PHPE 308M/PHIL 209F
Fairness

Eric Pacuit, University of Maryland

September 10, 2025

Preferences in the Ultimatum Game

- Two players: The Proposer (P) and the Responder (R)

Preferences in the Ultimatum Game

- Two players: The Proposer (P) and the Responder (R)

- An **outcome** of the game is (x_P, x_R) where x_P is the amount that player P receives and x_R is the amount that player R receives.

Preferences in the Ultimatum Game

- Two players: The Proposer (P) and the Responder (R)

- An **outcome** of the game is (x_P, x_R) where x_P is the amount that player P receives and x_R is the amount that player R receives.

- Players are assumed to have utility functions (a function that maps outcomes to real numbers) *representing* their preferences over the outcomes:

  The utilities for the outcome (x_P, x_R) are $u_P(x_P, x_R)$ and $u_R(x_P, x_R)$.

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

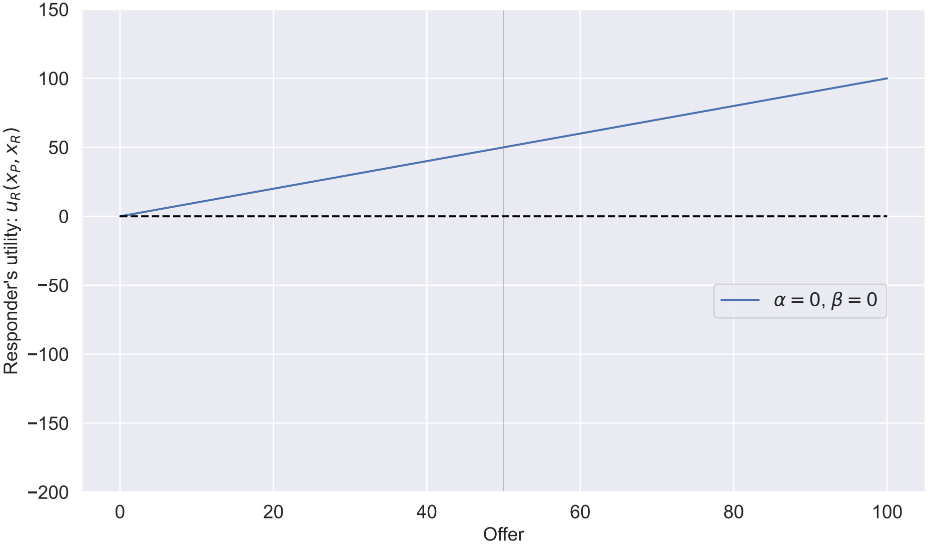

Responder's Utility

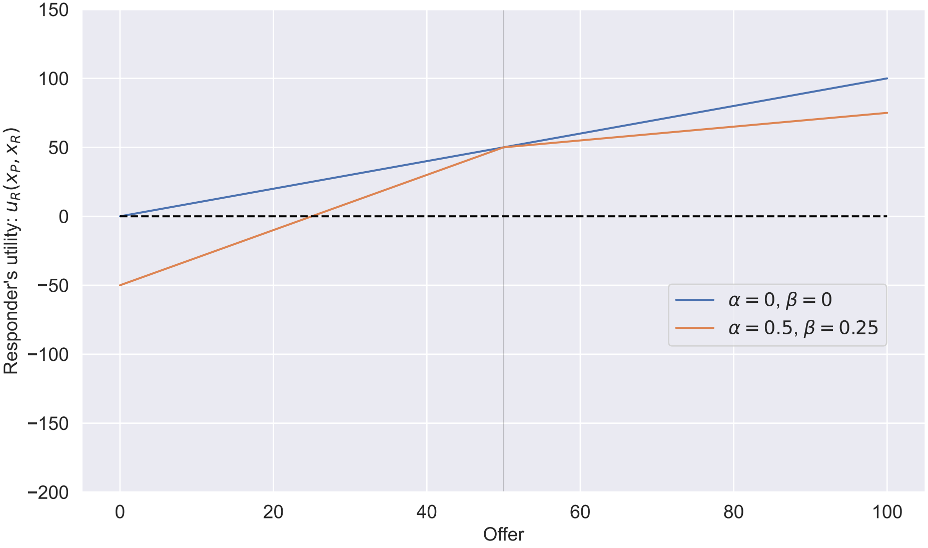

Responder's Utility

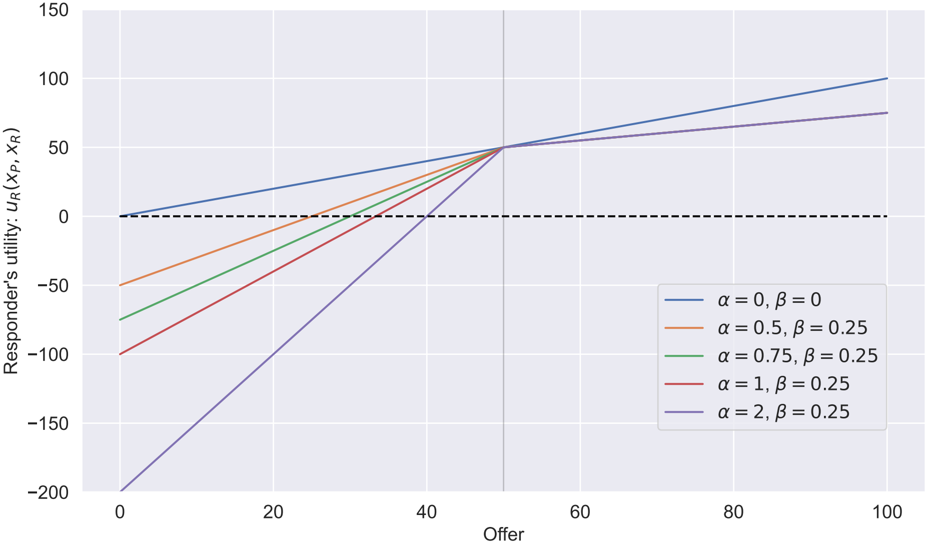

The Responder's Utility

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 $x_R = y$.

The Responder's Utility

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 $x_R = y$.

Simplifying Assumption: $x_P \geq x_R$ (i.e., $y \leq M/2$)

The Responder's Utility

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 $x_R = y$.

Simplifying Assumption: $x_P \geq x_R$ (i.e., $y \leq M/2$)

$$u_R(x_P, x_R) \quad = \quad x_R - \alpha_R \max(x_P - x_R, 0) - \beta_R \max(x_R - x_P, 0)$$

The Responder's Utility

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 $x_R = y$.

Simplifying Assumption: $x_P \geq x_R$ (i.e., $y \leq M/2$)

$$\begin{aligned} u_R(x_P, x_R) &= x_R - \alpha_R \max(x_P - x_R, 0) - \beta_R \max(x_R - x_P, 0) \\ &= x_R - \alpha_R(x_P - x_R) \end{aligned}$$

The Responder's Utility

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 $x_R = y$.

Simplifying Assumption: $x_P \geq x_R$ (i.e., $y \leq M/2$)

$$\begin{aligned} u_R(x_P, x_R) &= x_R - \alpha_R \max(x_P - x_R, 0) - \beta_R \max(x_R - x_P, 0) \\ &= x_R - \alpha_R(x_P - x_R) \\ &= y - \alpha_R((M - y) - y) \end{aligned}$$

The Responder's Utility

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 $x_R = y$.

Simplifying Assumption: $x_P \geq x_R$ (i.e., $y \leq M/2$)

$$\begin{aligned} u_R(x_P, x_R) &= x_R - \alpha_R \max(x_P - x_R, 0) - \beta_R \max(x_R - x_P, 0) \\ &= x_R - \alpha_R(x_P - x_R) \\ &= y - \alpha_R((M - y) - y) \\ &= y - \alpha_R M + \alpha_R 2y \end{aligned}$$

The Responder's Utility

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 $x_R = y$.

Simplifying Assumption: $x_P \geq x_R$ (i.e., $y \leq M/2$)

$$\begin{aligned} u_R(x_P, x_R) &= x_R - \alpha_R \max(x_P - x_R, 0) - \beta_R \max(x_R - x_P, 0) \\ &= x_R - \alpha_R(x_P - x_R) \\ &= y - \alpha_R((M - y) - y) \\ &= y - \alpha_R M + \alpha_R 2y \\ &= y(1 + 2\alpha_R) - \alpha_R M \end{aligned}$$

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 $x_R = y$:

$$u_R(x_P, x_R) = y(1 + 2\alpha_R) - \alpha_R M$$

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 $x_R = y$:

$$u_R(x_P, x_R) = y(1 + 2\alpha_R) - \alpha_R M$$

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

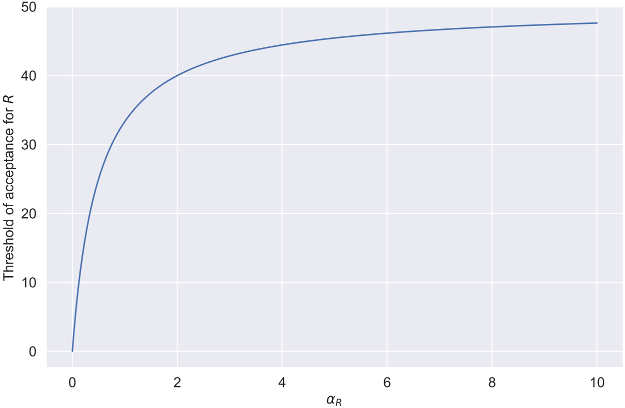

If α_R is close to zero — which indicates that R does not care much about being treated unfairly — the responder will accept very low offers. On the other hand, if α_R is sufficiently large, the offer has to be close to a half to be accepted.

The Proposer's Utility

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 $x_R = y$.

The Proposer's Utility

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 $x_R = y$.

Simplifying Assumption: $x_P \geq x_R$ (i.e., $y \leq M/2$)

The Proposer's Utility

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 $x_R = y$.

Simplifying Assumption: $x_P \geq x_R$ (i.e., $y \leq M/2$)

$$u_P(x_P, x_R) \quad = \quad x_P - \alpha_P \max(x_R - x_P, 0) - \beta_P \max(x_P - x_R, 0)$$

The Proposer's Utility

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 $x_R = y$.

Simplifying Assumption: $x_P \geq x_R$ (i.e., $y \leq M/2$)

$$\begin{aligned} u_P(x_P, x_R) &= x_P - \alpha_P \max(x_R - x_P, 0) - \beta_P \max(x_P - x_R, 0) \\ &= x_P - \beta_P(x_P - x_R) \end{aligned}$$

The Proposer's Utility

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 $x_R = y$.

Simplifying Assumption: $x_P \geq x_R$ (i.e., $y \leq M/2$)

$$\begin{aligned}
u_P(x_P, x_R) &= x_P - \alpha_P \max(x_R - x_P, 0) - \beta_P \max(x_P - x_R, 0) \\
&= x_P - \beta_P(x_P - x_R) \\
&= M - y - \beta_P((M - y) - y)
\end{aligned}$$

The Proposer's Utility

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 $x_R = y$.

Simplifying Assumption: $x_P \geq x_R$ (i.e., $y \leq M/2$)

$$\begin{aligned} u_P(x_P, x_R) &= x_P - \alpha_P \max(x_R - x_P, 0) - \beta_P \max(x_P - x_R, 0) \\ &= x_P - \beta_P(x_P - x_R) \\ &= M - y - \beta_P((M - y) - y) \\ &= M - y - \beta_P M + 2\beta_P y \end{aligned}$$

The Proposer's Utility

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 $x_R = y$.

Simplifying Assumption: $x_P \geq x_R$ (i.e., $y \leq M/2$)

$$\begin{aligned}
u_P(x_P, x_R) &= x_P - \alpha_P \max(x_R - x_P, 0) - \beta_P \max(x_P - x_R, 0) \\
&= x_P - \beta_P(x_P - x_R) \\
&= M - y - \beta_P((M - y) - y) \\
&= M - y - \beta_P M + 2\beta_P y \\
&= (1 - \beta_P)M - (1 - 2\beta_P)y
\end{aligned}$$

Proposer's Best Offer

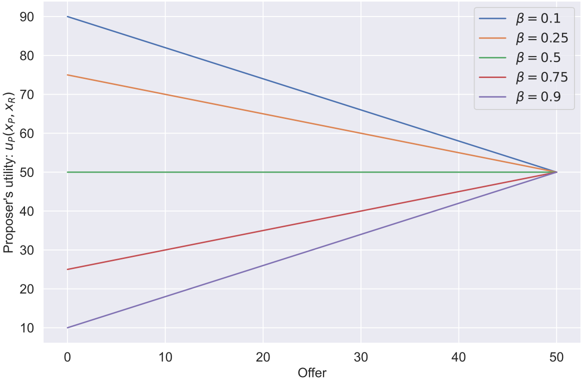

If $\beta_P > 1/2$ — if the proposer feels sufficiently guilty about treating others unfairly — then the best choice is to offer $M/2$. If $\beta < 1/2$, then the best offer is the minimum one that would be accepted, (a little bit more than $\alpha_R M/(1 + 2\alpha_R)$). If $\beta = 1/2$, then it any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

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

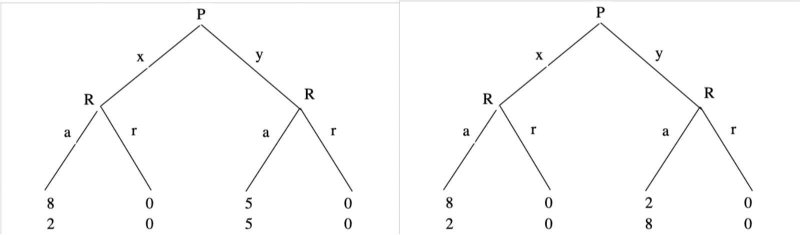

(a) (5/5)-game

(b) (2/8)-game

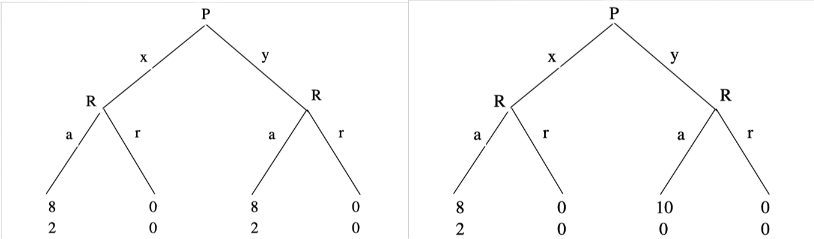

(c) (8/2)-game

(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

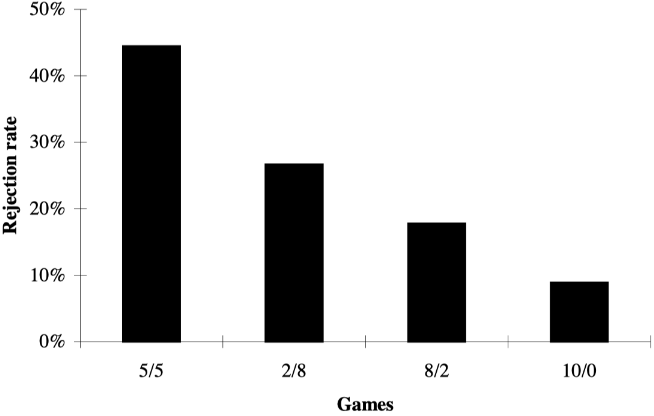

Figure 3:
Percentage and acceptance of (8/2) proposals

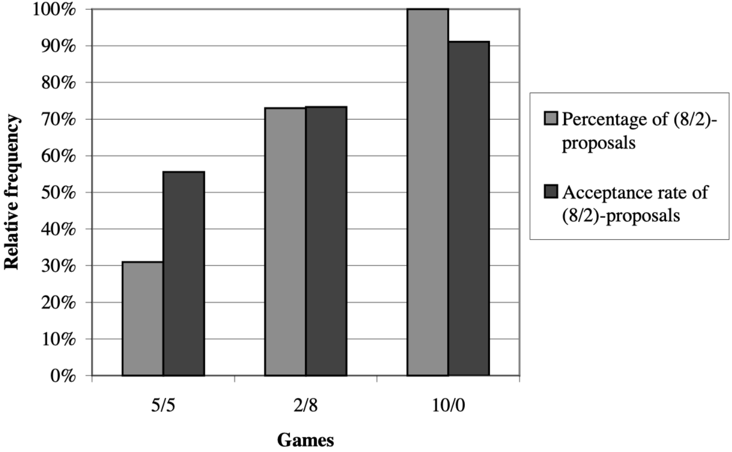

The results of our experiment clearly show that the same action by the proposer in a miniultimatum game triggers very different responses depending on the alternative action available to the proposer. **This suggests that responders do not only take into account the distributive consequences of the action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

Nash Bargaining Game

Brian Skyrms (2012). Chapters 1 & 2 in *Evolution of the Social Contract*. Cambridge University Press.

Nash Bargaining Game: Two players receive a windfall. Each player makes a demand, and if the two demands do not exceed the total good, both receive their demand. Otherwise, both receive nothing.

Nash Bargaining Game: Two players receive a windfall. Each player makes a demand, and if the two demands do not exceed the total good, both receive their demand. Otherwise, both receive nothing.

Simplified Nash Bargaining Game: Suppose that two players are dividing \$1.

Nash Bargaining Game: Two players receive a windfall. Each player makes a demand, and if the two demands do not exceed the total good, both receive their demand. Otherwise, both receive nothing.

Simplified Nash Bargaining Game: Suppose that two players are dividing \$1. The players write their demand on a sheet of paper and hand it to a referee. If the total is less than \$1, then each player gets what they demand. But, if the total is greater than \$1, then each player gets nothing.

Nash Bargaining Game: Two players receive a windfall. Each player makes a demand, and if the two demands do not exceed the total good, both receive their demand. Otherwise, both receive nothing.

Simplified Nash Bargaining Game: Suppose that two players are dividing $1. The players write their demand on a sheet of paper and hand it to a referee. If the total is less than $1, then each player gets what they demand. But, if the total is greater than $1, then each player gets nothing. Assume that each player has 5 available actions:

1. 0: Demand nothing
2. 1/3: Demand 1/3 of $1 (so 33 cents)
3. 1/2: Demand 1/2 of $1 (so 50 cents)
4. 2/3: Demand 2/3 of $1 (so 66 cents)
5. 1: Demand $1

Simplified Nash Bargaining Game

Player 2

Player 1

	0	1/3	1/2	2/3	1
0	$0, 0$	$0, \frac{1}{3}$	$0, \frac{1}{2}$	$0, \frac{2}{3}$	$0, 1$
1/3	$\frac{1}{3}, 0$	$\frac{1}{3}, \frac{1}{3}$	$\frac{1}{3}, \frac{1}{2}$	$\frac{1}{3}, \frac{2}{3}$	$0, 0$
1/2	$\frac{1}{2}, 0$	$\frac{1}{2}, \frac{1}{3}$	$\frac{1}{2}, \frac{1}{2}$	$0, 0$	$0, 0$
2/3	$\frac{2}{3}, 0$	$\frac{2}{3}, \frac{1}{3}$	$0, 0$	$0, 0$	$0, 0$
1	$1, 0$	$0, 0$	$0, 0$	$0, 0$	$0, 0$

Simplified Nash Bargaining Game

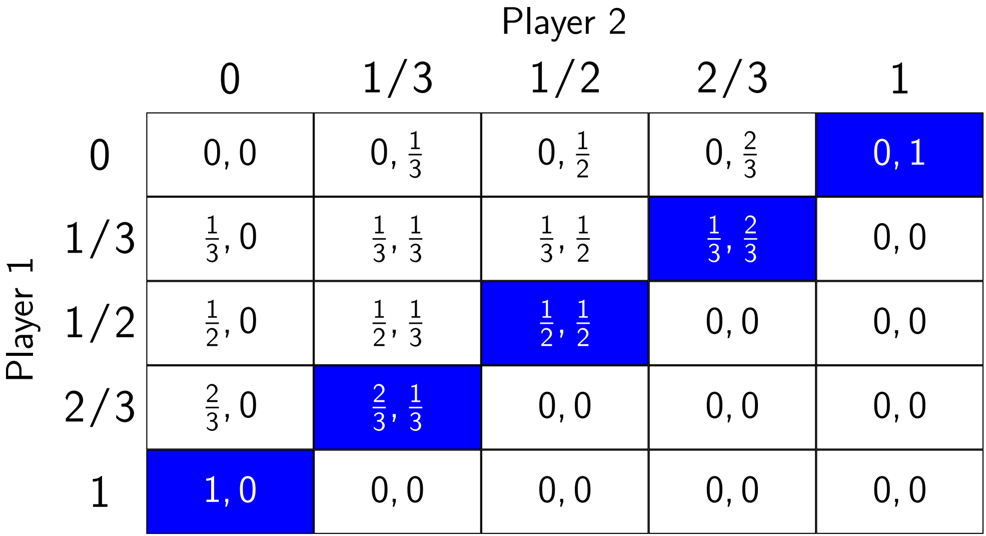

		Player 2				
		0	1/3	1/2	2/3	1
Player 1	0	$0, 0$	$0, \frac{1}{3}$	$0, \frac{1}{2}$	$0, \frac{2}{3}$	$0, 1$
	1/3	$\frac{1}{3}, 0$	$\frac{1}{3}, \frac{1}{3}$	$\frac{1}{3}, \frac{1}{2}$	$\frac{1}{3}, \frac{2}{3}$	$0, 0$
	1/2	$\frac{1}{2}, 0$	$\frac{1}{2}, \frac{1}{3}$	$\frac{1}{2}, \frac{1}{2}$	$0, 0$	$0, 0$
	2/3	$\frac{2}{3}, 0$	$\frac{2}{3}, \frac{1}{3}$	$0, 0$	$0, 0$	$0, 0$
	1	$1, 0$	$0, 0$	$0, 0$	$0, 0$	$0, 0$