

Minimally Acceptable Altruism and the Ultimatum Game

Julio J. Rotemberg*

April 4, 2004

Abstract

I suppose that people react with anger when others show themselves not to be minimally altruistic. I show that, with heterogeneous agents, this can account for the experimental results of ultimatum and dictator games. Moreover, it accounts for the surprisingly large fraction of individuals who offer an even split with parameter values that are more plausible than those that are required to explain outcomes in these experiments with the models of Levine (1998), Fehr and Schmidt (1999), Dickinson (2000) and Bolton and Ockenfels (2000).

*Harvard Business School, Soldiers Field, Boston, MA 02163, jrotemberg@hbs.edu. I wish to thank Rafael Di Tella and David Scharfstein for helpful conversations and the Harvard Business School Division of Research for research support.

This paper presents a model of individual preferences where individuals are mildly altruistic towards others and where they expect others to be mildly altruistic as well. When a person encounters evidence that another is less altruistic than he finds acceptable, he becomes angry and derives pleasure from harming the excessively selfish individual. The paper shows that these preferences can explain the experimental outcomes of the ultimatum game of Güth *et. al.* (1982) as well as those of an important variant, namely the dictator game of Forsythe *et. al.* (1994). Unlike models of social preferences that have been proposed in the past to explain these outcomes, the current model does not require extreme or unrealistic parameter values.

Because the basic results of these experiments are in such sharp conflict with the predictions of standard equilibrium models based on standard preferences, many experimental variations have been considered and this experimental literature is vast. Still, it is worth recalling the setting and some of the striking findings. Both games involve two players. The first player, who is called the proposer, offers to split a pie with the second, who is called the responder. In the ultimatum game, the responder can either accept (so that the pie is split in the way suggested by the proposer) or reject the offer, so that both players get nothing. In the dictator game, the responder must passively accept the proposer's offer.

With selfish preferences, standard game-theoretic equilibrium reasoning implies that the proposer ought to make a trivial (or zero) offer in the ultimatum game and a zero offer in the dictator game. In practice, the modal offer in the ultimatum game is to split the pie 50-50. The actual fraction of even splits varies somewhat from experiment to experiment, and sometimes across rounds of certain experiments. Using data from the late rounds of the experiments in Roth *et al.* (1991), for example, Levine (1998) computes that even splits constitute about 28% of offers.

Lower offers are rejected relatively often. Indeed, such rejections are sufficiently common that average earnings of proposers actually decline when they make lower offers. This is shown both in Figure 6 of Roth *et al.* (1991) and in Harrison and McCabe (1996), who show this finding to be quite robust. In the dictator game, expected monetary payoffs obviously

rise when lower offers are made. Not surprisingly, this implies that offers of even splits are observed less frequently. Still Forsythe *et. al.* (1994) report that about 20% of their proposers in the dictator game offered even splits.

The model I present is closely related to Levine (1998) who also supposes that agents' altruism for others depends on their assessment of how altruistic others are in return. However, the exact form of this dependence is different in the two papers. This allows me to avoid Levine's (1998) conclusion that the experimental outcomes in ultimatum games require that most people be spiteful. For his baseline parameters, 53% of the population is willing to give up more than 25 cents to ensure that a stranger loses a dollar. A further 20 % of the population stand ready to give up over 95 cents to bring about this outcome. In his model, this spitefulness ensures both that responders reject uneven offers and that proposers make uneven offers even though these lead to lower expected earnings.

One difficulty with these preferences is that, using variants of dictator games, Charness and Rabin (1999) show that most people would sacrifice some of their own resources to induce small gains by others whereas essentially no one is willing to sacrifice significantly to hurt others. They show, in particular, that 73% of their subjects are willing to give up 100 of their units to cause another agent to gain 400. Not only would all the agents in Levine (1998) turn this down (because even his altruists are not sufficiently altruistic for this) but 20% would also prefer giving nothing to either themselves and someone else to receiving 20 when an anonymous player receives 80. In the Charness and Rabin (1999) experiments, almost all the players choose the latter.¹

Unlike Levine (1998), other theoretical models of fairness do not (at least explicitly) make each agent's preferences depend on the parameters of other agents' utility functions. Instead, they focus on the interdependence that results when each agent cares about the payoffs to several agents. There are two related ways in which this has been done. In the literature

¹As I discuss below, the generous behavior observed in any dictator game, including those of Charness and Rabin (1999) could be due to fear of being found out and punished. However, relatively large punishments are required for the generosity observed by Charness and Rabin (1999) to be consistent with the spitefulness assumed by Levine (1998).

pioneered by Rabin (1993), agents care about how the payoffs of each agent compare to the payoffs these agents would have received if the players had taken different actions. In particular, Rabin (1993) supposes that an agent regards another as fair if he expects his actions to be “kind”, where kindness of the second player towards the first is defined to equal the difference between the payoffs that the first expects to receive from the second and the payoffs that it would have been “equitable” for the second to give to the first. He then defines equitable payoffs as the average of the highest and the lowest payoffs the second agent can give to the first under the assumption that the player acts efficiently. Lastly, he supposes that agents maximize the sum of their own material payoffs and the product of their own kindness times the kindness they expect to receive from the other agent.

This model has been extended by both Dickinson (2000) and Falk and Fischbacher (2000) to account for the results of ultimatum games, which is sequential unlike the game considered in Rabin (1993). One important modification in Dickinson (2000) and Falk and Fischbacher (2000) is that they allow an agent’s utility function to put lower weight on the agents own material payoffs than on ‘reciprocity’ *i.e.* on the product of the agent’s own kindness with the kindness they receive from the other player. It turns out that, in both papers, offers of even splits can be obtained as equilibria but only when agents care exclusively about reciprocity and ignore their own material payoffs. For obvious reasons, this limit is not particularly attractive as an empirical description of preferences.

A different way of letting agents care about the payoffs of other agents is to let them care about both about how their own payoff compares to that of others. This has been the approach taken in the widely cited work of Fehr and Schmidt (1999) and Bolton and Ockenfels (2000). These paper rationalize high offers in the ultimatum game by the unwillingness of responders to accept offers that give them payoffs that are low relative to those received by the proposer. Taken literally, these models imply that these rejections do not hinge on the actions to which the proposer had access, and depend only on the responder’s views of different allocations of resources. Falk and Fischbacher (2000) and Charness and Rabin (1999) provide some evidence against this hypothesis by showing that responders reactions

do in fact depend on to the initial choices available to the proposer. Even leaving this aside, these models seem incapable of explaining the outcomes in ultimatum games except under under implausible assumptions concerning parameters. As I show in section 2, the problems are very similar to those of Levine (1998) and the extensions of the Rabin (1993) model when these are used to explain these experiments. In effect, these models can predict even splits only when responders are unrealistically are mean or proposers are unrealistically kind hearted (or both).

I suppose instead that, unless they feel mistreated, most people are modestly altruistic towards those around them. They would be willing to incur small costs if this led to large gains in another person's well-being. Many individuals would, for example, offer some help to individuals in dire need if they felt that no one else is available to provide this help. At the same time, unlike what is required by the Fehr-Schmidt (1999) model to explain equal splits, essentially everyone values an additional dollar in their own pocket more highly than an additional dollar in the pocket of someone with a similar position in society than their own.

The main difference between the preferences I consider and Levine's (1998) proposal is that rather than letting the altruism of each agent depend linearly on his perception of the other's altruism, I let the response be highly nonlinear. I suppose in particular, that a reduction in a person's altruism triggers a much larger reduction in the altruism of those he interacts if his initial level of altruism is low. This captures the idea that people do not expect large levels of altruism from strangers (though family and friends are a different matter). Thus, reductions from a high to a more normal level of altruism do not have an appreciable effect on the behavior of casual acquaintances. By contrast, there is a minimal level of altruism that people do expect from those that they interact with, and reductions below this level are met with disapproval, and even anger. This fits with Pillutla and Murnighan (1996), who provide evidence that rejections in ultimatum games are indeed associated with anger.

One also observes rapid angry responses in the field when people feel mistreated by

strangers. These emotions seem particularly salient in “road rage”, the sometimes violent reaction of drivers who feel that other drivers have acted badly. As shown in the questionnaire responses of Parker *et. al.* (2002), road rage arises in a variety of settings, though some triggers are more reliable than others. When “someone cuts in and takes the parking spot you have been waiting for,” only about 26% of the respondents in Parker *et. al.* (2002) said they would not react at all.

Strong reactions of this sort lead proposers to try to signal that their altruism is not below par, and this can lead even mildly altruistic individuals to offer even splits. It would seem more difficult to rationalize the observation that proposers receive lower expected earnings when they make lower offers. One might imagine that, even if proposers are confused enough to make such offers at first, they ought to learn to stop making them. Nonetheless Levine (1998) shows that the proposers in Roth *et. al.* (1991) still make a large number of low offers in the 10th round and have low expected returns even then. These observations also raise questions about the even-handedness of responders, since it would appear that these are punishing proposers whose actions are simple mistakes.

It turns out, however, that both the making of low offers and the recurrent rejections of such offers can easily be rationalized once one takes into account how agents in other experimental settings react to risk. Whatever its significance for larger issues, agents appear to be risk loving in the case of small gambles offered in such settings. Conlisk (1993), for example, reports that when students are offered a choice between one dollar for sure and a fair coin toss that determines whether they receive 2 dollars or nothing, 63% prefer the coin toss. For a proposer, offers that are less favorable to the responder than an even split constitute a gamble in the sense that such offers are rejected with positive probability. Risk loving proposers with relatively low altruism would thus accept this gamble even if its expected return was slightly below the one obtained from offering an even split.

This still leaves one with the question of whether people who react to “bad acts” do so because of what these acts reveal about their perpetrators, as I implicitly assume, or whether their reaction is independent of the identity of those committing these acts. This question

seems difficult to settle empirically because, for a given pattern of rewards and punishments, people who carry out willful bad acts must differ in some way from people who do not. One small bit of evidence that the utility function parameters of perpetrators play a role in determining punishment is provided by the criminal justice system, which is admittedly quite far removed from the private justice realm that I study. Nonetheless, it seems relevant that sentencing for criminal violations in the United States as well as in other common law countries is affected by judge’s assessments of the personal character of convicted criminals.²

The paper proceeds as follows. In the next section, I briefly review the Fehr and Schmidt (1999) model and discuss outcome-based preferences more generally. I introduce the preferences I propose in section 3. To explain the fact that experimental outcomes are not all identical, I let agents differ both in their level of altruism and in the minimum altruism that they find acceptable. With this heterogeneity, the ultimatum game becomes a signaling game where proposers signal their altruism. As is common with signaling games there are many equilibria though I focus on the one where all possible offers, including the 50-50 division of the pie, are observed. Section 4 is devoted to the dictator game and Section 5 concludes.

1 Outcome-based Preferences

In this section, I briefly consider preferences where each individual cares both about their own allocation and about the way this allocation compares with that received by the other player. I focus in particular on the Fehr and Schmidt (1999) model where player i ’s utility is

$$x_i - \alpha \max(0, x_j - x_i) - \beta \max(0, x_i - x_j) \quad j \neq i. \quad (1)$$

In this expression, x_i is the level of resources in the hands of agent i while α and β are positive parameters. This utility function incorporates the idea that agents wish to hurt those whose income exceeds their own while they are altruistic towards those whose income is below their own. In an attempt to explain ultimatum outcomes without “fairness”, Burnell, Evans

²See Hill (1998) for evidence that this plays a role in the US as well as an argument that it should play an even more significant role in sentencing.

and Yao (1999) use exactly the same preferences except that they set $\beta = 0$. Bolton and Ockenfels (2000) consider a differentiable utility function which, like (1) imposes a penalty for differences between x_i and x_j .

Now suppose that the resources available to be split in the ultimatum game equal A . Denote the proposer's offer by y and suppose that, if the responder accepts the offer, he gets y while the proposer keeps $A - y$. If the players suppose that they have the same resources outside the experimental setting, one can ignore these outside resources, so that the x 's in (1) correspond to the amount that the agents receive in the game. It then follows from maximizing (1) and related utility functions that responders only accept offers in which y is above some threshold because this is needed to keep the x of the responder from being too far below the x of the proponent. In the case of (1) offers are accepted only if $y > \alpha A / (1 + 2\alpha)$. This gives an incentive to proposers to raise their offers above zero. An obvious question, however, is whether this incentive is sufficient to rationalize offers of $A/2$.

To consider this suppose first that $\beta = 0$ so that the proposer is selfish. If α is common knowledge, the proposer then offers the minimum acceptable offer of $\alpha A / (1 + 2\alpha)$, and this is strictly below $A/2$ unless α is infinite. An infinite α is absurd however, since it implies that individuals are willing to pay any price to achieve a one dollar reduction in the income of a single individual who has more than they do. This leads Fehr and Schmidt (1999) to suppose that α is finite and to use β to rationalize even splits.

For any $\beta < 1/2$, proposers prefer a dollar in their pocket to a dollar in the pocket of the responder. This means that increasing β above zero has no effect on equilibrium offers until $\beta = 1/2$. At that point, the proposer is indifferent to the allocation of A between the two parties so that even splits are possible, though there is no particular reason to suppose that they would emerge in equilibrium. Thus, Fehr and Schmidt (1999) rationalize even splits by supposing that $\beta > 1/2$ so that individuals strictly prefer a dollar in the pocket of someone who has a lower income than themselves to a dollar in their own pocket. Such a high level of altruism seems unreasonable.

Their model has the further disadvantage that individuals with this degree of altruism

would also offer even splits in the dictator game. Thus, Fehr and Schmidt (1999) predicts an equal proportion of even split offers in the ultimatum and dictator games, which is contrary to observation.³ This problem is avoided by the related model of Bolton and Ockenfels (2000).

Bolton and Ockenfels (2000) suppose that responders are heterogeneous and this has the obvious advantage of rationalizing the fact that offers of less than $A/2$ are sometimes accepted and sometimes rejected. Rather than using the piecewise linear utility function in (1), they consider a differentiable one. Nonetheless, each responder has a minimum level of y , which they call r , such that offers smaller than r are rejected. Since their theory is also one where rejections occur only to avoid differences in income, the parameter r can be interpreted much like α in (1). In particular, it can be interpreted as saying that a respondent is willing to incur a cost of up to r to eliminate a difference in income between the two agents of $A - 2r$. This means that, at least locally, he is willing to pay $\frac{1}{A/r-2}$ for each dollar by which he reduces the income difference between the proponent and himself.

One advantage of the heterogeneity in r is that it is not necessary for all respondents to have $r = A/2$, which would imply that respondents are willing to pay a dollar for an essentially trivial reduction in the proponent's income. To see this, suppose that the distribution of r 's among responders can be described by the distribution function $G(r)$. Suppose further that, if the proposal is accepted, the proposer's utility is $A - y + \lambda y$ so that the proposer is altruistic towards the respondent with altruism parameter λ . Then, the proposer's utility from making an offer $y \leq A/2$ is

$$G(y) (A - y(1 - \lambda)) \tag{2}$$

³Interestingly, Fehr and Schmidt (1999) contains a discussion of the implications for their setup for dictator games. What they argue is that utility functions that are not piecewise linear, but are curved instead, would yield offers in dictator games that are strictly between 0 and $A/2$. Curvature in the utility can only ensure that people who do not make offers of $A/2$ in ultimatum games make interior offers in dictator games. The right degree of curvature implies that proposers have increased marginal utility from a dollar in the respondents pocket as offers diminish in size, and thereby guarantees that the respondent gets more than zero. Interior offers are indeed observed in dictator experiments. However, with this curvature or without, the fraction of proposers that offer an equal split in the ultimatum game should be the same as the fraction that does so in the dictator game, and this is manifestly counterfactual.

An even split then satisfies the first order conditions for a maximum if

$$-G(A/2)(1 - \lambda) + G'(A/2)(1 + \lambda)\frac{A}{2} = 0 \quad (3)$$

For $\lambda < 1$, which is the reasonable case, this still requires that $G'(A/2) > 0$ so that there is positive density at an extreme level of meanness. If one is willing to suppose that G' is constant near $A/2$ (which is valid to first order) one can also use this equation to determine the fraction of people whose r must exceed any given threshold \bar{r} . This proportion is equal to

$$P(\bar{r}) = \left(\frac{A}{2} - \bar{r}\right) G' = \left(\frac{A}{2} - \bar{r}\right) \frac{2(1 - \lambda)}{A(1 + \lambda)} = \frac{(1 - \lambda)(1 - 2\bar{r}/A)}{(1 + \lambda)} \quad (4)$$

where the second equality is obtained from (3) taking into account that $G(A/2) = 1$ while the third equality is obtained by simplifying the earlier expression.

Now consider the responders whose r exceeds $2/3$ of $A/2$ so that the loss of a dollar reduces their utility by the same amount as having someone else gain a dollar relative to them. Interpreted broadly, this means that these people are willing to lose a dollar as long as they can cause two dollars of damage to someone richer than themselves, which seems like a degree of mean spiritedness that would lead to a great deal of malfeasance. Using $\bar{r} = A/3$ in (4) implies that $P(\bar{r}) = 1/3$ so that one third of responders must have a higher r for a selfish person to offer an even split. Even if one only requires people with $\lambda = .5$, who are quite altruistic, to offer an even split, it is still necessary for $1/9$ of responders to have an r larger than $A/3$.

As discussed above Charness and Rabin (2000) do not find their participants willing to incur even much lower costs to hurt people whose income is higher. It is important to stress that I am not suggesting that it is possible to explain rejections of ultimatum offers with smaller degrees of ill-will towards proposers. The difficulty, it seems to me, is with modelling this ill-will as emerging simply out of inequity.

What I do below is to allow this degree of ill-will to arise endogenously as a reaction to low offers. It is worth stressing, however, that simply stating that agent's utilities depend on other's intentions is not enough. The Levine (1998) model does this, but still requires

a large degree of "background" ill-will. That is, many individuals must be quite negatively predisposed towards others even without provocation. The variants of the Rabin (1993) model proposed by Davidson (2000) and Falk and Fischbacher (2000) do allow for sufficient negative reactions to low offers to induce even splits, but only if reciprocity becomes the only thing that individuals care about (so that agents do not care at all about their own material payoffs). The difficulty with these models seems to be that like the Levine (1998) model, the functional forms that they consider do not allow responders to be sufficiently influenced by low offers, which means that they require utility functions that are far from standard even under normal circumstances.

2 Anger at Insufficient Altruism

I continue to suppose that the two agents have material payoffs, *i.e.* utility functions leaving aside any feelings the agents have for one another, that equal their resources. I denote these by x_p and x_r for the proposer and responder respectively. The actual utility functions of the two players' are given by

$$W_p = E(x_p + \lambda^p x_r)^\gamma \quad (5)$$

$$W_r = x_r + (\lambda^r - \xi(\hat{\lambda}^p, \bar{\lambda}))x_p \quad (6)$$

where E is the expectations operator, γ represents the attitude towards risk of the proponent, while λ^p and λ^r represent the unconditional levels of altruism of the the proponent and the responder respectively. The variable $\hat{\lambda}$ represents the beliefs of the responder about λ^p while the function ξ takes a value $\bar{\xi}$ which is greater than λ^r if the responder can reject the hypothesis that $\lambda^p \geq \bar{\lambda}$ and equals zero otherwise. This means that the responder is willing to incur costs to inflict harm on the proposer if he can reject the hypothesis the latter's altruism is at least equal to $\bar{\lambda}$. I suppose that $\bar{\lambda}$ is randomly distributed among potential respondents with a distribution function $H(\bar{\lambda})$ so that different people demand a different level of benevolence from those they interact with in this game. The support for this distribution is the set $[\lambda_L, \lambda_H]$.

While the utility functions for the two players may appear dissimilar, it would make no difference to the analysis if I made them identical by incorporating the parameter γ in the utility function of the responder and by incorporating the reaction to the responder's altruism in the proposer's utility function. The latter is irrelevant for the analysis because the proponent has no information about the responder's altruism when he makes his proposal. At this point, the proponent is thus unable to reject the hypothesis that the respondent's altruism parameter equal at least $\bar{\lambda}$. Similarly, raising the expression on the right hand side of (6) to the power γ , does not affect responder behavior since the responder's actions determines the outcome with certainty so that monotone transformations of his utility function do not matter.

Another important aspect of (6) is that it allows agents to react to unfair acts even if they are not themselves the victims of these acts. In other words, the responder reacts to his assessment of the other player's social preferences, not to the effect of the other's acts on the responders material payoffs. These preferences are thus perfectly compatible with the qualitative findings of Fehr and Fischbacher (2004). They show that agents are willing to incur some costs if they can thereby reduce the payoffs of individuals that have made low offers in a dictatorship game that these individuals were playing with *third parties*. In my interpretation, these individuals revealed that they had a low λ^p and thereby incur the wrath of the agents who have access to a technology for punishing them.⁴

It is also worth noting that, except for allowing γ to differ from one, (5) is identical to (2). Having made an offer y , the expectations operator in (5) ensures that the proposer gets nothing if the offer gets rejected while he gets $A - y(1 - \lambda^p)$ if the responder accepts the offer. The net result is thus the expression in (2) with $G(y)$ representing the probability that y will be accepted. The reason it is worth drawing this parallel is that we saw that this

⁴In Fehr and Fischbacher (2004), this technology allows these agents to induce a loss of \$3 to an agent for each dollar they give up themselves. They claim that the Fehr and Schmidt (1999) model can explain these punishments because these punishments reduce the difference between the payoff to the person who acted greedily in the dictator game and the person who has access to the punishment technology. If this explanation were correct, these punishments would disappear if the person who has access to the punishment technology were given a larger endowment in the experimental setting. This seems unlikely.

objective function for the proposer has difficulty generating even splits when the responder cares only about his own resources and how these compare to the resources of the proposer. The difference in the results obtained in this section is thus mainly attributable to the preferences I assume for the responder, with a subsidiary role played by the parameter γ .

In supposing that responders are more altruistic towards proposers that they view as being altruistic in return, this model follows Levine (1998). The main difference between the two settings is that he supposes that ξ is a linear function of the proposer's altruism λ^p , which becomes known in equilibrium. In his setting, therefore, the responder's altruism responds smoothly to that of the proposer's. Here, instead, I suppose that the responders reacts discontinuously to his perceptions of the proposer's altruism. I do this for two reasons. First, as explained earlier, this allows the baseline levels of altruism of the two agents, λ^p and λ^r to be nonnegative while still rationalizing even splits. Second, it appears that people's response to what they perceive to be ill treatment often moves almost seamlessly from passive acceptance to violent outburst. In the context of industrial relations, for example, worker unhappiness sometimes erupts in sudden strikes. In the case of road rage, some drivers react with unusual force to what they see as inconsiderate driving by others.⁵

In principle, one can set $\bar{\xi}$ to be arbitrarily large so that respondents are willing to incur large costs to punish proposers whose λ^p is below $\bar{\lambda}$. In practice, however, people's willingness to incur costs to punish those that are insufficiently altruistic is bounded. A related observation is that proposers who offer $A/2$ in ultimatum experiments are essentially never punished. One simple condition on (6) which ensures that respondents are willing to accept such offers regardless of the λ^p that they imply is $(\bar{\xi} - \lambda^r) = 1$.⁶

With these preference parameters, respondents who face offers larger than $A/2$ accept them even if they believe that the proposer's λ^p is below $\bar{\lambda}$. I impose this condition on the parameters in what follows. I further simplify the analysis by supposing that responders

⁵While I suppose that these violent reactions are the result of a drastic changes in attitude, they may instead be due to the fact that the person who is reacting is incapable of milder reactions.

⁶A more subtle psychological reason for not rejecting such offers is that, while responders know that they receive high offers because of the threat of rejecting lower ones, they do not want to think of themselves as wielding this weapon so forcefully that they exploit the proposers.

accept offers when they are indifferent between accepting and rejecting them. Responders thus always accept offers of $A/2$.

The last parameter that deserves discussion is the risk attitude parameter γ . From an *a priori* point of view one might naturally suppose that the stakes in these games are sufficiently low that the agents ought to be effectively risk neutral. The fact that a majority of subjects in experimental settings often prefer risky to safe outcomes with the same expected reward, however, suggests that risk neutrality may not be a good description of the subjects attitude towards risk.

Battalio, Kagel and Jiranyakil (1990) display a large battery of experiments geared towards assessing various theories that explain risk attitudes. They find that most theories perform poorly, in part because it seems difficult to predict when subjects will act in a risk averse manner and when they will do the opposite. Risk loving behavior is fairly dominant, however. For example, 26 of 32 subjects preferred a 0.7 probability of receiving \$10 together with a 0.3 probability of receiving \$30 to receiving \$17 for sure. In a separate experiment, 21 of 35 subjects preferred an equal chance of winning and losing \$10 to receiving nothing with probability 1. These attitudes are consistent with the popularity of low-stakes gambling, particularly with slot machines, which offer unfair odds to the players.

The simplest way of capturing these attitudes towards experimental lotteries is to suppose that γ is greater than one in (5). Because the marginal utility of income is locally linear, it is more attractive to follow Conlisk (1993) and suppose instead that many people have a positive utility for gambling. However, for the purpose at hand such a formulation would have similar effects than assuming $\gamma > 1$. In particular, it would lead selfish people to prefer to make low offers that have positive probability of being turned down even though these offers lead to lower expected earnings than offering an even split.

I suppose that y takes values in a discrete grid and denote its equidistant values by y_0, y_1, \dots, y_n , where $y_0 = 0$ and $y_n = A/2$.⁷ Because this is a signaling model where the

⁷I have also studied the case where y can take a continuum of values. While the mass of proposers who offer even splits is more difficult to study in this case, the continuum formulation has the attractive feature that the resulting objective function is very similar to (2).

offer y signals the type λ^p , there are numerous equilibria, which are supported by different beliefs regarding off the equilibrium path behavior. I focus on equilibria with two properties. The first is that the highest offer be equal to $A/2$. This is not only essential for fitting the experimental results, but also has the attractive feature of being a conservative (or safe) offer for those proponents whose altruism parameter λ^p is greater than or equal to λ_H , so that their altruism is as large as that demanded by any responder. For these proposers, an offer of $A/2$ is optimal if these proponents are as pessimistic as possible about the beliefs of respondents, i.e. if the proposers believe that respondents associate lower offers with substantially lower values of λ so that they accept such a lower offer with substantially lower probability.

The second property I impose on equilibria is that each possible offer be made by some proponent in equilibrium. This, again, fits well with experimental findings. It also means that the equilibrium outcome is not caused by “unreasonable” off-the-equilibrium beliefs; respondents beliefs are reasonable because they correspond to actions that proponents with different λ^p 's actually take. This has the benefit of ensuring that the conservatism that I impose on proposers with $\lambda^p = \lambda_H$ is reasonable, since there is a good reason for offers below $A/2$ to lead to lower probabilities of acceptance.

Given that there are only a discrete set of offers available and that λ^p can take a continuum of values, many proposers with different λ^p 's must make the same offer. I let $\tilde{\lambda}_i$ represent the set of λ^p 's that make offer y_i , and let λ_i be the supremum of $\tilde{\lambda}_i$. Observing an offer of y_i , respondents cannot reject the hypothesis that the proposer's altruism parameter equals λ_i . This means that such an offer is accepted with probability $H(\lambda_i)$. A collection of disjoint sets $\tilde{\lambda}_i$ is thus an equilibrium if

$$\forall \lambda \in \tilde{\lambda}_i, \forall j \neq i, H(\lambda_i)(A - y_i(1 - \lambda))^\gamma \geq H(\lambda_j)(A - y_j(1 - \lambda))^\gamma \quad (7)$$

Simple inspection of the proposer's objective function makes it clear that, if the probability of acceptance H is increasing in the proposer's offer y , proposers with higher values of λ prefer higher values of y . By the same token, if proposers with higher λ^p tend to make

higher offers, higher offers lead fewer responders to reject the hypothesis that $\lambda^p \geq \bar{\lambda}$ so that offers are more likely to be accepted. The model thus satisfies Athey's (2000) single-crossing property and her analysis then implies that a weakly monotone equilibrium in pure strategies exists. In the equilibria I compute, this monotonicity is strong so that whenever $y_i < y_j$, $\lambda_i < \lambda_j$ and $H(\lambda_i) < H(\lambda_j)$.

Suppose that, in equilibrium, there exist both proposers that offer y_i and proposers that offer y_{i+1} . Then, a proposer with $\lambda^p = \lambda_i$ for $i < n$ must be indifferent between y_i and y_{i+1} . Since he chooses y_i , he cannot prefer y_{i+1} . And if he strictly preferred y_i , then a proposer with a slightly higher λ^p would prefer y_i as well, thereby violating the definition of λ_i . Thus,

$$H(\lambda_i)(A - y_i(1 - \lambda_i))^\gamma = H(\lambda_{i+1})(A - y_{i+1}(1 - \lambda_i))^\gamma \quad (8)$$

Note that once one fixes the probabilities $H(\lambda_i)$ and $H(\lambda_{i+1})$, the right hand side of (8) rises faster with λ_i than the left hand side. This means that, for given cutoffs λ_i and λ_{i+1} , all individuals with altruism parameters larger than λ_i prefer the higher offer.

For every offer to be made in equilibrium, the equilibrium cutoffs for the altruism parameter λ^p must satisfy the difference equation (8) for all $0 \leq i < n$. Thus, $\tilde{\lambda}_i = (\lambda_{i-1}, \lambda_i]$ for $i > 0$ while $\tilde{\lambda}_0 = [\lambda_L, \lambda_0]$. For the highest offer to be $A/2$, the relevant solution to this difference equation must have the boundary condition $\lambda_n = \lambda_H$. Such an equilibrium exists as long as the λ_0 one obtains from applying (8) is then consistent with $H(\lambda_0) \geq 0$. If this holds, no proposer whose λ^p is in $\tilde{\lambda}_i$ wishes to make any offer other than y_i . If it does not hold, the value of λ_0 one obtains does not make sense so that the other values of λ_i do not either.

Fortunately, it is immediately apparent that the solution of (8) with $\lambda_n = \lambda_H$ does satisfy $H(\lambda_0) \geq 0$. To see this, note first that $(A - y_{i+1}(1 - \lambda_i))^\gamma$ is positive if $0 \leq y_i \leq A/2$ and $\lambda_i < 1$, as I assume. This implies that whenever the right hand side of (8) is positive for a given i , $H(\lambda_i)$ is positive, implying that the right hand side of (8) is positive for $i - 1$. Since the right hand side of (8) is positive for $i = n - 1$, it follows that $H(\lambda_0)$ is positive as well.

One interesting feature of this equilibrium is that the actions of responders, and thus the

offers made by a proposer with a given λ^p , are independent of the distribution of λ^p itself.⁸ They depend only on the distribution of the $\bar{\lambda}$'s, the levels of altruism that people regard as minimally acceptable. If proposers and responders are drawn from the same population and people regard their own level of altruism as being the one that is minimally acceptable to them, the two distributions coincide, and this is thus an interesting special case. However, it seems more plausible to assume that the two distributions are distinct so that relatively altruistic individuals have $\bar{\lambda}$'s that are below their own λ 's. It is also quite possible that extremely selfish individuals have $\bar{\lambda}$'s that exceed their own λ 's.

By contrast, it does not seem plausible to suppose that all individuals are selfish while simultaneously supposing that all individuals demand that others have a minimum positive level of altruism. This would mean that every individual's altruism would be simulated, which immediately makes one wonder why and how people could all end up pretending something which is false. Such pretenses would seem more likely to flourish if some people were in fact altruistic. A second, and related, problem with the view that all generosity is the result of pretense is that it seems difficult to square this with the heterogeneity in the extent to which people behave generously. People do, in particular, make different offers in the ultimatum game. It might be possible to rationalize this by supposing that different people have different amounts to gain by convincing others that they are generous, but a simpler approach involves supposing that their generosity is in fact different.⁹

Offers made by proposers with different λ^p 's are illustrated together with their probability of being accepted in Figure 1. This Figure shows solutions to the difference equation in (8) for my base case where $\gamma = 1$ and where the distribution $H(\bar{\lambda})$ is uniform between 0 and 0.2.

8

In Levine (1998), by contrast, the distribution of λ 's does matter for the probability that an offer will be accepted by responders. In his model, the probability of accepting an offer depends on the altruism of responders and he supposes that the distribution of this altruism parameter is the same as the distribution of the altruism parameter among proponents. By contrast, the altruism of responders plays a more muted role in my analysis because this altruism is swamped by the anger that attends rejection of the hypothesis that a proponent's λ^p is at least equal to $\bar{\lambda}$.

⁹It is also worth mentioning that individuals that expect generosity from people they are close to feel extremely betrayed when such generosity is not forthcoming. This suggests that people do form strong beliefs about other people's λ 's. It would thus seem churlish to regard such beliefs as acts of delusion.

The first solution it displays is for the case where the grid size equals one tenth of the pie to be distributed (so that in a \$10 experiment, the proposers get to make integer offers. The second has a grid size of 5e-5, which is obviously much smaller. The figure fits a smooth line through the offers that are made in the two cases. Thus, it should be interpreted as saying that the probability of acceptance of offers that are common in both situations (such as 0.4 of the pie) are extremely similar. Similarly, of all the responders that accept a particular integer offer (say 0.4 of the pie) the one with the highest $\bar{\lambda}$ has a very similar $\bar{\lambda}$ in the two cases. In the case of offers of 0.4 of the pie, this highest $\bar{\lambda}$ equals 0.1753 in the case of the coarse grid and 0.1760 in the case of the fine one. On the other hand, the lowest $\bar{\lambda}$ that ends up accepting only a particular integer offer is smaller in the case of a coarse grid, because the fine grid ensures that there are smaller offers that are still acceptable to this more tolerant responder.

In my rational expectations equilibrium, the range of $\bar{\lambda}$'s of responders that find a particular offer y_i minimally acceptable corresponds to the range of λ 's of proposers that make this particular offer. This means, not surprisingly, that the range of proposers that make a particular integer offer is larger when the offer grids is coarser and there is smaller menu of offers available. Even splits, too ought to become less common when proposers have a finer grid to choose from.

This raises the question of whether this model can account for the fact that between 28 and 50% of proposers offer even splits. This turns out to be crucially dependent on how the distribution of λ^p 's compares to the distribution of $\bar{\lambda}$'s. If the two coincide, the parameters underlying Figure 1 cannot account for the high fraction of such offers, even with the coarse grid based on integer offers. With a uniform distribution between 0 and 0.2, only about 12% of responders have $\bar{\lambda}$'s between 0.175 and 0.2, so that this is the proportion of responders that only accepts an even split. If λ^p had the same distribution, only about 12% of proposers would offer even splits. On the other hand, if proposers with relatively high λ^p 's have $\bar{\lambda}$'s below their own λ^p , one would expect many more proposers to exist whose λ^p is above 0.175. Indeed, one would expect many of them to have λ^p 's above 0.2. These proposers would also

offer even splits, so that the proportion of even splits would easily exceed 12%.

Figure 2 maintains the coarse grid but varies some aspects of the distribution of $\bar{\lambda}$. It compares a uniform distribution between 0 and 0.2 to a uniform distribution between 0 and 0.1. It also considers a density whose range goes from 0 to 0.1 but where low values of $\bar{\lambda}$ are relatively more common. By letting the pdf of $\bar{\lambda}$ equal $35 - 500\bar{\lambda}$, the density of $\bar{\lambda}$ at zero becomes 25 times larger than the density at 0.1. Moving from the uniform distribution on the range $[0, .1]$ to this pdf has a similar effect than moving from the uniform distribution in the range $[0, .2]$ to the uniform distribution in the range $[0, .1]$. In both cases, the probability of acceptance falls somewhat more rapidly as offers decline, though the differences are modest.

Some intuition for these differences can be developed by considering a linearized version of the difference equation (8), where the linearization is carried out around the point λ_i, y_i . After dividing through by $(A - y_i(1 - \lambda_i))^{\gamma-1}$ this gives

$$h(\lambda_i)(A - y_i(1 - \lambda_i))(\lambda_{i-1} - \lambda_i) - \gamma H(\lambda_i)(1 - \lambda_i)(y_{i-1} - y_i) \approx 0 \quad (9)$$

where $h(x)$ is the density of H at x . The second term is the increase in the proposer's gain from offering a lower y holding constant the probability of acceptance H . The first term, meanwhile, is approximately equal to the proposer's loss from the reduction in the probability that his offer will be accepted, where $(A - y_i(1 - \lambda_i))$ is the amount he loses and $h(\lambda_i)(\lambda_{i-1} - \lambda_i)$ is approximately the reduction in the acceptance probability. This equation can be rearranged so that the change in the probability of acceptance that results from offering y_{i-1} rather than y_i is

$$h(\lambda_i)(\lambda_{i-1} - \lambda_i) \approx \gamma \frac{H(\lambda_i)(1 - \lambda_i)}{A - y_i(1 - \lambda_i)} \quad (10)$$

Since $\lambda_n = \lambda_H$ and $H(\lambda_H) = 1$, (10) implies that a higher value of λ_H leads to a slower decline in the probability of acceptance as offers are reduced from $A/2$. There are essentially two reasons for this. The first is that a higher λ_i lowers the extent to which a lower y raises utility, because the individual empathizes more with the responder's loss. This lowers the second terms of (9)) and implies that a smaller reduction in the probability of acceptance

is needed to keep the proposer indifferent to a reduction in y . The second reason is that a higher λ implies that the loss from having a given fall in the probability of acceptance (which is captured by the first term of (9)) is larger, precisely because the proposer is also losing the vicarious benefit he draws from the resources received by the responder. This also means that a smaller decline in the probability of acceptance is warranted to compensate for the gain that results from a given reduction in y .

These two effects explain the differences in acceptance probabilities displayed in Figure 2. When moving from having $\bar{\lambda}$ distributed uniformly on $[0, .2]$ to having it distributed uniformly on $[0, .1]$, λ^H clearly falls, and these two effects lead the probability of acceptance to fall more rapidly with y . Now consider the non-uniform distribution considered in the Figure. Equation (10) implies that the reduction in the probability of acceptance as y is reduced from $A/2$ depends only on the value of λ_H and is thus the same in the case of the uniform distribution over the range $[0, .1]$ and in the case of this non-uniform distribution. However, the density of $\bar{\lambda}$ at λ_H is much smaller when the pdf of $\bar{\lambda}$ is $35 - 500\bar{\lambda}$. This implies that, for a given reduction in the probability of acceptance, the distance between λ_n and λ_{n-1} must be greater and λ_{n-1} must be lower. The earlier argument then implies that the reduction in the probability of acceptance between y_{n-1} and y_{n-2} must be larger in the case of this non-uniform density.¹⁰

Even though I have shown that changes in the distribution of $\bar{\lambda}$ can change the speed at which the probability of acceptance falls with y , the parameters I have considered so far do not allow the probability to fall by enough to explain the evidence. Combining the results of various studies, Harrison and McCabe (1996) compute that the probability that offers of zero will be accepted equals 0.33. They also analyze data originally collected for Carter and Irons (1991) on how responders would react to various potential offers and show that only 15% of these respondents would accept such an offer. By contrast, Figures 1 and 2 show that the probability of accepting offers as low as zero always exceeds 0.5.

¹⁰The analysis above was carried out for $H(\lambda_i) = 1$ and $H(\lambda_{n-1})$ is a bit smaller. But, to first order, this effect is negligible.

For $\gamma = 1$, this turns out to be guaranteed if λ_L , the minimum level of altruism that anyone finds acceptable, is nonnegative. The reason is as follows. Any risk neutral selfish responder strictly prefers to offer an even split if this is accepted with probability one to having a zero offer rejected with a probability larger than 0.5. This means that anyone who makes a zero offer that has a probability lower than 0.5 of being accepted must be spiteful, *i.e.* must have a $\lambda^p < 0$. Since this is unacceptable to all responders, such an offer could not be accepted. This, in turn, contradicts the possibility that an offer of zero would be made if it had a probability less than 0.5 of being accepted.

This argument extends naturally to any offer that leads the proposer with expected earnings lower than $A/2$. Since a proposer can guarantee himself this level of earnings by offering $A/2$, a proposer with $\gamma = 1$ would only make a lower offer that led to lower expected earnings if he were spiteful, and this would make such an offer unacceptable. If, by contrast, proposers liked small gambles and responders knew this, such offers would become acceptable at least to some responders.

To see this, Figure 3 shows acceptance probabilities for γ equal to both 2 and 3, where the distribution H remains equal to the uniform distribution on $[0, .2]$. The Figure also displays the acceptance probabilities computed by Harrison and McCabe (1996) from a variety of previous studies as well as those they obtained analyzing the data of Carter and Irons (1991). The Figure shows that most of these data are consistent with the model with $\gamma = 2$. The exception is the probability of acceptance of zero offers obtained from the Carter and Irons (1991) data, which fits this model only when $\gamma = 3$. Leaving aside this observation, it seems like a moderate taste for small gambles can explain why proposers whose λ^p is smaller than λ_H offer y 's significantly lower than $A/2$ even though the acceptance probabilities that result from such offers are significantly lower than those that emerge in equilibrium when $\gamma = 1$.

3 The Dictator Game

If proposers in the dictator game could not be penalized at all for making small offers, and if their altruism parameter were smaller than one, their nonsatiation would imply that they would offer zero. In settings where receivers are anonymous, this is by far the most common offer (Burnham 2003). By contrast, in the original Forsythe *et. al.* (1994) setting, proponents and respondents were students at the same university and were placed in communicating rooms. It is thus more reasonable to think of this setting as one where penalties for making low offers were not eliminated altogether even though they were reduced relative to those of the ultimatum game. A more direct demonstration that proposers are worried about other people's reactions even in the dictator game is provided by Hoffman *et. al.* (1994). They conducted a dictator experiment where they increased the proposers anonymity vis-a-vis the experimenter and found that this reduced the offers made by proposers¹¹

In this section, I thus suppose that the proposer in the dictator game feels that there is a positive probability π that his offer will be discovered. If it is, and the person discovering it concludes that the dictator's altruism is below $\bar{\lambda}$, the dictator suffers a loss V . For this analysis, it does not seem important that the person who discovers a dictator's offer be the recipient of the offer. The results of Fehr and Fischbacher (2004) suggest instead that a variety of individuals are prepared to punish people whom they regard as having been insufficiently generous in the dictator game.

This mechanism of punishment may well have heterogenous effects across people. In particular, people might differ in their probability of being found out π . People might differ for example, in their ability to conceal their actions in response to direct questions. In addition, people could differ in the extent to which they like telling others about their experiences during the experiment. In addition, proposers might well differ in their V , particularly if the cost involved is the cost of being verbally assaulted (and this will often be

¹¹Their clever experimental design probably increased the extent to which proposers felt that their offer was anonymous with respect to everybody because the offer's only physical record was the currency that remained in the proposer's pocket (or was sent via envelope to the responder).

the maximum possible punishment in this game).

To simplify the discussion, I suppose that the dictator, whom I will continue to call the proposer, has a probability π of having his offer discovered by someone, whom I will call the responder. As before, the responder's $\bar{\lambda}$ is drawn from H . The responder then leads the proposer to incur the cost V if he can reject the hypothesis that the proposer's altruism parameter λ^p is equal to at least $\bar{\lambda}$. I also make an important additional assumption, namely that the responder knows the proposer's values for π and V . This rather strong assumption is meant to capture that this is a setting where the responder will typically have some personal knowledge of the proposer. Even if he does not, the responder might well glean information about the proposer's cost of being verbally mistreated by the responder and use this in crafting his response.

One advantage of this assumption is that it allows the earlier analysis to apply almost directly. In particular, for each offer y_i , let $\tilde{\lambda}_i(V, \pi)$ represent the set of λ^p 's that make offer y_i given their probability of detection π and their cost of punishment V . Further, let $\lambda_i(V, \pi)$ be once again the supremum of $\tilde{\lambda}_i(V, \pi)$. Respondents who know V and π and observe an offer of y_i cannot reject the hypothesis that the proposer's altruism parameter equals λ_i . This means that the probability that a responder would refrain from punishing the proposer is $1 - H(\lambda_i)$. The sets $\tilde{\lambda}_i$ are thus an equilibrium if

$$\forall \lambda \in \tilde{\lambda}_i, \forall j \neq i, (A - y_i(1 - \lambda))^\gamma - \pi(1 - H(\lambda_i))V \geq (A - y_j(1 - \lambda))^\gamma - \pi(1 - H(\lambda_j))V \quad (11)$$

As before, higher values of λ make higher values of y_i more valuable to proposers if responders are more likely to punish proposers whose y is lower. And similarly, responders who know that higher offers are made by individuals with higher λ^p 's are more likely to punish proposers whose offers are low. There is thus a separating equilibrium where higher offers are made by proposers with higher λ^p 's for given V and π . One immediate consequence of (11) is that, for given $\tilde{\lambda}$'s, the only individual-specific parameter that affects the behavior of the proposer is the product πV .

Suppose as before that some proposers with a given πV offer y_i while others offer y_{i+1} .

Then, proposers with this πV for whom $\lambda^p = \lambda_i$ must be indifferent between y_i and y_{i+1} . If they liked y_i less, (11) would be violated and, if they liked it more, proposers with $\lambda^p < \lambda_i$ would choose y_{i+1} thereby contradicting the definition of λ_i . Therefore

$$(A - y_i(1 - \lambda_i))^\gamma - \pi V(1 - H(\lambda_i)) = (A - y_{i+1}(1 - \lambda_i))^\gamma - \pi V(1 - H(\lambda_{i+1})) \quad (12)$$

and the sets $\tilde{\lambda}_i$ consist of the sets $(\lambda_{i-1}, \lambda_i]$. The single crossing property then ensures that proposers prefers offering y_i to offering y_{i+1} if and only if their λ^p exceeds λ_i . Equation (12) takes on a particularly simple form when $\gamma = 1$ and H is uniform in the range $[0, \lambda_H]$ so that $H(\lambda_i) = \lambda_i/\lambda_H$. The equation can then be rearranged to yield

$$\left(\frac{\pi V}{\lambda_H} - \delta_y\right) \lambda_i = \frac{\pi V}{\lambda_H} \lambda_{i+1} - \delta_y \quad (13)$$

where δ_y is the size of the offer grid $y_{i+1} - y_i$. Neither equation (12) nor (13) can constitute an equilibrium unless decreases in λ_{i+1} lead to reductions in λ_i . This requires that the gap between grid points δ_y be sufficiently small, and specifically in the case of (13), that it be smaller than $\pi V/\lambda_H$.

Before one can solve the difference equation (13) a boundary condition is needed. Unlike the case of the ultimatum game, proposers with altruism λ_H may no longer be willing to make an offer as high as $A/2$. An offer of zero would be preferable to such proposers even if it led to a guaranteed loss of πV if $A - \pi V > A(1 + \lambda_H)/2$ or if $\pi V < A(1 - \lambda_H)/2$. In the ultimatum game, the worst punishment for making a low offer is relatively costly for the proposer since he gets nothing. Here, by contrast, the loss of πV can be quite low and offers of $A/2$ are then unsustainable.

As before, I focus on the equilibrium where proposers are most pessimistic about the responders' reactions. This now means that the offer made by a proposer with $\lambda^p = \lambda_H$ equals either $A/2$ or, if this is unsustainable, is as large as possible. Let us suppose then that a proposer with $\lambda^p = \lambda_H$ offers y_m where $m \leq n$. According to (13), the most altruistic proposer who offers y_0 then has an altruism parameter equal to λ_0 where

$$\lambda_0 = \left(\frac{\pi V}{\pi V - \lambda_H \delta_y}\right)^m (\lambda_H - 1) + 1 \quad (14)$$

As long as the λ_0 that satisfies (14) is positive, having proposers with $\lambda^p = \lambda_H$ offer y_m is an equilibrium. Proposers with $\lambda^p \leq \lambda^0$ offer zero and (13) ensures that all proposers with higher λ^p prefer to make higher offers. By contrast, it is obviously not the case that y_m is offered in equilibrium if the λ_0 that solves (14) is negative. The resulting λ_i still satisfy (13) but high offers satisfy this equation only because low offers lead to a loss of πV with probability greater than one.

Therefore, having proposers with $\lambda^p = \lambda_H$ offer an even split is an equilibrium if

$$\left(\frac{\pi V}{\pi V - \lambda_H \delta_y} \right)^n (\lambda_H - 1) + 1 > 0 \quad (15)$$

Otherwise, the highest offer that can be made in equilibrium, y_m , has a value of m which is the highest value such that the right hand side of (14) is nonnegative. The immediate implication of this is that reductions in the expected punishment πV reduce the maximum sustainable value of y_m , and the same is true of reductions in the maximum altruism of proposers λ_H . Heterogeneity in πV across proposers can then explain why low offers are much more common in dictator games than in ultimatum games. Those proposers that have low values of πV cannot be induced to make higher ones.

4 Conclusions

This paper has presented a model of preferences that can account for the experimental findings of ultimatum and dictator games without imposing extreme parameter values. It supposes that people feel mildly altruistic towards one another in most circumstances. Their normal altruism is mild enough that they would not transfer a dollar from their pocket to someone with a similar marginal utility of income, though they would transfer resources to people whose marginal utility of income they perceive as being much higher. They would also be willing to give up a dollar if someone else thereby gained substantially more than a dollar, as in the experiments of Charness and Rabin (2002).

The reason why preferences that differ so little from the selfish ones that form the baseline of economic analysis can fit these experiments is that there is a trigger that leads people to

have very different preferences. In particular, people get upset with people who demonstrate extreme selfishness. One way of thinking about this is to think that people get angry at individuals who “behave badly” or even that have “poor manners”. Once this reaction has been triggered, people actually enjoy hurting the individuals they regard as excessively selfish. What makes the ultimatum and dictator different from more normal economic interactions is that they are fabulous litmus tests for the extent to which people are selfish. Indeed, the moves in the game are so simple that they signal essentially nothing else.

By contrast, the moves people make in many other economic interactions also signal their tastes for different commodity bundles, as opposed to the extent of their altruism. These interactions are thus less prone to trigger anger. To see this, consider two individuals who wish to buy the same good. Suppose first that the price is fixed, that there is only one unit of the good left and that the first manages to purchase the unit before the second. Even if the second covets the good as well, the purchase by the first does not necessarily prove that this individual is ungenerous. Thus, the second individual’s disappointment is unlikely to turn to anger.

Suppose that, instead, the two individuals are bidding for the single unit that they both desire. When the first individual raises his bid this raises the cost to the second of obtaining the unit. But, again, the second is not entitled to see this as purely reflecting the first individual’s selfishness. This is so not only because the first individual might desire the good intensely but also because the first individual might be equally altruistic towards the seller as towards the second individual, and such even handedness does not seem as subject to censure. This even handedness would imply that, if the second individual ends up with the good, the vicarious gain to the first individual from the resources gained by the seller equal the vicarious losses he experiences from the reduction in the resources available to the buyer. There is thus no altruistic gain to the first individual from keeping his bid low. This can explain why selling prices end up being close to the reservation price of buyers in the “market” experiment of Roth *et. al.* (1991).¹² This is true even though an ultimatum

¹²Note also that, as long as the seller’s altruism parameter is less than one, he would choose to sell to the

game with a similar structure, in which a single seller sells a good to a buyer, leads the seller to charge much less than the reservation price of the buyer (see Hoffman *et. al.* for an ultimatum game with this structure).

By contrast, the act of posting a price by a seller is more indicative of the seller's altruism for buyers. Still, unlike the case of the ultimatum game with full information, buyers do not typically know the cost conditions faced by sellers, and this means that a seller's price is a less accurate measure of the seller's altruism than is the proposer's offer in an ultimatum game. Thus, this setting is more similar to the ultimatum game with incomplete information introduced by Mitzkewitz and Nagel (1993). I consider a price-setting situation of this sort Rotemberg (2003) and show that, under plausible conditions, a single-shot price can be a quite poor signal of the seller's altruism. This would suggest that producers have a great deal of flexibility regarding the prices they charge. If, however, buyers do not believe that cost conditions change dramatically from period to period, Rotemberg (2003) shows that the ability of producers to change their prices without triggering anger is substantially lessened.

buyer that offers the largest price. Choosing this buyer therefore does not prove that the seller is excessively selfish.

References

- Battalio, Raymond C., John H. Kagel and Komain Jiranyakul, "Testing Between Alternative Models of Choice under Uncertainty: Some Initial Results," *Journal of Risk and Uncertainty*, 3, 1990, 25-50.
- Bolton, Gary E and Axel Ockenfels, "ERC: A Theory of Equity, Reciprocity, and Competition," *American Economic Review*, 90, March 2000, 166-93.
- Burnell, Stephen J., Lewis Evans and Shuntian Yao, "The Ultimatum Game: Optimal Strategies without Fairness," *Games and Economic Behavior*, 26, 1999, 221-252.
- Carter, John R., and Michael D. Irons, "Are Economists Different, and If So, Why?," *Journal of Economic Perspectives*, 5, Spring 1991, 171-77.
- Charness, Gary and Matthew Rabin, "Understanding Social Preferences with Simple Tests," *Quarterly Journal of Economics*, August 2002, 817-69.
- Conlisk, John, "The Utility of Gambling," *Journal of Risk and Uncertainty*, 6, 1993, 255-75.
- Dickinson, David L., "Ultimatum Decision Making: A Test of Reciprocal Kindness," *Theory and Decision*, 48, 2000, 151177.
- Falk, Armin and Urs Fischbacher, "A Theory of Reciprocity," *University of Zurich Institute for Empirical Research in Economics Working Paper No. 6*, July 2000.
- Fehr, Ernst and Urs Fischbacher, "Third Party Punishment and Social Norms," *Evolution and Human Behavior*,, 2004, forthcoming.
- Fehr, Ernst and Klaus M. Schmidt, "A Theory of Fairness, Competition and Cooperation," *Quarterly Journal of Economics*, 1999, 817-68.
- Harrison, Glenn W. and Kevin A. McCabe, "Expectations and Fairness in a Simple Bargaining Experiment," *International Journal of Game Theory*, 25, 1996, 303-27.
- Hill, Rachael A., "Character, Choice, and 'Aberrant Behavior': Aligning Criminal Sentencing with Concepts of Moral Blame," *University of Chicago Law Review*, 65, Summer 1998, 975.
- Forsythe, Robert, Joel L. Horowitz, N. E. Savin and Martin Sefton, "Fairness in Simple Bargaining Games," *Games and Economic Behavior*, 6, 1994, 347-69.
- Güth, Werner, Rolf Schmittberger, and Bernd Schwarze, An Experimental Analysis of

- Ultimatum Bargaining, *Journal of Economic Behavior and Organization*, 3, 1982, 367-388.
- Hoffman, Elizabeth, Kevin McCabe, Keith Shachat and Vernon Smith, "Preferences, Property Rights and Anonymity in Bargaining Games," *Games and Economic Behavior*, 7, 1994, 346-80.
- Levine, David K., "Modeling Altruism and Spitefulness in Experiments," *Review of Economic Dynamics*, 1: 593-622, 1998.
- Mitzkewitz, Michael and Rosemarie Nagel, "Experimental Results on Ultimatum Games with Incomplete Information," *International Journal of Game Theory*, 22, 1993, 171-98.
- Parker, Dianne, Timo Lajunen and Heikki Summala, "Anger and aggression among drivers in three European countries," *Accident Analysis and Prevention*, 34, 2002, 229-235.
- Pillutla, Madan M. and J. Keith Murnighan, "Unfairness, Anger, and Spite: Emotional Ultimatum Offers," *Organizational Behavior and Human Decision Processes*, 68, December 1996, pp. 208-224.
- Rotemberg, Julio J, "The Benevolence of the Baker: Fair Pricing under the Threat of Customer Anger," mimeo, November 2003.
- Roth, Alvin E., Vesna Prasnikar, Masahiro Okuno-Fujiwara and Shmuel Zamir, "Bargaining and Market Behavior in Jerusalem, Ljubljana, Pittsburgh and Tokyo: an Experimental Study," *American Economic Review*, 81, December 1991, 1068-95.

Figure 1: Probability of Acceptance: Variations in Grid Size




Figure 2: Probability of Acceptance: Variations in the distribution of $\bar{\lambda}$




Figure 3: Probability of Acceptance: Variations in γ

