PHPE 400
Individual and Group Decision Making

Eric Pacuit
University of Maryland
pacuit.org

Why play a Nash equilibrium?

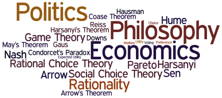

Self-Enforcing Agreements: Nash equilibria are recommended by being the only strategy combinations on which the players could make self-enforcing agreements, i.e., agreements that each has reason to respect, even without external enforcement mechanisms.

M. Risse (2000). *What is rational about Nash equilibria?*. Synthese, 124:3, pp. 361 - 384.

Column

Row		l	c	r
	u	4, 6	5, 4	0, 0
	m	5, 7	4, 8	0, 0
	d	0, 0	0, 0	1, 1

		Column		
		l	c	r
	u	$4, \underline{6}$	$\underline{5}, 4$	$0, 0$
Row	m	$\underline{5}, 7$	$4, \underline{8}$	$0, 0$
	d	$0, 0$	$0, 0$	$\underline{1}, \underline{1}$

(d, r) is a Nash equilibrium, but it is **not self-enforcing**

Row		Column	
		l	r
	u	$0, 0$	$4, 2$
	d	$2, 4$	$3, 3$

		Column	
		l	r
Row	u	$0, 0$	$\underline{4}, \underline{2}$
	d	$\underline{2}, \underline{4}$	$3, 3$

(d, r) is **not** a Nash equilibrium, but it is **self-enforcing**

Self-Enforcing Agreements: Nash equilibria are recommended by being the only strategy combinations on which the players could make self-enforcing agreements, i.e., agreements that each has reason to respect, even without external enforcement mechanisms.

- There are Nash equilibria that are not self-enforcing
- There are self-enforcing outcomes that are not Nash equilibria

Is a Nash equilibrium *guaranteed* by players that are rational and have *common knowledge* of each others' rationality?

Column

Row

	l	c	r
u	$3, 2$	$0, 0$	$2, 3$
m	$0, 0$	$1, 1$	$0, 0$
d	$2, 3$	$0, 0$	$3, 2$

		Column		
		l	c	r
Row	u	$\underline{3}, 2$	$0, 0$	$2, \underline{3}$
	m	$0, 0$	$\underline{1}, \underline{1}$	$0, 0$
	d	$2, \underline{3}$	$0, 0$	$\underline{3}, 2$

(m, c) is the unique Nash equilibrium

Column

Row	l	c	r
u	$\underline{3}, 2$	$0, 0$	$2, \underline{3}$
m	$0, 0$	$\underline{1}, \underline{1}$	$0, 0$
d	$2, \underline{3}$	$0, 0$	$\underline{3}, 2$

u, d, l, and r are all **rationalizable**

Column

Row		l	c	r
	u	$\underline{3}, 2$	$0, 0$	$2, \underline{3}$
	m	$0, 0$	$\underline{1}, \underline{1}$	$0, 0$
	d	$2, \underline{3}$	$0, 0$	$\underline{3}, 2$

Row plays d because she thought Column would play r

Column

Row	l	c	r
u	$\underline{3}, 2$	$0, 0$	$2, \underline{3}$
m	$0, 0$	$\underline{1}, \underline{1}$	$0, 0$
d	$2, \underline{3}$	$0, 0$	$\underline{3}, 2$

Column plays l because she thought Row would play d

Row \ Column	l	c	r
u	$\underline{3}, 2$	$0, 0$	$2, \underline{3}$
m	$0, 0$	$\underline{1}, \underline{1}$	$0, 0$
d	$2, \underline{3}$	$0, 0$	$\underline{3}, 2$

Column was correct, but Row was wrong.
Both players are **rational**.

Column

Row	l	c	r	x
u	$\underline{3}, 2$	$0, 0$	$2, \underline{3}$	$0, -5$
m	$0, 0$	$\underline{1}, \underline{1}$	$0, 0$	$\underline{100}, -5$
d	$2, \underline{3}$	$0, 0$	$\underline{3}, 2$	$1, -3$

Not every strategy is rationalizable

Column

Row	l	c	r	x
u	$\underline{3}, 2$	$0, 0$	$2, \underline{3}$	$0, -5$
m	$0, 0$	$\underline{1}, \underline{1}$	$0, 0$	$\underline{100}, -5$
d	$2, \underline{3}$	$0, 0$	$\underline{3}, 2$	$1, -3$

Not every strategy is rationalizable:
Row can't play m *because* she thinks Column will play x

An action a **strictly dominates** another action b for player i when i's utility is strictly better choosing a than choosing b no matter what actions are chosen by the other players.

Example

Row \ Column	l	r
u	$5, 5$	$-100, 4$
d	$0, 1$	$0, 0$

Example

Column

Row		l	r
	u	$5, 5$	$-100, 4$
	d	$0, 1$	$0, 0$

Since r is **strictly dominated** by l, Column will not play r.
Then, the best response for Row is u.

Important Games

Coordination

		Column	
		a	b
Row	a	$\underline{1}, \underline{1}$	$0, 0$
	b	$0, 0$	$\underline{1}, \underline{1}$

- Both (a, a) and (b, b) are Nash equilibria
- Both (a, a) and (b, b) are Pareto optimal
- The players want to coordinate by choosing the *same* action a or b.

Anti-Coordination

		Column	
		a	b
Row	a	$0, 0$	$\underline{1}, \underline{1}$
	b	$\underline{1}, \underline{1}$	$0, 0$

- Both (a, b) and (b, a) are Nash equilibria
- Both (a, b) and (b, a) are Pareto optimal
- The players want to mis-coordinate in which one player chooses a and the other chooses b.

Coordination and Competition

		Column	
		a	b
Row	a	$0, 0$	$\underline{2}, \underline{1}$
	b	$\underline{1}, \underline{2}$	$0, 0$

- Both (a, b) and (b, a) are Nash equilibria
- Both (a, b) and (b, a) are Pareto optimal
- Players want to mis-coordinate, and both prefer choosing b while the other chooses a.

Cailin O’Connor (2019). *The Origins of Unfairness: Social Categories and Cultural Evolution*. Oxford University Press.

Chicken

		Column	
		a	b
Row	a	$2, 2$	$\underline{1}, \underline{3}$
	b	$\underline{3}, \underline{1}$	$0, 0$

- Both (a, b) and (b, a) are Nash equilibria
- All profiles except (b, b) are Pareto optimal
- Also called the "hawk-dove game"

Stag-Hunt

		Column	
		a	b
Row	a	$\underline{3}, \underline{3}$	$0, 2$
	b	$2, 0$	$\underline{1}, \underline{1}$

- (a, a) and (b, b) are the Nash equilibria
- (a, a) Pareto dominates (b, b)
- Choosing a may lead to a better outcome, but it is riskier.

Stag-Hunt

B. Skyrms (2004). *The Stag Hunt and the Evolution of Social Structure*. Cambridge University Press.

Prisoner's Dilemma

Row \ Column	a	b
a	$3, 3$	$0, \underline{4}$
b	$\underline{4}, 0$	$\underline{1}, \underline{1}$

- (b, b) is the only Nash equilibrium
- (a, a) Pareto dominates (b, b)
- Typically, a is the "cooperate" action and b is the "defect" action.
- Often used to represent conflicts between individual rationality and cooperative behavior.

Prisoner's Dilemma

- Athletes using performance-enhancing drugs
- Two competing companies deciding advertising budgets
- Nation-states deciding to restrict $CO2$ emissions
- Two people meet and exchange closed bags, with the understanding that one of them contains money, and the other contains a purchase. Either player can choose to honor the deal by putting into his or her bag what he or she agreed, or he or she can defect by handing over an empty bag.
- `http://www.radiolab.org/story/golden-rule/`

Prisoner’s Dilemma

“Game theorists think it just plain wrong to claim that the Prisoners’ Dilemma embodies the essence of the problem of human cooperation. On the contrary, it represents a situation in which the dice are as loaded against the emergence of cooperation as they could possibly be. If the great game of life played by the human species were the Prisoner’s Dilemma, we wouldn’t have evolved as social animals!. . .. No paradox of rationality exists. Rational players don’t cooperate in the Prisoners’ Dilemma, because the conditions necessary for rational cooperation are absent in this game.” (Binmore, p. 63)

K. Binmore (2005). *Natural Justice*. Oxford University Press.

Prisoner's Dilemma

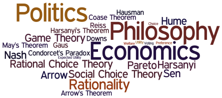

- S. Kuhn, Prisoner's Dilemma, Stanford Encyclopedia of Philosophy, plato.stanford.edu/entries/prisoner-dilemma/

- W. Poundstone, Prisoner's Dilemma, Anchor, 1993

Two Important Games

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.

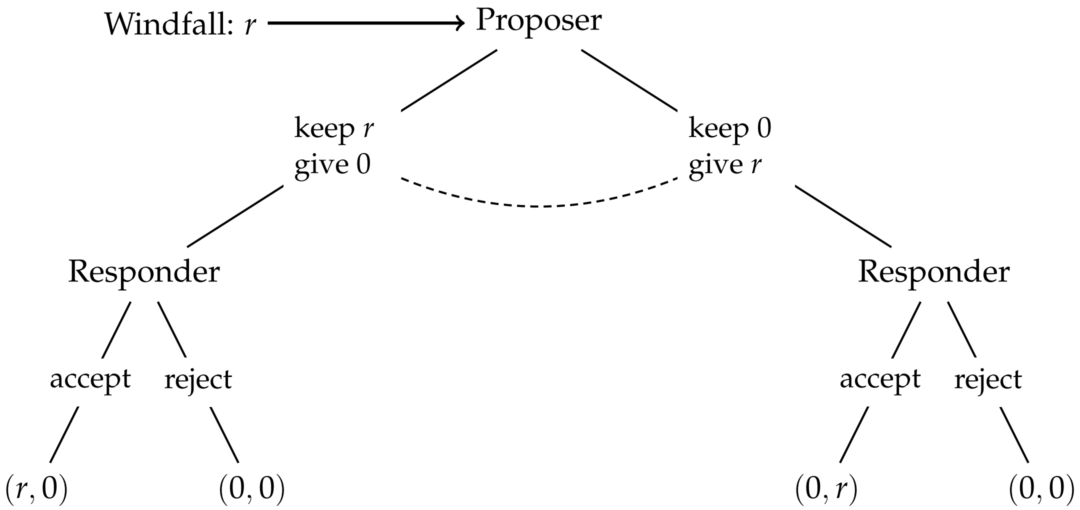
Windfall: r
Proposer
keep r
give 0
keep 0
give r
Responder
Responder
accept
reject
accept
reject
(r, 0)
(0, 0)
(0, r)
(0, 0)

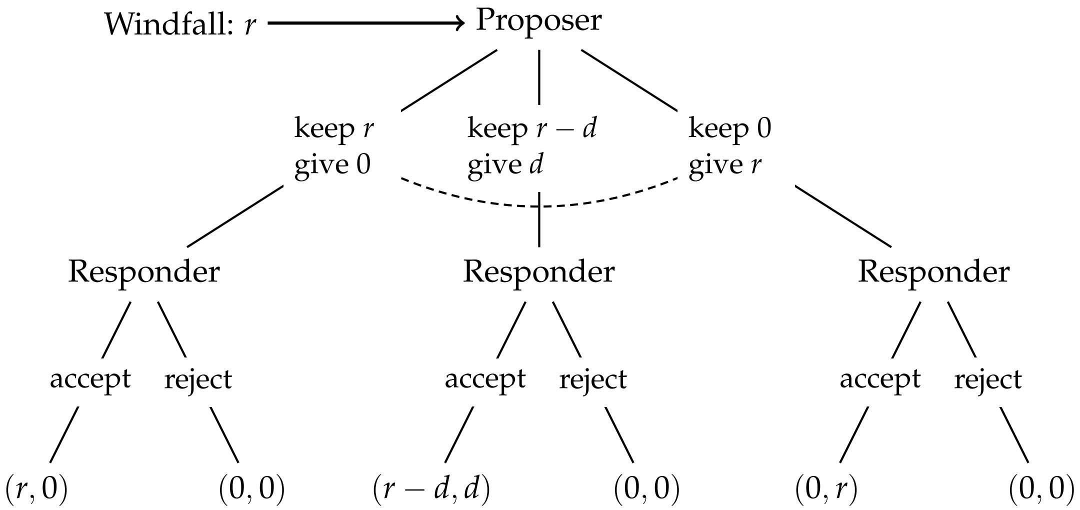
Windfall: r
Proposer
keep r
give 0
keep $r-d$
give d
keep 0
give r
Responder
Responder
Responder
accept
reject
accept
reject
accept
reject
$(r, 0)$
$(0, 0)$
$(r-d, d)$
$(0, 0)$
$(0, r)$
$(0, 0)$

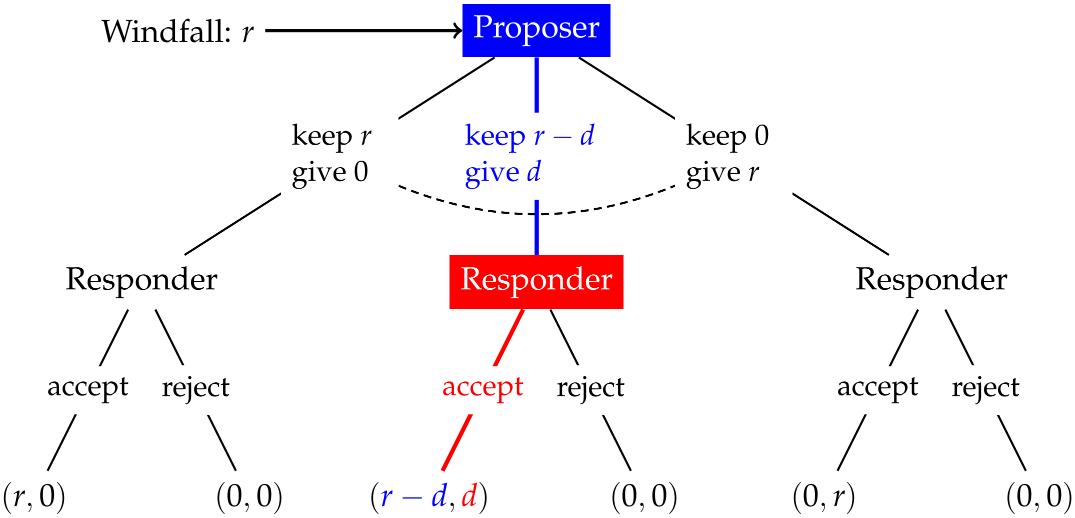

Proposer gets $r - d$ and Responder gets d

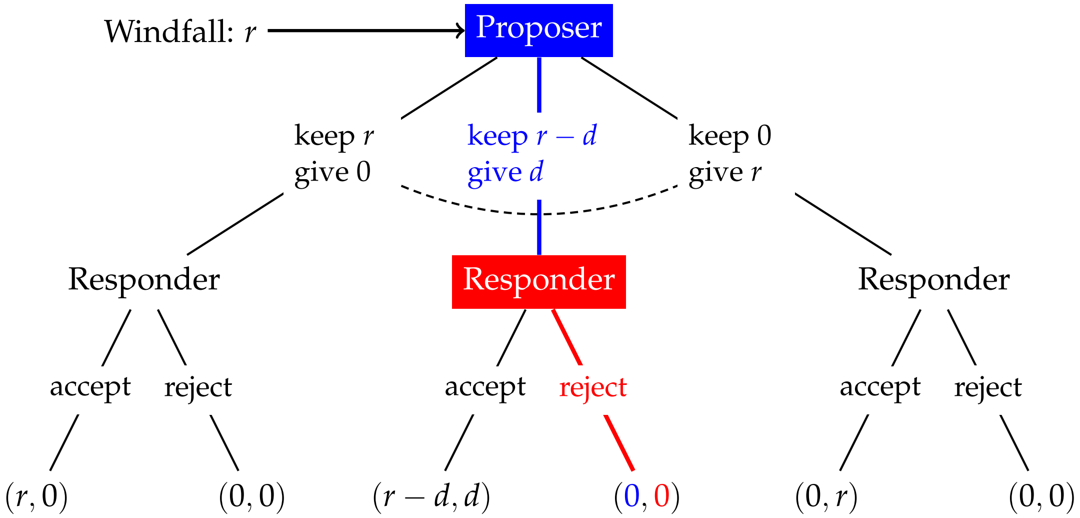

Proposer gets 0 and Responder gets 0

Sequential Rationality

If the proposer offers a split which gives the second any positive amount, the second does strictly worse by refusing the offer. So, no rejection strategies are sequentially rational.

Knowing this, the first player ought to offer the smallest amount possible to the second player.

This is not what is observed:

> ...offers typically average about 30-40 percent of the total, with a 50-50 split often the mode. Offers of less than 20 percent are frequently rejected. These facts are not now in question. What remains controversial is how to interpret the facts and how best to incorporate what we have learned into a more descriptive version of game theory.
>
> (p. 210, Camerer and Thaler)

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.

- Rejecting low offers is impossible to reconcile with a theory of *payoff maximization*.

- Rejecting low offers is impossible to reconcile with a theory of *payoff maximization*.

- Making a non-zero offer is consistent with payoff maximization, if a proposer believes that the responder will reject too low an offer.

- Rejecting low offers is impossible to reconcile with a theory of *payoff maximization*.

- Making a non-zero offer is consistent with payoff maximization, if a proposer believes that the responder will reject too low an offer.
  - However, offers are typically larger than the amount that proposers believe would result in acceptance.

    Joseph Henrich, Robert Boyd, Samuel Bowles, Colin Camerer, Ernst Fehr, Herbert Gintis, and Richard McElreath (2001). *In search of homo economicus: Behavioral experiments in 15 small-scale societies*. American Economic Review, 91(2), pp. 73–78.

Subjects in interpersonal experiments like the ultimatum game may be influenced by all kinds of factors: the wording of the instructions, the identity of the experimenters, whether the experiment is thought to be “economics” or “psychology,” and so forth. This means that initial results should be interpreted cautiously.

Subjects in interpersonal experiments like the ultimatum game may be influenced by all kinds of factors: the wording of the instructions, the identity of the experimenters, whether the experiment is thought to be "economics" or "psychology," and so forth. This means that initial results should be interpreted cautiously. At this point in ultimatum game research, enough independent studies have now been carried out with original designs and instructions to be confident that the basic phenomena are robust.

Subjects in interpersonal experiments like the ultimatum game may be influenced by all kinds of factors: the wording of the instructions, the identity of the experimenters, whether the experiment is thought to be "economics" or "psychology," and so forth. This means that initial results should be interpreted cautiously. At this point in ultimatum game research, enough independent studies have now been carried out with original designs and instructions to be confident that the basic phenomena are robust. The closely related "dictator game," however, turns out to be very sensitive to design issues.

(Camerer and Thaler, p. 213)

Dictator Game

In the dictator game, the first player, called the Allocator, makes a unilateral decision regarding the split of the pie. The second player, the Recipient, has no choice and receives only the amount that the dictator decides to give.

Since dictators have no monetary incentives to give, a payoff-maximizing dictator would keep the whole amount.

Dictator Game

Experimental Regularity: A significant number of Allocators give some money in the dictator game. Moreover, the distribution of donations tend to be bimodal, with peaks at zero and at half the total.

Daniel Kahneman, Jack L. Knetsch, and Richard Thaler (1986). *Fairness as a Constraint on Profit Seeking: Entitlements in the Market*. American Economic Review, 76, pp. 728 - 741.

Christoph Engel (2011). *Dictator games: A meta study*. Experimental Economics, 14(4), pp. 583 - 610.

The Dictator Game

The original dictator game experiments were used to help determine the extent to which generous offers in ultimatum games occurred because Proposers were fair-minded or because Proposers feared having low offers rejected.

The Dictator Game

The original dictator game experiments were used to help determine the extent to which generous offers in ultimatum games occurred because Proposers were fair-minded or because Proposers feared having low offers rejected.

- Offers in the dictator game are lower than in ultimatum games, but (in most variations) are still positive.

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:

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

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.
- 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 $u_R(x_P, x_R)$.

- The standard assumption is that players are *payoff maximizing*:
  - If $x_P < y_P$, then $u_P(x_P, x_R) < u_P(y_P, y_R)$ (and similarly for player R).
  - For simplicity we often identify money with utility, so $u_P(x_P, x_R) = x_P$ and $u_R(x_P, x_R) = x_R$; but this is not necessary.

Methodological Individualism

Traditional economic models presume that individuals do not take an interest in the interests of those with whom they interact. More particularly, the assumption of *non-tuism* implies that the utility function of each individual, as a measure of her preferences, is strictly independent of the utility functions of those with whom she interacts.

Methodological Individualism

Traditional economic models presume that individuals do not take an interest in the interests of those with whom they interact. More particularly, the assumption of *non-tuism* implies that the utility function of each individual, as a measure of her preferences, is strictly independent of the utility functions of those with whom she interacts. ...Interestingly, this idea is quite different from the usual egoistic assumption: a non-tuist may be a caring, altruistic human being, but when involved in an economic exchange, she must necessarily regard her own interest as paramount.

Methodological Individualism

Traditional economic models presume that individuals do not take an interest in the interests of those with whom they interact. More particularly, the assumption of *non-tuism* implies that the utility function of each individual, as a measure of her preferences, is strictly independent of the utility functions of those with whom she interacts. ...Interestingly, this idea is quite different from the usual egoistic assumption: a non-tuist may be a caring, altruistic human being, but when involved in an economic exchange, she must necessarily regard her own interest as paramount. Thus non-tuism is important insofar as it defines the scope of economic activities. When tuism to some degree motivates one's conduct, then it ceases to be wholly economic.

Cristina Bicchieri and Jiji Zhang (2012). *An Embarrassment of Riches: Modeling Social Preferences in Ultimatum Games*. Handbook of the Philosophy of Science, Volume 13: Philosophy of Economics.

Inequality Aversion: Fehr and Schmidt Utility

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

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

Inequality Aversion: Fehr and Schmidt Utility

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

- 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

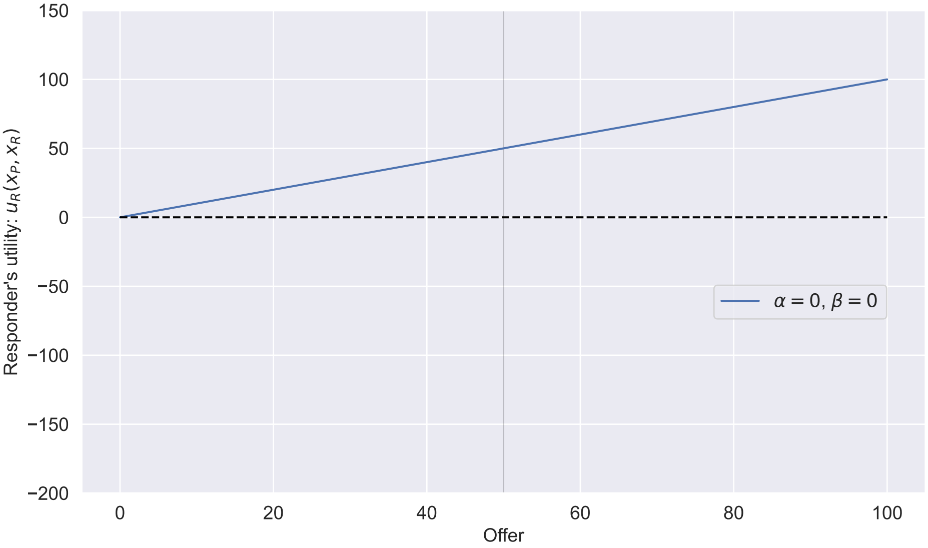

Responder's Utility

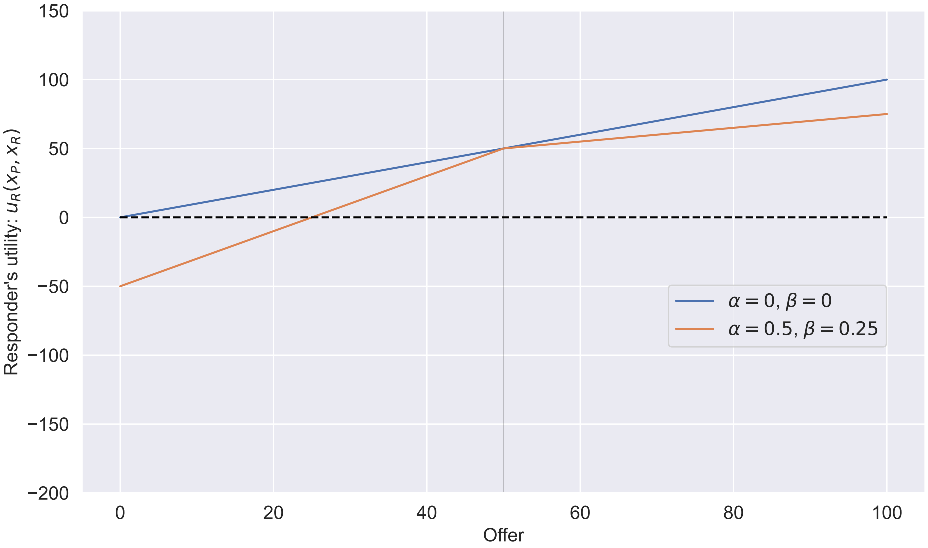

Responder’s Utility

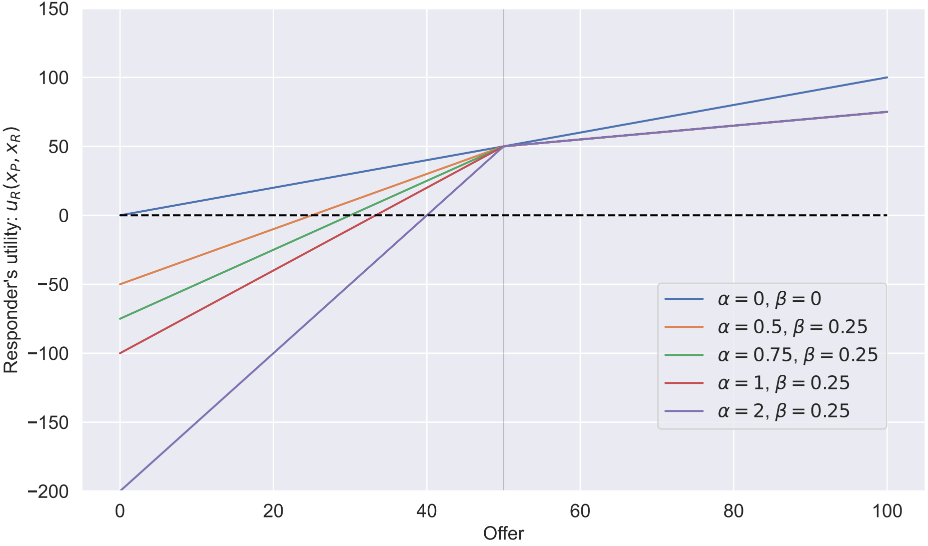

The Fehr-Schmidt Utility Function

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.

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.

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

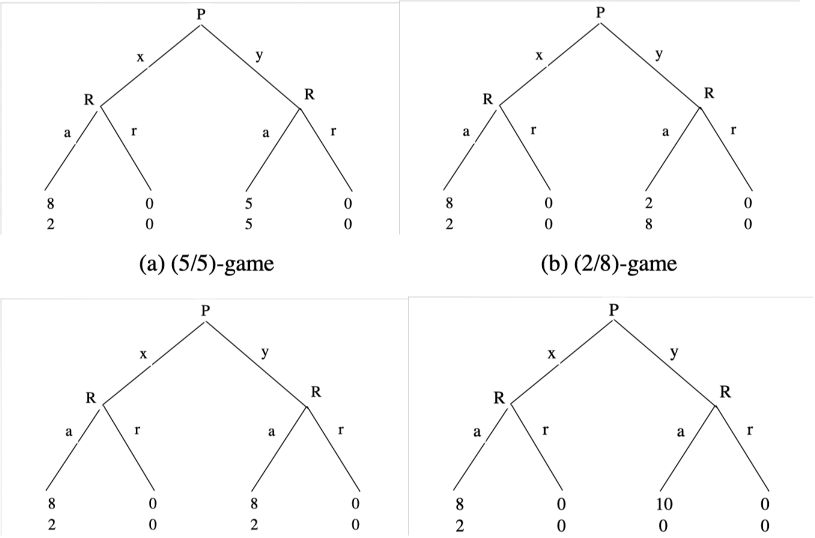

(a) (5/5)-game

(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

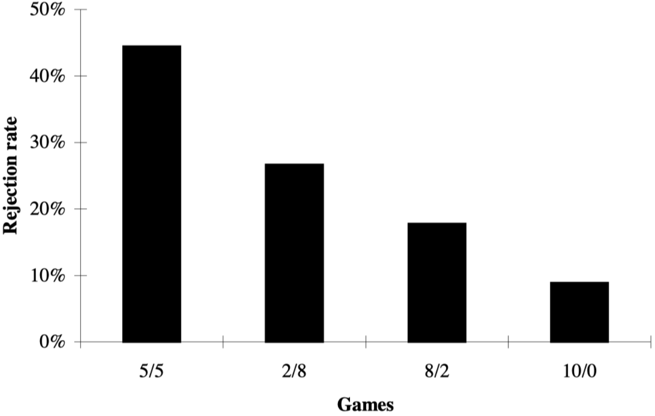

Figure 3:
Percentage and acceptance of (8/2) proposals

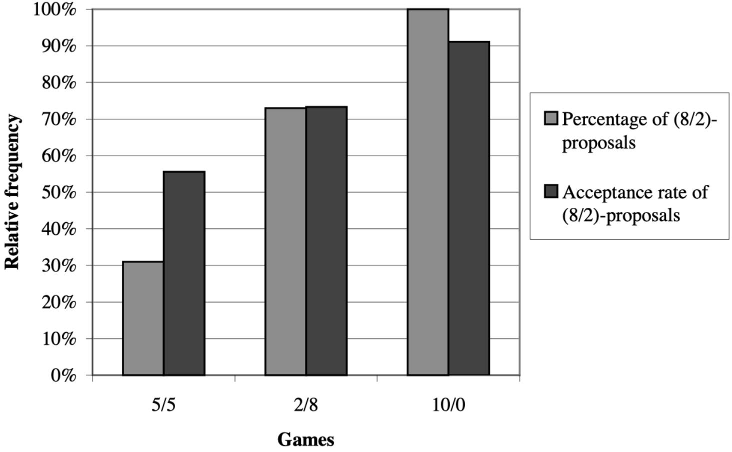

The results of our experiment clearly show that the same action by the proposer in a miniultimatum game triggers very different responses depending on the alternative action available to the proposer. This suggests that responders do not only take into account the distributive consequences of the action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

Concluding Remarks

"Rationality has a clear interpretation in individual decision making, but it does not transfer comfortably to interactive decisions, because interactive decision makers cannot maximize expected utility without strong assumptions about how the other participant(s) will behave. In game theory, common knowledge and rationality assumptions have therefore been introduced, but under these assumptions, rationality does not appear to be characteristic of social interaction in general." (Colman, 152)

A. Colman (2003). *Cooperation, psychological game theory, and limitations of rationality in social interaction*. Behavioral and Brain Sciences, 26, pp. 139 - 198.