

Modeling Cooperation Among Self-Interested Agents: A Critique

Herbert Gintis*

May 20, 2004

1 Introduction

Since Bernard de Mandeville (1924[1705]) and Adam Smith (1937[1776]), students of social life have stressed that cooperation occurs when it is in the personal interest of self-interested individuals to contribute to the collective good. We call this *mutualistic cooperation*. Mutualistic cooperation results from social rules that make prosocial behavior incentive-compatible for self-interested agents. When enforceable contracts regulate all relevant aspects of a process of exchange, for example, traders take account of the effects of their actions on others because these effects are written into the contract. Governments may simply enforce cooperation by sanctioning uncooperative acts. Even without the aid of contracts and governments, repeated interactions may sustain cooperation among actors with self-interested preferences, the anticipation of retaliation for an infraction deterring non-cooperative behavior.

The so-called *general equilibrium model*, first proposed by Léon Walras in the Nineteenth century and perfected by Kenneth Arrow, Gérard Debreu and others in the mid-Twentieth century, offers a comprehensive framework for explaining how a multitude of self-interested actors can succeed in instituting large-scale mutual cooperation based a finely gradated division of labor mediated by contractual exchange (Walras 1954 [1874], Arrow and Debreu 1954, Mas-Colell, Whinston and Green 1995). While immensely useful, the general equilibrium model does *not* in fact realize the dream of Mandeville and Smith. This is because the model assumes that contracts agreed upon by private agents can be enforced by third parties (e.g., the judiciary) costlessly to the contracting parties themselves. This is inaccurate for basic factors of production (labor and capital), and many consumer goods (Gintis 1976, Stiglitz 1987, Bowles and Gintis 2000, Gintis 2002).

*I would like to thank V. Bhaskar, Ken Binmore, Samuel Bowles, Drew Fudenberg, Ichiro Obara, Karl Schlag, and E. Somanathan for discussions that helped clarify the issues addressed in this paper.

Of course, one could argue that real world economies are just imperfect realizations of a fully methodologically individualist, internally coherent, model of cooperation based on self-interested behavior, in which contract enforcement is itself modeled in terms of incentive compatible interactions. However, no such model exists. Indeed, any such model would arguably suffer either from an infinite regress (second-level enforcers enforce contracts in the contract enforcement sector, third-level enforcers enforce contracts in the second-level enforcement sector, and so on) or have irremedial dynamic weakness (as is the case with dynamic programming models, some of which are discussed below).

To avoid these weaknesses, we posit a level of cooperation in which the enforcement of the norms of social interaction are carried out completely through the decentralized, voluntary, initiatives of individuals immersed in a social groups modeled as noncooperative games.¹ An acceptable model of sustained cooperation in such a framework, based on purely self-interested actors, would then complete the general equilibrium model's quest for a complete explanation of human cooperation based on self-interested agents. As we shall see, however, no such model currently exists, and for reasons detailed below, none is likely ever to be produced. A more plausible assumption is that cooperation in groups lacking enforceable contracts is ensured by the existence of a fraction of *strong reciprocators*, who cooperate, and punish defectors, without regard for the material losses they sustain by so doing. A modicum of strong reciprocity, we shall see, easily solves the puzzle of decentralized cooperation.²

2 Minimum Requirements for a Model of Cooperation

A *public goods game* is an n -agent game in which, by “cooperating,” each agent A adds more to the payoff of the other members than A 's cost of cooperating, but A 's share of the total gains he creates is less than his cost of cooperating. By “defecting,” the agent incurs no personal cost and produces no benefit for the group. For the remainder of this paper, we consider an extremely simplified example of a public goods game.

Consider a group of size n , where each member can work or shirk in each time period $t = 1, 2, \dots$. The cost of working is $c > 0$ and the benefit is $b > c$, shared

¹In effect, while the neoclassical general equilibrium model is often faulted by its critics for being *methodologically individualistic* (Granovetter 1985, Hodgson 1998), I stress here exactly the opposite weakness—its fails to model contract enforcement at the level of the individual.

²A number of papers have shown that strong reciprocity is a common human behavioral form (Fehr and Gächter 2000, Henrich, Boyd, Bowles, Camerer, Fehr and Gintis 2004, Gintis, Bowles, Boyd and Fehr 2004), and can persist in equilibrium despite the costs the behavior imposes on its bearers (Gintis 2000b, Bowles, Choi and Hopfensitz 2003, Bowles and Gintis 2004).

equally among the other group members. A self-interested member will shirk in a one-shot encounter because it costs c to cooperate and the benefits accrue only to others. However, suppose the encounter is repeated in each period, and all agents have discount factor δ , with $0 < \delta < 1$. In the interest of realism, we also assume that with probability $\epsilon > 0$, any agent who attempts to cooperate will fail to produce the benefit b , and will appear to the other members of the group to be shirking. If k other members work, a member receives $k \times b(1 - \epsilon)/(n - 1) = bv(1 - \epsilon)$ from their effort, where $v = k/(n - 1)$. The present value of working is then $v_c = bv(1 - \epsilon) - c + \delta v_c$, which gives

$$v_c = \frac{bv(1 - \epsilon) - c}{1 - \delta}. \quad (1)$$

An agent who does not work, by contrast earns $bv(1 - \epsilon)/(1 - \delta) > v_c$, so unless an additional assumption is made, each agent will defect, and total payoff will be zero. Repeated game models, discussed below, deal with this problem by noting that under appropriate conditions, the promise of future cooperation may induce self-interested agents to cooperate in the current period.

In this paper, I evaluate such models on the basis of the following criteria, which standard scientific practice would suggest are requirements for any successful explanatory theory.

- a. **Methodological Individualism.** Cooperation must be sustained without third-party contract enforcement or other unexplained social institutions. This requirement reflects not only the universal condition of humanity prior to some 10,000 years ago, but also the condition of most work groups in a modern economy that must solve its incentive and coordination problems without recourse to courts or police, except perhaps under extreme circumstances. In particular, this requirement implies that any payoff redistributions dictated by the rules for cooperation, as well as any punishments meted out, must be effected by agents who maximize utility subject only to the constraints and incentives imposed by the behavior of other group members. In particular, if agents are to be punished (by fine, social exclusion, or physical harm), the punishers must have an incentive to carry out the punishment.
- b. **Dynamic Stability.** Random fluctuations in costs and payoffs, as well as errors of commission (e.g., attempting, but failing to cooperate), signal transmission (e.g., appearing to defect, while in fact having cooperated, and *vice-versa*), and mutation (the periodic introduction of novel strategies) do not disrupt cooperation or entail excessive efficiency losses. The model must also be *evolutionarily stable* with respect to standard monotonic dynamics (i.e., more successful strategies tend to grow at the expense of less successful strategies), under a variety of

plausible conditions, in the sense of being impervious to invasions by mutant agents or by groups that deploy competing forms of social organization. Finally, the proposed game structure must be capable of emerging and expanding its presence in a hostile or indifferent economic environment.

- c. **Empirically Relevance.** The organizational forms and incentive mechanisms deployed must reflect the types of strategic interaction and incentives widely observed in human groups. In particular, the model should work well with group sizes on the order of ten to twenty, and the incentive to punish defectors should reflect those deployed in real-world public goods game settings.
- d. **Plausible Informational Requirements.** Signals concerning the behavior of group members should be *imperfect* in the sense that the signal indicating the cooperation of particular agent may be incorrectly received with moderate probability (say, 5% or 10% for any agent in any period), and *private* in the sense that the signals indicating the cooperation of particular agent are uncorrelated across other agents in the group, conditional on the actual performance of the signaler. This requirement follows from the fact that in observed public goods game setting, group members observe only a subset of other group members in any one period. In this paper I discuss public information models, in which all group members receive the *same* imperfect signals, but only because there may be some (as yet unknown) means whereby groups aggregate private signals into relatively accurate public signals, and agents cannot gain by conditioning their behavior on their private signals.
- e. **Plausible Discount Factors.** It is reasonable to suppose that within a group faced by a public goods game, there will be a distribution of discount factors among members, some of whom will have quite long time horizons, while others are concerned predominantly with more immediate payoffs. For various reasons, the probability of future interactions will differ across group members, and there is individual variation in discount rates across individuals, and within the same individual, across time. An acceptable model must therefore function effectively in the face of plausible statistical distribution of discount factors.

As I shall show, all models of cooperation based on self-interested agents violate one or more of these conditions, and hence fail to solve the problem of cooperation among unrelated agents.

3 Modeling Cooperation among Self-interested Agents: Trigger Strategies

If agents are self-interested, according to the celebrated *folk theorem* of repeated game theory, cooperation can be sustained if all agents use the following *trigger strategy*: cooperate as long as all other team members cooperate. Whenever a member defects (or appears to defect), defect for a sufficient number of periods, say k , so as to render the defector worse off than if he had cooperated. Cooperation is a best response for all players provided they are sufficiently patient, and if a defection is observed, defecting for the required number of rounds is also a best response.

Suppose cooperate/defect decision of each member is an imperfect public signal, with probability $\epsilon > 0$ that an intended cooperation will appear to be a defection. The value of cooperating when all other members cooperate is now given by the recursion

$$v_c = b(1 - \epsilon) - c + \delta(1 - \epsilon)^n v_c + (1 - (1 - \epsilon)^n) \delta^k v_c,$$

which gives

$$v_c = \frac{b(1 - \epsilon) - c}{1 - \delta^k - \delta(1 - \delta^{k-1})(1 - \epsilon)^n}. \quad (2)$$

The present value of defecting is $v_d = b(1 - \epsilon) + v_c \delta^k$. By the one-shot deviation principle,³ cooperation is Nash subgame perfect if and only if $v_c \geq v_d$, which simplifies to

$$\frac{c}{b} \leq \delta(1 - \epsilon)^{n+1} \left(\frac{1 - \delta^{k-1}}{1 - \delta^k} \right). \quad (3)$$

It is easy to check that the right hand side is increasing in k , but if

$$c > b(1 - \epsilon)^{n+1}, \quad (4)$$

the cooperative equilibrium cannot be sustained, no matter how patient the group members (i.e., no matter how close δ is to unity). Indeed, (3) shows that *an error rate of ϵ in a group of size n reduces the effective discount factor by a factor of $(1 - \epsilon)^n$* . Thus, no matter how small the probability ϵ , if the group is sufficiently large, cooperation cannot be sustained.

It is also easy to check that, assuming k is chosen so that (3) is an equality, the total discounted payoff to members is

$$v_c = \frac{b(1 - \epsilon)}{1 - \delta^k}, \quad (5)$$

³For a justification of this *one-shot deviation principle* as sufficient to prove subgame perfection in a repeated game, see Fudenberg and Tirole (1991) or Gintis (2000a).

which is close to the one-shot payoff to cooperation for even moderate values of k and δ . By contrast, if cooperation could be sustained costlessly, the total discounted payoff would be larger by a factor of about $1/(1 - \delta)$.

For example, suppose $n = 15$, $b = 1.5$, $c = 1.0$, and $\delta = 0.95$. Then the highest error rate that can be sustained is $\epsilon \approx 2.2\%$. At $\epsilon = 2.0\%$, we must set $k = 20$ for a cooperative equilibrium, the punishment stage occurs with 20% probability after each production state, and the present value of the game is 2.30, whereas if cooperation could be costlessly enforced, the present value of the game would be 9.40.

The inefficiency in this model stems from the fact that, given the self-interested nature of the actors, the only way to punish a defector is to cease cooperation completely for a sufficiently long period of time. When the group is large, this is clearly an inefficient form of punishment—guilty and innocent suffer equally and with a plausible error rate, most of the groups' time is spent idling.

4 Modeling Cooperation among Self-interested Agents: Directed Punishment

Suppose we reinstitute the direct punishment of defectors, but players agree that any agent who is detected not punishing a defector is himself subject to punishment by the other players. Without such an incentive, self-interested agents would of course refrain from punishing. Suppose with probability ϵ a agent either who intends to punish fails to do so, or is perceived publicly by the other members to have failed. We choose ϵ to be the same as the error rate of cooperation. If all agents cooperate and punish, the rate of defection observed will be ϵn , so the mean number of punishment events per period per agent will be ϵn , and hence the mean number of punishment events per period will be ϵn^2 . The probability that no agent will be punished for nonpunishing is thus $(1 - \epsilon)^{\epsilon n^2}$. The recursion equation for complete cooperation is thus

$$v_c = b(1 - \epsilon) - c - \epsilon(p + c_p) + \delta(1 - \epsilon)^{n^2\epsilon} v_c + \left(1 - (1 - \epsilon)^{n^2\epsilon}\right) \delta^k v_c,$$

where k is the number of periods of universal defection needed to ensure that nonpunishing is unprofitable. Solving, we get

$$v_c = \frac{b(1 - \epsilon) - c - (c_p + p)\epsilon}{1 - \delta^k - \delta(1 - \delta^{k-1})(1 - \epsilon)^{n^2\epsilon}}. \quad (6)$$

It is easy to show that v_c is decreasing in ϵ and n , and increasing in δ . The gain from not punishing for one period and then returning to cooperation is $b(1 - \epsilon) -$

$c - \epsilon p + \delta^k v_c$. The gain from punishing over nonpunishing is thus positive when

$$b(1 - \epsilon) - c > p\epsilon + \frac{1 - \delta^k}{1 - \delta^{k-1}} \frac{\epsilon c_p}{\delta(1 - \epsilon)^{n^2\epsilon}}. \quad (7)$$

A necessary condition for a cooperative equilibrium is then

$$b - c > (b + p)\epsilon + \frac{\epsilon c_p}{\delta(1 - \epsilon)^{n^2\epsilon}}. \quad (8)$$

This will clearly be violated for sufficiently large ϵ , for any given $\delta < 1$.

For example, suppose the parameters of the model are as in the previous numerical example except $\epsilon = 8\%$ and, in addition, $c_p = c/3$ and $p = 3c$ (i.e., it is not very costly to punish nonpunishers, but it is costly to be punished for nonpunishing). Then it is easy to calculate $k = 9$. The probability of completing a cooperative period without the need for punishing nonpunishers is $(1 - \epsilon)^{n^2} \approx 22\%$. The inefficiency of the “cooperative” equilibrium is underlined by the fact that $v_c \approx 0.3797$ is less than the one-shot with costless enforcement, $b(1 - \epsilon) - c = 0.38$.

5 Trigger Strategies: Empirical Evidence

Given the prominent position of trigger strategies in repeated game theory, it might be thought that serious attempts would have been made to ascertain the prevalence of such behaviors in real human groups. However, I do not know of a single study of this type, except in the area of tacit collusion among firms in the same industry (Porter 1985). I cannot even think of an anecdotal example of a large team of cooperators that deters shirking by ceasing all cooperation for a limited period of time when a defection is detected. Of course, problems with defection frequently lead to the dissolution of groups involved in a public goods situation. Dissolution, however, marks the failure of the group to cooperate, and is not a strategic tool in providing incentives for cooperation. While there is no doubt but that the threat of dissolution is an element in the array of incentives that induce agents to cooperate, in general groups deploy less drastic incentives as well, such as the shunning, ostracism, and directed punishment of offenders.

6 Second Order Punishing: Empirical Evidence

It may be thought that punishing non-punishers (so-called *second order punishing*) is a viable alternative to strong reciprocity. There are several problems with this alternative. First, it is difficult to find examples of second order punishing in real human groups. Second, No one has proposed a model of second order punishing

with private information, and when errors are admitted into the model, second order punishing with imperfect public information is highly inefficient. Finally, even with public information, models of second order punishment generally have implausible informational requirements and/or are not dynamically robust. I will take these issues in turn.

Second order punishment takes the form of group members punishing those who shirk in the duty to punish shirkers. It is difficult to find examples of second order punishing in the literature on group dynamics. The closest we have come is the widespread behavior of groups in disciplining members who befriend or otherwise associate with “outsiders.” However, such “outsiders” are not group members who are subject to group discipline, and the punishing involved is more accurately described as the first-order punishing of violators of the group norm prohibiting associations with “outsiders.” There is, to our knowledge, no case of punishing group members who fail to punish violators of this norm.

Wiessner (2003) used data from 308 conversations among the Ju/’hoansi (!Kung) Bushmen of North West Botswana collected in 1974 and 1996-8 to examine norm enforcement through reward and punishment. Wiessner investigated the dynamics of punishment among the Ju/’hoansi, asking the following questions:

- a. What are the roles of reward and punishment respectively in norm enforcement?
- b. Which behaviors elicit punishment by individuals and by groups?
- c. Who punishes whom? d. What are the different forms of punishment applied and what are their outcomes?

Her second objective was to determine if the Ju/’hoansi data provide evidence for strong reciprocity or whether other hypotheses that center on individual self-interest provide sufficient explanation. She finds that the hypothesis of second-order punishment was refuted; “not a single case of second order punishment occurred during conversations noted.” Wiessner continues by saying, “I have also not observed any second order punishment in all of my years working with the Ju/’hoansi. This is not surprising for autonomy of action is a privilege of an egalitarian society: no single person can tell another what to do. It is unlikely that second order punishment would be found in egalitarian societies...” Wiessner goes on to say

Does punishing give high social regard and make a person a desirable group member? While ability to mediate is highly valued among the Ju/’hoansi, as it is in most societies, willingness to punish is not. As mentioned earlier, those who punish frequently or harshly are considered to be “tchi n!ai, angry, sharp or biting things.” They are frequently told to be quiet and stop causing trouble, and they are not considered

desirable group members. This point is further underscored by the fact that the most influential member of a group, the recognized “owner of the land” refrained from punishment in 44 out of 87 cases (46%) of group punishment in which he or she was present.

It is frequently suggested that the case of “insiders” punishing their members for fraternizing with “outsiders” is a case of second-order punishment. But, in an insider/outsider situation, the outsiders are not being punished—they are simply members of other groups. If the norm of a group is to refrain from fraternizing with members of other groups, defectors from this norm may be punished. But this is clearly a case of first-, not second-order punishment. Second-order punishment in this case would be to punish those who refuse to punish those who fraternize with outsiders.

7 Cooperation of Self-interested Agents with Public Information: The Classical Game-Theoretic Model

Shubik (1959) appears to be the first author to suggest that noncooperative game theory be applied to the problem of using game repetition and the threat of retaliation against defectors to ensure a high level of cooperation. The first formal statement of this principle was Friedman (1971). A complete theory for the case of perfect public information was given in the seminal paper Fudenberg and Maskin (1986). A quite different approach was needed to extend this to imperfect public information. The key step here was provided by the dynamic programming method of Abreu, Pearce and Stacchetti (1990), culminating in the definitive paper by Fudenberg, Levine and Maskin (1994) (hereafter FLM), who show that close to full cooperation can be attained for sufficiently patient agents, provided only that the information structure is sufficiently rich to detect accurately which agents appear to have defected (I say “appear” because the signal may be inaccurate).

In this section I will outline the FLM analysis as applied to the model of cooperation laid out above and show why it does not solve the problem of cooperation. It is convenient to represent the decision of an agent to cooperate by 1 and defect by 0. Let A be the set of strings of n zeros and ones, representing the possible strategy profiles of the n players, the k th entry representing the choice of the k th player. Let $\tau(a)$ be the number of ones in $a \in A$, and write a_i for the i th entry in $a \in A$. For any $a \in A$, the random variable $y \in A$ represents the imperfect public information concerning $a \in A$. We assume defections are signaled correctly, but intended cooperation fails and appears as defection with probability $\epsilon > 0$. Let $\pi(y|a)$ be the probability that signal $y \in A$ is received by players when the actual strategy profile

is $a \in A$. Clearly, if $y_i > a_i$ for some i , then $\pi(y|a) = 0$. Otherwise

$$\pi(y|a) = \epsilon^{\tau(a)-\tau(y)}(1-\epsilon)^{\tau(y)}.$$

The payoff to player i who chooses a_i and receives signal y is given by $r_i(a_i, y|a) = b\tau(y)(1-\epsilon) - a_i c$. The expected payoff to player i is just

$$g_i(a) = \sum_{y \in A} \pi(y|a) r_i(a_i, y) = b\tau(a)(1-\epsilon) - a_i c.$$

We will show that payoffs close to the full cooperation payoff $b^* = b(1-\epsilon) - c$ can be supported by a sequential Nash equilibrium. We write the full cooperation vector as $\mathbf{v}^* = (b^*, \dots, b^*) \in \mathbf{R}^n$. The n vectors \mathbf{v}_i representing the payoffs when all agents cooperate except one, who defects, are given by

$$\begin{aligned} \mathbf{v}_0 &= (0, \dots, 0) \\ \mathbf{v}_1 &= \left(b, \frac{(n-2)b}{n-1}, \dots, \frac{(n-2)b}{n-1} \right) \\ \mathbf{v}_2 &= \left(\frac{(n-2)b}{n-1}, b, \frac{(n-2)b}{n-1}, \dots \right) \\ &\dots \\ \mathbf{v}_n &= \left(\frac{(n-2)b}{n-1}, \dots, \frac{(n-2)b}{n-1}, b \right). \end{aligned}$$

We shall show that the convex hull V^* of $\{\mathbf{v}^*, \mathbf{v}_0, \mathbf{v}_1, \dots, \mathbf{v}_n\}$ include an open set of which \mathbf{v}^* is a boundary point. Let $\mathbf{v}_\gamma = (\gamma_1, \dots, \gamma_n)$ be a vector of perturbations, and define $\gamma_0 = \sum_{i=1}^n \gamma_i$. Let

$$\alpha_0 = 1 - \frac{\gamma_0}{b-c}$$

and

$$\alpha_i = \frac{b(\gamma_0 - (n-1)\gamma_i) + 2c\gamma_i}{(b-c)(b + (n-1)c)}$$

for $i = 1, \dots, n$. It is easy to check that $\sum_{i=0}^n \alpha_i = 1$ and the set of vectors \mathbf{v}_γ for which $\alpha_0, \alpha_1, \dots, \alpha_n > 0$ is an open set in \mathbf{R}^n . Moreover, a straightforward calculation shows that

$$a_0 \mathbf{v}^* + \sum_{i=1}^n \alpha_i \mathbf{v}_i = \mathbf{v}^* - \mathbf{v}_\gamma,$$

which proves the assertion.

For any choice of strategies a_- for the other players, let $a^+ = (1, a_-)$ and $a^- = (0, a_-)$. Let $w_i(a^+) = v_i$ and $w_i(a^-) = w_d$, to be determined. We interpret

$w_i(a^+)$ (resp. $w_i(a^-)$) as the expected per-period payoff (EPPP) to i over all future periods if he cooperates (resp. defects), given the choice a_{-i} by the other players. Let V_i^* be the interior of the projection of V^* on the i th axis, which is thus the interval $(0, b^*) \subset \mathbf{R}_+$. For any $v_i \in V_i^*$

$$v_i = (1 - \delta)b^* + \delta(1 - \epsilon)v_i + \delta\epsilon w_d \quad (9)$$

is the EPPP to cooperating, and we must have

$$v_i \geq (1 - \delta)(b^* + c) + \delta w_d \quad (10)$$

to ensure that cooperating is a best response. It is easy to check that (9) implies

$$w_d = b^* - \gamma(1 - \delta(1 - \epsilon))/\delta\epsilon.$$

For δ sufficiently close to unity, we have $w_d \in V$. Also, (10) is equivalent to

$$\gamma > \frac{\epsilon(1 - \delta)c}{\delta - \epsilon(1 - \delta)},$$

which holds for δ sufficiently close to unity.

If player i signals defection, then we can repeat the above construction, replacing v_i by w_d , constructing a new w'_d . We continue this process recursively, deriving a set of continuation payoffs for each history of signals y_1, y_2, \dots in periods $t = 1, 2, \dots$. FLM prove that this process indeed represents a sequential equilibrium.

This model is highly attractive in that it shows that for any number n of sufficiently patient agents (δ sufficiently close to unity), and any error rate $\epsilon > 0$, the repeated game can support close to a Pareto optimum, using trigger strategies alone (i.e., by reacting to perceived defections by some pattern of defection of other members, plus cooperation of the guilty parties, of sufficient duration to render purposeful defection unprofitable). This is in sharp contrast to the highly error-sensitive existence condition (4) and payoff schedule (5).

Section 2 presented several attractive properties of a successful model of cooperative behavior. I propose to evaluate the classical game-theoretic model with public information according to these criteria:

- a. **Methodological Individualism.** The FLM model fully satisfies the conditions for methodological individualism.
- b. **Dynamic Stability.** The fact that a pattern of activity is a sequential Nash equilibrium says nothing whatever about its dynamic properties. Sequential Nash equilibria are impervious to single agents playing alternative strategies, whether on or off the game path. Dynamic stability, by contrast, requires that

the model be capable of recovering from small, simultaneous, perturbations in the strategies of *all agents*. This is certainly not the case for the FLM model, simply because there is an open set of sequential equilibria (corresponding to the various specifications of \mathbf{v}_γ), implying that in every neighborhood of such equilibrium, there are uncountable other such equilibria.

- c. **Empirically Relevance.** The FLM model uses trigger strategies, which, as we have seen, have not been validated empirically. It would be easy to reformulate the model to use directed punishment, in which case second order punishment of non-punishers is used to secure compliance. However, second order punishment has no more empirical validation than does the use of trigger strategies.
- d. **Weak Informational Requirements.** The FLM model has extremely strong, indeed quite implausible, informational requirements. Consider, for instance, the well-known sufficient conditions for Nash equilibrium of Aumann and Brandenburger (1995), which can be stated as follows:⁴ *Let \mathcal{G} be a game with n players and let σ be an n -dimensional strategy profile. Suppose the players have a common prior, which assigns positive probability to it being mutually known that \mathcal{G} is the game being played, mutually known that all players are rational, and commonly known that σ is the strategy profile being played. Then σ is a Nash equilibrium.*⁵ The implausibility in the case of the FLM model is that the equilibrium deviation \mathbf{v}_γ from the \mathbf{v}^* point, as well as all the continuation payoffs w_d , one for each current strategy profile and each pattern of observed defection from this profile, must be common knowledge. Since the n players have conflicting interests concerning \mathbf{v}_γ , and there is an open set of sequential equilibria, not even mutual knowledge, much less common knowledge, it plausible.
- e. **Plausible Discount Factors.** While it is of great interest to know that a certain Nash equilibrium exists for sufficiently patient agents, we cannot expect a real social group to have extremely patient members. Indeed, there is likely to be a statistical distribution of discount rates among members, reflecting such considerations as age, health, reproductive status, as well as idiosyncratic factors (Hansson and Stuart 1990, Rogers 1994, Becker and Mulligan 1997). The lower

⁴While these conditions are not necessary, they have been shown to be strict (i.e., with a violation of any one, a non-Nash equilibrium counterexample can be found). There can be no substantive necessary conditions for Nash equilibrium, since a random strategy profile can, by accident, represent a Nash equilibrium.

⁵We say $\sigma = \{\sigma_1, \dots, \sigma_n\}$ is *mutually known* if each agent knows σ , and σ is *commonly known* if each agent knows σ , each knows that the others know σ , and so on. A player is *rational* if he maximizes his payoff subject to beliefs.

tail of such a distribution determines the degree of patience of the group, and so long as the distribution of discount factors has positive variance, the practice of taking the limit as the discount factor goes to unity will not be acceptable.

8 Cooperation in Repeated Games with Private Information

In the models discussed to this point, each member of the groups receives the same, perhaps imperfect, signal of the behavior of each other member. In most empirically relevant cases, however, different group members will receive different, and perhaps conflicting, signals concerning other members. For instance, A might see B sleeping under a tree when he should be hunting, but no other group member may be in the vicinity to witness the scene. To illustrate the problems that arise with private signals, let us see what happens to a very robust public information self-interested agent model when the information becomes private even to a relatively small degree. The reader will note that our model of cooperation with strongly reciprocal preferences (Section 2) uses a private information model and exhibit a high level of cooperation.

One of the most important developments in economic theory in the past decade is a rigorous analysis of cooperation with private signaling. It may be thought that simply pooling information would transform private into public information. This is the case, however, only if we assume that agents report their private information truthfully. However, without the proper incentives for truth-telling, we cannot plausibly make this assumption. The situation is especially serious because the private information is also imperfect. For, in this case, observing a defection, one does not know whether others observed it, and hence if C observes the “failure” of B to punish the defection of C could occur when A did not defect, or when A did defect but B did not witness the defection, or when A defected, B witness the defection, but C mistakenly failed to witness B’s punishment of A.

The obvious next step is consider more general ways to use private information efficiently. This is in fact the tack taken in recent years by several economists (see Kandori, 2002 for an overview). Important contributions to this research agenda include Sekiguchi (1997), who was the first to propose a Nash equilibrium that approximately sustains the cooperative payoff in the two-person prisoner’s dilemma, assuming that private monitoring is nearly perfect. Following this, contributions by Piccione (2002), Ely and Välimäki (2002), and Bhaskar and Obara (2002), and Matsushima (2000) considerably deepened the approach.

The technical problems involved in developing an equilibrium with a high level of cooperation assuming private information and self-interested agents are extreme when the private signals are noisy. If punishing an observed breach of cooperative

norms is costly, and if team members generally do not know which members observed which breeches, costly first-order punishment will not occur because those who see the defection know that they will not be punished for failing to punish. Therefore, first-order punishment must fail to be costly. There are various ways of achieving this result involving the use of mixed strategy sequential equilibria, so these models are vulnerable to the critique that the mechanisms involved are not seen empirically and they have very poor stability properties.

I shall sketch two quite distinct models in this tradition. These are the most powerful of their respective types, yet each has irremedial weakness as an explanation of cooperation in public goods games. The first, analyzed by Bhaskar and Obara (2002), is a private signal version of Section 3. We revise this model by assuming that when an agent moves in each period, he sends signals to each of the other $n - 1$ agents. These signals are independently distributed, given the move, but each is in error with probability $\epsilon > 0$. We will consider only symmetric equilibria, in which all agents employ the same strategy, and we assume that after a defection, a player defects forever (this is called the *grim trigger* strategy).

The first complication of the private signal assumption is that no sequential equilibrium can support full cooperation in any period. To see this, consider the first period. If each player uses the full cooperation strategy, then if a player receives a defection signal from another player, with probability one this represents a bad signal rather than an intentional defection. Thus, with very high probability, no other member received a defection signal. Therefore no player will react to a defect signal by defecting, and hence the always defect strategy will have a higher payoff than the always cooperate strategy. To deal with this problem, we must have all players defect with positive probability in the first period. A similar analysis applies to all future periods.

Now, in any Nash equilibrium, the payoff to any two pure strategies that are used with positive probability by a player must have equal payoffs against the equilibrium strategies of the other players. Therefore, the probability of defecting must be chosen so *each player is indifferent between cooperating and defecting on each round*. The analysis of the implications of this fact by Bhaskar and Obara is subtle and ingenious. They prove that under plausible conditions, for any n there is an error rate ϵ^* sufficiently small that, when errors are of frequency ϵ^* or less, there is a sequential equilibrium in which the discounted payoff to agents is approximately that of the fully cooperative solution.

There are two problems with this solution. First, empirical relevance suggests that the error rate will be given by social and technical conditions, and will be a function of group size: $\epsilon = f(n)$. Thus, it is not reasonable to choose n first, then freely choose ϵ . Indeed, we might expect ϵ to increase with group size. This will occur, for instance, if each member has a fixed frequency of interacting with

other members, so as group size increases, the probability of detecting a defection of any particular member decreases. As we have seen in Section 3, when we hold ϵ constant, there is a finite limit to the group size that can be supported, and using plausible parameter values, this limit is likely to be quite small. Moreover, as we have seen, even when a Nash equilibrium exists, it is likely to be quite inefficient.

The second problem is that, as we have stressed, there is no reason to believe that a sequential Nash equilibrium will have any particularly valuable dynamic stability properties. For empirical relevance, one must explicitly show that a solution of the form suggested by Bhaskar and Obara has acceptable dynamic properties. To illustrate this problem, I have constructed an agent-based simulation of the Bhaskar-Obara model in the Pascal programming language, as implemented by Borland Delphi 6.0. The stage game is as above, with $b = 3$ and $c = 1$. Agents are randomly assigned to groups of size n in each of 100,000 periods. In each period, each group plays the stage game repeatedly, the game terminating with probability 0.05 at the end of each round, thus implementing a discount factor of 0.95. The simulation begins by creating 210 agents, each endowed at time of creation with two parameters. The first, *DefectRound*, indicates at which round the agent will voluntarily defect. If this is very large, the agent never defects. Since we wish to assess the stability of equilibrium rather than whether it is globally stable, the program initially assigns 80% of agents with *DefectRound* = 100, which effectively means they never defect. The other 20% of agents are randomly assigned *DefectRound* values between 1 and 10. The second parameter is *Tolerance*, which indicates how many defections an agent who voluntarily cooperates must see before beginning to defect. All agents are assigned *Tolerance* = 0, so they defect at the first defection signal they receive (this is the equilibrium value for the Bhaskar-Obara model).

In each round, for each group, each member sends a signal indicating whether he cooperated or defected, with error rate ϵ , to every other group members. On the basis of this signal, all agents then update their willingness to cooperate in the next round. As soon as the round hits or exceeds an agent's *DefectRound*, or he accumulates more than *Tolerance* defect signals, the agent defects from that point on with this particular group.

At the end of every 100 periods, the simulation implements a reproduction phase, using the relative fitness of the agents, as measured by their accumulated score over the 100 periods, and replacing 5% of poorly performing agents by copies of better performing agents. We implement this by a simple imitation process that has the same dynamic properties as the replicator dynamic (Taylor and Jonker 1978). For each replacement, we randomly choose two agents, and the agent with the higher score is copied into the agent with the lower score.

At the completion of each reproduction phase, the simulation implements a mutation phase, in which each agent's parameters are increased or decreased by

one unit (except if so doing would lead to negative values) with probability 0.001.

As might be expected, when we set $n = 2$, the dynamic process exhibits a high level of efficiency (about 90% of full cooperation), as well as a high level of tolerance (agents defect after about seven defect signals, on average) even with the quite high error rate of $\epsilon = 10\%$, after 100,000 rounds.

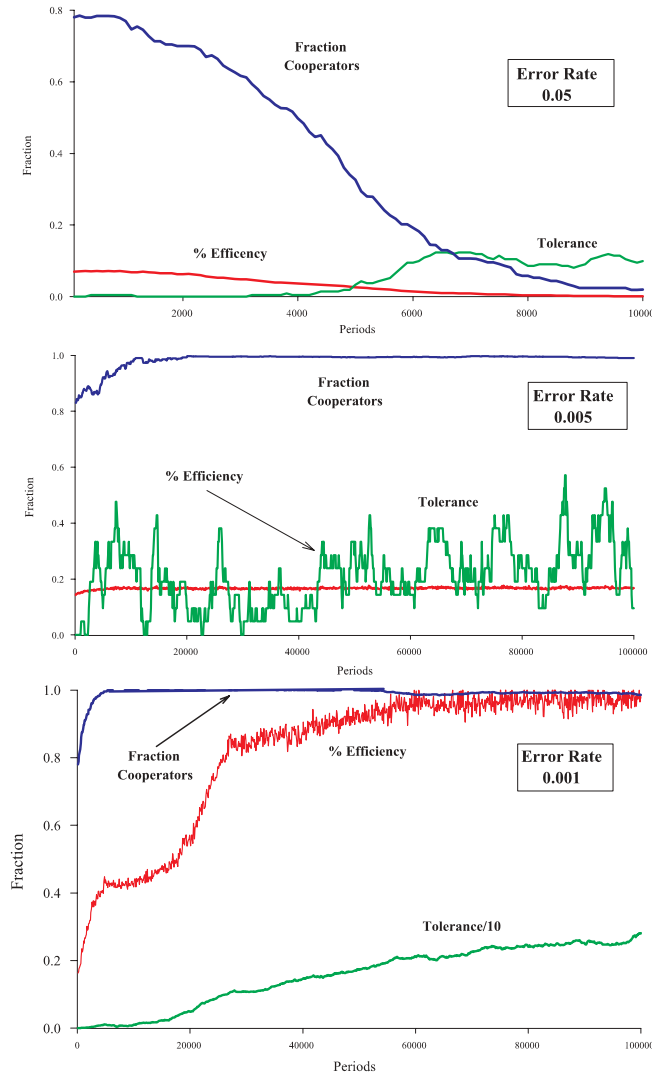


Figure 1: Simulation of Bhaskar-Obara model with group size $n = 10$, with model parameters as described in the text.

When we raise group size to $n = 10$, however the picture is quite different. The first graph in Figure 1 illustrates the case with error rate $\epsilon = 5\%$. Note that even with this relatively small group size, the level of cooperation falls to very low levels. Lowering the error rate to $\epsilon = 0.5\%$, as in the second graph in Figure 1, we see that the level of cooperation becomes high, but the efficiency of cooperation is only about 17%. This is because cooperation is signaled as defection between some pair of agents with probability close to 40%. Only when we set the error level to $\epsilon = 0.1\%$, as in Figure 1, do we achieve a high level of efficiency, the probability of an agent receiving a defection signal when in fact all are cooperating now falling below 10%. Since this low error level also allows a high level of tolerance, defections become quite rare. However, a 0.1% error rate is implausibly low.

Ely and Vålímäki (2002) have developed a quite different approach to the problem, following the lead of Piccione (2002), who showed how to achieve coordination in a repeated game with private information without the need for complex belief updating of the sort required by Bhaskar and Obara (2002). They achieve this by constructing a sequential equilibrium in which at every stage, each player is indifferent between cooperating and defecting no matter what his fellow members do. Such an agent is thus willing to follow an arbitrary mixed strategy in each period, and the authors show that there exists such a strategy for each player that ensures close to perfect cooperation, provided agents are sufficiently patient and the errors are small.

The weakness of this approach in explaining real-world cooperation, which can be stated and supported without explicitly presenting the Ely-Vålímäki model, is one shared by the mixed strategy Nash equilibria of many games. In a general one-shot game, if a player's mixed strategy is part of a Nash equilibrium, then the payoffs to all the pure strategies used with positive probability must be equal. Hence no player has an incentive to calculate and use the mixed strategy at all, since he does equally well by simply choosing among the pure strategies occurring in the support of the mixed strategy in question. If there are costs to computing and randomizing, however small, choosing the most convenient pure strategy will be strictly preferred to computing and playing the mixed strategy.

One means of avoiding this difficulty is the epistemological approach of Harsanyi (1973) and Aumann (1987), which treats the best response mixed strategy a *conjecture* of the other players, representing the limitations of knowledge they face in playing the game. While useful for some purposes, this approach does nothing to explain why one might observe players actually using the mixed strategy, in the sense that many observations of randomly paired agents would conclude that the historical average for each pure strategy observed is close to the computed Nash equilibrium distribution of pure strategies.

Harsanyi (1973) provided more pertinent solution in a altogether remarkable

demonstration of the following *purification principle*: Suppose $\sigma = \{\sigma_1, \dots, \sigma_n\}$ is a completely mixed Nash equilibrium of finite game \mathcal{G} . Then, if a certain regularity conditions is satisfied (the condition is generic in the game payoffs), and if the payoffs to each player are subject to small stochastic shocks, then with probability one each player will play a strict Nash equilibrium involving only one pure strategy, and as the size of the shocks goes to zero, the distribution of pure strategies used converges to the mixed strategy equilibrium.

	L	R		L	R
U	2,1	0,0	U	$\frac{2 + \epsilon\theta_1}{1 + \epsilon\theta_2}$	$\epsilon\theta_1, 0$
D	0,0	1,2	D	$0, \epsilon\theta_2$	0,0

Figure 2: The Battle of the Sexes. The left-hand pane is the unperturbed game, while the right-hand pane is the game with payoffs perturbed as described in the text.

For example, consider the two player game known as the *Battle of the Sexes*, with payoffs shown in Figure 2. The left-hand pane shows the payoffs in the unperturbed game. In the right-hand pane we see the game where player 1 receives additional payoff $\epsilon\theta_1$ for playing U, and player 2 receives additional payoff $\epsilon\theta_2$ for playing L. Suppose player 1 plays U with probability α and player 2 plays L with probability β . We assume θ_1 and θ_2 are uniformly distributed on $[0,1]$, although any other choices with finite support would do as well. The payoffs to the four pure strategies are then given by

$$\{\pi_U, \pi_D, \pi_L, \pi_R\} = \{2\alpha + \epsilon\theta_1, 1 - \alpha, \beta + \epsilon\theta_2, 2 - 2\beta\}.$$

We then have

$$\begin{aligned} \beta &= P[U > D] = P[\pi_U > \pi_D] = P\left[\theta_1 > \frac{1 - 3\alpha}{\epsilon}\right] = 1 - \frac{1 - 3\alpha}{\epsilon}, \\ \alpha &= P[L > R] = P[\pi_L > \pi_R] = P\left[\theta_2 > \frac{2 - 3\beta}{\epsilon}\right] = 1 - \frac{2 - 3\beta}{\epsilon}, \end{aligned}$$

Solving these equations simultaneously, we get

$$\alpha = \frac{3 - \epsilon - \epsilon^2}{9 - \epsilon^2}, \quad \beta = \frac{6 - 2\epsilon - \epsilon^2}{9 - \epsilon^2}.$$

Clearly, this approaches $\{1/3, 2/3\}$, which is the unique mixed strategy Nash equilibrium of this game.

The problem is that the Harsanyi Purification Theorem does not apply to infinite games, of which the repeated Prisoner's Dilemma is an example. Indeed, Bhaskar, Mailath and Morris (2003) have shown that their construction cannot be purified in the two-player case with perfect information, provided the players have finite memory. It is not likely, then, that purification would hold with imperfect signals and many players.

The intuition behind this result is straightforward. The Ely-Välimäki construction has the explicit goal of making agents indifferent to the actual moves of other members, thus admitting a particular constellation of mixed strategies that are specified precisely to implement an efficient, cooperative equilibrium. But, if payoffs are perturbed, players are no longer indifferent and do not have the proper incentives to implement the near-efficient solution.

9 Strong Reciprocity: An Alternative Model of Cooperation

Consider a group involved in a public goods game. Suppose a fraction $f > 0$ of the group consists of *strong reciprocators* (Gintis 2000b, Bowles and Gintis 2004) who impose a total punishment p upon a self-interested type who chooses to defect. Strong reciprocators are altruists, in the sense that they contribute to the well-being of group members at a cost to themselves that is not shared by self-interested group members.

Let v_c be the present value of cooperating for a self-interested member of the group. Since such a member defects by accident with probability ϵ , he receives expected punishment ϵp . The value of this new game for such an agent is thus

$$v_c = \frac{b(1 - \epsilon) - c - \epsilon p}{1 - \delta}. \quad (11)$$

The return to defecting for one period and then returning to cooperation is $b(1 - \epsilon) - p + \delta v_c$. It is easy to check that this quantity is negative precisely when

$$c \leq p(1 - \epsilon). \quad (12)$$

Note that this is *independent of both group size and the discount factor*. In the interests of minimizing the cost of compliance, we may assume p is chosen so that (12) is an equality.

To assess the cost of punishing for strong reciprocators, suppose that the signal indicating that an agent has defected is *private* and *imperfect*. Specifically, suppose a defectors is detected by each group member with probability q . The number of strong reciprocators who detect an intentional defector is thus $f q(n - 1)$, so each must supply a level $c/(1 - \epsilon) f q(n - 1)$ of punishment per detected infraction. If

the per unit cost of punishment is r , each strong reciprocator will incur a cost of $rc/(1-\epsilon)f q(n-1)$ per detected infraction. Since the expected number of detected infractions by other group members is $\epsilon(n-1)q$, the strong reciprocator incurs a cost $rc\epsilon/(1-\epsilon)f$ of punishing. Since this cost is not shared by self-interested types, we are obliged to exhibit a mechanism whereby strong reciprocators can be sustained in equilibrium despite their altruistic behavior. This is carried out in Gintis (2000b), Boyd, Gintis, Bowles and Richerson (2003), Gintis (2003), and Bowles and Gintis (2004), and will not be addressed here.

A model of cooperation with a sufficient fraction of strong reciprocators satisfies all the requirements listed in Section 2. It is clearly methodologically individualistic, its dynamic stability is shown in the various papers cited in the previous paragraph, its empirical relevance is supported by many experimental studies that demonstrate the existence of strong reciprocators (Fehr, Gächter and Kirchsteiger 1997, Fehr and Gächter 2000, Gintis et al. 2004). In addition, the informational requirements are such that private signals are completely adequate to ensure proper incentives, and the incentive to punish does not depend on the distribution of discount factors.

10 Conclusion

A major thrust of economic theory since the second half of the Twentieth century has been to show the plausibility of wide-scale cooperation among self-interested agents. The first thrust in this endeavor involved the rigorous specification of Walras' general equilibrium model. Despite the stunning success of this endeavor (Arrow and Debreu 1954, Debreu 1959, Arrow and Hahn 1971), the assumption that contracts could be costlessly written and costlessly enforced by third parties has never been shown feasible in a methodologically individualist model. Not surprisingly, we find that economic institutions (e.g., firms and state agencies) depend strongly on incentive mechanisms involving strategic interaction rather than explicit contracts (Gintis 1976).

The second major thrust involved the development of sophisticated repeated game-theoretic models of strategic interaction, which began with Shubik (1959), culminating in the major contributions of Fudenberg and Maskin (1986), Abreu et al. (1990), Fudenberg et al. (1994), Piccione (2002), Ely and Välimäki (2002), Bhaskar and Obara (2002) and others. While there is no question but these models have strongly advanced our understanding of the theory of social cooperation, the very fact these contributions prove "folk theorems" that sustain full-dimensional open sets of sequential equilibria virtually assure that these models will have poor dynamic qualities. Even if it can be successfully argued that the equilibria will be confined to the Pareto frontier of feasible payoffs (which is difficult to do, since this

frontier is on the boundary rather than the interior of the set of feasible payoffs), there remains an $n - 1$ dimensional manifold of sequential equilibria. The relevance of these models is further compromised by a tendency to look only at the limiting characteristics of solutions as discount rates approximate unity and error rates approach zero. Finally, not attempt has ever been made to show that such models apply to the major forms of non-market-mediated economic cooperation. There is considerable doubt that they could so apply, since neither trigger strategies nor second order punishment is commonly observed in socially efficient cooperative groups.

By contrast, strong reciprocity models have excellent dynamic properties, do not depend on discount factors near unity, and support relatively high error rates in private signals. Dynamic robustness is important because social institutions must both emerge and remain stable against adverse shocks. High error rates must be supported because we observe quite imperfect signals of cooperation in the real world, and it is often difficult to separate the failure of an agent to cooperate with the failure of the agent to succeed through forces beyond his control. Discount factors well below unity must be supported because there is evidence that humans have very high short-term discount rates (Ainslie and Haslam 1992, Laibson 1997), the variance of discount rates across human individuals at any point in time is considerable, if only because of age- and health-related differential life expectancy (Kaplan, Gurven, Hill and Hurtado 2003, Robson and Kaplan 2003). Indeed, individual discount factors are likely to have been high throughout most of human history, both because of the riskiness of life and the fragility of group ties. For instance, hunter-gather groups typically experience periodic threats to their existence, in the form of pestilence, famine, and war, at which time the discount factor is quite low, since the probability of group dissolution is high. Self-interested cooperation models predict the dissolution of such groups, whereas our models predict that such conditions may favor the emergence of agents who cooperate and punish without regard to the discount factor. Behavioral experiments show that such agents do exist in large numbers, lending credence to the gene-culture coevolution models of which strong reciprocity models are a prime example (Gintis et al. 2004).

REFERENCES

- Abreu, Dilip, David Pearce, and Ennio Stacchetti, "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring," *Econometrica* 58,5 (September 1990):1041–1063.
- Ainslie, George and Nick Haslam, "Hyperbolic Discounting," in George Loewenstein and Jon Elster (eds.) *Choice Over Time* (New York: Russell Sage, 1992)

pp. 57–92.

Arrow, Kenneth J. and Frank Hahn, *General Competitive Analysis* (San Francisco: Holden-Day, 1971).

— and Gerard Debreu, “Existence of an Equilibrium for a Competitive Economy,” *Econometrica* (1954):265–290.

Aumann, Robert and Adam Brandenburger, “Epistemic Conditions for Nash Equilibrium,” *Econometrica* 65,5 (September 1995):1161–80.

Aumann, Robert J., “Correlated Equilibrium and an Expression of Bayesian Rationality,” *Econometrica* 55 (1987):1–18.

Becker, Gary S. and Casey B. Mulligan, “The Endogenous Determination of Time Preference,” *Quarterly Journal of Economics* 112,3 (August 1997):729–759.

Bhaskar, V. and Ichiro Obara, “Communication in the Repeated Prisoner’s Dilemma with Private Monitoring,” *Journal of Economic Theory* 102 (2002):40–69.

—, George J. Mailath, and Stephen Morris, “Purification in the Infinitely Repeated Prisoner’s Dilemma,” 2003. University of Essex.

Bowles, Samuel and Herbert Gintis, “Walrasian Economics in Retrospect,” *Quarterly Journal of Economics* (November 2000):1411–1439.

— and —, “The Evolution of Strong Reciprocity: Cooperation in Heterogeneous Populations,” *Theoretical Population Biology* 65 (2004):17–28.

—, Jung-kyoo Choi, and Astrid Hopfensitz, “The Co-evolution of Individual Behaviors and Social Institutions,” *Journal of Theoretical Biology* 223 (2003):135–147.

Boyd, Robert, Herbert Gintis, Samuel Bowles, and Peter J. Richerson, “Evolution of Altruistic Punishment,” *Proceedings of the National Academy of Sciences* 100,6 (March 2003):3531–3535.

Debreu, Gérard, *Theory of Value* (New York: Wiley, 1959).

Ely, Jeffrey C. and Juuso Välimäki, “A Robust Folk Theorem for the Prisoner’s Dilemma,” *Journal of Economic Theory* 102 (2002):84–105.

Fehr, Ernst and Simon Gächter, “Cooperation and Punishment,” *American Economic Review* 90,4 (September 2000):980–994.

—, —, and Georg Kirchsteiger, “Reciprocity as a Contract Enforcement Device: Experimental Evidence,” *Econometrica* 65,4 (July 1997):833–860.

Friedman, James W., “A Non-cooperative Equilibrium for Supergames,” *Review of Economic Studies* 38,113 (January 1971):1–12.

Fudenberg, Drew and Eric Maskin, “The Folk Theorem in Repeated Games with Discounting or with Incomplete Information,” *Econometrica* 54,3 (May 1986):533–554.

— and Jean Tirole, *Game Theory* (Cambridge, MA: MIT Press, 1991).

- , David K. Levine, and Eric Maskin, “The Folk Theorem with Imperfect Public Information,” *Econometrica* 62 (1994):997–1039.
- Gintis, Herbert, “The Nature of the Labor Exchange and the Theory of Capitalist Production,” *Review of Radical Political Economics* 8,2 (Summer 1976):36–54.
- , *Game Theory Evolving* (Princeton, NJ: Princeton University Press, 2000).
- , “Strong Reciprocity and Human Sociality,” *Journal of Theoretical Biology* 206 (2000):169–179.
- , “Some Implications of Endogenous Contract Enforcement for General Equilibrium Theory,” in Fabio Petri and Frank Hahn (eds.) *General Equilibrium: Problems and Prospects* (London: Routledge, 2002) pp. 176–205.
- , “The Hitchhiker’s Guide to Altruism: Genes, Culture, and the Internalization of Norms,” *Journal of Theoretical Biology* 220,4 (2003):407–418.
- , Samuel Bowles, Robert Boyd, and Ernst Fehr, *Moral Sentiments and Material Interests: On the Foundations of Cooperation in Economic Life* (MIT Press, 2004).
- Granovetter, Mark, “Economic Action and Social Structure: The Problem of Embeddedness,” *American Journal of Sociology* 91,3 (November 1985):481–510.
- Hansson, Ingemar and Charles Stuart, “Malthusian Selection of Preferences,” *American Economic Review* 80,3 (June 1990).
- Harsanyi, John C., “Games with Randomly Distributed Payoffs: A New Rationale for Mixed-Strategy Equilibrium Points,” *International Journal of Game Theory* 2 (1973):1–23.
- Henrich, Joe, Robert Boyd, Samuel Bowles, Colin Camerer, Ernst Fehr, and Herbert Gintis, *Foundations of Human Sociality: Economic Experiments and Ethnographic Evidence from Fifteen Small-scale Societies* (Oxford: Oxford University Press, 2004).
- Hodgson, Geoffrey M., “The Approach of Institutional Economics,” *Journal of Economic Literature* 36,1 (1998):166–192.
- Kandori, Michihiro, “Introduction to Repeated Games with Private Monitoring,” *Journal of Economic Theory* 102 (2002):1–15.
- Kaplan, Hillard, Michael Gurven, Kim Hill, and A. Magdalena Hurtado, “The Natural History of Human Food Sharing and Cooperation: A Review and a New Multi-Individual Approach to the Negotiation of Norms,” in Herbert Gintis, Samuel Bowles, Robert Boyd, and Ernst Fehr (eds.) *The Moral Sentiments: Origins, Evidence, and Policy* (???: ??, 2003) pp. —.
- Laibson, David, “Golden Eggs and Hyperbolic Discounting,” *Quarterly Journal of Economics* 112,2 (May 1997):443–477.

- Mandeville, Bernard, *The Fable of the Bees, or Private Vices, Publick Benefits* (Oxford: Clarendon Press, 1924[1705]).
- Mas-Colell, Andreu, Michael D. Whinston, and Jerry R. Green, *Microeconomic Theory* (New York: Oxford University Press, 1995).
- Matsushima, H., “The Folk Theorem with Private Monitoring and Uniform Sustainability,” 2000. CIRJE Discussion Paper F-84, University of Tokyo.
- Piccione, Michele, “The Repeated Prisoner’s Dilemma with Imperfect Private Monitoring,” *Journal of Economic Theory* 102 (2002):70–83.
- Porter, R. H., “On the Incidence and Duration of Price Wars,” *Journal of Industrial Economics* 33 (1985):415–426.
- Robson, Arthur and Hillard Kaplan, “The Evolution of Human Life Expectancy and Intelligence in Hunter-Gatherer Economies,” *American Economic Review* 93,1 (March 2003):150–169.
- Rogers, Alan, “Evolution of Time Preference by Natural Selection,” *American Economic Review* 84,3 (June 1994):460–481.
- Sekiguchi, Tadashi, “Efficiency in Repeated Prisoner’s Dilemma with Private Monitoring,” *Journal of Economic Theory* 76 (1997):345–361.
- Shubik, Martin, *Strategy and Market Structure: Competition, Oligopoly, and the Theory of Games* (New York: Wiley, 1959).
- Smith, Adam, *The Wealth of Nations* (New York: Modern Library, 1937[1776]).
- Stiglitz, Joseph, “The Causes and Consequences of the Dependence of Quality on Price,” *Journal of Economic Literature* 25 (March 1987):1–48.
- Taylor, P. and L. Jonker, “Evolutionarily Stable Strategies and Game Dynamics,” *Mathematical Biosciences* 40 (1978):145–156.
- Walras, Leon, *Elements of Pure Economics* (London: George Allen and Unwin, 1954 [1874]).
- Wiessner, Polly, “The Pleasures and Perils of Reward and Punishment: Strong Reciprocity Among the Ju/’hoansi Bushmen,” 2003. Department of Anthropology, University of Utah.