ELSEVIER

Journal of Public Economics 89 (2005) 1457–1468

JOURNAL OF PUBLIC ECONOMICS

www.elsevier.com/locate/econbase

Cooperation among strangers with limited information about reputation

Gary E. Bolton[a,*], Elena Katok[a], Axel Ockenfels[b]

[a]*Penn State University, Smeal College of Business, University Park, PA, USA*
[b]*Department of Economics, University of Cologne, Germany*

Received 27 February 2003; received in revised form 17 March 2004; accepted 24 March 2004
Available online 13 November 2004

Abstract

The amount of institutional intervention necessary to secure efficiency-enhancing cooperation in markets and organizations, in circumstances where interactions take place among essentially strangers, depends critically on the amount of information informal reputation mechanisms need transmit. Models based on subgame perfection find that the information necessary to support cooperation is recursive in nature and thus information generating and processing requirements are quite demanding. Models that do not rely on subgame perfection, on the other hand, suggest that the information demands may be quite modest. The experiment we present indicates that even without any reputation information there is a non-negligible amount of cooperation that is, however, quite sensitive to the cooperation costs. For high costs, providing information about a partner's immediate past action increases cooperation. Recursive information about the partners' previous partners' reputation further promotes cooperation, regardless of the cooperation costs.

JEL classification: C7; C91; L14
Keywords: Experimental economics; Reputation; Cooperation

* Corresponding author. Tel.: +1 814 865 0611.
E-mail address: geb3@psu.edu (G.E. Bolton).

0047-2727/$ - see front matter
doi:10.1016/j.jpubeco.2004.03.008

1. Introduction: reputation and the role of information

Economists have long recognized reputation as an effective means of enforcing cooperation when an institution exists to track and disseminate such information (e.g., credit agencies; Milgrom et al., 1990), or within a small group where people are intimately familiar with of one another's history (Fudenberg and Maskin, 1986). In contrast, the effectiveness of reputation in circumstances where players are essentially strangers, knowing about one another only through word-of-mouth, is far less certain. The issue is important because word-of-mouth mechanisms are typically less costly than formal institutional interventions such as legal contracts.

We report here on a laboratory investigation of the information about reputation necessary to support cooperative effort among strangers. In practice, only partial information about a stranger's reputation is typically available. Models based on subgame perfection imply that enforcing cooperation requires reputational information that is recursive in nature; one needs to know not only one's partner's past action but also one's partner's partners' past actions, and so on. Among strangers, and absent a formal tracking institution, available information is unlikely to be this extensive, implying a rather limited domain on which informal reputation systems can be effective. Models that relax the subgame perfection requirement, however, suggest that cooperation can be sustained on quite modest amounts of information; perhaps solely on information about a partner's recent past, information that is relatively easy to disseminate.

Our experiment involves the image scoring game (Nowak and Sigmund, 1998a). In each period, players are partnered and one given the chance to take a costly action that helps the other. Cooperating in this manner is socially efficient, but the only way to monitor free riding is through the image score (reputation) which, in this game comes to an accounting of a player's past helping actions. Our experiment excludes any direct reciprocal benefit from reputation building. Thus, our reputation mechanisms are systems of indirect reciprocity, in which actions taken with one group of partners are reciprocated by strangers.

The critical manipulation in the experiment has to do with the information available in the image score: we study no information, immediate past action (first-order information) and one step of recursive information (second-order information). We look at each information condition for two cost-to-benefit ratios for cooperation.

Our main result is that informal strictures that punish cheaters along with those who fail to punish cheaters generate substantial cooperation even when the cost of cooperating is relatively high. Before giving a detailed accounting of our results we review the relevant theory and lay out our experimental design and related experimental work.

2. The image scoring game and theories of indirect reciprocity

The image scoring game is constructed out of a series of asymmetric stage games; each stage features an interaction between two players who have never met before and who will never meet again. One player is randomly assigned the role of mover, and the other the role of receiver. The mover chooses to either keep or give. Keep pays the mover c and the receiver nothing, give pays the receiver b and the mover nothing, $b>c>0$. Prior to choosing, the mover observes the receiver's image score, the information available to the mover about the receiver's past actions. The efficient outcome of the game requires all movers to play give, in spite of the incentive to play keep. Consequently, the information content in the image score must be sufficient to support a system that, in the long run, rewards giving and punishes keeping.

Symmetric games, such as the prisoner's dilemma, involve direct reciprocal behaviors based on expectations of what the partner will do at the current stage (e.g., Croson, 2000; Clark and Sefton, 2001; Güth et al., 2003). The asymmetry of the image scoring game excludes direct reciprocal relationships, permitting a tight focus on the influence of reputation information.

Absent information about reputation, there is little reason to suppose that cooperation can be sustained: The image scoring game becomes essentially a set of one-shot encounters in which the mover has a dominant strategy to keep.[1] Models of norm observance in a series of one-shot interactions incorporating subgame perfection, however, show that indirect reciprocity can be sustained if a complete history of the game is available to all players, or if a (local) mechanism or (central) institution is available to process and to provide the necessary information honestly (Milgrom et al., 1990; Kandori, 1992; Okuno-Fujiwara and Postlewaite, 1995). However, to make cooperation the outcome of subgame perfect equilibria, reputation generally need convey both information about the partner's past actions and whether these actions were in line with the prevailing strictures; one needs to be able to judge whether a partner's past actions were justified given the past actions of their own partners, etc. Theoretically, the amount of recursive information necessary can be enormous, making it is unlikely that indirect reciprocal systems would be stable outside of special circumstances (ex., small, non-anonymous groups) that lessen the information gathering and processing demands.

One of the signature implications of backward induction is that cooperation cannot be sustained in equilibrium unless the stop time for the game is indefinite (at least for complete information models, a point we return to in the final section).[2] Our experiment is conducted with a finite, publicly known stop time, and we observe substantial cooperation across almost all treatments. Apparently the necessary conditions for cooperation laid down by backward induction are too strong. This finding is consistent with a large empirical literature on finitely repeated prisoners' dilemmas where cooperation has been shown to be quite robust (ex., Andreoni and Miller, 1993; Selten and Stöcker, 1986; Ledyard, 1995). For this reason, we consider models that assume people have a very limited ability to do backward induction. Nowak and Sigmund (1998a,b) exhibit a model in which there are only two types of players: those who play keep every time and those that discriminate on the basis of what their partner did the last time he or she was a mover. Under these circumstances, and in a finitely repeated game, discriminating is an

[1] In the infinite horizon version of the game, depending on the matching scheme there may also be a "contagious" strategy in which each player cooperates until they have evidence another has defected. So if one person defects, eventually everyone defects. While there are no defections along the equilibrium path, the fact that one defection collapses the system forever makes this scenario a very fragile one (Kandori, 1992; see also Ellison, 1994).

[2] With a definite stop, not cooperating is a dominant strategy in the last round and independent of the image score, and rolling back, cooperation unravels.

evolutionarily stable strategy. Hence the only information the image score need contain to sustain cooperation is the action a player took last time as a mover—what we will refer to as first-order information. Nowak and Sigmund also show that the strategy's success is sensitive to the cost-to-benefit ratio of giving, something our experimental design enables us to explore.[3]

Discriminating on the basis of a partner's last action alone is quite myopic. The threat to punish someone who plays keep is not in the interest of the punisher since it leads, by the discriminating strategy, to the punisher being punished when next in the role of receiver. A player might therefore deviate from discriminating, and play keep—assuming he is confident that the next mover he meets will also recognize that it is not in his interest to punish. Supposing this is so, cooperation eventually collapses. So consider adding a second-order of information; the receiver's image score includes not only what he did last time as a mover, but also what the receiver he faced did last time as a mover. For example, the image score might state that the receiver last played keep with a player who last played give. To see the effect, consider a mover who, for the first time, encounters a receiver who played keep on a giver. To support his punishment, keeping on a keeper would have to be rewarded, meaning that there needs to be giving to someone who gives to a keeper—which is not consistent with self-interest since keeping on a keeper pays more. So now players would have to think two steps ahead, and be confident others do so as well, before cooperation would unravel.

Our experiment investigates indirect reciprocity on human subjects with limited ability to think sufficiently ahead to do backward induction (studies analyzing and confirming this limitation include Johnson et al. (2002) and Selten and Stöcker (1986)). This allows us to formulate hypotheses with respect to the cooperation dynamics and the subjects' response to the reputation information. In particular, since all our games have a commonly known finite end, cooperation may emerge early in the games but should be close to zero at the final games regardless of the reputation information available to the subjects; in fact, this is what we find. Furthermore, in our context, a failure to do backward induction also suggests that information about a person's past action is sufficient to generate cooperative behavior. Comparing this case against a no information baseline game, this is what we find when the cost-to-benefit ratio is high, the situation where inducing cooperative behavior appears more difficult—although the amounts of cooperation generated are quite substantially less than 100%. It then seems plausible that, adding an additional layer of recursive information might increase the amount of cooperation, and this is the second hypothesis we investigate. Indeed, this is what we find for both cost–benefit ratios.

3. Design of the experiment

The design manipulates two factors: the amount of information available in the image score, and the cost of giving, c. We examine three information conditions—no

[3] Specifically, they show that discriminating is evolutionarily stable if the probability of observing an image score, q, is greater than the cost-to-benefit ratio, c/b. Lotem et al. (1999) demonstrate that discriminating in this manner is robust to a variety of factors.

information, first-order information and second-order information—each for two different costs of giving.

Subjects in the experiment are matched in pairs for each of 14 rounds. No two subjects are matched together more than once, and subjects are told this prior to play. So any reciprocal behavior is necessarily indirect, minimizing influences other than the reputation information. Partners interface with one another via computers. Identities are anonymous, making opportunities for direct reciprocity outside the experiment unlikely. (The computers are in three-sided cubicles and neither other subjects nor the experiment's monitor can watch someone make choices.) Upon being paired, one subject is selected mover and the other receiver. The mover chooses between two actions. He can keep, in which case he gets c, and the receiver gets nothing. Or the mover can give, in which case he gets nothing and the receiver gets b. Subjects know that they will be in each role for half the trials (seven times) and that for the most part they will rotate roles between rounds.[4]

The experiment has a 2×3 (six treatment) design. In half of the treatments, c is set high (US$0.75), and in half it is set low (US$0.25). The value of b is the same for all cases (US$1.25). There are three information conditions differentiated by what the image score contains. For *no information*, the mover chooses while knowing nothing about the receiver's history. For *first-order information*, the mover knows whether the receiver played give or keep when last in the role of mover. Nowak and Sigmund's benchmark for when cooperation is possible, $q>c/b$, implies we should see giving in our first-order treatments (q=1) but not in the no information treatments (q=0).

For *second-order information*, the mover knows not only if the receiver last played give or keep, but also whether the receiver last played give or keep with someone who last played give or keep. A mover can then condition his behavior according to whether a receiver who kept the money the last time he was a mover 'justifiably punished a keeper', or whether he 'exploited a giver.' As mentioned earlier, this information improves prospects for cooperation: if one cannot distinguish between justified and unjustified giving and keeping, the only way to sustain a positive image score is to give, independent of the image score of the game partner.

In all games, it was publicly stated that the number of rounds to be played is 14. Consequently, if we expect backward induction to prevail, we cannot expect to see any giving, regardless of the information condition, since all experimental games end with probability one.[5] This is clearly rejected by the data, as we will see in the next section. If on the other hand the players' capability to think ahead is limited, we would expect more giving in the early rounds of the game than in the later rounds.

Each treatment consisted of two sessions, with 16 subjects per session (32 per treatment) for a total of 192 participants. Subjects were Penn State University students,

[4] Specifically, subjects were told that they might, on occasion, have the same role for two rounds running (the exact instructions can be found along with data sets at http://lema.smeal.psu.edu/lema under the data archive link). Given the capacity of the lab available to us at the time, it is not possible to have strict rotation while at the same time observing the one-play per pair rule.

[5] In fact, all the models mentioned above that are based on backward induction are cousins of standard Folk theorems (e.g., Fudenberg and Maskin, 1986), meaning that they are dealing with infinitely repeated matching games.

mostly undergraduates, recruited by fliers posted around campus. Cash was the only incentive. Subjects were paid their total individual earnings from the game plus a US$7 show-up fee.

Two independently conducted experiments demonstrate that cooperation can arise in image scoring games.[6] Wedekind and Milinski (2000) demonstrate that cooperation can arise when movers know the sum number of times the receiver has given and kept in the past, and that movers are more likely to give to those that are generous. Seinen and Schram (2001) present their subjects with similar information and also compare to the no information case, sometimes varying the cost of giving. It is not clear, however, to what extent their results are attributable to indirect reciprocity. In principle, cooperation could be explained by the kind of Folk theorem results associated with direct reciprocity.[7] A recent follow-up study by Engelmann and Fischbacher (2002) attempts to disentangle the relative importance of strategic and non-strategic motives to give in image scoring games, by allowing only half of the subjects to build a reputation at any time. Their data give additional evidence that strategic considerations drive giving, though they also find non-trivial non-strategic indirect reciprocity. Unlike these experiments, we focus on the effects of differing types of (recursive) information. We examine this issue in the context of the cost-to-benefit ratio, in a fully crossed design.

4. The impact of cost, information, and future horizon on the cooperation level

Fig. 1 compares giving rates by cost. Giving is higher in all low cost treatments than in any of the high cost treatments. Recall that each treatment involves two separate sessions (shown aggregated in Fig. 1). For all three comparisons, the giving in either of the low cost sessions was higher than in either of the high cost sessions. The probability of this happening by chance is quite small, p=0.005.

Fig. 1 also demonstrates that the willingness to give tends to be lower in the last rounds. In particular, the lowest cooperation rate in 10 of the 12 separate sessions occurred in the last round.[8] Finally, we also see from Fig. 1 that there is substantial giving even in the no information treatments, and so it is immediately apparent that not all of the giving is attributable to image scoring. We return to this observation below.

[6] A related line of empirical research deals with the effect of reputation in eBay on price and probability of transaction (Lucking-Reiley et al., 1999; Resnick and Zeckhauser, 2002; Bolton et al., in press (a,b)).

[7] In their experiment, subjects interacted repeatedly with the same opponents (in random order). To offset this, subjects were led to believe that the group size was twice the actual size. So the extent to which the results reflect indirect reciprocity depends to a large degree on what subjects believed about the rematch probability.

[8] Within each cost treatment, last round cooperation rates do not differ significantly. One might think that, when subjects are repeatedly exposed to the game and learn that the endgame effect is coming, unraveling would begin. However, sequential equilibrium theory suggests that this need not be the case (ex., Kreps et al., 1982). Furthermore, empirically, Andreoni and Miller (1993) found in a finitely repeated prisoner's dilemma supergame, that with experience people wait longer before their first defection, while Selten and Stöcker (1986) found some unraveling across games. The latter authors discuss the possibility that behavior will not fully unravel, both from the perspective of a simple learning model and from their data. Indeed, to our knowledge there are no experiments on finitely repeated dilemma games in which cooperation is completely eliminated with game repetition.

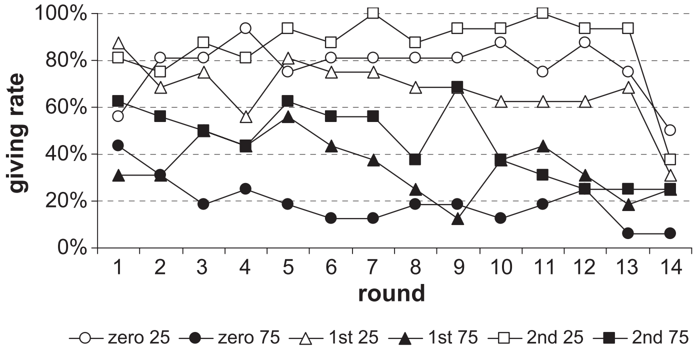

Fig. 1. Comparison of giving by cost.

Fig. 2 compares giving rates by information level. That adding information increases the rate of giving is unambiguously so for high cost treatments. The rate of giving by session rank exactly as predicted: 2nd–2nd–1st–1st–zero–zero. The probability of this by chance is 0.011. For the low cost treatments, the second-order information sessions still rank highest, but the evidence is somewhat ambiguous for the ranking of first-order and no information sessions: 2nd–2nd–zero–1st–zero–1st. The probability of at least this extreme an outcome by chance is 0.100; here 'at least this extreme' includes all outcomes exactly consistent with the image scoring prediction plus all those where no more than two of the six observations are out of place. Overall, however, the null hypothesis that information has no effect in the experiment is rejected in favor of the hypothesis that more information increases giving, p=0.003.[9]

In sum, all three strategic factors we examine—cost, information and the future horizon—influence giving in the direction anticipated by the indirect reciprocity hypothesis.

5. How history affects cooperation

Table 1 tabulates the effects of the receiver's last action as a mover, and of the mover's last experience as a receiver, on the mover's probability to give. The averages in the last column show that the tendency to give is strongly influenced by both how the receiver treated others in the past: The giving rate is 71.1% when the receiver chose give the last time he was a mover, and only 37.4% when he chose keep (t-test $p<0.01$). Conditioning behavior on the recipients' history is, of course, anticipated by theories of indirect reciprocity via image scoring. But Table 1 also shows that the tendency to give is

[9] Even though our null hypothesis concerning the effect of information can be statistically rejected, we caution that when costs are low our results do not strongly confirm that first-order information affects giving compared to no information. Two hypotheses suggest themselves. First, low costs of giving may create more indifference and thus more noise in the behavior, so that more independent observations are needed to confirm the information effect. Second, and related, if giving is cheap, self-sacrificing norm observance—regardless of the image score of the receiver—may become relatively more important.

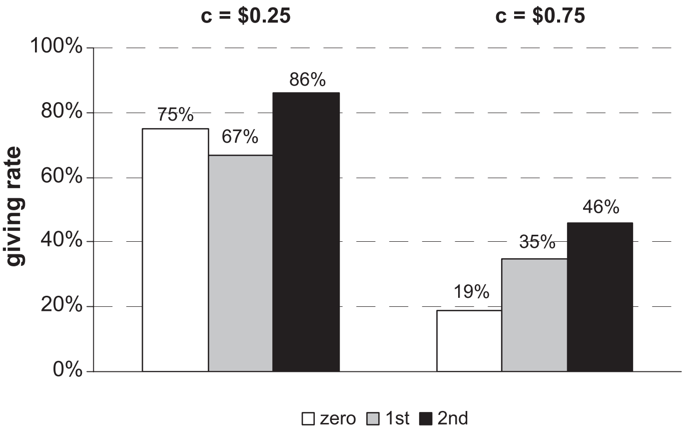

Fig. 2. Comparison of giving by amount of information.

influenced by how the mover was treated before by others: The giving rate is 71.0% if the mover received a gift the last time he was a receiver and only 32.9% otherwise (t-test $p<0.01$).[10]

Conditioning behavior on one's own history might be interpreted as evidence for 'contagious' equilibrium principles (these call for a player to cooperate until deviation is observed) or perhaps a preference for fairness (Fehr and Schmidt, 1999; Bolton and Ockenfels, 2000). Both rationalizations predict that cooperation is partly triggered by observations unrelated to the past behavior of the receiver. Such giving may also explain cooperation even in environments in which information about image scores is not available.[11] In fact, the overall rate of giving for the no information–low cost treatment (75%) is higher than for any of the high cost treatments. And while the rate for no information–high cost is quite a bit lower, it is not negligible (19%) (see Fig. 2).

Turning back to the impact of the image score, there is a remarkable difference between the first-order and second-order information condition: while the probabilities of giving to a receiver with a 'keep-history' are almost identical across these information conditions and for each cost condition, second-order information appears to provoke much more giving to receivers with 'give-histories' than first-order information does. Averaged over both cost treatments, giving increases the probability of receiving a reciprocal gift by 54% (from 39% to 60%) in the first-order condition, but

[10] Testing for each cost condition separately within each information condition yields analogous conclusions. Also, while the probability of giving after having received a gift almost triples in the no information condition, it less than doubles in the other information conditions suggesting that movers rely less on their own history when the receiver's image score is made available (see Table 1). Due to space constraints, we dropped a probit analysis that account for all available history information (not just the last round), and better controls for round effects, individual heterogeneity, etc. The results, available from the authors upon request, imply the findings are robust.

[11] This sort of giving is reminiscent of data from the dictator game, in which one player decides how to split a monetary pie between self and one other. Numerous studies find that many dictators give some money (e.g., Roth, 1995; Andreoni and Miller, 2002). This, with data from a variety of games studied in the economics and psychology literature, suggests that cooperation is based, at least in part, on considerations of fairness.

Table 1
The impact of receiver and mover histories on the probability of giving (%)

Information condition		No			First order			Second order			Row
Cost condition		Low		High	Low		High	Low		High	Avg.
Receiver's last move	keep				58.3		30.3	57.1		31.0	
(first-order info)	avg.					39.0			35.0		37.4
	give				68.9		43.4	89.8		60.0	
	avg.					60.3			79.8		71.1
Receiver's last move	keep/keep							75.0		32.0	36.3
(second-order info)	keep/keep							50.0		28.6	33.0
	give/keep							80.0		48.6	58.0
	give/give							91.2		66.0	85.2
Mover last received	keep	73.3		12.0	54.1		28.8	89.5		31.5	
	avg.		25.0			36.8			40.2		32.9
	give	81.0		39.0	70.8		46.1	86.2		58.0	
	avg.		72.6			62.3			76.5		71.0

The 'receiver's last move' is keep (give) if he chose keep (give) the last time he was a mover. It is keep/give if he did not give to a giver the last time he was a mover. The 'mover last received' is keep (give) if he received zero (a gift) the last time he was a receiver. All numbers include only observations where the respective history information was available.

by more than twice this amount, 128% (from 35% to 80%), in the second-order information condition.

It appears then that players make use of recursive information in a way that stabilizes indirect reciprocal giving. In fact, Table 1 shows that movers distinguish whether the recipient's last partner was a giver or a keeper. In particular, movers are more likely to give to givers when the receivers' giving was 'justifiable': giving rates are 91% (80%) for give/give in contrast to 66% (49%) for give/keep in the low (high) cost condition (t-test p=0.08 and p=0.06, respectively). And movers are somewhat more likely to give to keepers when keeping was 'justifiable', though this effect is not statistically significant at the 10% level: giving probabilities are 75% (32%) for keep/keep in contrast to 50% (29%) for keep/give in the low (high) cost condition (t-test p=0.15 and p=0.41, respectively). This suggests that, in line with the theoretical arguments outlined in Section 2, it is the opportunity to justify a receiver's action in the second-order information treatment that promotes giving beyond cooperation rates in the first-order information conditions.

6. Conclusions

Consistent with both backward induction and action discrimination models, reputation information, the cost-to-benefit ratio for cooperating, and the length of the game horizon, all play important roles in the decision of strangers to cooperate. Where the data differs with both models is with regard to the information necessary to sustain cooperation. Compared to no information, information about a player's immediate past action increases cooperation when the cost of cooperation is high, and somewhat decreases cooperation when it is low. Yet, adding one layer of recursive information (information about a partner's last partner's behavior) is unambiguously more effective than no or only first-order information—regardless of the cost-to-benefit ratio. In short, backward induction models overstate the information necessary to sustain cooperation, whereas action discrimination models overstate the effectiveness of action information alone.

We find that informal strictures that punish both cheaters and those who fail to punish cheaters are sufficient to generate substantial efficiency gains, even when the benefits from cooperation are relatively modest.[12] In this sense, reputation information mechanisms take advantage of a naturally occurring willingness to cooperate—conditional on a good, if incomplete, reputation—without relying on more costly institutions such as contracts.

Greif (1989) presents a well documented historical example of an essentially second-order information system: A group of eleventh-century Mediterranean traders, the Maghribi, commonly used overseas agents to complete trades. The moral hazard problem implicit in this arrangement was controlled through informal reputation. Essentially, traders wrote one another letters evaluating the competence and efficiency of the various agents (these letters, in large volume, have been passed down). Grief describes the system this way: "[A]ll coalition merchants agree never to employ an agent that cheated while operating for a coalition member. Furthermore, if an agent who was caught cheating operates as a merchant, coalition agents who cheated in their dealings with him will not be considered by other coalition members to have cheated." Our data suggests that the stability of this system is not narrowly historically or culturally specific.

A modern example of the use of reputation to enforce trust in a market setting is eBay auction markets. The market is geographically disperse; recourse to legal remedies, if at all available, is expensive. For each transaction, buyers and sellers are permitted to post evaluations of one another's trustworthiness. The eBay reputation system differs in several ways from the one we study here, but as in our information treatments, the reputation signal is noisy, and seller and buyer sometimes blame one another for a transaction that has gone bad. To date, there are but a few studies of the effectiveness of this system (see Dellarocas, in press, for a survey), but the volume of trade is large and eBay is profitable, suggesting that the system is effective to a substantial degree. This would seem to buttress our finding that a noisy signal can be enough to support a meaningful level of cooperation.

Finally, we suspect that our results could be captured by an incomplete information model that supposes some people are more prone to cooperate than others.[13] Partial information about reputation then acts as a (possibly noisy) signal of a player's type. Such a model would have some of the strategic characteristics apparent in our data (in particular,

[12] Overall, players earned on average 2.65 times more than their payoff with no indirect reciprocity (when movers always play keep), although a minority of about 12% earned less than the payoff they received when nobody gives. On the individual level, there is, in all treatments, a negative correlation between numbers of gifts given and total payoffs. The losses are not particularly high, however. Averaging across all treatments and all levels of giving, an additional gift costs about 9 cents as opposed to an average of 50 cents when gifts are not reciprocated. This loss plausibly stems from the individual tendency to reciprocate gifts received (independent of receiver image scoring).

[13] A different approach would be to theoretically investigate reputation building as a function of the information available in finite games assuming incomplete information about preferences following the research agenda outlined by Wilson (1985); see Bolton and Ockenfels (2004) for a study along these lines.

incomplete information models predict cooperation can arise even in finite horizon games), and might also begin to explain the importance of a player's own history in his or her decision to cooperate, since one's own history is potentially informative of the proportion of cooperatives types in the population.

Acknowledgements

The authors thank Steffen Huck, Jim Andreoni, John List, an anonymous referee, and seminar participants at Berlin, Bielefeld, Bonn, Chicago, Harvard, Jena, and Magdeburg for helpful comments, and Pamela Sodini for assistance in running the experiments. Bolton and Katok gratefully acknowledge the support from the U.S. National Science Foundation. Ockenfels gratefully acknowledges the support of the Deutsche Forschungsgemeinschaft. Special thanks from all the authors for the warm hospitality of the Making Choices project, Center for Interdisciplinary Research, University of Bielefeld, summer of 2000.

References

Andreoni, J., Miller, J.H., 1993. Rational cooperation in the finitely repeated prisoner's dilemma: experimental evidence. Economic Journal 103, 570–585.

Andreoni, J., Miller, J.H., 2002. Giving according to GARP: an experimental test of the consistency of preferences for altruism. Econometrica 70, 737–753.

Bolton, G.E, Ockenfels, A., 2000. ERC: a theory of equity, reciprocity and competition. American Economic Review 90, 166–193.

Bolton, G.E, Ockenfels, A., 2004. Reputation, information and matching: theory and evidence. Mimeo.

Bolton, G.E, Katok, E., Ockenfels, A., in press (a). How Effective are Online Reputation Mechanisms? An Experimental Investigation. Management Science.

Bolton, G.E, Katok, E., Ockenfels, A., in press (b). Bridging the Trust Gap in Electronic Markets. A Strategic Framework for Empirical Study. In: Akcali, J. Geunes, P.M. Pardalos, H.E. Romeijn and Z.J. Shen (Eds.), Applications of Supply Chain Management and E-Commerce Research in Industry. Kluwer Academic Publishers.

Clark, K., Sefton, M., 2001. The sequential prisoner's dilemma: evidence on reciprocation. Economic Journal 111, 51–68.

Croson, R., 2000. Thinking like a game theorist: factors affecting the frequency of equilibrium play. Journal of Economic Behavior and Organization 41, 299–314.

Dellarocas, C., 2003. The Digitization of Word-of-Mouth: Promise and Challenges of Online Reputation Mechanisms. Management Science, vol. 49 (10), pp. 1407–1424.

Ellison, G., 1994. Cooperation in the prisoner's dilemma with anonymous random matching. Review of Economic Studies 61, 567–588.

Engelmann, D., Fischbacher, U., 2002. Indirect reciprocity and strategic reputation building in an experimental helping game. Working Paper, Humboldt University, Berlin.

Fehr, E., Schmidt, K., 1999. A theory of fairness competition, and cooperation. Quarterly Journal of Economics 114, 817–868.

Fudenberg, D., Maskin, E., 1986. The folk theorem in repeated games with discounting or with incomplete information. Econometrica 50, 533–554.

Greif, A., 1989. Reputation and coalitions in medieval trade: evidence on the Maghribi traders. Journal of Economic History 49, 857–882.

Güth, W., Kliemt, H., Ockenfels, A., 2003. Fairness versus efficiency: an experimental study of (mutual) gift giving. Journal of Economic Behavior and Organization 50, 465–475.

Johnson, E.J., Camerer, C., Sen, S., Rymon, T., 2002. Detecting failures of backward induction: monitoring information search in sequential bargaining. Journal of Economic Theory 104 (1), 16–47.

Kandori, M., 1992. Social norms and community enforcement. Review of Economic Studies 59, 63–80.

Kreps, D.M., Milgrom, P., Roberts, J., Wilson, R., 1982. Rational cooperation in the finitely repeated prisoners' dilemma. Journal of Economic Theory 27, 245–252.

Ledyard, J., 1995. Public goods: a survey of experimental research. In: Kagel, H., Roth, A.E. (Eds.), Handbook of Experimental Economics. Princeton University, Princeton, NJ, pp. 111–194.

Lucking-Reiley, D., Bryan, D., Prasad, N., Reeves, D., 1999. Pennies from eBay: the Determinants of Price in Online Auctions. Working paper, Vanderbilt University.

Lotem, A., Fishman, M.A., Stone, L., 1999. Evolution of cooperation between individuals. Nature 400, 226–227.

Milgrom, P., North, D., Weingast, B., 1990. The role of institutions in the revival of trade: the law merchant, private judges, and the champagne fairs. Economics and Politics 2, 1–23.

Nowak, M.A., Sigmund, K., 1998a. Evolution of indirect reciprocity by image scoring. Nature 393, 573–577.

Nowak, M.A., Sigmund, K., 1998b. The dynamics of indirect reciprocity. Journal of Theoretical Biology 194, 561–574.

Okuno-Fujiwara, M., Postlewaite, A., 1995. Social norms and random matching games. Games and Economic Behavior 9, 79–109.

Resnick, P., Zeckhauser, R., 2002. Trust among strangers in internet transactions: empirical analysis of eBay's reputation system. In: Baye, M.R. (Ed.), The economics of the internet and e-commerce, Advances in Applied Microeconomics, vol. 11. Amsterdam, pp. 127–157.

Roth, A.E., 1995. Bargaining experiments. In: Kagel, H., Roth, A.E. (Eds.), Handbook of Experimental Economics. Princeton University, Princeton, NJ, pp. 253–348.

Seinen, I. and Schram, A., 2001. Social Status and Group Norms: Indirect Reciprocity in a Helping Experiment. Working paper.

Selten, R., Stöcker, R., 1986. End behavior in sequences of finite prisoner's dilemma supergames. Journal of Economic Behavior and Organization 7, 47–70.

Wedekind, C., Milinski, M., 2000. Cooperation through image scoring in humans. Science 288, 850–852.

Wilson, R., 1985. Reputations in games and markets. In: Roth, A. (Ed.), Game-Theoretic Models of Bargaining. Cambridge University Press, pp. 27–62.