise definitions of both dynamics are provided in Section V as special cases of the generalized fictitious play model.

There have been two theoretical papers studying the dynamic properties of these Nash mechanisms. Muench and Walker (1983) studied the convergence condition of the basic quadratic mechanism using Cournot best-response dynamics in a parameterized quadratic quasi-linear environment and found that the process converged when $\gamma > I$. In Section VI a generalization of this result to a general quasi-linear environment and a much wider class of learning dynamics will be provided. Kim (1986) proved that in certain quadratic non-quasi-linear environments, game forms that implement Lindahl allocations, including the paired-difference mechanism, are unstable under any decentralized adjustment process. To our best knowledge there has been no published theoretical study of the dynamic stability properties of the paired-difference mechanism in quasi-linear environments.

III. Experimental Design

The experimental design reflects both technical and theoretical considerations. The economic environment and experimental procedures are discussed in the subsections below.

A. *The Economic Environment*

We are interested in an environment in which theoretically the voluntary contribution mechanism predicts zero provision, whereas the

TABLE 1
PARAMETERS, EQUILIBRIUM VALUES, AND PAYOFFS (in Fictitious Currency)

SUBJECT	A_i (1)	B_i (2)	α_i (3)	x_i^e			π_i^e		
				BQ1 (4)	BQ100 (5)	PD (6)	BQ1 (7)	BQ100 (8)	PD (9)
1	26	1	200	-3	1	1.8	202	205	225
2	104	8	10	5	1	3.8	224	230	210
3	38	2	160	-1	1	1.8	204	200	210
4	82	6	40	3	1	-2.2	204	200	190
5	60	4	100	1	1	-.2	207	200	200
Total	310	21	510	5	5	5			

basic quadratic mechanism and paired-difference mechanism predict Pareto-efficient provision of the public good. A second consideration is the influence of the punishment parameter in the basic quadratic mechanism on the convergence of the dynamic processes.

The parameters chosen for the experiments involve five individuals, $I = 5$. In all experiments a simple constant unit cost, b , is used to produce the public good, which is set to 100. Preferences are induced on units of the abstract public good by an individually specified value function, $V_i(X)$, which indicates the amount of money an individual will receive if the group choice of the public good is X and if the individual pays nothing for it. For simplicity and for comparison of our results with previous experiments, the valuation functions are set to be quadratic:

$$V_i(X) = A_iX - B_iX^2 + \alpha_i.$$

Therefore, individual i 's payoff per period is $\pi_i = V_i(X) - T_i^m$, where $m \in \{\text{BQ}, \text{PD}\}$.

Given the size of the economy, the punishment parameter, γ , defines a family of basic quadratic mechanisms. To study the effects of the punishment parameter on the dynamics and learning processes of the basic quadratic mechanism, we set $\gamma = 1$ and 100.

In implementation of the paired-difference mechanism, one problem is the selection of one possible mechanism from an entire family. Given an economy with I individuals, there are $|I|!/(2I)$ different possible circles and hence $|I|!/(2I)$ corresponding equilibria. Of all the 12 circles that correspond to different paired-difference mechanisms in this environment, we let the computer randomly pick a circle, 1-2-4-3-5, to implement.

Table 1 lists the parameters of individual subjects' valuation functions (A_i , B_i , and α_i), their equilibrium proposals (x_i^e), and payoffs (π_i^e) under the three different mechanisms. The particular values of

the preference parameters, A_i , B_i , and α_i , are chosen for the following reasons: (1) In a voluntary contribution mechanism, the theoretical equilibrium is a zero contribution for all subjects, but the three incentive-compatible mechanisms predict the Pareto-efficient level of public goods, $X = 5$, as shown in the last row. (2) The parameter $\{B_i\}_i$ varies among subjects to induce diverse tastes for the public good. (3) Initial endowments $\{\alpha_i\}_i$ were set such that the equilibrium payoffs of all subjects, π_i^e , tabulated in columns 7–9 of table 1, are approximately the same in all three mechanisms. (4) The equilibrium contributions for BQ1 and BQ100 are all integers, and the paired-difference equilibria multiplied by five are also integers. To avoid fractions, the subjects actually chose $5x$, and all formulas were adjusted accordingly. Therefore, in all three mechanisms, subjects could choose any integer from -20 to 30 , which includes all the stage game equilibria.

B. Experimental Procedures

Seven independent computerized sessions for each mechanism were conducted in March, April, and May 1996: four at Caltech and three at the University of Amsterdam. All sessions were conducted in English by the first author. Thus the problems with the experimenter effect and the language effect (Roth et al. 1991) are circumvented. Our subjects were students from the two universities. No subject was used in more than one session. This gives us a total of 105 subjects and 21 independent sessions. Each session consisted of 100 rounds with no practice round; the sessions lasted between 1 and 2 hours, with the first half hour being used for instructions.

Table 2 summarizes session numbers, dates and places in which experiments were conducted, and the conversion rates of these experiments. The conversion rates were set such that the expected average earning per hour was approximately the same as that of other experiments in each lab. The paired-difference mechanism took longer than the basic quadratic mechanisms; therefore, the conversion rate was set lower.

Subjects who participated in an experimental session randomly drew an identification number. Then each of them was seated in front of the corresponding terminal, with a folder containing a set of instructions, payoff chart(s), and record sheets. After the instructions were read aloud, subjects were required to finish the review questions, which were designed to test their understanding of the instructions.² Afterward, the experimenter checked answers individ-

² In session 16 (04/05/96 [glcr]), one subject did not seem to understand the instructions and answered all review questions incorrectly; this was discovered by

TABLE 2
FEATURES OF EXPERIMENTAL SESSIONS

Session	Date (Code) (1)	Place (2)	Conversion Rate (3)
PD Mechanism			
1	03/31/96 (wb)	Caltech	400
2	03/31/96 (wr)	Caltech	400
3	04/01/96 (wb)	Caltech	400
4	04/01/96 (wr)	Caltech	400
5	05/08/96 (wb)	Amsterdam	400
6	05/08/96 (wr)	Amsterdam	400
7	05/08/96 (wy)	Amsterdam	400
BQ1 Mechanism			
8	04/04/96 (glib)	Caltech	800
9	04/04/96 (glir)	Caltech	800
10	04/05/96 (glib)	Caltech	800
11	04/05/96 (glir)	Caltech	800
12	05/06/96 (glib)	Amsterdam	700
13	05/06/96 (glir)	Amsterdam	700
14	05/06/96 (gliy)	Amsterdam	700
BQ100 Mechanism			
15	04/08/96 (glcb)	Caltech	800
16	04/05/96 (glcr)	Caltech	800
17	04/06/96 (glcb)	Caltech	800
18	04/06/96 (glcr)	Caltech	800
19	05/07/96 (glcb)	Amsterdam	700
20	05/07/96 (glcr)	Amsterdam	700
21	05/07/96 (glcy)	Amsterdam	700

ually and answered questions. After this, subjects signed the financial agreement.³ Then the experimenter read the computer instructions.

The mechanisms were implemented as finitely repeated games of incomplete information. At round t , a subject submitted her proposed addition, $x_i(t)$. After everyone submitted his or her proposal, the following information appeared on i 's screen:

$$\left\{ x_i(t), S_{-i}(t), \sigma_{-i}^2(t), \pi_i(t), \sum_{s=1}^t \pi_i(s) \right\} \quad (\text{in BQ1 and BQ100})$$

the experimenter only after the session. This was the only session in which a subject gave wrong answers to the review questions. We decided that it was not comparable to the rest of the sessions and therefore excluded it from the analysis.

³ Of the 105 subjects who participated in the experiments, all but three (subject 5 in session 20 and subjects 2 and 4 in session 7) signed the financial agreement.

or

$$\left\{ x_i(t), S_{-i}(t), d_{-i}(t), \pi_i(t), \sum_{s=1}^t \pi_i(s) \right\} \quad (\text{in PD}).$$

Subjects were then required to record $x_i(t)$ and $\pi_i(t)$ on a record sheet. Space was also provided for them to record $S_{-i}(t)$ and $d_{-i}(t)$, but this information was optional.⁴

During the experiment, subjects could access past history at any time by hitting the history key. When subject i used the history key at round t , the following information was available on i 's screen:

$$\{x_i(s), S_{-i}(s), \sigma_{-i}^2(s), \pi_i(s)\}_{s=1}^t \quad (\text{in BQ1 and BQ100})$$

or

$$\{x_i(s), S_{-i}(s), d_{-i}(s), \pi_i(s)\}_{s=1}^t \quad (\text{in PD}).$$

Some subjects did use the history page. Since most subjects had this information recorded, they did not use the history page often.

The process was repeated for 100 rounds, as announced at the beginning of the instructions. At the end of a session, the subjects recorded their total earnings (in fictitious currency) for all rounds and converted them to dollar (or guilder) payments. The conversion rate was announced in the instructions and was written on the board for their attention.

To summarize the information conditions, apart from the information on her screen, each subject knew her own valuation function, the basic quadratic cost function, or the paired-difference cost function. The subjects knew that other subjects in the same group might have different valuation functions, but everyone faced the same cost function. They did not know the distribution of preferences.

IV. Group-Level Results

The instructions and the complete data are available from the authors on request. Results on the aggregate performance of the mechanisms are summarized in results 1–4. Two questions are of overrid-

⁴ We required the subjects to record $x_i(t)$ and $\pi_i(t)$ to prevent the loss of data in case of a computer crash. In the pilot experiments, $x_i(t)$ and $\pi_i(t)$ were the only two columns on the record sheets, but most subjects also recorded $S_{-i}(t)$ and $d_{-i}(t)$ on the margin; we subsequently added this information to the record sheets with a mark of [optional] in all the formal sessions. We found that all subjects recorded $x_i(t)$ and $\pi_i(t)$; about 80 percent also recorded $S_{-i}(t)$ and $d_{-i}(t)$.

TABLE 3
STATISTICS FOR EXPERIMENTAL SESSIONS

Session	Efficiency (1)	Mean (2)	Standard Deviation (3)	Average Absolute Deviation (4)	Violation of Individual Rationality Constraint (5)
PD Mechanism					
1	.741	4.596	3.566	2.708	101
2	.787	4.840	3.250	2.344	62
3	.879	4.206	2.329	1.710	36
4	.746	4.188	3.460	2.624	125
5	.628	4.790	4.302	3.226	117
6	.878	5.322	2.448	1.294	126
7	.577	5.576	4.552	3.708	150
BQ1 Mechanism					
8	.968	4.906	1.266	.826	14
9	.987	4.938	.839	.330	7
10	.945	4.670	1.628	1.026	26
11	.965	5.082	1.338	.642	15
12	.961	4.710	1.411	.474	8
13	.791	4.870	3.242	2.590	79
14	.965	4.976	1.343	.684	13
BQ100 Mechanism					
15	.997	4.982	.390	.182	5
16	.995	5.010	.510	.206	5
17	.979	4.684	.978	.688	29
18	.985	4.946	.857	.490	15
19	.986	4.836	.825	.256	6
20	.973	4.910	1.152	.630	17
21	.983	5.088	.911	.492	10

ing importance. The first is related to the actual performance of the basic quadratic and paired-difference mechanisms in general. The second is related to the underlying principles of individual behavior. We address the first question in this section. A more detailed examination of individual behavioral models is reserved for the next section.

Group efficiency, the mean and standard deviation of the level of the public good, the average absolute deviation from the optimal level of the public good, and the number of violations of the individual rationality constraint for each session are tabulated in table 3.

Group efficiency is calculated by taking the ratio of the sum of the actual earnings of all subjects in a session and the Pareto-optimal

earnings of the group. As a benchmark case, if no public good is produced, the system efficiency is

$$E_0 = \frac{\text{total initial endowment in the private good}}{\text{Pareto-optimal value of the group}} \\ = \frac{525}{1,035} = 50.73\%.$$

RESULT 1. The ranking of group efficiency is highly significant: BQ100 > BQ1 > PD.

Support. Column 1 in table 3 lists the sessional group efficiency under the three mechanisms. Permutation tests (see, e.g., Siegel and Castellan 1988, pp. 95–100) show that BQ100 > BQ1 at a significance level of 0.23 percent (one-tailed), BQ100 > PD at a significance level of 0.03 percent (one-tailed), and BQ1 > PD at a significance level of 0.20 percent (one-tailed).

Result 1 shows that BQ100 generates the highest group efficiency, followed by BQ1, and then by PD.

As can be seen from column 2 of table 3 and confirmed by permutation tests, the average levels of public good provision are not significantly different across mechanisms since the mean averages out the over- and underprovision of the public good across different rounds. However, the standard deviations from the average levels of the public good are significantly different.

RESULT 2. The ranking of the standard deviation from the average level of the public good is highly significant: BQ100 < BQ1 < PD.

Support. Column 3 of table 3 shows the standard deviation of the average level of the public good. Permutation tests show that BQ100 < BQ1 at a significance level of 0.17 percent (one-tailed), BQ100 < PD at a significance level of 0.03 percent (one-tailed), and BQ1 < PD at a significance level of 0.12 percent (one-tailed).

Result 2 shows that BQ100 induces the least amount of dispersion in the level of the public good provided from period to period. It is followed by BQ1 and then by PD.

To assess how successful each mechanism is in providing close to the Pareto-optimal level of the public good, we define a measure

$$D = \sum_{t=1}^{100} \frac{|X(t) - 5|}{100}$$

as the average absolute deviation of the total level of the public good each round from the Pareto-optimal level of five.

RESULT 3. The ranking of the average absolute deviation from

the Pareto-optimal level of the public good is highly significant: $BQ100 < BQ1 < PD$.

Support. Column 4 of table 3 shows the average absolute deviation from the Pareto-optimal level of the public good. Permutation tests show that $BQ100 < BQ1$ at a significance level of 2.10 percent (one-tailed), $BQ100 < PD$ at a significance level of 0.03 percent (one-tailed), and $BQ1 < PD$ at a significance level of 0.26 percent (one-tailed).

Result 3 states that BQ100 produces the closest to Pareto-efficient level of the public good, followed by BQ1 and then by PD.

One advantage of the paired-difference mechanism over the basic quadratic mechanism is that it is individually rational in equilibrium. However, if the equilibrium is not reached, the individual rationality constraint can be violated.

RESULT 4. The ranking of the number of violations of the individual rationality constraints is highly significant: $PD > BQ1$ and $PD > BQ100$.

Support. Column 5 of table 3 shows the total number of violations of individual rationality constraints in each session. Permutation tests show that $PD > BQ1$ at a significance level of 0.12 percent (one-tailed), $PD > BQ100$ at a significance level of 0.03 percent (one-tailed), and $BQ1 > BQ100$ at a significance level of 19.30 percent (one-tailed), which is not significant at the usual 5 percent level.

Result 4 is striking in that there are significantly more violations of the individual rationality constraint in the paired-difference sessions than in the basic quadratic sessions, even though theoretically the paired-difference mechanism is supposed to be individually rational in equilibrium. This result demonstrates the importance of the dynamic properties of the mechanisms, which had been largely ignored in the literature.

The aggregate results indicate that the performance of BQ100 is far better than that of BQ1, and both are better than PD. Since the three mechanisms have very similar static properties, it is clear that individual behavior is important in understanding the dynamics that lead to the results above. In the next section, we evaluate several learning models in an attempt to understand individual behavior.

V. Learning

“Learning” can be viewed as any systematic change of behavior due to an accumulation of experience. A learning model, following the probabilistic approach of Bush and Mosteller (1955), is a mathematical system that predicts the probabilities of available choices or feasible actions at the next occurrence. There are many learning models

attempting to capture the principles of human learning behavior (Tang 1995). Since it might be misleading to claim which model is the “true” description, we evaluate two major classes of models to see which one tracks the data better under different mechanisms.

To evaluate the accordance between model predictions and the experimental data, one can measure the deviation of the model predictions from the actual choices by the quadratic deviation measure (QDM), which is a proper scoring rule.⁵ We also evaluated all models by two other scoring rules, the absolute deviation measure⁶ and the proportion of inaccuracy (POI) scores.⁷ All qualitative results hold under all three scoring rules, although only the QDM scores are reported here. Other scores are available from the authors.

Recall that subjects can choose any integer, $5x_i \in \{-20, \dots, 30\}$; namely, each has 51 stage game strategies under each mechanism. We reduce the 51 strategies to 11 choice intervals by dividing a choice number by five and rounding it up to the nearest integer in order to have multiple observations for each strategy interval (see, e.g., Roth and Erev 1995). Note that under this treatment, all equilibria in each mechanism are still treated equally, even if they are not integers, since choices in the neighborhood of radius 0.5 of an integer are also given credit.

Let $j = 1, \dots, 11$ correspond to the strategies of choosing the number $\{-4, -3, \dots, 5, 6\}$.

Let $\mathbf{c}_i(t) = (c_{i1}(t), c_{i2}(t), \dots, c_{i11}(t))$ denote the indicator vector of subject i 's contribution at round t :

$$c_{ij}(t) = \begin{cases} 1 & \text{if alternative } j \text{ is chosen in round } t \\ 0 & \text{otherwise.} \end{cases}$$

⁵ A scoring rule is “proper” if it does not give the forecaster any incentive to “ignore the verification system” or, even worse, to “play the system.” See Yates (1990) for a recent survey and Selten (1995) for the axiomatization of the quadratic scoring rule.

⁶ The absolute deviation measure is calculated as

$$\text{ADM} = \sum_{i=1}^5 \sum_{t=1}^{100} \sum_{j=1}^{11} \frac{|c_{ij}(t) - p_{ij}(t)|}{5}.$$

⁷ According to Erev and Roth (in press), the POI score “returns the value of 0 if the subject made the most likely choice under the model, the value of 1 if the subject chose a strategy that differs from the most likely prediction, and $1 - 1/b$ if the model predicts that b strategies are equally likely and the subject chose one of them. (Thus the POI score judges all the models on the basis of their ‘deterministic’ predictions, which should facilitate comparison of the deterministic models and the stochastic models.)”

Let $\mathbf{p}_i(t) = (p_{i1}(t), p_{i2}(t), \dots, p_{i11}(t))$ denote the predicted choice probability vector for subject i at round t . Then the quadratic deviation for subject i at round t is

$$\text{QDM}_i(t) = \sum_{j=1}^{11} [c_{ij}(t) - p_{ij}(t)]^2.$$

It follows that the average quadratic deviation for an entire session is

$$\text{QDM} = \sum_{i=1}^5 \sum_{t=1}^{100} \frac{\text{QDM}_i(t)}{5}.$$

And the overall average quadratic deviation measure for each mechanism is the average of the QDM scores over all seven sessions. Apparently, the smaller the QDM score a model produces, the better its prediction is.

Three different classes of models are evaluated on this data set: a static benchmark, an exponentialized relative payoff sum (RPS) model, and a class of population learning models, called a generalized fictitious play model. We focus on the last two categories of learning models for two reasons: The RPS-type models emerge as the top performers among the 18 different models evaluated in Tang (1995); however, the fictitious play model, which revived a lot of theoretical attention lately, is too influential to be ignored. In the following subsections, we present the most important parts of the learning models we have tested and discuss the implications of the results.

A. *Static Benchmark: The Equilibrium Model*

The equilibrium model uses the stage game equilibrium as the prediction:

$$p_{ij}(t) = \begin{cases} 1 & \text{if alternative } j \text{ is a stage game equilibrium strategy for } i \\ 0 & \text{otherwise.} \end{cases}$$

Note that while the QDM score of any learning model is an indicator of how well the learning model performs, only the QDM score of the equilibrium model is an indicator of how well a mechanism performs. One important measure of the performance of a mechanism is whether it induces convergence to its stage game equilibrium.

TABLE 4

QUADRATIC DEVIATION MEASURE SCORES OF LEARNING MODELS: AVERAGE OVER ALL SESSIONS

Mechanism	Initial Values	Parameters	Parameter Values	QDM
Equilibrium Model				
PD	Stage game equilibrium	183.09
BQ1	Stage game equilibrium	169.71
BQ100	Stage game equilibrium	13.03
Exponentialized RPS Model				
PD	200	(λ, q)	$(.007, .70)$	62.15
BQ1	200	(λ, q)	$(.007, .70)$	33.05
BQ100	200	(λ, q)	$(.006, .80)$	12.46
Generalized Fictitious Play Model				
PD	$(\frac{1}{11}, \dots, \frac{1}{11})$	δ	.90	134.39
BQ1	$(\frac{1}{11}, \dots, \frac{1}{11})$	δ	.72	83.02
BQ100	$(\frac{1}{11}, \dots, \frac{1}{11})$	δ	$ [.49, .52] \cup [.57, .64] \cup [.67, .78]$	13.42

Tables 4 and 5 summarize the numerical results from evaluating all three learning models; the “best” or minimum QDM scores are tabulated for each model. Table 4 presents the initial values, the estimated parameter values, and the average QDM scores over all seven sessions for PD, BQ1, and BQ100, respectively. Table 5 gives the QDM scores of each individual session using the same estimated parameter values as those listed in table 4. Recall that a QDM score is the sum of 100 rounds of quadratic deviations between model predictions and actual choices. Therefore, the smallest possible score is zero if a model gives completely correct predictions, and the largest possible score is 200 if every prediction is wrong.

One striking result is that the static equilibrium model produces extraordinarily small QDM scores under BQ100 but very large QDM scores under the other two mechanisms.

RESULT 5. Individual players under BQ100 followed their stage game equilibria at an extraordinarily high frequency, much higher than under either BQ1 or PD. Individual players under BQ1 followed their stage game equilibria at a higher frequency than under PD.

Support. The equilibrium model segment from table 4 shows that the sessional average QDM scores for PD, BQ1, and BQ100 are 183.09, 169.71, and 13.03, respectively. On the sessional level, the equilibrium model segment from table 5 shows that $QDM(BQ100)$

TABLE 5
QUADRATIC DEVIATION MEASURE SCORES OF LEARNING MODELS: INDIVIDUAL SESSIONS

Mechanism	QDM _{1s}	QDM _{2s}	QDM _{3s}	QDM _{4s}	QDM _{5s}	QDM _{6s}	QDM _{7s}
Equilibrium Model							
PD	182.80	174.70	179.20	189.20	183.60	197.60	174.80
BQ1	156.80	194.00	170.00	150.40	198.80	180.40	137.60
BQ100	6.40	4.80	26.00	10.40	8.80	20.80	14.00
Exponentialized RPS Model							
PD	79.37	67.68	53.90	61.09	76.19	25.89	70.90
BQ1	34.27	19.84	43.67	17.44	14.22	81.00	20.93
BQ100	7.87	5.80	24.53	10.44	7.64	16.97	13.95
Generalized Fictitious Play Model							
PD	135.71	121.31	120.51	158.51	142.51	99.31	162.91
BQ1	79.71	83.71	94.51	62.51	64.11	134.91	61.71
BQ100	6.91	4.51	26.51	10.91	8.51	21.71	14.91

NOTE.—Parameter values are the same as in table 4. QDM_{is} is the QDM score of session $i \times s$, where $s = 1$ for PD, $s = 2$ for BQ1, and $s = 3$ for BQ100.

$< \text{QDM}(\text{BQ1})$ or $\text{QDM}(\text{BQ100}) < \text{QDM}(\text{PD})$ is so obvious that any statistical test is superfluous. The permutation test shows that $\text{QDM}(\text{BQ1}) < \text{QDM}(\text{PD})$ at a weak significance level of 8.4 percent (one-sided).

As an extension of this result, we would like to see whether a mechanism induced convergence to its stage game equilibrium. Theoretically, convergence implies that no deviation will ever be observed once the system equilibrates. In an experimental setting with long iterations, even after the system equilibrates, subjects sometimes experiment by occasional deviation. Therefore, it is necessary to have some behavioral definition of convergence: a system converges to an equilibrium at round t if $x_i(s) = x_i^e$, for all i and for all $s \geq t$, except for a maximum of n rounds of deviation for $s > t$, where n is small. For our experiments of 100 rounds, we let $n \leq 5$; that is, there could be a total of up to five rounds of experimentation or mistakes after the system converged. Admittedly, the requirement of $n \leq 5$ is to some extent arbitrary. However, it is necessary to have some criterion in order to distinguish between sessions that converged and those that did not converge and to have a measure of the speed of convergence.

RESULT 6. Every session of BQ100 converged to its stage game equilibrium; most of the sessions converged fairly quickly. The ses-

sions of BQ1 and PD never converged to their stage game equilibrium.

Support. The seven sessions of BQ100 (sessions 15–21) converged to their stage game equilibria in the following rounds: 22, 9, 76, 44, 9, 44, and 60. In all seven sessions, every deviation after convergence was made by a single subject whereas all other subjects still chose their stage game equilibrium strategy. The other two mechanisms never converged to their stage game equilibria. Moreover, stage game equilibrium under BQ1 and PD was not even reached by all subjects simultaneously at any round in any session.

Results 5 and 6 provide further evidence in ranking the performance of the mechanisms. Although the equilibrium model is not really a dynamic learning model, it provides a baseline for the comparison of genuinely dynamic learning models. In the next subsection we shall present results that rank the performance of two dynamic learning models.

B. Dynamic Learning Models

This subsection contains an analysis of the exponentialized RPS model and the generalized fictitious play model.

The basic idea of the RPS-type model is that an individual is more likely to choose a strategy that has yielded relatively higher payoffs to her in the past. Therefore, it is also called a reinforcement learning model or stimulus response model (Fudenberg and Levine 1996). Learning models in this spirit have a long history in biology and psychology, but their systematic application in experimental economics seems to start from Roth and Erev (1995). In that paper, several variants of the basic linear form were used to construct computer simulations at the group or population level to track the ultimatum bargaining, best-shot, and market game experimental data from a comparative study in Jerusalem, Ljubljana (Slovenia), Pittsburgh, and Tokyo. They did not generalize their simulations to nonlinear functional forms.

Since each of our experimental sessions consists of 100 rounds, which is long enough for performing a more detailed analysis than the possibility available to Roth and Erev (1995), we can compare the performance of various learning models in tracking the data down to the individual level.

We use a nonlinear variant of the basic RPS model called the exponentialized RPS model. Define $M_{ij}(t)$ as the discounted payoff sum of individual i to choose strategy j :

$$M_{ij}(t) = qM_{ij}(t-1) + c_{ij}(t)\pi_{ik}(t),$$

where $q \in [0, 1]$ is the time/memory discount factor. Then the predicted probability for subject i at round $t + 1$ is

$$p_{ij}(t + 1) = \frac{e^{\lambda \cdot M_{ij}(t)}}{\sum_{k=1}^{11} e^{\lambda \cdot M_{ik}(t)}} \quad \forall i, j,$$

where $\lambda \geq 0$ helps to scale up (when $\lambda > 1$) or scale down (when $\lambda < 1$) the relative weights of the discounted payoff sums. When $\lambda = 0$, the model degenerates into a random choice model.

This model is also called the quantal response learning model (Mookherjee and Sopher 1996), which is a dynamic learning version of the quantal response equilibrium model of McKelvey and Palfrey (1995). This approach originated in the multinomial logit framework used in the econometric models of discrete choice (see, e.g., McFadden 1984).

One advantage of the exponentialized RPS model is that negative payoffs can be treated the same as positive payoffs, since the exponential function gives a positive number whether the discounted payoff sum is positive or negative.

The initial value we used for the exponentialized RPS model is $M_{ij}(0) = 200$, for all i, j , since the first-round payoffs for most of our subjects were around 200. We have also tried various other initial values, ranging from 10 to 500, which produced little difference. It seems that because of the long sequence of play, as long as the initial values are not set too large or too small, the performance of the model is hardly affected.⁸ Furthermore, these initial values result in probability predictions around the centroid, $(1/11, \dots, 1/11)$, a somewhat “natural” starting point for the first round when no history information is available. For this model we have searched the discount factor $q \in [0, 1]$ at a grid size of 0.05 and the λ parameter at a grid size of 0.001 until the minimum QDM scores are obtained.

Compared with the exponentialized RPS model, where an individual subject bases her decision on her own past payoff information only, population learning models allow an individual to base her decision on some summary statistics of the population as well. We use a generalized fictitious play model (Cheung and Friedman 1995) to analyze the data.

It is straightforward to calculate the best response for both mecha-

⁸ Roth and Erev (1995) found that the performance of their model was robust to various initial values for both the best-shot and market games, although initial conditions seemed to be sensitive for the ultimatum bargaining game.

nisms. For the basic quadratic mechanism, a player's best response to some predicted population characteristics is

$$x_i = a_i S_{-i} + b_i,$$

where

$$a_i = \frac{(\gamma/I) - 2B_i}{[\gamma(I-1)/I] + 2B_i},$$

$$b_i = \frac{A_i - (b/I)}{[\gamma(I-1)/I] + 2B_i}.$$

For the paired-difference mechanism, the best-response function is

$$x_i = m_i - n_i d_{-i} - S_{-i},$$

where

$$m_i = \frac{A_i - (b/I)}{2B_i}, \quad n_i = \frac{1}{2B_i}.$$

The dynamics of the decision process is specified according to a retrospective learning rule. For some discount factor, δ , we assume that players predict the S_{-i} and d_{-i} at round $t+1$ according to

$$S_{-i}(t+1) = \frac{S_{-i}(t) + \sum_{u=1}^{t-1} \delta^u S_{-i}(t-u)}{1 + \sum_{u=1}^{t-1} \delta^u},$$

$$d_{-i}(t+1) = \frac{d_{-i}(t) + \sum_{u=1}^{t-1} \delta^u d_{-i}(t-u)}{1 + \sum_{u=1}^{t-1} \delta^u}.$$

Note that this model is quite general. When $\delta = 0$, it yields the *Cournot best-response* model, $S_{-i}(t+1) = S_{-i}(t)$ and $d_{-i}(t+1) = d_{-i}(t)$. When $\delta = 1$, it yields the *fictitious play* model. The usual *adaptive learning model* assumes $0 < \delta < 1$, so all observations influence the expected state but more recent observations have greater weight.

For the same reason as with the exponentialized RPS model, we have used the centroid, $(1/11, \dots, 1/11)$, as the starting probability

prediction vector for the first round for the generalized fictitious play model. The discount factor, $\delta \in [0, 1]$, was searched at a grid size of 0.01.

Since BQ100 induces fast and stable convergence to its stage game equilibria, all learning models perform very well under this mechanism. The relatively volatile dynamic paths of BQ1 and PD provide a sharp separation of the performance of the exponentialized RPS model and the generalized fictitious play model.

RESULT 7. The exponentialized RPS model fits the BQ1 and PD data much better than the generalized fictitious play model.

Support. Tables 4 and 5 show that the generalized fictitious play model produces much larger (almost double) QDM scores than the exponentialized RPS model, not only at overall averages but also at independent sessional averages. The difference is so obvious that statistical tests are superfluous. Either permutation test or Wilcoxon test can give a clear-cut statistic separation at the 1 percent significance level (one-tailed).

One might argue that these two types of learning models are not entirely comparable since the generalized fictitious play model is a deterministic model that makes extreme predictions of zero or one, whereas the exponentialized RPS model makes stochastic predictions. To correct for this bias, we also evaluated both learning models under the absolute deviation measure and POI scores. All the results still hold under these two scoring rules.

VI. Incentives to Learn, Deviation Sensitivity, and Stability

Since implementation of a static mechanism usually starts somewhere off the equilibrium path, disequilibrium aspects of a mechanism are especially important in inducing convergence to the equilibrium. We define *deviation cost*, DC^ϵ , as a subject's net utility loss when she deviates ϵ from the equilibrium strategy, that is,

$$DC^\epsilon = \pi(x^e) - \pi(x^e + \epsilon),$$

where x^e is the equilibrium strategy of a player. Subscripts are suppressed for simplicity. It is straightforward to calculate the deviation costs for the two mechanisms:

$$\begin{aligned} DC^\epsilon(\text{BQ}) &= [V(S + x^e) - T^{\text{BQ}}(x^e|\mu, \sigma^2)] \\ &\quad - [V(S + x^e + \epsilon) - T^{\text{BQ}}(x^e + \epsilon|\mu, \sigma^2)] \\ &= \left(B + \frac{I-1}{2I} \gamma \right) \epsilon^2 \end{aligned}$$

$$\begin{aligned} \text{DC}^\epsilon(\text{PD}) &= [V(S + x^\epsilon) - T^{\text{PD}}(x^\epsilon | S, d)] \\ &\quad - [V(S + x^\epsilon + \epsilon) - T^{\text{PD}}(x^\epsilon + \epsilon | S, d)] \\ &= B\epsilon^2. \end{aligned}$$

The measure DC^ϵ captures the incentive a mechanism gives a subject to learn to play equilibrium strategies. When a subject is away from the equilibrium, the higher the punishment is, the higher an incentive she has to learn to play an equilibrium strategy. This incentive is captured by the possible increase in utility (or monetary pay-offs). Since

$$\text{DC}^\epsilon(\text{BQ100}) \gg \text{DC}^\epsilon(\text{BQ1}) > \text{DC}^\epsilon(\text{PD}),$$

the punishment for deviation from equilibrium strategies is much higher in BQ100 than in either BQ1 or PD. This partially explains why convergence was so fast in BQ100 and why the frequencies with which the subjects play their stage game equilibrium strategies follow the same ordering. When a mechanism is implemented and a player is not playing her equilibrium strategy, she should “know” that she is not doing her best. Under BQ100 it can really result in big losses if one is not doing one’s best, but not so much under BQ1 or PD.

A practical measure of system stability in actual implementation is how sensitive the system is to deviation. When a system reaches equilibrium, if one person deviates from equilibrium, what are the effects on the rest of the subjects? Does the noise get diminished or amplified? For simplicity, the following analysis assumes best-response dynamics. One could easily carry out the same analysis with other models, for example, generalized fictitious play.

Suppose that player j deviates ϵ from her equilibrium message at time t , $x_j(t) = x_j^\epsilon + \epsilon$. Then for everyone else at time $t + 1$, the summary statistics are changed to

$$S_{-i}(t + 1) = S_{-i}(t) + \epsilon \quad \forall i \neq j$$

and

$$d_{-i}(t + 1) = \begin{cases} d_{-i}(t) - \epsilon & \text{if } i = j + 1 \\ d_{-i}(t) + \epsilon & \text{if } i = j - 1 \\ d_{-i}(t) & \text{if } i \neq j - 1, j + 1. \end{cases}$$

Under the basic quadratic mechanism, player i ’s ($i \neq j$) best response at time $t + 1$ is

$$x_i(t+1) = x_i(t) + \left[\frac{(\gamma/I) - 2B_i}{\gamma - (\gamma/I) + 2B_i} \right] \epsilon.$$

If the planner knows the distribution of preferences, then she can pick γ such that the *deviation sensitivity* coefficient

$$DS^{BQ} \equiv \left| \frac{(\gamma/I) - 2B_i}{\gamma - (\gamma/I) + 2B_i} \right| < 1.$$

With the parameters of our experiments, this is satisfied for both $\gamma = 1$ and 100. Therefore, any noise in the system due to deviation or a mistake of some player gets diminished under the basic quadratic mechanism. Furthermore, notice that $\partial DS^{BQ}/\partial I < 0$; that is, the noise gets diminished more, the larger the population is. The reason is that with the basic quadratic mechanism, players react to the mean of everyone else's message. In a large population, noise created by deviation gets averaged out.

This is not true with the paired-difference mechanism. Under the paired-difference mechanism,

$$x_i(t+1) = \begin{cases} x_i(t) - \epsilon - \frac{\epsilon}{2B_i} & \text{if } i = j+1 \\ x_i(t) - \epsilon + \frac{\epsilon}{2B_i} & \text{if } i = j-1 \\ x_i(t) - \epsilon & \text{if } i \neq j-1, j+1. \end{cases}$$

Therefore, the deviation sensitivity coefficient for the paired-difference mechanism is

$$DS^{PD} = \begin{cases} \left| 1 + \frac{1}{2B_i} \right| & \text{if } i = j+1 \\ \left| 1 - \frac{1}{2B_i} \right| & \text{if } i = j-1 \\ 1 & \text{if } i \neq j-1, j+1. \end{cases}$$

So noise in the system from someone's deviation or mistake does not get diminished except possibly for $i = j-1$; rather, it either remains the same ($i \neq j-1, j+1$) or gets amplified ($i = j+1$), which can cause the system to unravel.

For a full stability analysis of the two families of mechanisms in a general quasi-linear environment, a sufficient condition for convergence under a wide class of learning dynamics is provided below for

the basic quadratic mechanism, and an observation is made for the paired-difference mechanism.

For a general quasi-linear utility function, $V_i = u_i(X) + \alpha_i$, where $u_i(\cdot)$ is C^2 and concave, the payoff to individual i is $\pi_i^m = u_i(X) + \alpha_i - T_i^m$, where $m \in \{\text{BQ}, \text{PD}\}$. From Milgrom and Roberts (1990), we know that supermodular games⁹ converge to their unique Nash equilibrium under a wide class of interesting learning dynamics, including Bayesian learning,¹⁰ fictitious play, adaptive learning, Cournot best-response, and many others. Therefore, supermodularity is a very robust stability criterion for public goods mechanisms.

PROPOSITION 1. The basic quadratic mechanism is a supermodular game in a quasi-linear environment if $\gamma \in [-\min_{i \in N} \{\partial^2 u_i / \partial X^2\} I, +\infty)$ and if the strategy space is bounded.

Proof. Since π_i is C^2 , it has increasing differences in (x_i, x_{-i}) if and only if $\partial^2 \pi_i^{\text{BQ}} / \partial x_i \partial x_j \geq 0$, for all $i \neq j$ (Milgrom and Roberts 1990). Since

$$\frac{\partial^2 \pi_i^{\text{BQ}}}{\partial x_i \partial x_j} = \frac{\partial^2 u_i}{\partial X^2} + \gamma \frac{1}{I},$$

if

$$\gamma \in \left[-\min_{i \in N} \left\{ \frac{\partial^2 u_i}{\partial X^2} \right\} I, +\infty \right),$$

then π_i has increasing differences or strategic complementarity between players' strategies. Note that the strategy space is one-dimensional, so π_i is automatically supermodular in x_i . These two conditions, together with bounded strategy space in R^1 , give us a supermodular game. Q.E.D.

Proposition 1 generalizes Muench and Walker's (1983) convergence result to a wider class of learning dynamics over a more general set of preferences. It also explains why BQ100, which is a supermodular game, converges so fast and remains stable under all learning models evaluated. On the other hand, BQ1 is not a supermodular game.

⁹ A supermodular game is one in which, for each player i , her strategy space is a subset of a finite Euclidean space, π_i has increasing differences in (x_i, x_{-i}) , and π_i is supermodular in x_i .

¹⁰ By Bayesian learning, we mean that each player has a prior belief about her opponents' types or possible payoff functions, which is updated according to Bayes's rule after each round of repeated play. In each period players play a Bayesian Nash equilibrium for their current probability beliefs about their opponents' types (see, e.g., Fudenberg and Levine 1996).

For the paired-difference mechanism, we have

$$\frac{\partial^2 \pi_i^{\text{PD}}}{\partial x_i \partial x_j} = \begin{cases} \frac{\partial^2 u_i}{\partial X^2} + 1 & \text{if } i = j + 1 \\ \frac{\partial^2 u_i}{\partial X^2} - 1 & \text{if } i = j - 1 \\ \frac{\partial^2 u_i}{\partial X^2} & \text{if } i \neq j - 1, j + 1. \end{cases}$$

Therefore, with general convex preferences the paired-difference mechanism is not a supermodular game.

The analysis above suggests that the success of a mechanism depends not only on its properties in equilibrium but also on its disequilibrium properties. The comparative performance of the basic quadratic mechanism and the paired-difference mechanism, as well as their disequilibrium properties, provides some lessons for mechanism design. Two aspects are identified: the incentives to learn and deviation sensitivity. The deviation cost, DC^ϵ , imposes incentives for subjects to learn to play their equilibrium strategies by punishing deviations. With proper incentives, such as that of BQ100, a mechanism can successfully induce a subject to play equilibrium strategies. The deviation sensitivity coefficient affects whether noise in a system gets diminished or amplified. A mechanism that uses population characteristics, such as the mean of others' messages, can be designed in such a way that the noise gets diminished in the system. On the other hand, a mechanism that uses individual players' characteristics, such as the difference of one's two neighbors' messages, tends to get unstable because idiosyncrasies or mistakes of a single player can cause the entire system to unravel. A supermodular mechanism has a robust stability property since a wide class of learning dynamics converge to its Nash equilibrium.

VII. Comparison with Previous Work

There has been little experimental work on the paired-difference mechanism. Robin Hanson (School of Public Health, University of California at Berkeley) ran a pilot experiment testing the paired-difference mechanism with the Smith process and found nonconvergence. Results of his pilot experiments have not been published or summarized in a working paper. This paper reports the first systematic experimental study of the paired-difference mechanism.

There have been three groups of experiments with mechanisms motivated by the basic quadratic mechanism.

First, Smith (1979) conducted two sets of experiments using a simplified version of the basic quadratic mechanism, which only balanced the budget in equilibrium. The process used was the Smith process, where all the subjects need to repeat the same choice three times in a row to finalize the production of public goods, and they were paid only when agreements were reached.

Second, Harstad and Marrese (1982) had run experiments using the complete version of the basic quadratic mechanism. The serial-tim process used by Harstad and Marrese also requires unanimity of the subjects to produce the public good, but it differs from the Smith process in that subjects need to repeat their messages only once for an iteration to end.

Neither Smith nor Harstad and Marrese studied the effects of the punishment parameter, γ , on the performance of the mechanism. More recently, Chen and Plott (1996) did the first set of experiments to assess the performance of the basic quadratic mechanisms under different punishment parameters, $\gamma = 1$ and 100. The periodic process used by Chen and Plott implemented the public goods game as a finitely repeated game, where subjects proposed a contribution each period and were paid for each decision they made. Unanimity was not required to produce the public good.

Our experimental design for the basic quadratic mechanism resembles the Chen and Plott experiments in that we also consider two treatments, $\gamma = 1$ and $\gamma = 100$; but it differs significantly in the experimental procedures, which lead to much sharper statistical comparisons of different treatments. The individual behavioral models analyzed in this paper are much richer than those in Chen and Plott. Our theoretical result on the sufficient conditions for convergence of the basic quadratic mechanism is new. These differences will be explained in detail below.

The Chen and Plott experiments consist of four sequential sessions: two sessions with the order of $\gamma = 1$ following $\gamma = 100$ and the other two with the reversed order. Since the two treatments are not independent, it is almost impossible to disentangle the learning effects from the incentive effects provided by the mechanism. In contrast, we have run seven independent sessions for each treatment, which allows us to perform an analysis that requires statistical independence. Note that one session is only one independent observation due to the intrinsic strategic interaction among subjects within each session. To compare the dynamic paths induced by the mechanisms, we used much longer iterations, 100 rounds per session with no practice rounds, whereas the Chen and Plott experiments have only 25 rounds per mechanism with five practice rounds before each trial. Practice rounds, as well as the sequential treatments that mixed up different mechanisms, interfere with the experimenters'

control for learning effects; and 25 rounds are not enough to study the learning dynamics.

On the aggregate level, because of the lack of independent treatments, Chen and Plott compared the efficiency and public goods levels of BQ1 and BQ100 by comparing the average. No statistical results were presented. In comparison with all previous experiments on the basic quadratic mechanism, this is the first time some clear and highly significant statistical results are presented.

On the individual behavioral level, Chen and Plott examined the Cournot and fictitious play learning models, both of which are special cases of our generalized fictitious play model. Our estimation of the generalized fictitious play model suggests that the best discount factor lies generally between .5 and .9, which yields neither the Cournot nor the fictitious play model. Apart from that, we also present some highly significant statistical results about the comparative performance of the exponentialized RPS model, the generalized fictitious play model, and the static equilibrium model. Our new results show that the exponentialized RPS model fits the BQ1 and PD data much better than the generalized fictitious play model and that individual players under BQ100 followed their stage game equilibria at a much higher frequency than under either BQ1 or PD.

On the theoretical level, Chen and Plott did not provide a convincing dynamic theory to explain why BQ100 performed so much better than BQ1; therefore, they did not answer the question about the range of γ that ensures stability of the mechanism. Proposition 1 in this paper provides a sufficient condition for the convergence of the basic quadratic mechanism in a class of general quasi-linear environments, thus giving the precise range of γ that induces stability under a wide class of learning dynamics. This result also generalizes theoretical work on the dynamic stability of the basic quadratic mechanism by Muench and Walker (1983).

Therefore, from the perspectives of the experimental design, aggregate and individual-level analysis and results, and theoretical findings, this study is a major substantive advance over Chen and Plott's study and other previous experimental and theoretical studies of the basic quadratic mechanism.

VIII. Concluding Remarks

The free-rider problem has been the cornerstone of the problem of public goods provision. Many mechanisms promise a solution. Two of the most famous ones are the basic quadratic mechanism and the paired-difference mechanism. Both have very similar static properties: Nash efficiency and a balanced budget with the same dimension of message space. The paired-difference mechanism also

satisfies the individual rationality constraint. However, our experiments show that they induce very different dynamics. Despite all the perfect static theoretical properties of the paired-difference mechanism, the empirical evidence from our experiments indicates that in a simple quasi-linear environment the basic quadratic mechanism with a properly chosen punishment parameter has much better dynamic properties.

Comparing the performance of the basic quadratic mechanism under a high punishment parameter ($\gamma = 100$), a low parameter ($\gamma = 1$), and the paired-difference mechanism, we conclude that the performance of BQ100 is far better than that of BQ1, which, in turn, is better than that of PD, in terms of system efficiency, close to Pareto-optimal level of public goods provision, less violation of the individual rationality constraint, and convergence to the stage game equilibrium. All rankings are statistically highly significant.

These results suggest that when we design a mechanism, standard considerations, such as incentive compatibility, individual rationality, and balanced budget, are not enough to guarantee that these desirable properties can actually be obtained in a dynamic process with real human subjects. Other disequilibrium aspects, such as *deviation costs*, which impose incentives for subjects to learn to play their equilibrium strategies, and *deviation sensitivity*, which can either amplify or diminish noise in a system, are also important to induce good dynamics and stability of a mechanism. We present a sufficient condition for the basic quadratic mechanism to converge under a wide class of learning dynamics, which provides a rigorous theoretical explanation for the good dynamic properties of BQ100.

Individual learning rules are important for us to understand the dynamic properties of incentive-compatible mechanisms. In an attempt to understand the principles of individual learning behavior, we estimate three individual behavioral models. The exponentialized relative payoff sum model outperforms the population model of generalized fictitious play on this data set.

To abstract aspects of mechanisms that induce boundedly rational individuals to play equilibrium strategies is an important but difficult task, which requires experimental studies of many mechanisms. This study begins to give us some intuition from comparing two interesting mechanisms. Further experimental study of other mechanisms is needed to confirm the intuition obtained from this study.

References

- Bagnoli, Mark, and Lipman, Barton L. "Provision of Public Goods: Fully Implementing the Core through Private Contributions." *Rev. Econ. Studies* 56 (October 1989): 583–601.

- Bergstrom, Theodore C.; Simon, Carl P.; and Titus, Charles J. "Counting Groves-Ledyard Equilibria via Degree Theory." *J. Math. Econ.* 12 (October 1983): 167-84.
- Bush, Robert P., and Mosteller, Frederick. *Stochastic Models for Learning*. New York: Wiley, 1955.
- Chen, Yan, and Plott, Charles. "The Groves-Ledyard Mechanism: An Experimental Study of Institutional Design." *J. Public Econ.* 59 (March 1996): 335-64.
- Cheung, Yin-Wong, and Friedman, Daniel. "Individual Learning in Games: Some Laboratory Results." Manuscript. Santa Cruz: Univ. California, Dept. Econ., 1995.
- Crawford, Vincent P. "Adaptive Dynamics in Coordination Games." *Econometrica* 63 (January 1995): 103-43.
- Erev, Ido, and Roth, Alvin. "Modeling How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria." *A.E.R.* (in press).
- Fudenberg, Drew, and Levine, David. "Theory of Learning in Games." Manuscript. Los Angeles: Univ. California, Dept. Econ., 1996.
- Green, Jerry, and Laffont, Jean-Jacques. "Characterization of Satisfactory Mechanisms for the Revelation of Preferences for Public Goods." *Econometrica* 45 (March 1977): 427-38.
- Groves, Theodore, and Ledyard, John O. "Optimal Allocation of Public Goods: A Solution to the 'Free Rider' Problem." *Econometrica* 45 (May 1977): 783-809.
- Harstad, Ronald M., and Marrese, Michael. "Behavioral Explanations of Efficient Public Good Allocations." *J. Public Econ.* 19 (December 1982): 367-83.
- Hurwicz, Leonid. "On Informationally Decentralized Systems." In *Decision and Organization: A Volume in Honor of Jacob Marschak*, edited by C. B. McGuire and Roy Radner. Amsterdam: North-Holland, 1972.
- . "Outcome Functions Yielding Walrasian and Lindahl Allocations at Nash Equilibrium Points." *Rev. Econ. Studies* 46 (April 1979): 217-25.
- Jackson, Matthew O., and Moulin, Herve. "Implementing a Public Project and Distributing Its Cost." *J. Econ. Theory* 57 (June 1992): 125-40.
- Kim, Taesung. "On the Nonexistence of a Stable Nash Mechanism Implementing Lindahl Allocations." Manuscript. Minneapolis: Univ. Minnesota, Dept. Econ., 1986.
- Ledyard, John O. "Public Goods: A Survey of Experimental Research." In *Handbook of Experimental Economics*, edited by John H. Kagel and Alvin E. Roth. Princeton, N.J.: Princeton Univ. Press, 1995.
- McFadden, Daniel L. "Econometric Analysis of Qualitative Response Models." In *Handbook of Econometrics*, vol. 2, edited by Zvi Griliches and Michael D. Intriligator. Amsterdam: North-Holland, 1984.
- McKelvey, Richard D., and Palfrey, Thomas R. "Quantal Response Equilibria for Normal Form Games." *Games and Econ. Behavior* 10 (July 1995): 6-38.
- Milgrom, Paul, and Roberts, John. "Rationalizability, Learning, and Equilibrium in Games with Strategic Complementarities." *Econometrica* 58 (November 1990): 1255-77.
- Mookherjee, Dilip, and Sopher, Barry. "Learning and Decision Costs in Experimental Constant Sum Games." Manuscript. New Brunswick, N.J.: Rutgers Univ., Dept. Econ., 1996.
- Muench, Thomas, and Walker, Mark. "Are Groves-Ledyard Equilibria Attainable?" *Rev. Econ. Studies* 50 (April 1983): 393-96.

- Roberts, John. "Incentives in Planning Procedures for the Provision of Public Goods." *Rev. Econ. Studies* 46 (April 1979): 283-92.
- Roth, Alvin E., and Erev, Ido. "Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term." *Games and Econ. Behavior* 8 (January 1995): 164-212.
- Roth, Alvin E.; Prasnikar, Vesna; Okuno-Fujiwara, Misahiro; and Zamir, Shmuel. "Bargaining and Market Behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo: An Experimental Study." *A.E.R.* 81 (December 1991): 1068-95.
- Selten, Reinhard. "Axiomatic Characterization of the Quadratic Scoring Rule." Manuscript. Bonn: Univ. Bonn, Dept. Econ., 1995.
- Siegel, Sidney, and Castellan, N. John, Jr. *Nonparametric Statistics for the Behavioral Sciences*. 2d ed. New York: McGraw-Hill, 1988.
- Smith, Vernon L. "Incentive Compatible Experimental Processes for the Provision of Public Goods." In *Research in Experimental Economics*, vol. 1, edited by Vernon L. Smith. Greenwich, Conn.: JAI, 1979.
- . "Experimental Economics: Behavioral Lessons for Microeconomic Theory and Applications." Nancy L. Schwartz Memorial Lecture. Evanston, Ill.: Northwestern Univ., Kellogg Grad. School Management, 1990.
- Tang, Fang-Fang. *Anticipatory Learning in Two-Person Games: An Experimental Study*. Lecture Notes in Mathematical and Economic Systems. Berlin: Springer-Verlag, 1998 (in press).
- Walker, Mark. "On the Nonexistence of a Dominant Strategy Mechanism for Making Optimal Public Decisions." *Econometrica* 48 (September 1980): 1521-40.
- . "A Simple Incentive Compatible Scheme for Attaining Lindahl Allocations." *Econometrica* 49 (January 1981): 65-71.
- Yates, Frank. *Judgement and Decision Making*. Englewood Cliffs, N.J.: Prentice-Hall, 1990.