

Behaving as Expected: Public Information and Fairness Norms

CRISTINA BICCHIERI^{1*} and ALEX CHAVEZ²

¹*Department of Philosophy, University of Pennsylvania, USA*

²*Department of Psychology, University of Michigan, USA*

ABSTRACT

What is considered to be fair depends on context-dependent expectations. Using a modified version of the Ultimatum Game, we demonstrate that both fair behavior and perceptions of fairness depend upon beliefs about what one *ought* to do in a situation—that is, upon *normative expectations*. We manipulate such expectations by creating informational asymmetries about the offer choices available to the Proposer, and find that behavior varies accordingly. Proposers and Responders show a remarkable degree of agreement in their beliefs about which choices are considered fair. We discuss how these results fit into a theory of social norms. Copyright © 2009 John Wiley & Sons, Ltd.

KEY WORDS game theory; fairness; social norms; normative expectations

INTRODUCTION

Social norms are often invoked in explanations of pro-social behavior. The relative dearth of direct explanations (and predictions) in terms of social norms in behavioral economics is due to the fact that they are usually only vaguely defined, and thus are difficult to operationalize and draw predictions from. As a result, social norms are usually incorporated into research as ex-post interpretations for behavior or outcomes that are otherwise difficult to explain. However, without an operational definition of a social norm, it is impossible to design experiments that tease apart the behavioral effects of a social norm from those of personal values or generic social preferences. Furthermore, the absence of a clear definition makes it impossible to assess when a norm exists, and to make testable predictions about when it will be followed. In what follows we adopt a definition of norms that is grounded upon individuals' preferences and expectations (Bicchieri, 2006, p. 11). Such a definition makes claims about norms' influence on behavior easily testable.

For a social norm to exist and be followed by a specific population, three conditions must be present. First, it is necessary that the individuals involved *believe* it exists and know the class of situations to which the norm pertains. This condition implies that individuals must be aware they are in a situation in which a particular

* Correspondence to: Cristina Bicchieri, Department of Philosophy, University of Pennsylvania, 433 Logan Hall, 249 S. 36th Street, Philadelphia, PA 19104, USA. E-mail: cb36@sas.upenn.edu

norm applies, since lack of awareness may lead to non-compliance. We thus hypothesize that making a norm salient will lead, *ceteris paribus*, to greater compliance (Cialdini, Kallgren, & Reno, 1990). The second condition is that individuals must have a *conditional preference* for following the norm. Specifically, an individual will prefer to obey a given norm if she (a) expects others in the population to comply with it (empirical expectations) and (b) believes that others in the population expect her to obey the norm and may sanction transgressions (normative expectations). The third condition is the actual presence of the empirical and normative expectations. It follows that an individual may not obey a norm she knows applies to a given situation if she fails to have the right kind of expectations. Transgressions may occur because one observes non-compliance, or alternatively normative expectations are absent, or they are present but one can violate them without being observed.

The conditional preference condition distinguishes social norms from personal values. In the latter case, one usually has an unconditional preference for following a certain rule, as expectations about others' compliance play little or no role in one's decision. Having a conditional preference for conformity also implies that one might follow a norm in the presence of the relevant expectations, but disregard it in their absence. We thus hypothesize that manipulation of expectations will produce major shifts in norm-abiding behavior. To test this hypothesis, we focused on a simple version of the Ultimatum game, in which one of the parties proposes a division of a fixed amount of money to another party who can then accept or reject the offer. If the offer is rejected, both parties get nothing. Experimental results show that participants' modal and mean offers are 40–50% of the total amount, and offers below 20% are rejected about half of the time (Camerer, 2003). These results are generally interpreted as showing that subjects have a preference for fairness, and have been formalized in inequity aversion models¹ (Bolton & Ockenfels, 2000; Fehr & Schmidt, 1999).

A drawback of the inequity aversion model is that it assumes that what matters to an agent is the final payoff distribution, not the way the distribution came about. This consequentialist assumption has been experimentally challenged. Several authors have shown that, contrary to the consequentialism of the inequity aversion model, intentions matter (Blount, 1995; Falk, Fehr, & Fischbacher, 2003). Moreover, it has also been shown that fair procedures can be as acceptable as fair outcomes (Bolton, Brandts, & Ockenfels, 2005). However, there has been no comprehensive explanation of why subjects care not only about consequences but also the process through which such consequences occur. If intentions matter, it must be recognized that they only matter against a background of expectations. Thus, to judge a Proposer's intentions as normatively acceptable or unacceptable, we must have a relatively clear idea of what Responders believe a Proposer ought to do in a given situation. What is socially appropriate, in turn, is defined by our shared social norms. An operational definition of social norms, it must be added, not only helps us understand Responders' reactions to unfair offers but also helps in predicting when, or under which conditions, Proposers will choose to be fair.

Our hypothesis is that individuals have a conditional preference for following fairness norms, and that manipulating participants' expectations across conditions will lead to different behavior. These predictions cannot be explained by a fairness preference hypothesis, since the material consequences are the same across conditions, and the only difference between conditions is the level of information (and thus the expectations) of the players. Furthermore, even if we were to modify the fairness preference hypothesis by making preferences conditional, it would remain to be explained under which conditions preferences would change. The theory of norms we adopt provides such a testable explanation.

To show that “fair” behavior is dictated by norms that are conditionally followed, we manipulated both salience and expectations. In the salience treatment, we asked Proposers which of the options they thought Responders believed to be fair, thus focusing them on Responders' normative expectations. We hypothesized that making a fairness norm more salient would induce greater fairness on the part of Proposers. In both the salience and non-salience conditions, we asked Responders which of the Proposers' choice options they

¹Such models typically specify that an individual suffers disutility from outcomes in which others earn more or less than they do.

believed to be fair. The goal of this assessment was twofold. First, we wanted to check whether there was agreement in Responders' normative expectations. Agreement strongly indicates the presence of a shared norm of fairness. Second, we wanted to assess the agreement of Responders' normative expectations with Proposers' beliefs about them. It is significant that Proposers' answers (in the salience conditions) agreed overall with Responders' answers, a fact that further indicates the presence of a shared norm. The direct measurement of first and second-order expectations about what choices are fair is lacking in most behavioral studies to date. By measuring expectations, we provide evidence of the actual process that governs choices in situations involving social norms.

What counts as fair may vary by context, and in a given context more than one fairness criterion might apply. Equity and equality considerations may coexist, and randomization is often perceived as a fair allocation mechanism (Bolton et al., 2005), as when we use a lottery to allocate transplant organs. In our experiment, we added a choice option that randomized with equal probability between an equal and a very unequal share. Remarkably, a majority of Responders found this option fair, a finding also reflected in the Proposers' beliefs about Responders' assessments.² Adding this option allowed us to check for *norm manipulation* (Bicchieri, 2008). That is, when a norm can have several interpretations, individuals will tend to adopt the one most favorable to them. This form of self-serving bias has been studied in the context of equity versus equality interpretations of fairness (Frey & Bohnet, 1995; Hoffman & Spitzer, 1985), and we show that we can elicit the same bias with a suitable random option, provided this choice option is common knowledge among participants.

To manipulate expectations, we changed the information conditions in three versions of the same division problem. We expected a change in normative expectations to affect Proposers' choices. We also predicted the occurrence of *norm evasion* (Bicchieri, 2008): Whenever it was possible for Proposers to defy normative expectations without Responders' knowledge, the proportion of unfair offers would be significantly higher. Other authors have shown that participants tend to skirt offering a fair share when their actions cannot be detected by the other parties (Kagel, Kim, & Moser, 1996; Straub & Murnighan, 1995), because of asymmetric information about chip values or ignorance about the amount being divided. A theory of conditional norm compliance predicts these effects. Due to the conditional nature of norm compliance, the more ambiguous the choice situation, the higher will be the proportion of individuals that flaunt the norm, since the outcome of their behavior cannot be clearly interpreted as intentional, and thus no sanctioning is expected to occur.

Our paper is a contribution to a field that may be labeled *behavioral ethics*. It should be seen as a subfield of behavioral decision theory, taking into account of social and psychological considerations in an attempt to understand pro-social, "moral" choices. In this regard, we explore whether and why a pro-social norm will be followed, trying to gain insights about decision-making and motivation that are useful for further developing a behavioral account of moral choices.

METHOD

Participants

One hundred and six college-age subjects participated in our study across 11 experimental sessions. Advertisements specified that participants would earn 5 USD in addition to an amount that would depend on decisions made during the experiment.

²Easy acceptance of the randomization option may be a relative effect, namely of comparing it with the direct proposal of (8,2). Another possibility is that determining allocations by random devices is usually seen as fair and acceptable, as participants focus on the fair procedure and not on the expected utilities.

Game paradigm

Our experimental design employed a variant of the Ultimatum Game (Guth, Schmittberger, & Schwarze, 1982), in which one participant, the *Proposer*, provisionally received a sum of 10 USD—provided by the experimenter—and then proposed a division of that money with a *Responder*. The Responder subsequently decided to accept or reject the proposal. If the Responder accepted, both players received the amounts specified in the proposal. If the Responder rejected, both players received \$0. The Proposer chose from one of the following options:

- (5,5)—to propose \$5 for the Proposer and \$5 for the Responder;
- (8,2)—to propose \$8 for the Proposer and \$2 for the Responder; and
- Coin—to let the outcome of a fair coin flip determine the proposal: Heads corresponded to (5,5) and tails to (8,2).

Procedure

Upon their arrival, participants randomly drew ID codes labeled as A_1, A_2, \dots, A_n and B_1, B_2, \dots, B_n , where n is twice the number of participants in the session, and proceeded either to Room A or B based on their codes. We distributed and read aloud a set of instructions, which explained the following:

- Based on random assignment, the participants in Room A were Proposers, and those in Room B were Responders.
- Participants would play three games,³ each with a different person in the other room.
- By means of a public randomization device, we would select two of the three games at the end of the study to determine participants' cash payments.
- For each game, we would post on the blackboard the Proposer–Responder pairings using their ID codes, distribute instructions specific to that game, and then administer a short quiz to test their understanding of the instructions. By posting pairings publically, we maximized transparency and the likelihood that each participant would believe he or she actually was interacting with three different people in the other room.
- Subsequently, Proposers would be given a proposal form on which they would write their ID codes and choose one of the available proposals.
- After all Proposers completed their proposal forms, an assistant would take them to Room B, where the Responders would mark their decision to accept or reject the proposal.

Information

The proposal form that was used for each of the three games was determined by the information condition which we assigned to that game. Appendix B contains the proposal forms which differed slightly by information condition, as described below.

In the *full information* condition, the proposal form listed all three options: (5,5), (8,2), and Coin. After Proposers completed their proposal forms, the experimenter in Room B flipped a coin in front of the Responders. On any forms on which the Proposer chose Coin, the experimenter then marked (5,5) or (8,2) based on the coin flip outcome. The instructions explained this, so that prior to making their decisions, all participants understood that Responders would know if the Proposers with whom they were paired chose Coin.

In the *private information* condition, the Responder did not know that Coin was available to Proposers, and Proposers knew this fact. To create this informational asymmetry, we left Coin out of the proposal form. The experimenter in Room A explained to Proposers that they could indicate a Coin choice by leaving both (5,5) and (8,2) unmarked; subsequently, the experimenter would flip a coin in Room A, and on any such forms,

³We described the games using the language of Hoffman et al. (1994).

mark (5,5) if the outcome was heads, and (8,2) if the outcome was tails. Thus, Proposers understood that Responders were unaware that Coin was an available option.

In the *limited information* condition, both Proposers and Responders knew that Coin was available, but also that Responders would not be able to infer whether the Proposer chose Coin or marked one of (5,5) and (8,2) directly. Specifically, the instructions explained that Proposers could indicate a Coin choice by leaving both (5,5) and (8,2) unmarked on their proposal forms. The experimenter in Room B would flip a coin and mark (5,5) or (8,2) accordingly on any such forms. Crucially, he would flip the coin behind a small screen, so that no participant could see the outcome. Thus, all participants understood that the Responder would be unable to distinguish forms on which the Proposer chose (5,5) or (8,2) directly from forms on which the Proposer chose Coin.

Finally, because Proposers did not receive feedback between conditions, and to minimize participants' confusion that occurred during a pilot study, we fixed the order of the information conditions as (1) full, (2) private, and (3) limited.⁴ It might be argued that the within-subjects design may have led to leakage or learning without feedback (Weber, 2003) between experimental conditions. In general, however, learning takes many trials, and in the context of different information conditions, leakage would be expected to be minimal. We only have three trials, and the size of the information effect is very large, thus suggesting there was no such learning.

Saliency

In each information condition, prior to making their choices, Responders completed a questionnaire that measured their normative expectations (see Appendix C).

The questionnaire asked whether the responder found each of the choice options fair. The questionnaire was aimed at assessing whether there was an agreement in Responders' normative expectations, an indicator of (as well as a necessary condition for) the existence of a social norm. In addition, roughly half of our experimental sessions included an incentive-based questionnaire for both Responders and Proposers, which they completed in each information condition, following the quiz but prior to making their choices (See Appendices D and E).⁵

These questionnaires asked participants about their beliefs about the percentage of Responders who indicated (5,5), (8,2), and Coin, respectively, as fair options. The questionnaires were designed to (1) make fairness norms more salient, and (2) test for an agreement between Responders' normative expectations and Proposers' beliefs about them. The other half of the sessions included no such questionnaires; they just included the first Responders' questionnaire (Appendix C). We distinguish salient sessions which included these two extra questionnaires from non-salient sessions which did not.

Design and analyses

The experiment used a 3×2 design, crossing three levels of a within-participant variable, information (full, private, and limited), with two levels of a between-participant variable, saliency (non-salient and salient). The primary dependent variable was the Proposer's choice: (5,5), (8,2), or Coin.

We used a mixed-effects multinomial logit model to estimate choice probabilities. To determine the significance of information, saliency, their interaction, and participant-level random effects, we used nested model comparisons based on the likelihood ratio test statistic (LRT). Because small cell sizes made it

⁴In the pilot sequence, we put the private information condition first, followed by the full and limited information conditions. We did that since we thought that responders might not be unaware of the COIN option after the full information condition. Yet the pilot results were not qualitatively different from the results we report with the present ordering of conditions.

⁵For Responders in the saliency condition, the second questionnaire in Appendix D was appended to the one in Appendix C.

inadvisable to base inference about the LRT on the χ^2 distribution, we generated p -values for model comparisons using the non-parametric bootstrap. Bootstrap p -values lead to more accurate inference when cell counts are sparse. For the same reason, we generated p -values for the logit coefficients using the bootstrap. Appendix A contains additional details on the bootstrap methods we employed.

Hypotheses

We first hypothesized that, if there is a social norm, there must be agreement between Responders' normative expectations and Proposers' beliefs about them. We also hypothesized that choices would depend on information condition and salience. Based on the theory presented in the Introduction, we made several directional predictions by considering possible scenarios involving relevant fairness norms.

Firstly, we predicted that the proportion of Coin choices would be higher in the full information condition than in the private or limited conditions. This follows because in the private information condition, there are no normative expectations for Coin (as the availability of Coin is unknown to Responders), and there should therefore be very few Coin choices, since an outcome of (8,2) is likely to be rejected as unfair. In the limited condition, Responders could not determine whether the Proposer chose Coin. Therefore the expected utility of Coin is a combination of the probability of getting (5,5) or (8,2), and the probability that (8,2) will be accepted.⁶ If the Proposer assesses a probability greater than 5/8 that (8,2) will be accepted, she will choose (8,2). If the probability is less than 5/8, the Proposer will choose (5,5). Only if the probability is exactly 5/8, the Proposer will be indifferent among the three options. Therefore, Proposers who believed that less than 62.5% of Responders would accept (8,2) would choose (5,5), since the expected value of choosing (8,2) would be less than 5. Proposers who believed that more than 62.5% of Responders would accept (8,2) would choose (8,2), since the expected value of choosing (8,2) would be more than 5. In either case, there will be few or no Coin choices. Finally, in the full information condition, as long as there were Proposers who thought that choosing Coin was perceived to be fair by Responders, the proportion of Coin choices would be positive. This is because these Proposers believe that the expected utility of choosing Coin is 6.5, while the expected utilities of choosing (5,5) or (8,2) are, respectively, 5 and 0. This is an example of *norm manipulation* i.e., participants would adopt the (allegedly shared) interpretation of fairness that would benefit them the most. Indeed, Coin maximizes expected monetary payoff without violating fairness.

Secondly, we predicted that there would be more (8,2) choices in the limited than in the full or private information conditions, because some Proposers would take advantage of the ambiguity of their choice. Specifically, if some Proposers believed that a large enough fraction of Responders thought that choosing Coin was fair and that Proposers were playing fair by choosing Coin, then these Proposers could evade the norm and maximize their monetary expected value by choosing (8,2) in the limited information condition.⁷ Thus, the proportion of (8,2) choices should be higher in the limited information condition than in the full information condition. In the private information condition, we expected the fraction of (8,2) choices to be low for reasons discussed below.

Thirdly, we predicted that there would be more (5,5) choices in the private than in the full or limited information conditions. Previous work has indicated that the proportion of (5,5) choices is around 70% when (5,5) and (8,2) are the only available options (Falk et al., 2003). Because the private information condition is most similar to this situation, and we expected (5,5) to be almost universally considered fair, we expected similar proportions of (5,5) choices. In the full information condition, because we expected a relatively larger number of Coin choices (due to norm manipulation), we expected fewer (5,5) choices. In the limited

⁶The expected utility of coin is equal to $\frac{1}{2}(5) + \frac{1}{2}(8)p$. p is the probability that the Responder will accept (8,2) because he believes that the Proposer is 'playing fair' and thinks coin is a fair choice. The expected utility of (5,5) is 5, and the expected utility of (8,2) is $8p$. The value of p that makes one indifferent between the three options is 5/8, or .625.

⁷That is, p must be greater than .625 for a Proposer to choose (8,2) in the limited condition.

information condition, because we expected a relatively larger number of (8,2) choices (as the ambiguity of choices led to norm evasion), we expected fewer (5,5) choices.⁸

Finally, because we expected the questionnaire in the salient condition to focus Proposers on fairness—i.e., on choosing either (5,5) or Coin, and not (8,2)—we predicted that the above effects would be amplified.

RESULTS

Choices by information and salience

Figure 1 shows the choice proportions for each level of Information \times Salience. As salience effects appeared to be small, we averaged across salience conditions in reporting the choice proportions below. 37.7% (20/53) of Proposers chose Coin in the full information condition, compared to 11.3% (6/53) in the private condition and 5.7% (3/53) in the limited condition, consistent with our first hypothesis. Consistent with our second hypothesis, more Proposers chose (8,2) in the limited condition (58.5% [31/53]) than in the full (24.5% [13/53]) or private (37.7% [20/53]) conditions. We observed the highest frequency of (5,5) choices in the private condition (50.9% [27/53]) relative to the full (37.7% [20/53]) and limited (35.8% [19/53]) conditions, consistent with our third hypothesis.

The primarily additive effect of salience combined with differences in choice proportions across information conditions suggests the presence of both main effects, but not their interaction. Indeed, nested model comparisons (see Table 1) confirmed that a model with main effects of information and salience—but not their interaction—fit the choice data the best. Participant-level random effects were not significant for any of the models in Table 1.

Consistent with a norm-based explanation of choices, when normative expectations for Coin were either absent (in the private condition) or could be defied without consequence (in the limited condition), the predicted probabilities of (5,5) and (8,2), respectively, were considerably higher than those of Coin. This is indicated by the significantly positive coefficients in Table 2, which shows the estimates for the model with main effects of information and salience. On the other hand, when normative expectations for Coin were present and choices were transparent (in the full information condition), (5,5) was no more likely than Coin ($p = .963$), whereas (8,2) was in fact less likely ($p = .038$); see Table 2.⁹

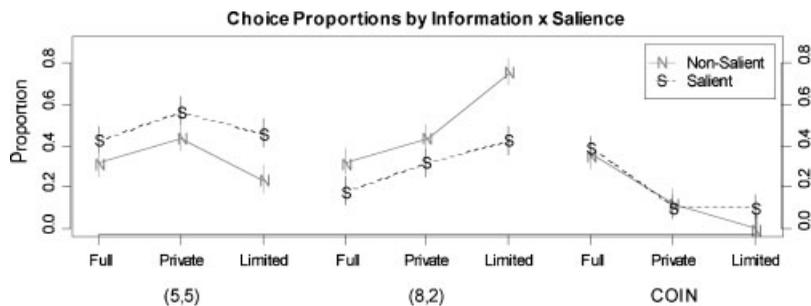


Figure 1. Choice proportions of (5,5), (8,2), and Coin in each level of information \times salience. Error bars are bootstrap estimates of one standard error of the choice proportion.

⁸Our results in the limited condition are similar to those obtained by Kagel, Kim and Moser (1996). They show how participants “skirt” giving a fair share when their actions cannot be detected by the other parties: They offer a 50/50 share of chips when fairness would dictate a 1/3, 2/3 one.

⁹The substantial discrepancy between some of the maximum likelihood (ML) and bootstrap estimates of the standard error were due to right-skewed coefficients distributions; accordingly, we used percentiles of the bootstrap distribution to calculate p -values, instead of relying on the asymptotic normality of the coefficients.

Table 1. Model comparisons for information and salience. *P*-values are based on bootstrap quantiles (see Appendix A for details), and not the χ^2 distribution*

Null	Alternative	df	LRT	<i>p</i>
—	Null (N)	—	—	—
N	Information (I)	4	25.73	<.0001
N	Salience (S)	2	6.54	.05
I	I + S	2	7.05	.04
S	I + S	4	26.23	<.0001
I + S	I × S	4	3.67	.52

*The χ^2 would have yielded the following less conservative *p*-values: From top to bottom (see Table 2), *p* < .0001, *p* = .04, *p* = .03, *p* < .0001, and *p* = .45, respectively.

Table 2. Multinomial logit model of information and salience as predictors of choices. Maximum likelihood (ML) and bootstrap standard error estimates are included. *P*-values are based on bootstrap quantiles (see Appendix A for details). The reference level for choice is Coin, for information is the full condition, and for salience is the salient condition*

	Log-odds	ML SE	Bootstrap SE	<i>p</i>
(5,5)				
Intercept	0.02	0.37	0.39	.963
Limited	1.84	0.69	1.72	.0007
Private	1.50	0.55	0.72	.0009
Non-salient	−0.07	0.48	0.48	.913
(8,2)				
Intercept	−0.89	0.45	0.47	.0380
Limited	2.83	0.71	1.73	.0001
Private	1.66	0.59	0.76	.0011
Non-salient	0.86	0.50	0.53	.0777

*Thus, for example, the predicted probability of choosing (5,5) in the limited condition and salient treatment is $\exp(.02 + 1.84) / (\exp(0) + \exp(.02 + 1.84) + \exp(-.89 + 2.83)) = .45$.

Using the model with both main effects, we performed formal tests corresponding to our first three hypotheses, and found support for the first two. Firstly, as suggested by Figure 1, the probability of Coin was significantly higher in the full information condition than in the private or limited information conditions (*p* < .001, for both the salient and non-salient conditions). Secondly, (8,2) was more likely to be chosen in the limited condition than in the full or private information conditions (*p* < .02, for both the salient and non-salient conditions). Lastly, (5,5) was not more likely to be chosen in the private condition than in the other two conditions (*p* = .1 and *p* = .2, respectively, for the salient and non-salient conditions).

We found mixed support for our final set of hypotheses, that Proposers would be more focused on fairness in the salient condition and would therefore be more likely to choose (5,5) and Coin, and less likely to choose (8,2). In the salient and non-salient conditions, there were 41 and 25 (5,5) choices (*p* = .003); 17 and 12 Coin choices (*p* = .2); and 26 and 38 (8,2) choices (*p* = .016), respectively, where again we generated bootstrapped *p*-values using the model with main effects of salience and information. Because the interaction between salience and information was not significant (see Table 1), the norm-focusing effects reported above were limited to across information conditions.

Table 3. Normative expectations of Responders. Each cell contains the proportion (fraction) of Responders who indicated that the choice was fair

Condition	Choice					
	5,5		8,2		Coin	
Full	96.4%	27/28	14.3%	4/28	64.3%	18/28
Private	96.4%	27/28	17.9%	5/28		
Limited	96.4%	27/28	14.3%	4/28	57.1%	16/28

Normative expectations and beliefs about them

Table 3 reports the normative expectations of Responders. Clearly, almost all Responders considered (5,5) to be fair in all information conditions, and a majority of them also thought that Coin was fair. This may be surprising because the expected utility of Coin is only 3.5 USD for the Responder, whereas it is 6.5 USD for the Proposer. A possible explanation is that using a random device is perceived as a fair way to choose between alternatives. Responders might have compared Coin to the temptation of (8,2) and found the Proposer who refused to choose between (5,5) and (8,2) as one making a fair choice. Alternatively, some might have thought the Proposer, as the first mover, had an entitlement to a greater share, and Coin seemed a fair compromise between \$5 and \$8.

In order for a norm to be shared, second-order beliefs should be the same across Responders and Participants. Figures 2–4 show participants’ beliefs about Responders’ normative expectations for each information condition. Indeed, there is a remarkable degree of agreement between Responders’ and Proposers’ beliefs about the normative expectations of Responders. Moreover, a comparison of Table 3 with Figures 2–4 shows that participants’ beliefs about normative expectations are in agreement with the normative expectations themselves.

Such high degree of agreement is the strongest possible indication that there is a shared norm of fairness. Not only is (5,5) universally perceived as fair, but also Coin is thought to be fair by a majority of participants. This agreement explains the tendency to choose a self-serving interpretation of fairness (norm manipulation) in the full information condition, as well as the pattern of choices across information conditions, as we show in the following section.

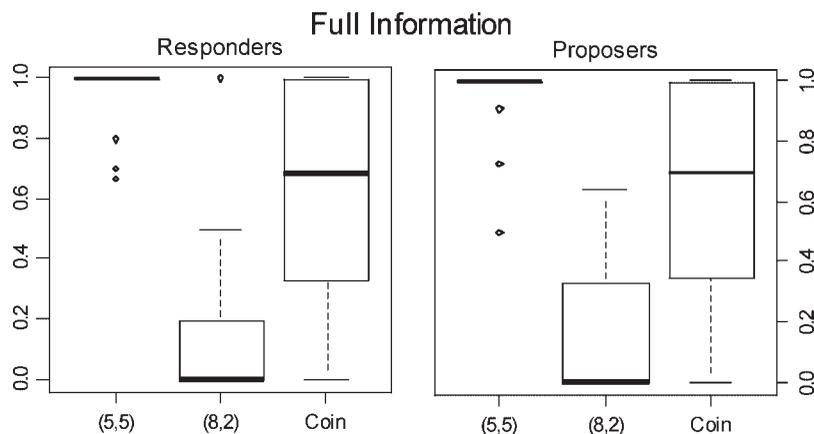


Figure 2. Boxplots of beliefs about normative expectations in the full information condition. Responders’ mean beliefs were 97.0, 12.6, and 63.6%, for (5,5), (8,2), and Coin, respectively. Proposers’ mean beliefs were 96.6, 14.9, and 65.0%.

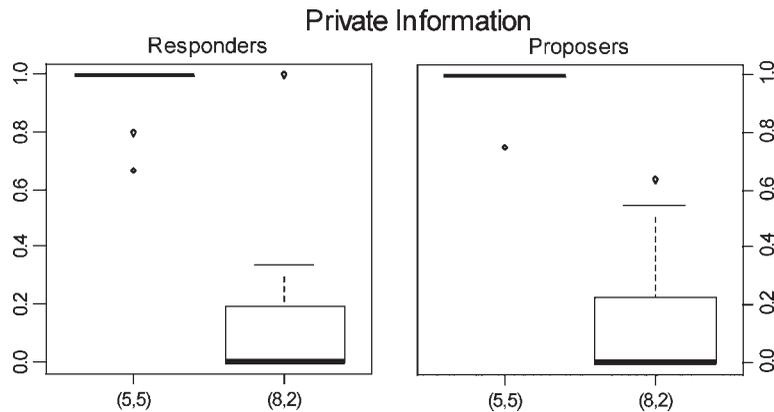


Figure 3. Boxplots of beliefs about normative expectations in the private information condition. Responders' mean beliefs were 98.1 and 16.0%, for (5,5) and (8,2), respectively. Proposers' mean beliefs were 99.1 and 12.5%.

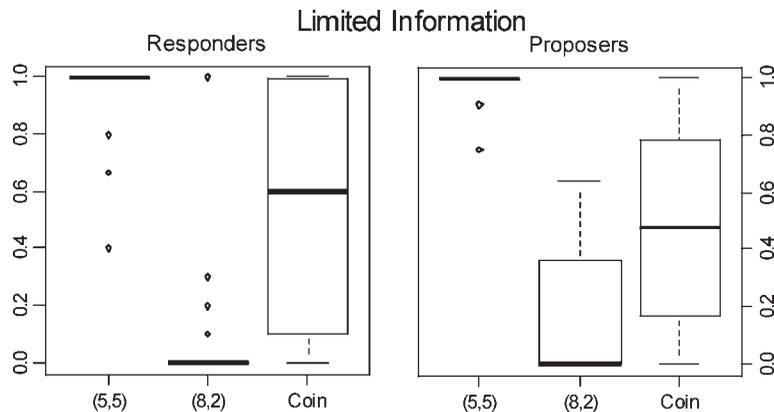


Figure 4. Boxplots of beliefs about normative expectations in the limited information condition. Responders' mean beliefs were 96.0, 10.0, and 54.4%, for (5,5), (8,2), and Coin, respectively. Proposers' mean beliefs were 98.8, 17.6, and 49.3%.

Choices as a function of beliefs about normative expectations

A social norm explanation presupposes consistency between beliefs and behavior. In particular, it requires consistency between participants' beliefs about normative expectations and their subsequent choices. As we show below, the data show a high degree of consistency between beliefs and behavior.

We denote the Proposers' belief about the proportion of Responders who consider (5,5), (8,2), and Coin, respectively, as being fair by $\varphi(5,5)$, $\varphi(8,2)$, and $\varphi(\text{Coin})$. The distributions of these variables are summarized in Figures 2–4. To determine whether these beliefs could explain the variance in Proposer's choices, we fit a multinomial logit model with information condition and the three questionnaire variables as predictors, based on a stepwise search.¹⁰ The signs of the coefficients for the questionnaire variables were in the appropriate

¹⁰We used a stepwise search that minimized Akaike's information criterion (AIC), a function of the likelihood which penalizes larger models ($AIC = -2(\log\text{-likelihood}) + 2k$, where k is the number of parameters in the model that are being estimated). The scope of the model search was information, $\varphi(5,5)$, $\varphi(8,2)$, and $\varphi(\text{coin})$. When we expanded the scope to search amongst all two-way interactions as well, AIC selected a model which included $\varphi(8,2) \times \varphi(\text{coin})$ in addition to the variables in Table 4 ($AIC = 158.61$). However, we chose to present the simpler model without this interaction term, as it did not change our substantive results, and Figures 5 and 6 looked the same with or without it.

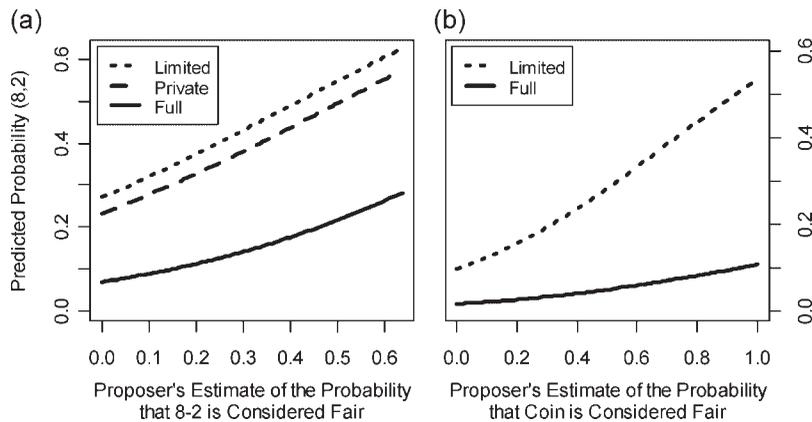


Figure 5. Predicted probability of (8,2) as a function of the Proposer's estimate of the proportion of Responders who considered (a) (8,2) fair, and (b) Coin fair, by condition. For each graph, the other beliefs are held fixed at their median values.

directions (see Table 4). As the Proposer's belief that the Responder considered (5,5) to be fair increased, (8,2) became less likely relative to the reference choice of Coin ($p = .068$). As $\varphi(\text{Coin})$ increased, the odds of choosing Coin over (5,5) increased as well ($p = .048$). For higher levels of $\varphi(8,2)$, Coin was less likely than (8,2) ($p = .068$).

Figures 5a–b show the predicted probabilities of (8,2) choices by condition as functions of $\varphi(8,2)$ and $\varphi(\text{Coin})$, respectively. In the private information condition, for example, a Proposer who believed that no Responders found (8,2) to be fair (i.e., $\varphi(8,2) = 0$) had a predicted probability of .23 of choosing (8,2)—see Figure 5a.¹¹ However, a Proposer who believed that roughly 60% of Responder found (8,2) to be fair was more than twice as likely to choose (8,2), with a predicted probability of .55.

The strong dependency of (8,2) choices on the Proposer's belief that Coin was considered to be fair in the limited information condition (see Figure 5b) is also noteworthy. A Proposer who believed that no Responders found Coin to be fair had a predicted probability of only .10 of choosing (8,2). However, a Proposer who believed that all Responders found Coin to be fair had a predicted probability of .53 of choosing Coin, a five-fold increase. This occurs because the Proposer could choose (8,2) with impunity when the Proposer's estimate of the probability that Coin was considered fair by Responders was high. If the Proposer believed that (a) the Responder thought Coin was fair and (b) the Responder believed that the Proposer was playing fair by choosing Coin, then (8,2) is expected to be accepted with high probability because it would be interpreted as the unlucky outcome of a coin flip. Moreover, as Figure 5b shows, the predicted probability of (8,2) choices in the limited information condition increased with the Proposer's estimate of the probability that Coin was considered fair by Responders. Whereas the existence of more than one interpretation of fairness led to norm manipulation (i.e., the Proposer chose Coin instead of (5,5) in the full information condition), the presence of ambiguity led to norm evasion (i.e., the Proposer chose (8,2) in the limited information condition because the source of the offer is not identifiable as being intentional or due to chance). Indeed, 76% of Proposers chose (8,2) in the limited information condition. And even when Proposers were focused on fairness in the salience condition, the proportion of (8,2) choices is relatively high.

Finally, Figure 6 shows the predicted probability of Coin choices as a function of $\varphi(\text{Coin})$ by condition. Clearly, the probability of choosing Coin increased with the Proposer's estimate of the proportion of

¹¹The other belief variables are held fixed at their median values for these predictions.

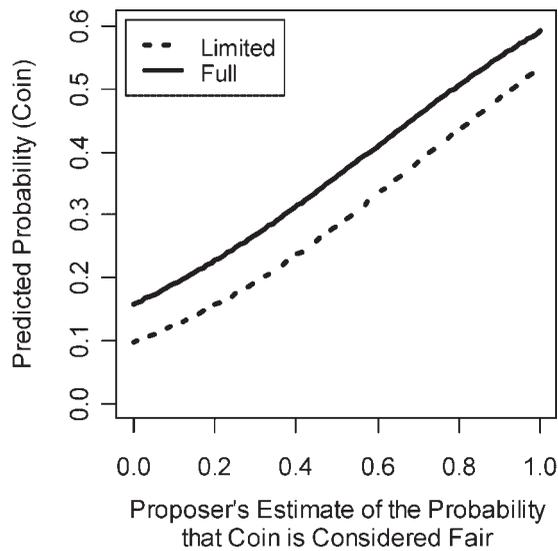


Figure 6. Predicted probability of Coin as a function of the Proposer's belief about Responders' normative expectations by condition. The other beliefs are held fixed at their median values.

Responders who considered Coin to be fair. As mentioned earlier, the Proposer's self-serving bias would lead to a greater frequency of Coin choices whenever it was clear that Responders knew that Coin had been chosen and that they considered Coin to be fair. This can also be seen in Figure 1, where the proportion of Coin choices in the full information condition was substantially larger relative to the other information conditions.

Responder behavior

Table 5 shows that Responders discriminated between intentional offers and offers that were generated by a chance mechanism. In the full information condition, no (8,2) offers resulting from a Coin choice were

Table 4. Multinomial logit model of information and beliefs about Responders' normative expectations as predictors of choices. The model was selected via a stepwise search (AIC = 166.83)

	Log-odds	<i>p</i>
(5,5)		
Intercept	-13.34	.144
Limited	0.86	.318
Private	-0.03	.958
$\varphi(5,5)$	14.99	.131
$\varphi(\text{Coin})$	-2.33	.048
$\varphi(8,2)$	0.89	.587
(8,2)		
Intercept	4.95	.176
Limited	2.64	.002
Private	2.87	.05
$\varphi(5,5)$	-7.21	.068
$\varphi(\text{Coin})$	0.55	.663
$\varphi(8,2)$	3.11	.068

Table 5. Rejection rates and frequencies by offer source, offer, and condition

	Offer			
	(5,5)		(8,2)	
Full				
Direct	0.0%	0/20	23.1%	3/13
Coin	0.0%	0/16	0.0%	0/4
Private	0.0%	0/28	40.0%	10/25
Limited	0.0%	0/20	18.2%	6/33

rejected, whereas intentional (8,2) offers were rejected 23% of the time. Responders' rejection rates of (8,2) were also sensitive to the presence of an intermediate choice. In the private information condition, in which Responders believe the only choices are (5,5) and (8,2), the rejection rate is much higher than in the full information condition (40% vs. 23%), in which Responders know that the Coin choice is available. When there is ambiguity as to the source of the choice, as in the limited information condition, the rejection rate of (8,2) was low (16.1%). This suggests that the large majority of Responders who receive an (8,2) offer in the limited information condition believe both that it is the result of a Coin choice and that choosing Coin is fair.¹² The large number of Proposers who choose (8,2) in the limited information condition seem to expect this low rate of rejections.

CONCLUSIONS

Having a well-defined, testable theory of social norms allows us to explain what would *prima facie* appear to be inconsistencies in individual behavior. Individuals choose to be fair on occasion, but revert to selfish behavior on others. The theory of social norms we adopt explains these apparent inconsistencies. Compliance with a norm is conditional upon having the right kind of empirical and normative expectations. It is also important that individuals focus on the relevant norm in order for them to comply with it. The theory we adopt (Bicchieri, 2006) predicts that making a norm salient will tend to increase compliance and, even more important, that the presence of the appropriate expectations is crucial for attaining conformity to the norm. We thus expected that manipulating expectations by changing the information available to individuals would result in large shifts in behavior.

The data we presented highlight two important phenomena connected to norm compliance. The first is norm manipulation (Bicchieri, 2008): When a norm can have several interpretations, individuals will tend to choose the interpretation that best serves their interests. This effect is evident in the choice of Coin in the full information condition. In this condition, the participants have common knowledge of the Proposer's access to the three choices of (5,5), (8,2), and Coin, as well as the Responder's ability to differentiate between an intentional unfair offer and an unfair offer resulting from a chance event. In the full information condition, the proportion of Coin choices is the same as the proportion of (5,5) choices. Moreover, this manipulation of the fairness norm is made possible by the implicit and strong agreement among the Responder's normative expectations about the fairness of Coin choices, and both Proposers' and Responders' implicit agreement

¹²In private conversation, Werner Guth suggested that Responders who, in the limited condition, accept the outcome (8,2) display let-down aversion, i.e., they *want* to believe that COIN was chosen.

about such normative expectations. As we mentioned at the outset, such an agreement is a strong indication of the presence of a shared norm.

The second phenomenon is norm evasion. Norm evasion differs from norm avoidance, in which an individual avoids a specific situation to which a norm applies. A vivid example of norm avoidance is the behavior of the Iks described by Turnbull (1972). The Iks repaired their huts in the middle of the night so as to avoid their neighbors' offers of help, as such offers had to be accepted, and that involved incurring an obligation that the beneficiary wanted to avoid. Norm evasion, on the contrary, is the deliberate, private flouting of a norm even if one knows the normative expectations of the relevant parties. For many individuals, the presence of normative expectations without the sanctioning element weakens the grip of the norm. Such expectations can be violated at no cost, as the victim will not be able to distinguish an intentional action from a chance event. Norm evasion explains Proposers' behavior in the limited information condition, in which (8,2) was the most frequent choice.

Our work introduces the field of behavioral ethics and builds on the seminal work of Guth et al. (1982) and Hoffman, McCabe, Keith, and Smith (1994), who showed the importance of context in games requiring the division of resources. We add to that work by showing that the effect of context is mediated by the role that normative expectations and shared norms hold in explaining behavior. Secondly, our measurement of first-order (Responders' fairness judgments) and second-order (Proposers' and Responders' beliefs about Responders' fairness judgments) beliefs is a useful but underused method in experiments on strategic interactions. Thirdly, our results potentially allow us to distinguish between different types of individuals. Future work will explore how a more fine-grained account of individuals' sensitivity to specific norms explains their choices, and whether there are correlations between sensitivities to different norms, such as those of cooperation, fairness, and reciprocity. Strong correlations would indicate the existence of a general disposition to follow social norms, whereas a low correlation would indicate that norm compliance is a local, norm-specific phenomenon.

ACKNOWLEDGEMENTS

We thank the Goldstone Research Unit at the University of Pennsylvania for financial support, Werner Güth, Jonathon Baron, Robert Sugden, and seminar participants at the Universities of Michigan, CUNY, Yale, Milano, British Columbia, Alabama, Duke, Madrid, Texas at Austin, Arizona, and the Wharton School for many useful comments and suggestions. We are grateful to Ryan Muldoon, Doug Paletta, and Giacomo Sillari for their assistance in data collection. Alex Chavez's work was supported by a National Science Foundation Graduate Research Fellowship.

APPENDIX A: STATISTICAL METHODS

Rationale for statistical methods

To evaluate our hypotheses and predictions, traditional approaches using the general linear model or the χ^2 test for independence are inadequate for at least two reasons. Firstly, our data contain multiple responses from each participant. If these responses were substantially correlated, the standard error estimates produced by traditional methods would be too small. Secondly, four of the eighteen cells (3 choices \times 3 information conditions \times 2 salience treatments) in our data had counts of less than five, making inference based on asymptotic results unadvisable.

The bootstrap

Asymptotically, logit coefficients and the LRT follow known distributions (under general regularity conditions). However, this is not true for small samples. The bootstrap distribution of a statistic can be used to check the validity of basing inference on asymptotic results. Roughly half of the coefficients in the multinomial logit models we estimated had right-skew (and hence, non-normal) bootstrap distributions, and the majority of the bootstrap distributions of the LRTs were not χ^2 . Had we relied on traditional analyses based on the normal and χ^2 distributions, our standard error estimates and p -values would have been too small. We provide details of our bootstrap methods below.

Let $y_{\text{obs}} = (y_1, y_2, \dots, y_n)$ denote the vector of observed data, let $T(y_{\text{obs}})$ denote the *statistic* of interest, and let F denote the underlying distribution that generated the data. The bootstrap estimates F using y_{obs} , thereby producing an estimate, F^\wedge , of the distribution of $T(y_{\text{obs}})$. The *non-parametric bootstrap* estimates F^\wedge as the distribution generated by random sampling with replacement from the set of observed data points, $\{y_1, y_2, \dots, y_n\}$. The *parametric bootstrap* estimates F^\wedge by estimating the model's free parameters and then sampling from the assumed distribution conditional on its parameters being fixed at their estimated values. Using either nonparametric or parametric bootstrap, one generates a bootstrap data set, y , by taking n draws from F^\wedge . The quantity $T(y)$ is the *bootstrap estimate* of $T(y_{\text{obs}})$, and the distribution of $T(y)$ is its *bootstrap distribution*. Letting y_1, y_2, \dots, y_n be n bootstrap data sets, the standard error estimate of $T(y_{\text{obs}})$ is the standard deviation of $T(y_1), T(y_2), \dots, T(y_n)$.

We produced each parametric bootstrap estimate based on the multinomial logit model by (1) fitting the multinomial logit model using maximum likelihood, (2) generating 159 draws from the multinomial distribution with its parameters fixed at their maximum likelihood estimates from step 1, and (3) computing and recording the bootstrap estimates of the statistic of interest. We then repeated the above steps 9,999 times to determine the distribution of the statistic. The procedure for generating non-parametric bootstrap estimates is similar: (1) randomly select 25 Proposers in the control and 28 Proposers in the salient condition, with replacement, (2) compute the statistic of interest, and (3) repeat the previous two steps 9,999 times.

To determine the p -value for each multinomial logit coefficient, we computed the smallest value of α for which zero was not contained in the confidence interval formed by the $(n * \alpha/2)$ th and $(n * (1 - \alpha/2))$ th quantiles (sorted values) of the parametric bootstrap estimates of the coefficients, $T(y_1), T(y_2), \dots, T(y_n)$.¹³ The estimate of the p -value is the proportion of bootstrap estimates that fell strictly outside this interval, which estimates the probability of drawing a coefficient more extreme in either tail of its distribution than the observed value. Although the distribution of the coefficient divided by its estimated standard error converges in probability to the standard normal, this is not necessary true for small samples. Thus, computing a p -value based on quantiles of the bootstrap distribution is preferable. We determined the p -value for the LRT in the same way, except that we constructed a one-sided confidence interval by omitting the division of α by two, as we expected LRTs to be positive under the null hypothesis.

We also used the bootstrap to generate p -values for our hypothesis tests. For our first three hypotheses, we formed our test statistics by taking the predicted probability for the choice and information condition of interest, and subtracting from it the maximum of the predicted probabilities for the two other information conditions. For example, for the first hypothesis, that Coin choices would be more likely in the full information condition than in the private or limited conditions, we subtract the maximum of the predicted probabilities of Coin in the private and limited conditions from the predicted probability of Coin in the full information condition. Because we expected this test statistic to be positive, the p -value was the proportion of bootstrap replicates that were negative. Test statistics for the last three hypotheses were formed analogously.

¹³The decimal precision of this value is determined by the number of bootstrap replicates.

APPENDIX B: PROPOSAL FORMS

Proposal form for the first game (full information condition)

(1) Identification number (Proposer fills this out): _____
(2) Paired with (Proposer fills this out): _____
(3) Proposer's choices (Proposer check one): <input type="checkbox"/> \$5 for Proposer and \$5 for Responder <input type="checkbox"/> \$8 for Proposer and \$2 for Responder <input type="checkbox"/> Let a coin flip decide which of the above choices will be made.
(4) Responder's decision (Responder check one): <input type="checkbox"/> Accept <input type="checkbox"/> Reject

Proposal form for second game (private information condition)

(1) Identification number (Proposer fills this out): _____
(2) Paired with (Proposer fills this out): _____
(3) Proposer's choices (Proposer check one): <input type="checkbox"/> \$5 for Proposer and \$5 for Responder <input type="checkbox"/> \$8 for Proposer and \$2 for Responder
(4) Responder's decision (Responder check one): <input type="checkbox"/> Accept <input type="checkbox"/> Reject

Proposal form for third game (limited information condition)

(1) Identification number (Proposer fills this out): _____
(2) Paired with (Proposer fills this out): _____
(3) Proposer's choices (Proposer check one): <input type="checkbox"/> \$5 for Proposer and \$5 for Responder <input type="checkbox"/> \$8 for Proposer and \$2 for Responder Let a coin flip decide which of the above choices will be made
(4) Responder's decision (Responder check one): <input type="checkbox"/> Accept <input type="checkbox"/> Reject

APPENDIX C: QUESTIONNAIRE 1 FOR RESPONDERS

Item (3) of each block of questions was omitted for the private condition, because Responders did not know that the Coin option was available to Proposers.

Questionnaire

Your identification number: _____.

Please guess how many Proposers will choose:

- (1) \$5 for Proposer and \$5 for Responder: _____.
- (2) \$8 for Proposer and \$2 for Responder: _____.
- (3) Let a coin flip decide: _____.

If your guess is correct, you will earn a \$1 bonus.

Please mark any options you believe are fair options. You are free to choose none of the options, one, or more than one option. Your answer will not affect your payment.

- (1) \$5 for Proposer and \$5 for Responder [].
- (2) \$8 for Proposer and \$2 for Responder [].
- (3) Let a coin flip decide [].

APPENDIX D: QUESTIONNAIRE 2 FOR RESPONDERS

Please guess how many Responders (excluding you) will select each of the options in the above question as fair options. For each line on which your guess is correct, you will earn a \$1 bonus.

- (1) \$5 for Proposer and \$5 for Responder _____.
- (2) \$8 for Proposer and \$2 for Responder _____.
- (3) Let a coin flip decide _____.

APPENDIX E: QUESTIONNAIRE FOR PROPOSERS

The questionnaire item regarding the Coin option was omitted in the private condition, as Proposers understood that Responders thought the Coin option was unavailable in that condition.

Questionnaire

Your identification number: _____.

Each Responder was asked to decide whether the “\$5 and \$5” option is fair.

Please guess how many Responders selected this option as fair: _____.

Each Responder was asked to decide whether the “\$8 and \$2” option is fair.

Please guess how many Responders selected this option as fair: _____.

Each Responder was asked to decide whether the coin flip option is fair.

Please guess how many Responders selected this option as fair: _____.

Note that Responders may answer no to all questions or yes to one or more questions.

For each line on which your guess is correct, you will earn a \$1 bonus.

REFERENCES

- Bicchieri, C. (2006). *The grammar of society: The nature and dynamics of social norms*. Cambridge: Cambridge University Press.
- Bicchieri, C. (2008). The fragility of fairness: An experimental investigation on the conditional status of pro-social norms. *Philosophical Issues 18 Interdisciplinary Core Philosophy*, 227–246.
- Blount, S. (1995). When social outcomes aren't fair: The effect of causal attributions on preferences. *Organizational Behavior and Human Decision Processes*, 63, 131–144.
- Bolton, G. E., & Ockenfels, A. (2000). ERC: A theory of equity reciprocity and competition. *American Economic Review*, 90, 166–193.
- Bolton, G., Brandts, J., & Ockenfels, A. (2005). Fair procedures: Evidence from games involving lotteries. *Economic Journal*, 115, 1054–1076.
- Camerer, C. (2003). *Behavioral game theory: Experiments on strategic interaction*. Princeton, NJ: Princeton University Press.
- Cialdini, R., Kallgren, C., & Reno, R. (1990). A focus theory of normative conduct: A theoretical refinement and reevaluation of the role of norms in human behavior. *Advances in Experimental Social Psychology*, 24, 201–234.

- Falk, A., Fehr, E., & Fischbacher, U. (2003). On the nature of fair behavior. *Economic Inquiry* 41, 20–26.
- Fehr, E., & Schmidt, K. (1999). A theory of fairness, competition, and cooperation. *The Quarterly Journal of Economics*, 114, 817–868.
- Frey, B., & Bohnet, I. (1995). Institutions affect fairness: Experimental investigations. *Journal of Institutional and Theoretical Economics*, 151, 286–303.
- Guth, W., Schmittberger, R., & Schwarze, B. (1982). An experimental analysis of ultimatum bargaining. *Journal of Economic Behavior and Organization*, 3, 367–388.
- Hoffman, E., & Spitzer, M. (1985). Entitlements, rights, and fairness: An experimental examination of subjects' concept of distributive justice. *Journal of Legal Studies*, 2, 259–297.
- Hoffman, E., McCabe, K. A., Keith, S., & Smith, V. (1994). Preferences, property rights, and anonymity in bargaining games. *Games and Economic Behavior*, 7, 346–380.
- Kagel, J. H. K., Kim, K., & Moser, D. (1996). Fairness in ultimatum games with asymmetric information and asymmetric payoffs. *Games and Economic Behavior*, 13, 100–110.
- Straub, P. G., & Murnighan, J. K. (1995). An experimental investigation of ultimatum games: Information, fairness, expectations, and lowest acceptable offers. *Journal of Economic Behavior and Organization*, 27, 345–364.
- Turnbull, C. (1972). *The mountain people*. New York: Simon and Schuster.
- Weber, R. A. (2003). “Learning” with no feedback in a competitive guessing game. *Games and Economic Behavior*, 44, 134–144.

Authors' biographies:

Cristina Bicchieri is the Carol and Michael Lowenstein Professor of Philosophy and Legal Studies at the University of Pennsylvania, where she is also a member of the Psychology graduate group. Her research interests lie at the intersection between philosophy, game theory, and psychology. Her primary research focus is on judgment and decision making with special interest in decisions about fairness, trust, and cooperation, and how expectations affect behavior.

Alex Chavez is a graduate student in the Department of Psychology at the University of Michigan. His research interests are in behavioral game theory, the modeling of social phenomena, and consumer behavior.

Authors' addresses:

Cristina Bicchieri, Department of Philosophy, University of Pennsylvania, 433 Logan Hall, 249 S. 36th Street, Philadelphia, PA 19104, USA.

Alex Chavez, Department of Psychology, University of Michigan, 4436 East Hall, 530 Church Street, Ann Arbor, MI 48109-1043, USA.