PHIL 408Q/PHPE 308D
Fairness

Eric Pacuit, University of Maryland

January 30, 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.

- 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.

Robustness of the Experiments

Does the overall amount of money matter?

Robustness of the Experiments

Does the overall amount of money matter?

- Passing up \$0.01 (10% of \$1) or \$1 (10% of \$10) because it is unfair is very different than passing up \$10 (10% of \$100), \$100 (10% of \$1,000) or \$100,000 (10% of \$1,000,000), ...

Robustness of the Experiments

Does the overall amount of money matter?

- Passing up \$0.01 (10% of \$1) or \$1 (10% of \$10) because it is unfair is very different than passing up \$10 (10% of \$100), \$100 (10% of \$1,000) or \$100,000 (10% of \$1,000,000), ...

- Telser (1993) offers a carefully reasoned argument along the following lines: Think of fairness as a commodity such that consumption is maximized when the split is equal. It is reasonable to assume that fairness varies inversely with its price. Therefore, if we raise the stakes the price of fairness increases; i.e. offers of some small percentage of the total get larger, and it is not implausible that such amounts will be accepted for large M.

  (Hoffman et al., p. 292)

Robustness of the Experiments

Does the overall amount of money matter?

- Although Telser's example involves raising the stakes to millions of dollars, we propose (consistent with our budget limitations) experiments in which the stakes are raised from \$10 to \$100. (Hoffman et al., p. 292)

Elizabeth Hoffman, Kevin McCabe, and Vernon Smith (1996). *On Expectations and the Monetary Stakes in Ultimatum Games*. International Journal of Game Theory, 25, pp. 289 - 301.

Robustness of the Experiments

Does the overall amount of money matter?

- No

  The results of this expensive experiment show that it is safe to go back to more affordable ultimatum experiments, since they generate data that are indistinguishable from highstakes data. (Camerer and Thaler, p. 211)

  (although would you *really* pass up 10% of $1,000,000 because it is unfair??)

Robustness of the Experiments

What about cultural differences? Does the nationality of the subjects make a difference?

Robustness of the Experiments

What about cultural differences? Does the nationality of the subjects make a difference?

- While there are some intriguing differences, results are generally the same across a wide-range of cultures.

Alvin E. Roth, Vesna Prasnikar, Shmuel Zamir, and Masahiro Okuno-Fujiwar (1991). *Bargaining and Market Behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo: An Experimental Study*. American Economic Review, 81, pp. 1068 - 1095.

H. Oosterbeek, R. Sloof, and G. van de Kuilen (2004). *Cultural Differences in Ultimatum Game Experiments: Evidence from a Meta-Analysis*. Experimental Economics, 7, pp. 171 - 188.

If we accept that the basic empirical data on the ultimatum game is capturing a real phenomenon, then we can ask what is going on.

If Responders reject small offers because they deem them to be unfair, then their willingness to reject should depend on what they think the Proposers are keeping for themselves.

John H. Kagel, Chung Kim, Donald Mose (1996). *Fairness in Ultimatum Games with Asymmetric Information and Asymmetric Payoffs*. Games and Economic Behavior, 13(1), pp. 100 - 110.

Asymmetric Information

- The players are asked to divide 100 chips. The experiments manipulate both the value of the chips (each chip is worth 10 cents or 30 cents) and who knows what about their relative value.

Asymmetric Information

- The players are asked to divide 100 chips. The experiments manipulate both the value of the chips (each chip is worth 10 cents or 30 cents) and who knows what about their relative value.
- Players always know their own chip valuation, but in some conditions they do not know the other player's value. Who knows what is common knowledge.

Asymmetric Information

- The players are asked to divide 100 chips. The experiments manipulate both the value of the chips (each chip is worth 10 cents or 30 cents) and who knows what about their relative value.
- Players always know their own chip valuation, but in some conditions they do not know the other player's value. Who knows what is common knowledge.
- There are ten repetitions of the game, always under the same informational conditions, against ten different opponents. (One round is picked at random to count for the actual monetary payoff.)

Asymmetric Information

- The players are asked to divide 100 chips. The experiments manipulate both the value of the chips (each chip is worth 10 cents or 30 cents) and who knows what about their relative value.
- Players always know their own chip valuation, but in some conditions they do not know the other player's value. Who knows what is common knowledge.
- There are ten repetitions of the game, always under the same informational conditions, against ten different opponents. (One round is picked at random to count for the actual monetary payoff.)
- In the case when both players have the same payoff, and this in known to both, we have a standard ultimatum game, and offers tend to converge to 50 percent.

Asymmetric Information

Do Proposers want to *be* fair or *seem* fair?

Asymmetric Information

Do Proposers want to *be* fair or *seem* fair?

- The Proposer knows that the chips are worth 30 cents to him but only worth 10 cents to the Responder, an equal division of money would imply an offer of 75 percent of the chips to the Responder.

Asymmetric Information

Do Proposers want to *be* fair or *seem* fair?

- The Proposer knows that the chips are worth 30 cents to him but only worth 10 cents to the Responder, an equal division of money would imply an offer of 75 percent of the chips to the Responder.
- If only the Proposer knows that he has the higher valuation, he can offer 50 percent and still seem fair.

Asymmetric Information

Do Proposers want to *be* fair or *seem* fair?

- The Proposer knows that the chips are worth 30 cents to him but only worth 10 cents to the Responder, an equal division of money would imply an offer of 75 percent of the chips to the Responder.
- If only the Proposer knows that he has the higher valuation, he can offer 50 percent and still seem fair.
- The data imply that the appearance of fairness is enough; offers in this condition are close to 50 percent, and since Responders have no reason to think that the 50 percent offer is anything but fair, rejections are rare.

Asymmetric Information

Do Proposers want to *be* fair or *seem* fair?

- The Proposer knows that the chips are worth 30 cents to him but only worth 10 cents to the Responder, an equal division of money would imply an offer of 75 percent of the chips to the Responder.
- If only the Proposer knows that he has the higher valuation, he can offer 50 percent and still seem fair.
- The data imply that the appearance of fairness is enough; offers in this condition are close to 50 percent, and since Responders have no reason to think that the 50 percent offer is anything but fair, rejections are rare.
- These results serve as an important reminder that self-interested behavior is alive and well, even in ultimatum games.

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.

Manipulating the Social Distance

Elizabeth Hoffman, Kevin McCabe and Vernon L. Smith (1996). *Social Distance and Other-Regarding Behavior in Dictator Games*. The American Economic Review, 86(3), pp. 653660.

In laboratory experiments we cannot assume that subjects behave as if the world is completely defined by the experimenter. Past experience is important in so far as beliefs are based on experience....

Manipulating the Social Distance

Elizabeth Hoffman, Kevin McCabe and Vernon L. Smith (1996). *Social Distance and Other-Regarding Behavior in Dictator Games*. The American Economic Review, 86(3), pp. 653660.

In laboratory experiments we cannot assume that subjects behave as if the world is completely defined by the experimenter. Past experience is important in so far as beliefs are based on experience....In short, subjects bring their ongoing repeated game experience and reputations from the world into the laboratory, and the instructional language, especially in single-play sensitive experiments like the dictator game, can subtly suggest more or less isolation from that interactive experience. (p. 655)

Manipulating the Social Distance

When the instructions state that the subject and his or her counterpart "has been provisionally allocated \$10," and suggesting that the task is to "divide" the \$10 may imply that the objective is to share the money with someone, who, though anonymous, is socially relatively near to the decision maker.

Manipulating the Social Distance

When the instructions state that the subject and his or her counterpart "has been provisionally allocated $10," and suggesting that the task is to "divide" the $10 may imply that the objective is to share the money with someone, who, though anonymous, is socially relatively near to the decision maker. At the other pole, defining the greatest social distance, is our Double Blind procedure, which goes to some pains to guarantee the decision maker absolute privacy, and isolation from any social consequence or association with the person's decision.

(Hoffman et al., p. 655)

Manipulating the Social Distance

When the instructions state that the subject and his or her counterpart "has been provisionally allocated $10," and suggesting that the task is to "divide" the $10 may imply that the objective is to share the money with someone, who, though anonymous, is socially relatively near to the decision maker. At the other pole, defining the greatest social distance, is our Double Blind procedure, which goes to some pains to guarantee the decision maker absolute privacy, and isolation from any social consequence or association with the person's decision.

(Hoffman et al., p. 655)

There are 6 variations of the experiment that vary the "social distance" between the Allocator and the Recipient, labeled DB1, DB2, SB1, SB2, FHSS V, FHSS R

Manipulating the Social Distance

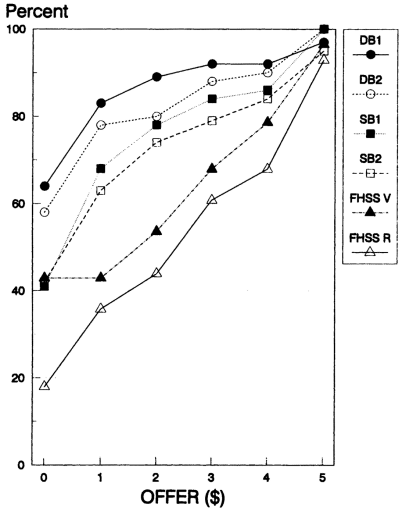

FIGURE 1. CUMULATIVE DISTRIBUTIONS FOR DICTATOR EXPERIMENTS

Manipulating the Social Distance

[W]hat is it that is being consumed when someone rejects an offer in an ultimatum game, or when someone gives money away in either the ultimatum or the dictator experiments. From the perspective of this experiment the answer, which we will call reputation (or image), is largely explained as self-regarding, that is, people act as if they are other regarding because they are better off with the resulting reputation. Only under conditions of social isolation are these reputational concerns of little force. (Hoffman et al., p. 659)

Manipulating the Social Distance

Subjects are handed $10 in manna from experimental heaven and asked whether they would like to share some of it with a stranger who is in the same room. Many do. However, if the first player is made to feel as if he earned the right to the $10, or the relationship with the other player is made less personal, then sharing shrinks. Etiquette may require you to share a windfall with a friend, but it certainly does not require you to give up some of your hard-earned year-end bonus to a stranger. (Camerer and Thaler, p. 216)

Why does game theory fail as a predictive model in ultimatum and dictator games? These games are so simple that we can rule out rationality as the source of any problem, so the difficulty presumably has something to do with the assumption that the players are income maximizers.

(Camerer and Thaler, p. 216)

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 $u_R(x_P, x_R)$.

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, Selfish, Egoistic Preferences

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, Selfish, Egoistic Preferences

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, Selfish, Egoistic Preferences

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.

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.