PHIL 408Q/PHPE 308D
Fairness

Eric Pacuit, University of Maryland

February 1, 2024

Ultimatum Game: Two players receive a windfall. One of the players suggests a division. After learning of the first player's proposal, the second must either accept or reject it. If the second accepts, both receive the amounts suggested by the first, otherwise they receive nothing.

Experimental Regularity: In the ultimatum game, a substantial proportion of responders reject non-zero offers and a significant number of proposers offer an equal split.

Social Preferences

Social preferences share the underlying assumption that the utility of an individual depends not only on the individual's monetary payoff, but also on the monetary payoff of the other players involved in the interaction.

Social Preferences

Social preferences share the underlying assumption that the utility of an individual depends not only on the individual's monetary payoff, but also on the monetary payoff of the other players involved in the interaction.

Social preferences are examples of **outcome-based preferences**: utility functions that depend only on:

1. the individuals involved in the interaction and
2. the monetary payoffs associated with each strategy profile.

Inequality Aversion: Fehr and Schmidt Utility Function

$$u_P(x_P, x_R) = x_P - \alpha_P \max(x_R - x_P, 0) - \beta_P \max(x_P - x_R, 0)$$
$$u_R(x_P, x_R) = x_R - \alpha_R \max(x_P - x_R, 0) - \beta_R \max(x_R - x_P, 0)$$

Inequality Aversion: Fehr and Schmidt Utility Function

$$u_P(x_P, x_R) = x_P - \alpha_P \max(x_R - x_P, 0) - \beta_P \max(x_P - x_R, 0)$$
$$u_R(x_P, x_R) = x_R - \alpha_R \max(x_P - x_R, 0) - \beta_R \max(x_R - x_P, 0)$$

- α_i is i's 'envy' weight and β_i is i's 'guilt' weight
- $0 < \beta_i < \alpha_i$: indicates that people dislike inequality against them more than they do inequality favoring them.
- $\beta_i < 1$: agents do not suffer terrible guilt when she is in a relatively good position. For example, a player would prefer getting more without affecting other people's payoff even though that increases inequality.

Ernst Fehr and Klaus M. Schmidt (1999). *A theory of fairness, competition, and cooperation*. The Quarterly Journal of Economics, 114(3), pp. 817 - 868.

Responder's Utility

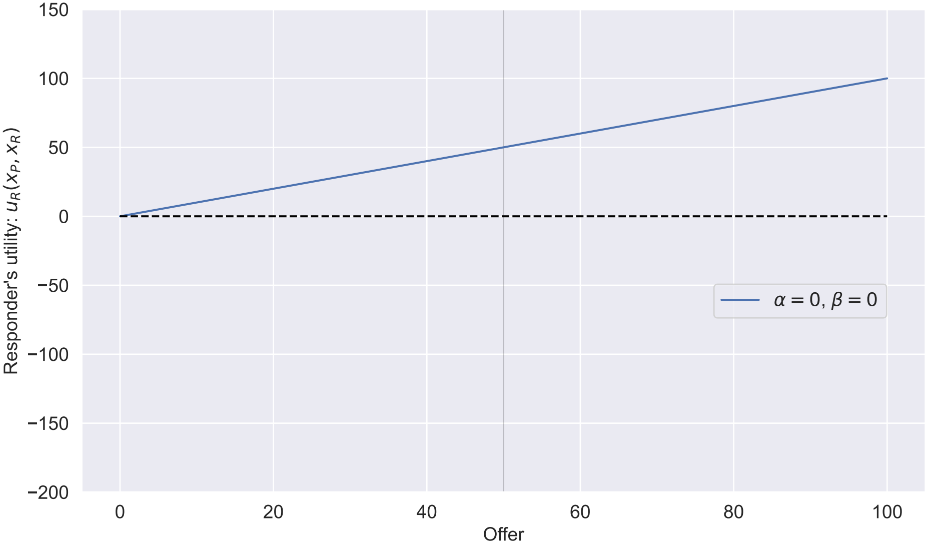

Responder's Utility

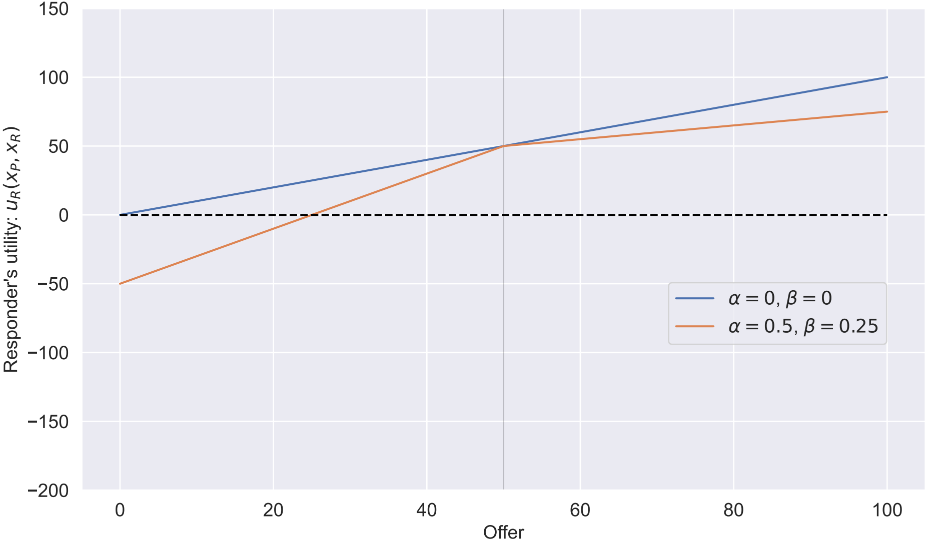

Responder's Utility

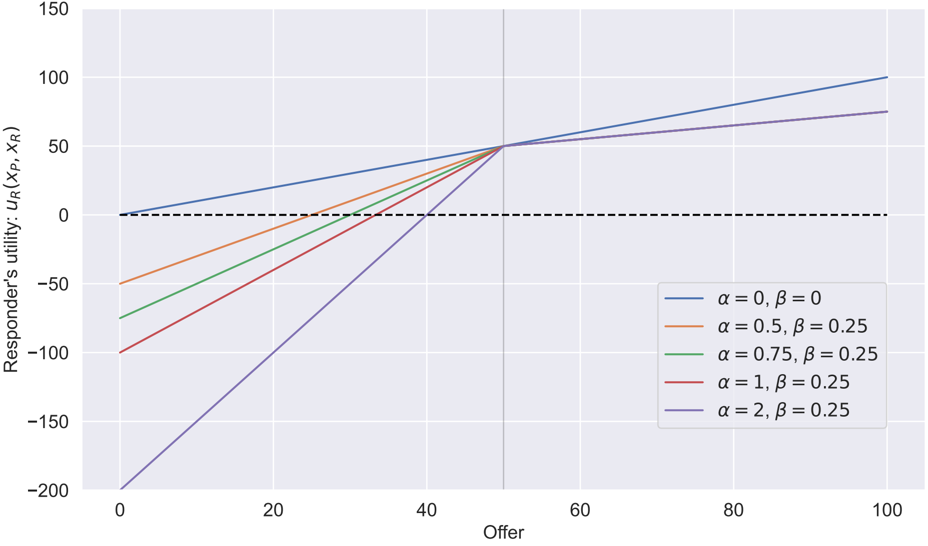

The Responder's Utility

$$u_R(x_P, x_R) = x_R - \begin{cases} \alpha_R(x_P - x_R) & x_P \geq x_R \\ \beta_R(x_R - x_P) & x_R > x_P \end{cases}$$

Suppose that the total amount to be distributed is M and y is how much P offers to R. So, $x_P = M - y$ and $x_R = y$. If R accepts, then:

The Responder's Utility

$$u_R(x_P, x_R) = x_R - \begin{cases} \alpha_R(x_P - x_R) & x_P \geq x_R \\ \beta_R(x_R - x_P) & x_R > x_P \end{cases}$$

Suppose that the total amount to be distributed is M and y is how much P offers to R. So, $x_P = M - y$ and $x_R = y$. If R accepts, then:

$$u_R(x_P, x_R) = \begin{cases} y - \alpha_R((M - y) - y) & y < M/2 \\ y - \beta_R(y - (M - y)) & y \geq M/2 \end{cases}$$

The Responder's Utility

$$u_R(x_P, x_R) = x_R - \begin{cases} \alpha_R(x_P - x_R) & x_P \geq x_R \\ \beta_R(x_R - x_P) & x_R > x_P \end{cases}$$

Suppose that the total amount to be distributed is M and y is how much P offers to R. So, $x_P = M - y$ and $x_R = y$. If R accepts, then:

$$u_R(x_P, x_R) = \begin{cases} y - \alpha_R((M - y) - y) & y < M/2 \\ y - \beta_R(y - (M - y)) & y \geq M/2 \end{cases}$$

Simplifying:

$$u_R(x_P, x_R) = \begin{cases} (1 + 2\alpha_R)y - \alpha_R M & y < M/2 \\ (1 - 2\beta_R)y + \beta_R M & y \geq M/2 \end{cases}$$

Suppose that the total amount to be distributed is M and y is how much P offers to R. So, $x_P = M - y$ and $x_R = y$:

$$u_R(x_P, x_R) = \begin{cases} (1 + 2\alpha_R)y - \alpha_R M & y < M/2 \\ (1 - 2\beta_R)y + \beta_R M & y \geq M/2 \end{cases}$$

Suppose that the total amount to be distributed is M and y is how much P offers to R. So, $x_P = M - y$ and $x_R = y$:

$$u_R(x_P, x_R) = \begin{cases} (1 + 2\alpha_R)y - \alpha_R M & y < M/2 \\ (1 - 2\beta_R)y + \beta_R M & y \geq M/2 \end{cases}$$

Then, R should accept provided that $u_R(x_P, x_R) > 0$. Solving, for y, we get:

$$y > \frac{\alpha_R M}{1 + 2\alpha_R}$$

Acceptance Threshold for R

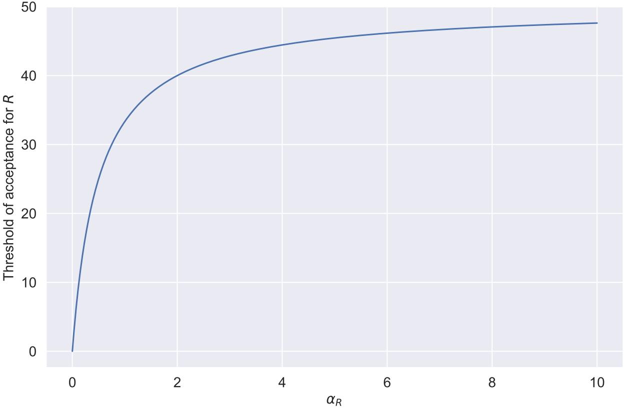

If α_R is close to zero — which indicates that R does not care much about being treated unfairly — the responder will accept very low offers. On the other hand, if α_R is sufficiently large, the offer has to be close to a half to be accepted.

The Proposer's Utility

$$u_P(x_P, x_R) = x_P - \begin{cases} \alpha_P(x_R - x_P) & x_R \geq x_P \\ \beta_P(x_P - x_R) & x_P > x_R \end{cases}$$

Suppose that the total amount to be distributed is M and y is how much P offers to R. So, $x_P = M - y$ and $x_R = y$. If R accepts, then:

The Proposer's Utility

$$u_P(x_P, x_R) = x_P - \begin{cases} \alpha_P(x_R - x_P) & x_R \geq x_P \\ \beta_P(x_P - x_R) & x_P > x_R \end{cases}$$

Suppose that the total amount to be distributed is M and y is how much P offers to R. So, $x_P = M - y$ and $x_R = y$. If R accepts, then:

$$u_P(x_P, x_R) = \begin{cases} (M - y) - \alpha_P(y - (M - y)) & y \geq M/2 \\ (M - y) - \beta_P((M - y) - y) & y < M/2 \end{cases}$$

The Proposer's Utility

$$u_P(x_P, x_R) = x_P - \begin{cases} \alpha_P(x_R - x_P) & x_R \geq x_P \\ \beta_P(x_P - x_R) & x_P > x_R \end{cases}$$

Suppose that the total amount to be distributed is M and y is how much P offers to R. So, $x_P = M - y$ and $x_R = y$. If R accepts, then:

$$u_P(x_P, x_R) = \begin{cases} (M - y) - \alpha_P(y - (M - y)) & y \geq M/2 \\ (M - y) - \beta_P((M - y) - y) & y < M/2 \end{cases}$$

Simplifying:

$$u_P(x_P, x_R) = \begin{cases} (1 + \alpha_P)M - (1 + 2\alpha_P)y & y \geq M/2 \\ (1 - \beta_P)M - (1 - 2\beta_P)y & y < M/2 \end{cases}$$

Proposer's Best Offer

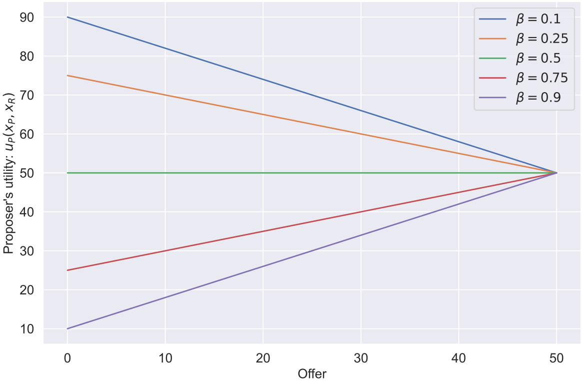

If $\beta_P > 1/2$ — if the proposer feels sufficiently guilty about treating others unfairly — then the best choice is to offer $M/2$. If $\beta < 1/2$, then the best offer is the minimum one that would be accepted, (a little bit more than $\alpha_R M/(1 + 2\alpha_R)$). If $\beta = 1/2$, then it any (acceptable) offer is best.

The Fehr-Schmidt Utility Function

As noted by Fehr and Schmidt, the model allows for the fact that individuals are heterogeneous. Different αs and βs correspond to different types of people. Although the utility functions are common knowledge, the exact values of the parameters are not. The proposer, in most cases, is not sure what type of responder she is facing.

The Fehr-Schmidt Utility Function

As noted by Fehr and Schmidt, the model allows for the fact that individuals are heterogeneous. Different αs and βs correspond to different types of people. Although the utility functions are common knowledge, the exact values of the parameters are not. The proposer, in most cases, is not sure what type of responder she is facing.

$$EU(y) = Pr(\alpha_R M/(1+2\alpha_R) < y) \times ((1-\beta_P)M - (1-2\beta_P)y)$$

The Fehr-Schmidt Utility Function

As noted by Fehr and Schmidt, the model allows for the fact that individuals are heterogeneous. Different αs and βs correspond to different types of people. Although the utility functions are common knowledge, the exact values of the parameters are not. The proposer, in most cases, is not sure what type of responder she is facing.

$$EU(y) = Pr(\alpha_R M/(1+2\alpha_R) < y) \times ((1-\beta_P)M - (1-2\beta_P)y)$$

The experimental data suggest that for many proposers, either β_P is large ($\beta_P > 1/2$) or they estimate the responder's α_R to be large.

The Fehr-Schmidt Utility Function

The advantages of the Fehr-Schmidt utility function are that it can rationalize both positive and negative outcomes, and that it can explain the observed variability in outcomes with heterogeneous types.

The Fehr-Schmidt Utility Function

The advantages of the Fehr-Schmidt utility function are that it can rationalize both positive and negative outcomes, and that it can explain the observed variability in outcomes with heterogeneous types.

One of the major weaknesses of this model, however, is that it has a consequentialist bias: players only care about final distributions of outcomes, not about how such distributions come about.

Shortly after the explosion of inequity aversion models, several economists observed that some decision-makers appear to act in a way that *increases* inequity, if this increase results in an increase in the total payoff of the participants.... This observation is hard to reconcile with inequity-aversion models, and suggests that people not only prefer to minimize inequity, but also prefer to maximize social welfare.

James Andreoni and John Miller (2002). *Giving According to GARP: An Experimental Test of the Consistency of Preferences for Altruism*. Econometrica, 70(2), pp. 737-753.

Andreoni and Miller conducted an experiment in which participants made decisions in a series of modified dictator games where the cost of giving is in the set $\{0.25, 0.5, 1, 2, 3\}$.

TABLE I

ALLOCATION CHOICES

Budget	Token Endowment	Hold Value	Pass Value	Relative Price of Giving	Average Tokens Passed
1	40	3	1	3	8.0
2	40	1	3	0.33	12.8
3	60	2	1	2	12.7
4	60	1	2	0.5	19.4
5	75	2	1	2	15.5
6	75	1	2	0.5	22.7
7	60	1	1	1	14.6
8	100	1	1	1	23.0
9[a]	80	1	1	1	13.5
10[a]	40	4	1	4	3.4
11[a]	40	1	4	0.25	14.8

[a] Were only used in session 5, others used in all sessions.

Results

- 22.7% of the dictators were perfectly selfish $(u_P(x_P, x_R) = x_P)$

Results

- 22.7% of the dictators were perfectly selfish $(u_P(x_P, x_R) = x_P)$
- 14.2% of dictators split the monetary payoff equally with the recipient (rationalized by the *Rawlsian utility function* $u_P(x_P, x_R) = \min(x_P, x_R)$)

Results

- 22.7% of the dictators were perfectly selfish ($u_P(x_P, x_R) = x_P$)
- 14.2% of dictators split the monetary payoff equally with the recipient (rationalized by the *Rawlsian utility function* $u_P(x_P, x_R) = \min(x_P, x_R)$)
- 6.2% of the dictators gave to the recipient only when the price of giving was smaller than 1 (rationalized by the *utilitarian utility function* $u_P(x_P, x_R) = 1/2 * x_P + 1/2 * x_R$)

Results

- 22.7% of the dictators were perfectly selfish ($u_P(x_P, x_R) = x_P$)
- 14.2% of dictators split the monetary payoff equally with the recipient (rationalized by the *Rawlsian utility function* $u_P(x_P, x_R) = \min(x_P, x_R)$)
- 6.2% of the dictators gave to the recipient only when the price of giving was smaller than 1 (rationalized by the *utilitarian utility function* $u_P(x_P, x_R) = 1/2 * x_P + 1/2 * x_R$))
- To rationalize the behavior of the remaining 57% of the dictators, Andreoni and Miller fit their data to a utility function of the form

$$u_P(x_P, x_R) = (\alpha x_P^\rho + (1-\alpha)x_R^\rho)^{1/\rho}$$

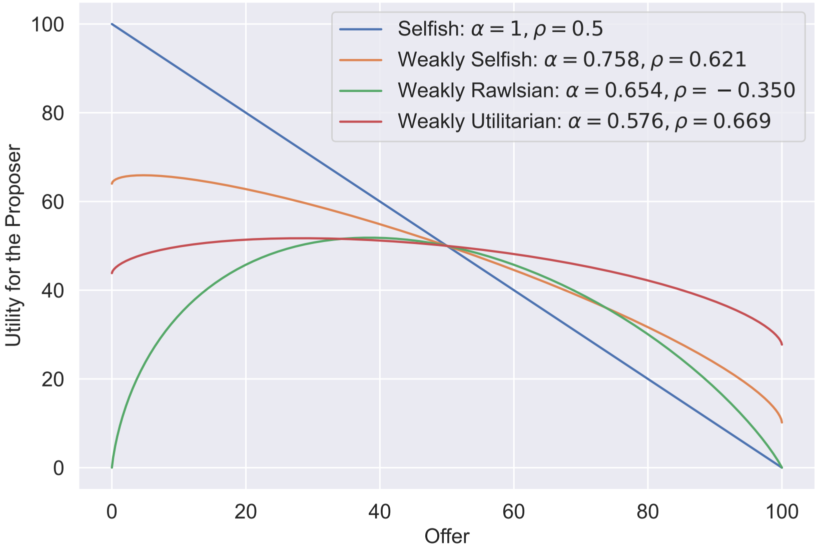

Selfish: $\alpha = 1, \rho = 0.5$
Weakly Selfish: $\alpha = 0.758, \rho = 0.621$
Weakly Rawlsian: $\alpha = 0.654, \rho = -0.350$
Weakly Utilitarian: $\alpha = 0.576, \rho = 0.669$
Utility for the Proposer
Offer
0
20
40
60
80
100

What light can our findings shed on efforts to suggest utility functions for fairness and altruism?

What light can our findings shed on efforts to suggest utility functions for fairness and altruism? One essential observation from our study is that individuals are heterogeneous. There is clearly not one notion of fairness or inequality-aversion that all people follow preferences range from Utilitarian to Rawlsian to perfectly selfish....

What light can our findings shed on efforts to suggest utility functions for fairness and altruism? One essential observation from our study is that individuals are heterogeneous. There is clearly not one notion of fairness or inequality-aversion that all people follow preferences range from Utilitarian to Rawlsian to perfectly selfish.... A second critical observation is that fairness must be addressed and analyzed on an individual level.....Capturing the variety of choices among individuals and then aggregating their behavior will lead to better understanding of both individuals and markets when altruism matters

What light can our findings shed on efforts to suggest utility functions for fairness and altruism? One essential observation from our study is that individuals are heterogeneous. There is clearly not one notion of fairness or inequality-aversion that all people follow preferences range from Utilitarian to Rawlsian to perfectly selfish.... A second critical observation is that fairness must be addressed and analyzed on an individual level.....Capturing the variety of choices among individuals and then aggregating their behavior will lead to better understanding of both individuals and markets when altruism matters.... [O]ur results beyond simple dictator games suggests that many things other than the final allocation of money are likely to matter to subjects. Theories may need to include some variables from the game and the context in which the game is played if we are to understand the subtle influences on moral behavior like altruism. (p. 751-2)

A. Falk, E. Fehr, and U. Fishbacher (2003). *On the nature of fair behavior*. Economic Inquiry, 41(1), pp. 20 - 26.

A. Festré (2019). *On the Nature of Fair Behaviour: Further Evidence*. Homo Oeconomicus, 36, pp. 193 - 207.

...identical offers in an ultimatum game trigger vastly different rejection rates depending on the other offers available to the proposer. In particular, a given offer with an unequal distribution of material payoffs is much more likely to be rejected if the proposer could have proposed a more equitable offer than if the proposer could have proposed only more unequal offers.

...identical offers in an ultimatum game trigger vastly different rejection rates depending on the other offers available to the proposer. In particular, a given offer with an unequal distribution of material payoffs is much more likely to be rejected if the proposer could have proposed a more equitable offer than if the proposer could have proposed only more unequal offers. ... This result not only casts serious doubt on the consequentialist practice in standard economic theory that defines the utility of an action solely in terms of the consequences of this action. It also shows that the recently developed models of fairness...are incomplete to the extent that they neglect "nonconsequentialist" reasons for reciprocally fair actions.

A. Falk, E. Fehr, and U. Fishbacher (2003). *On the nature of fair behavior*. Economic Inquiry, 41(1), pp. 20 - 26.

Each of 90 experimental subjects participated in four different games. In all games the proposer P is asked to divide 10 points between himself and the responder R, who can either accept or reject the offer. Accepting the offer leads to a payoff distribution according to the proposer's offer. A rejection implies zero payoffs for both players.

Figure 1: The mini ultimatum games

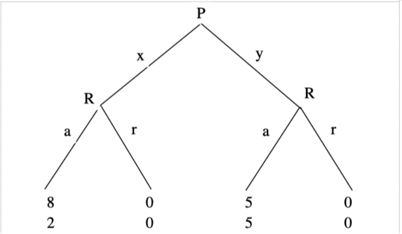

(a) (5/5)-game

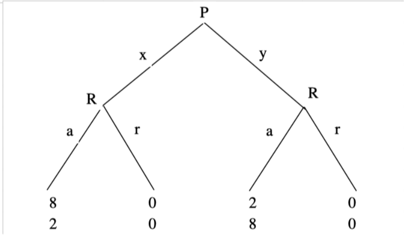

(b) (2/8)-game

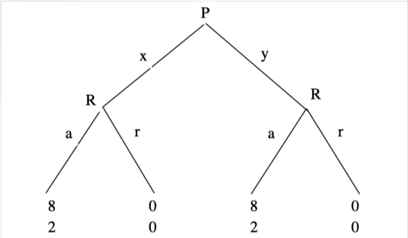

(c) (8/2)-game

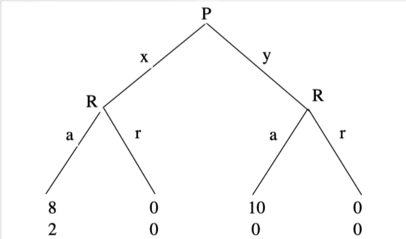

(d) (10/0)-game

Every responder had to indicate his action at both decision nodes, i.e., for the case of an x- and for the case of a y-offer, without knowing what P had proposed.

Every responder had to indicate his action at both decision nodes, i.e., for the case of an x- and for the case of a y-offer, without knowing what P had proposed.

At the beginning subjects were randomly assigned the P- or the R-role and they kept this role in all four games.

Every responder had to indicate his action at both decision nodes, i.e., for the case of an x- and for the case of a y-offer, without knowing what P had proposed.

At the beginning subjects were randomly assigned the P- or the R-role and they kept this role in all four games.

Subjects faced the games in a random order and in each game they played against a different anonymous opponent. They were informed about the outcome of all four games, i.e., about the choice of their opponents, only after they had made their decision in all games.

Every responder had to indicate his action at both decision nodes, i.e., for the case of an x- and for the case of a y-offer, without knowing what P had proposed.

At the beginning subjects were randomly assigned the P- or the R-role and they kept this role in all four games.

Subjects faced the games in a random order and in each game they played against a different anonymous opponent. They were informed about the outcome of all four games, i.e., about the choice of their opponents, only after they had made their decision in all games.

After the end of the fourth game subjects received a show-up fee of plus their earnings from the experiment (about \$23 was at stake).

Figure 2
Rejection rate of the (8/2)-offer across games

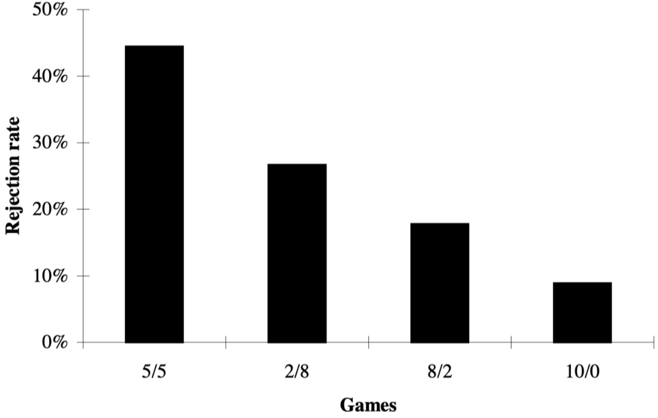

Figure 3:
Percentage and acceptance of (8/2) proposals

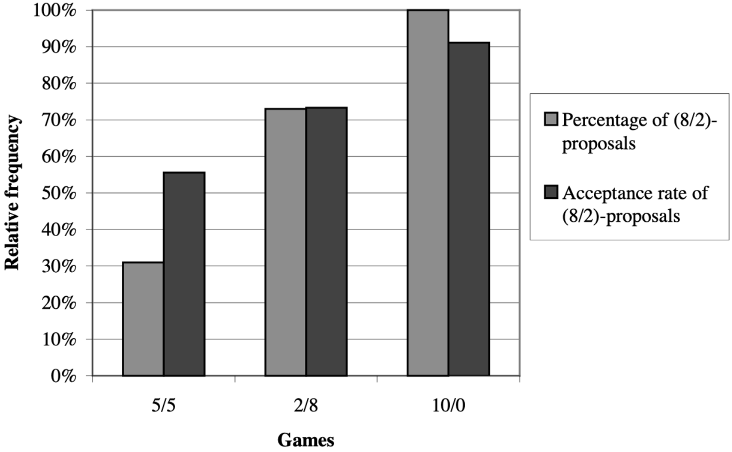

The results of our experiment clearly show that the same action by the proposer in a miniultimatum game triggers very different responses depending on the alternative action available to the proposer. This suggests that responders do not only take into account the distributive consequences of the action by the proposer but also the intention that is signaled by the action.

- ✓ Social Preferences: Outcome-based Preferences
- ▶ Moral Preferences
- ▶ Reasoning based on Norms

J. Halpern, V. Capraro and M. Perc (2022). *From outcome-based to language-based preferences*. Journal of Economic Literature.