Reputation and Distribution in a Gift Giving Game[1]

David K. Levine[2]

March 18, 1995

Abstract: The folk theorem allows a very unequal division between players. In non-repeated experimental games with many equilibria, such as ultimatum, observed play involves a relatively equal division between players. In a two-player repeated game setting there is a simple intuition about this: a poor player has little to lose by deviating from his equilibrium strategy. So a rich player ought to be willing to concede a reasonable amount of pie. We investigate whether reputation effects lead to this conclusion in a simple two person gift giving game. When types are known equilibria are socially feasible and pairwise individually rational. In the reputational case the set of equilibria is smaller than the set of socially feasible, average individually rational and incentive compatible payoffs, and larger than the set of socially feasible, pairwise individually rational and incentive compatible points. In general the set of equilibrium payoffs need not be smaller or more equitable in the reputational case. However, in sensible example, we can how the incentive constraints do have the desired effect of reducing the set of equilibrium payoffs and eliminating many inequitable equilibria.

[1] The author is grateful for financial support from NSF grant SBR-9409180 and the UCLA Academic Senate. This paper benefited from conversations with Drew Fudenberg, David Kreps and Ed Prescott.

[2] Department of Economics, UCLA, Los Angeles, CA 90024; e-mail: dlevine@ucla.edu

1. Introduction

Roughly speaking, the folk theorem for repeated games says that with sufficiently patient players any assignment of utility to players that is both socially feasible and individually rational arises in equilibrium. This means that there will typically be equilibria that involve a very unequal division between players. In non-repeated experimental games with many equilibria, such as the ultimatum game[3], the equilibria actually observed typically involve a relatively equal division between players. Evidence from the dictator game[4] suggests that experimental subjects are suspicious that their experimental play may be revealed to others after the end of the experiment. This means that a repeated game or reputational model may be appropriate, despite the apparently non-repeated nature of the experimental game. In a two-player repeated game setting, there is a simple intuition about why we might expect a relatively equal division between players: if one player, the poor player, is relatively close to the level of utility that is individually rational, and his opponent, the rich player, is well above that level, then the poor player has little to lose by deviating from his equilibrium strategy, to, say, his minmax strategy, and this will typically impose large costs on the rich player. The rich player, recognizing this threat ought to be willing to concede a larger part of the pie.

The goal of this paper is to investigate this intuition in the context of a very simple two person gift giving game. Each player has one unit of a divisible consumption good that may either be consumed or given as a gift to the other player. Utility is linear and non-decreasing in the amount of the player's own good consumed, and in the amount that his opponent gives him to consume. An economic example where this type of gift giving occurs is the exchange of services between family members. There are different types of players with different preferences between the two goods. We consider this game under two different information conditions: the case in which types are known, and the case in which each player knows only his own type. We refer to these as the *perfect information game*, and *reputational game*, respectively. In both cases we look at Bayesian subgame perfect equilibrium payoffs as the common discount factor goes to one.

In the case of the perfect information game, the set of limit equilibrium payoffs is a trivial consequence of the folk theorem: it is the set of socially feasible, pairwise individually rational payoffs. That is, the payoffs must satisfy individual rationality for each pair of players. We also give a partial characterization of the equilibrium set in the

[3] A good discussion of ultimatum experiments can be found in Roth et al [1991].
[4] See Hoffman et al [in press].

reputational case: we show that it is smaller than the set of socially feasible, average individually rational and incentive compatible payoffs, and larger than the set of socially feasible, pairwise individually rational and incentive compatible points. Here average individual rationality means that each type on average gets his individually rational payoff, while incentive compatibility means that no type prefers to switch the the strategy of a different type. The point where all players donate their entire endowment to the other is a limiting equilibrium of both the perfect information and reputational game. For a wide variety of preferences, it is Pareto efficient as well. We should point out that this type of corner solution where one partner provides all of the services of a particular type are relatively common in the case of families exchanging services. Naturally another equilibrium under both informational conditions is the static Nash equilibrium of the stage game in which there is no gift giving at all. These two equilibria establish a range of possibilities from complete autarky to full efficiency for both informational conditions. This establishes a sense in which the difference between equilibria in the two games is that they differ in the possible distribution of equilibrium payoffs.

We cannot say that in general the set of equilibrium payoffs is smaller or more equitable in the reputational case. We give an example showing how average individual rationality, a weaker condition than pairwise individual rationality, may lead to the equilibrium payoffs in the reputational case being larger than in the perfect information case. However, in a more sensible example, we show how the incentive constraints do have the desired effect of reducing the set of equilibrium payoffs and eliminating many inequitable equilibria.

This type of reputational model differs from those that have been studied previously in several respects. Most of the literature on reputation studies two players one of whom is relatively more patient than the other: a recent paper that discusses this literature in greater detail is Celentani, Fudenberg, Levine and Pesendorfer[1993]. Typically this literature assumes that there are some types committed to playing particular strategies. The usual result is that the more patient player can get the result he most prefers subject to the individual rationality of the less patient player. This can be interpreted as a type of efficiency, that is on the boundary of the utility possibility set. It may also be interpreted as one player being able to prevent his opponent from forcing him to low levels of utility, an interpretation more consistent with the model in this paper. There is a limited literature on reputation with equally patient players, including Kreps, Milgrom, Roberts and Wilson [1982], Fudenberg and Maskin [1986] and Aumann and Sorin [1989] but the focus in these paper is on types committed to playing particular punishment strategies. In this paper types do not play punishment strategies, and indeed no type is committed to playing any particular strategy. Different types simply have different preferences in the static game. This is most similar to the work of Schmidt [1989] on reputation in a bargaining game.

In addition to Schmidt's [1989] study of a repeated bargaining game, Abreu and Gul [1994] have studied the role of reputation in a non-repeated game bargaining setting. Their work also shows how reputation can narrow down the set of possible equilibrium distributions. Indeed, their results are considerably stronger than the results here: they show that the division of pie is essentially unique, and independent of details of the bargaining procedure. The stronger results seem to be largely a consequence of the fact that in the non-reputational game bargaining game the folk theorem does not hold, and this significantly reduces the possibilities once one player has revealed his type.

One important difference between other work on reputation and the work here should be emphasized. Traditionally the reputational literature has focused on a game that is slightly perturbed by the addition of types that occur with small probability. This is not the case here: it is crucial that the strong types occur with relatively high probability. The reason is due to the folk theorem. One possibility is to offer the other player terrible terms of trade. A weak opponent will accept this and a strong opponent will not. However, the fact that this is inefficient does not prevent it from being an equilibrium. If the probability of a strong opponent is very small this is a very good strategy to follow, since it extracts a lot of surplus from the weak opponent at the expense of a small probability of losing when meeting the strong opponent. However, when the probability of a strong opponent is appreciable, it is necessary to weight the gain to extracting surplus from a weak opponent against the gain to trade with a strong opponent. To make this point we do not compare an "unperturbed" with a "perturbed" game. Rather with examine a game with fixed type probabilities under two different information conditions: the case in which types are public information and the case in which they are private information.

From a technical point of view, the current analysis is closely related to mechanism design. The major idea here is that the special nature of linear preferences in a gift giving game makes it possible to reduce the analysis of reputational equilibria to the study of incentive compatibility and individual rationality constraints in a static mechanism design setting. Although it is not always emphasized, it is well known that the size of the Pareto frontier in static mechanism design problems is reduced by the presence of incentive constraints. A good example is the signaling problem studied in Prescott and Townsend [1984] where the incentive constraints reduce the Pareto frontier to a single point. From this point of view, the results here should not be terribly surprising: the main point is that reputational equilibria have much in common with mechanism design problems since different types can potentially imitate each other.

2. The Model

We first describe the stage game. There are two players $i = 1,2$. Each has one unit of a distinct infinitely divisible consumption good which may either be consumed or any portion of which may be given to the other player as a gift. We denote by g_i the amount of the gift given by player *i* to player *-i*. Each player's utility is linear in the amount of consumption of their own good and the other player's good. Specifically, the utility of player *i* is given by

$$u_i = \omega_i^0 g_{-i} - \omega_i^1 g_i$$

We refer to the vector $\omega_i = (\omega_i^0, \omega_i^1)$ as player *i*'s type. If

$$\omega_i^0 - \omega_{-i}^1 \geq 0,$$

for all pairs of types, with strict inequality for one pair, then we say that there are *universal gains to trade*. In this case the sum of the two players payoffs is never decreased by having them exchange gifts, regardless of which pair of types is playing.

Types are drawn from a finite set of types Ω_i with type ω_i having probability $\mu_i(\omega_i)$. The types of the two players are drawn independently of one another. These probability distributions over types are common knowledge. We consider two cases. In the *perfect information game* each player knows their own type and the type of opponent. In the *reputational game* each player knows only their own type and the distribution μ_{-i} of the others type. We denote player *i*'s information by $\hat{\Omega}_i = \Omega_i \times \Omega_{-i}$ or Ω_i as the game is the perfect information or reputational game.

In the repeated game, both players have a common discount factor δ and act to maximize expected average present value. Histories are $h = (g_1(1), g_2(1), \ldots, g_1(t), g_2(t))$. The length of a history is denoted $t(h)$. The set of all histories is denoted by H, and by notational convention, this includes a null history h_0. Behavior strategies are maps $\sigma_i : \hat{\Omega}_i \times H \to \Delta([0,1])$. Our solution concept is *perfect Bayesian equilibrium* by which we mean that no player has any incentive to deviate from the equilibrium strategies following any history, and that in the reputational game beliefs about opponents types on the equilibrium path are determined by Bayes law, and off the equilibrium path are arbitrary.

As the discount factor converges to one, we may consider average present value vectors that arise as limits of equilibrium payoffs. We denote these limiting sets by

V^R, V^P in the reputational and perfect information games respectively.

3. Allocations

If we denote by $\pi(h)$ the probability of the history *h*, the we may write average present value as

$$\begin{aligned} V_i &= (1-\delta)\sum_h \delta^{t(h)-1}\left(\omega_i^0 g_{-i} - \omega_i^1 g_i\right)\pi(h) \\ &= \omega_i^0(1-\delta)\sum_h \delta^{t(h)-1}\pi(h) g_{-i} - \omega_i^0(1-\delta)\sum_h \delta^{t(h)-1}\pi(h) g_i \\ &= \omega_i^0 G_{-i} - \omega_i^0 G_i \end{aligned}$$

In other words, utility in the repeated game depends only on the average present value of the gifts $G_i = \sum_h \delta^{t(h)-1}\pi(h) g_i$. This leads us to define an allocation as a map $G{:}\Omega \to \Re^2$.

Several different allocations play a role in characterizing equilibrium. The set of *socially feasible allocations* is *SF* and satisfies $1 \geq G_i(\omega) \geq 0$. There are several notions of individual rationality: notice that each player always has the option of giving nothing, and guaranteeing himself zero. The *pairwise individually rational allocations* are *PIR* and are allocations that satisfy are individually rationality restrictions conditional on each pair of types

$$\omega_i^0 G_{-i}(\omega) - \omega_i^1 G_i(\omega) \geq 0, \forall \omega .$$

The *average individually rational allocations* are *AIR* and are allocations that satisfy the individual rationality restrictions conditional on the player's own type.

$$\sum_{\omega_{-i} \in \Omega_{-i}} \left(\omega_i^0 G_{-i}(\omega) - \omega_i^1 G_i(\omega)\right)\mu_{-i}(\omega_{-i}) \geq 0, \forall \omega_i .$$

Finally, *individual rational allocations* are *IR* simply require no player get less than zero in expected utility

$$\sum_{\omega \in \Omega}\left(\omega_i^0 G_{-i}(\omega) - \omega_i^1 G_i(\omega)\right)\mu_i(\omega_i)\mu_{-i}(\omega_{-i}) \geq 0 .$$

It is obvious that $PIR \subseteq AIR \subseteq IR$, but $PIR \neq IR$ in general, as we shall see below.

Since players value their opponents good more than their own, a particularly important allocation is the allocation in which $G_i(\omega) = 1$. We refer to this as the *full exchange allocation*, denoted $\hat{G}$.

Finally, we have the notion of *incentive compatible allocations*, denoted *IC*, which require that no type wish to receive the allocation given another type.

$$\sum_{\omega_{-i} \in \Omega_{-i}} \left(\omega_i^0 G_{-i}(\omega) - \omega_i^1 G_i(\omega)\right)\mu_{-i}(\omega_{-i}) \geq \sum_{\omega_{-i} \in \Omega_{-i}} \left(\omega_i^0 G_{-i}(\tilde{\omega}_i, \omega_{-i}) - \omega_i^1 G_i(\tilde{\omega}_i, \omega_{-i})\right)\mu_{-i}(\omega_{-i}), \forall \omega_i, \tilde{\omega}_i$$

We also need to convert allocations into utilities. If $\Im$ is a set of allocations define

$$V(\Im) \equiv \left\{ \left(\omega_1^0 G_2(\omega) - \omega_1^1 G_1(\omega), (\omega_2^0 G_1(\omega) - \omega_2^1 G_2(\omega) \right) \middle| G(\omega) \in \Im \right\}.$$

4. Results

We now give some general results characterizing equilibrium payoffs. In the perfect information case, conditional on each pair of types the Fudenberg-Maskin[1986] folk theorem applies, and we may immediately conclude

Folk Theorem: $V^P = V(SF \cap PIR)$

We can also relatively easily give a partial characterization of equilibrium payoffs in the reputational case

Reputational Theorem: $V^R \subseteq V(SF \cap AIR \cap IC)$; if $SF \cap PIR \cap IC$ has non-empty interior, then $V(SF \cap PIR \cap IC) \subseteq V^R$.

proof: Since each type has the option of playing the strategy of any other type, it is obvious that $V^R \subseteq V(SF \cap IC)$. Moreover, each type must get at least the minmax for that type after learning his type, so this establishes $V^R \subseteq V(SF \cap AIR \cap IC)$. Supposing conversely that we have an allocation in the interior of $SF \cap PIR \cap IC$ we assign each type a unique gift. In the first period each type reveals itself by giving the unique gift associated with that type. In the second period, each pair of types moves to an equilibrium in which the utility when added to the utility from the first period signal is exactly the utility from the target allocation. This is possible by the interiority condition, which enables us to adjust for the first period signals, and by the folk theorem which applies once the types are revealed since *PIR* is satisfied. Since *IC* is satisfied, no type has any incentive to deviate in the first period either.

❑

There are two fairly obvious candidates for equilibrium, autarky and full exchange. Autarky, is the unique equilibrium of the stage game, so is obviously an equilibrium of the repeated game. We have

Theorem: $0, V(\hat{G}) \in V(SF \cap PIR \cap IC)$; if there are universal gains to trade $V(\hat{G})$ is Pareto efficient.

proof: We discussed the case of autarky where all players get zero above. Obviously $\hat{G}$ is socially feasible, and since each type prefers the opponents good to his own satisfies

PIR. Moreover, since every type both gives and receives the same gift, it makes no difference what type is announced. Finally, with universal gains to trade, exchange of gifts never decreases the sum of the two players utility, and the sum is increased for at least one pair.

❑

5. Example: Pairwise IR vs. Average IR

One might expect that due to the incentive constraints the set of reputational equilibria is generally smaller than the set of perfect information equilibria. However, since *AIR* is weaker than *PIR* it is possible that there are actually reputational equilibria that are not perfect information equilibria. We begin with an example that shows that this is indeed the case.

We suppose that player 1 has only one type (1,1); while player 2 has two equi-probable types (1,1) and (0,0). We consider the pairwise case first. In this case if types (1,1) and (1,1) meet there is no gain from trade, so both players must get zero by *PIR.* If type (1,1) meets type (0,0), player 2 is indifferent, and always gets zero, while player 1 can get any utility from zero to one, depending on how much player 2 gives. Since there is a .5 chance of meeting type (0,0) the expected utility in the best case is only 1/2. It follows that in the perfect information game player 1 gets between 0 and 1/2, while player 2 gets zero. Notice that in this game $V(\hat{G}) = (0,0)$ so that full exchange is not efficient.

We now consider the reputational game. Let T be such that the average present value of a unit of consumption on or before T is 1/6, and after that date the weight is 5/6. (This can be achieved with an arbitrarily high degree of accuracy for discount factors near one.) Consider the following strategies: Player 1 gives one through time T, then donate zero forever after. Type (1,1) player 2 never donates. Type (0,0) player 2 donates zero through time T then donates one forever. If play is ever observed that is inconsistent with these strategies revert to autarky forever.

Notice that player 2 is clearly optimizes in every contingency: type (0,0) is indifferent, and type (1,1) is never called on to make a gift and his opponent gives him no reason to do so. Player 1 is also optimizing. He does not learn his opponents type until time T, at which point he stops donating anyway. If he tries to cut his donation prior to time T he gains at most 1/6, and looses a .5 chance of 5/6. So it is best to wait until T and hope that the other type is the indifferent type.

In this equilibrium, player 1 gets 1/4. More important, however, is the fact that player 2 gets 1/12. By contrast in the perfect information game he never got anything at all.

Although this example is special, it may straightforwardly be extended to the case in which the type (0,0) player 2 is replaced by a type $(\varepsilon_1, \varepsilon_2)$ player, for example. Notice that in both the original example, and in this extension there are not universal gains to trade. Indeed, efficiency demands that player 1 make gifts to the type (1,1) player 2 and receive gifts from the type $(\varepsilon_1, \varepsilon_2)$ player 2. Such trades cannot be accomplished once the types are known, but as the example shows, can be conducted in the reputational game. An open question is whether there are limit reputational equilibria that fail to be limit perfect information equilibria when there are universal gains to trade. In the next example where there are universal gains to trade, we can at least show that there are relatively few such points.

6. Example: The Effect of Incentive Compatibility

We now consider a more symmetric example, and use this to show how the incentive constraints serve to rule out relatively asymmetric outcomes.

We now suppose that each player has two equi-probable types (1,1) and $(1, \varepsilon)$, referred to as the strong and weak types respectively. Notice that in this example there are universal gains to trade.

We start by considering the $SF \cap PIR$ set. First we maximize player 2's payoff holding each type of player 1 at zero.

player 1 type	player 2 type	player 1 gift	player 2 gift	player 2 payoff
strong	strong	=p2 gift	=p1 gift	0
strong	weak	1	1	$1-\varepsilon$
weak	strong	1	ε	$1-\varepsilon$
weak	weak	1	ε	$1-\varepsilon^2$

This yields a payoff vector of $\left(0, \frac{3}{4} - \frac{\varepsilon}{2} - \frac{\varepsilon^2}{4}\right)$. This marks the upper left corner of the equilibrium set indicated with the solid line in the figure below.

We now wish to move along the efficiency frontier by transferring utility from player 2 to player 1. To stay on the frontier we should look for the least cost transfer, which is to have the weak player 2 facing a weak player 1 increase the size of his gift. When the weak player 2 gives the maximum gift, player 1's utility is increased to $\frac{1}{4} - \frac{\varepsilon}{4}$, while player 2's utility is decreased to $\frac{3}{4}-\frac{3\varepsilon}{4}$. Consequently the point $\left(\frac{1}{4}-\frac{\varepsilon}{4},\frac{3}{4}-\frac{3\varepsilon}{4}\right)$ marks the next extremal point on the frontier in the figure. We can now move further along the efficiency frontier only by having either the strong player 1 reduce his gift to the weak player 2, or having the strong player 2 increase his gift to the weak player 1. In either case, utility is transferred 1 for 1 until we reach the full exchange payoff of $\left(\frac{1}{2}-\frac{\varepsilon}{2},\frac{1}{2}-\frac{\varepsilon}{2}\right)$.

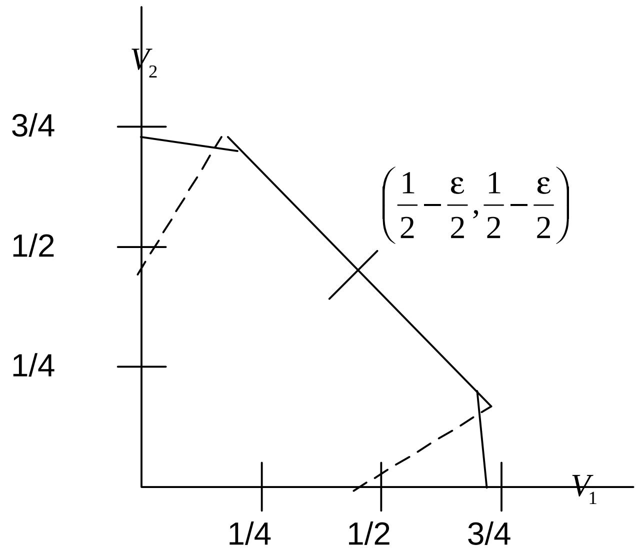

Next we take up the reputational game. Here the goal is to show that the set of equilibrium payoffs is small, so we consider $V(SF \cap AIR \cap IC)$ which is an upper bound on how large the set of equilibrium payoffs might be. As in the perfect information case, we begin by maximizing player 2's payoff subject to *AIR* and each type of player 1 getting 0. In addition, we must now consider *IC*.

Let G_2^ω be 2's expected donation to type ω; let $\tilde{G}_1^\omega$ be the expected donation by type 1 (note that 2 does not care how it is allocated between his types). We have

AIR $\qquad \tilde{G}_1^s \le G_2^s,\ \varepsilon\tilde{G}_1^w \le G_2^w$

IC $\qquad G_2^w - \varepsilon\tilde{G}_1^w \ge G_2^s - \varepsilon\tilde{G}_1^s$.

We ignore the incentive constraint that the strong type does not wish to imitate the weak type, and it may easily be shown that this constraint does not bind on the pareto frontier anyway. If player 1 is to get 0 then the *AIR* conditions hold with exact equality, and by *IC* this implies $0 \ge G_2^s(1-\varepsilon)$ and nothing may be given to the strong type. Utility will then be maximal for player 2 when his weak type gives 2ε to the weak type player 1 in exchange for the player 1 weak type giving 1 all the time. This gives 2 a utility of $\frac{1}{2}-\frac{\varepsilon^2}{2}$, and establishes the upper left corner of our set that bounds the reputational equilibrium payoffs at $\left(0,\frac{1}{2}-\frac{\varepsilon^2}{2}\right)$.

As we move along the upper boundary player 2 can increase his utility by having his weak type trade with the strong type player 1. Note that to maximize player 2's payoff holding fixed player 1's utility *AIR* must hold with equality for the strong type of player 1, so that $\tilde{G}_1^s = G_2^s$; Therefore *IC* therefore implies that the weak type player 1 gets $G_2^s(1-\varepsilon)$ and since he is going to give 1 this means

$$G_2^w - \varepsilon = G_2^s(1-\varepsilon) \text{ or } G_2^w = G_2^s(1-\varepsilon)+\varepsilon .$$

Total giving by the weak type player 2 is twice the average of the donation to weak and strong types of player 1, for a total of

$$2\left(G_2^s + G_2^s(1-\varepsilon)+\varepsilon\right)/2 = 2G_2^s + (1-G_2^s)\varepsilon .$$

The largest this can be is one, in which case $G_2^s = (1-\varepsilon)/(2-\varepsilon)$. In exchange, player 2 gets $G_2^s/2$ from the strong type player 1 plus the 1/2 he was already getting from the weak type player 1. Total utility for 2 is therefore

$$\frac{1}{2}+\frac{G_2^s}{2}-\frac{1}{2}\varepsilon\left(2G_2^s+(1-G_2^s)\varepsilon\right).$$

In particular, when the weak type donates the maximum, the utilities of the two players are

$$\left(\frac{G_2^s(1-\varepsilon)}{2},\frac{1}{2}+\frac{G_2^s}{2}-\frac{1}{2}\varepsilon\left(2G_2^s+(1-G_2^s)\varepsilon\right)\right)$$
$$=\left(\frac{\frac{1-\varepsilon}{2-\varepsilon}(1-\varepsilon)}{2},\frac{1}{2}+\frac{\frac{1-\varepsilon}{2-\varepsilon}}{2}-\frac{1}{2}\varepsilon\left(2\frac{1-\varepsilon}{2-\varepsilon}+(1-\frac{1-\varepsilon}{2-\varepsilon})\varepsilon\right)\right)$$
$$=\left(\frac{(1-\varepsilon)^2}{2(2-\varepsilon)},\frac{3-4\varepsilon+\varepsilon^2}{2(2-\varepsilon)}\right)$$

From this point, payoffs along the upper frontier decline linearly to the full exchange point. This is illustrated in the figure above. Note that

$$\frac{(1-\varepsilon)^2}{2(2-\varepsilon)} \le \frac{1-\varepsilon}{4}.$$

This implies that the efficiency frontier of upper bound on the reputational equilibrium lies outside the set of equilibrium payoffs in perfect information case for a small range as shown in the figure. This is because we are allowing the weak type of player 2 to make a donation of ε to the strong type of player 1 on behalf of the strong type of player 2, while in the perfect information case no type can donate on behalf of another. Since this is simply an upper bound on what player 2 can get in equilibrium, it does not actually imply that player 2 can get this extra amount in the reputational equilibrium. In any case the difference is quite small.

7. Conclusion

The overall conclusion of the paper is found in the previous example and the figure that goes with it. In this particular gift giving game, reputation substantially narrows the range of utility distributions. For ε small, a player can get about 3/4 in with or without reputation, but in the reputational game, the player must cede about 1/4 to his opponent in order to get this amount. To push his opponents utility below this level requires that he make a sacrifice as well. If we subscribe to the notion that efficient arrangements are more likely than inefficient arrangements, in this game with reputation efficiency must necessarily lead also to a higher degree of equity.

References

Abreu, D. and F. Gul [1994]: “Bargaining and Reputation,” Princeton.

Aumann, R. and S. Sorin [1989]: “Cooperation and Bounded Recall”, *Games and Economic Behavior,* 1, 5-39.

Celentani, M., D. Fudenberg, D. Levine and W. Pesendorfer [1993]: “Maintaining a Reputation Against a Patient Opponent,” Northwestern.

Fudenberg, D. and E. Maskin [1986]: “The Folk Theorem for repeated games with discounting and incomplete information”, *Econometrica,* 54, 533-554.

Hoffman, E., K. McCabe, K. Shachat, and V. Smith [in press]: “Preferences, Property Rights and Anonymity in Bargaining Games,” *Games and Economic Behavior.*

Kreps, D.M., P. Milgrom, J. Roberts and C. Wilson [1982]: "Rational Cooperation in the Finitely Repeated Prisoners' Dilemma," *Journal of Economic Theory*, 27, 245-252, 486-502.

Prescott, E. and R. Townsend [1984]: “Pareto Optima and Competitive Equilibria with Adverse Selection and Moral Hazard”, *Econometrica*, 52, 21-46.

Roth, Alvin E., V. Prasnikar, M. Okuno-Fujiwara and S. Zamir [1991]: “Bargaining and market behavior in Jerusalem, Liubljana, Pittsburgh and Tokyo: an experimental study,” *American Economic Review*, 81:1068-1095.

Schmidt, K. M. [1989]: “Commitment Through Incomplete Information in a Simple Bargaining Model,” London School of Economics.