Fairness and Social Preferences

Behavioral Economics

Patrick McAlvanah

Fairness

What is the standard economic model? Pure self-interest.

From Adam Smith's Wealth of Nations:

> **"It is not from the benevolence of the butcher, the brewer, or the baker that we expect our dinner, but from their regard for their own self-interest. We address ourselves not to their humanity, but to their self-love, and never talk to them of our necessities, but of their advantage."**

However, fairness matters. In particular, people don't just care about final outcomes, they also care about the intent behind those outcomes.

Fairness

Kahneman, Knetsch, Thaler – phone surveys

Two goals: 1) Identify community standards of fairness
2) Consider implications for market outcomes

A football team normally sells tickets on the day of their games. Recently, interest in the next game has increased greatly, and tickets are in great demand. The team owners can distribute the tickets in one of three ways.

1) **By auction: the tickets are sold to the highest bidders.**
2) **By lottery: the tickets are sold to the people whose names are drawn**
3) **By queue: the tickets are sold on a first-come first-served basis**

Fairness

Rank these 3 in order of what you feel is the most fair and least fair:

TABLE 1 Ranking of Allocation Methods

Allocation Method	Most Fair (%)	Least Fair (%)
Auction	4	75
Lottery	28	18
Queue	68	7

Notice that in this case, people's fairness judgment rankings are the exact inverse of the economic efficiency rankings

Q1) A hardware store has been selling snow shovels for $15. The morning after a snow storm, the store raises the price to $20.

Acceptable 18% Unfair 82%

Fairness

Q 2A) A small photocopying shop has one employee who has worked in the shop for 6 months and earns $9 an hour. Business continues to be satisfactory but unemployment has increased in the area. Other firms have now hired similar workers at $7. The owner of the photoshop reduces the employee's wage to $7. **Acceptable 17%** **Unfair 83%**

Q 2B) A small photocopying shop … The current employee leaves, and the owner decides to pay a replacement $7 an hour.

Acceptable 73% **Unfair 27%**

The basic finding is that standards of fairness are determined by a reference transaction, which yields dual entitlement: a relevant reference price or wage for consumers, and a reference profit level for the firm.

Fairness

Firms are not allowed to increase profits by violating the entitlements of the reference transaction, as in the snowstorm example.

In the photoshop wage example, the current wage of the employee serves as a reference for evaluating the fairness of wage adjustments.

However, that same reference wage does not apply to new employees, or even the same employees in new industries:

Fairness

Q3) A house painter employs two assistants and pays them $9 an hour. The painter decides to quit painting and go into landscaping, where the going wage is lower. He reduces the workers' wages to $7 for landscaping work.

Acceptable 63% Unfair 37%

The reference transaction provides a basis for fairness judgments because it is normal, not because it is just and fair.

Recall the power of adaptation. Terms of exchange that might initially seem unfair can easily become the reference transaction over time.

Coding

How people code outcomes will impact perceptions of fairness. If people exhibit loss aversion around a reference transaction, then framings that make the loss salient should have an impact

Q4A) A company is making a small profit. It is located in a community experiencing a recession with substantial unemployment and no inflation. There are many workers anxious to work at the company. The company decides to decrease wages and salaries 7%.

Acceptable 38% Unfair 62%

Q4B) … with substantial unemployment and inflation of 12% … the company decides to increase salaries only 5% this year.

Acceptable 78% Unfair 22%

Coding

We've seen money illusion before, affecting happiness, demand, and risk preferences. Now it's affecting fairness judgments.

A nominal wage cut is coded as loss and thus readily judged as unfair.

A nominal raise which does not cover inflation is more acceptable because it is coded as a gain relative to the reference wage.

Thus, a little inflation might be good because it "greases the wheels."

If there were no inflation, firms could never adjust wages downward without worker revolt.

However, with inflation, firms can effectively lower wages by simply under-compensating for inflation.

Coding

Q5A) A shortage has developed for a popular model of automobile, and customers must now wait two months for delivery. A dealer who has been selling these cars at list price raises the price to $200 above list price.

Acceptable 29% Unfair 71%

Q5B) …A dealer who has been selling these cars at a discount of $200 below list price now lists this model only at list price.

Acceptable 58% Unfair 42%

In 5A the reference point is unambiguously the list price. In 5B, some must view the current price as the reference and others the list price.

Coding

Similar example, but with wages:

Q6A) A small company employs several people. The workers' incomes have been about average for the community. In recent months, business for the company has not increased as it had before. The owners reduce the workers' wages by 10%.

Acceptable 39% Unfair 61%

Q6B) … The workers have been receiving a 10 percent annual bonus each year and their total incomes have been about average for the community… The owners eliminate the workers' bonus for the year.

Acceptable 80% Unfair 20%

Profit

Q7) Suppose that, due to a transportation mixup, there is a local shortage of lettuce and the wholesale price has increased. A local grocer has bought the usual quantity of lettuce at a price that is 30 cents per head higher than normal. The grocer raises the price of lettuce to customers by 30 cents per head.

Acceptable 79% Unfair 21%

Q8) A landlord owns and rents out a single small house to a tenant living on a fixed income. A higher rent would mean that the tenant would have to move. Other small rental houses are available. The landlord's costs have increased substantially over the past year and the landlord raises the rent when the lease is due for renewal.

Acceptable 75% Unfair 25%

Profit

Even though a random sample of adults will contain many more customers, tenants, and employees than store owners, landlords, or employers, most people state that the firm is entitled to its reference profit.

Firms that are threatened by a reduction in profits below the reference level are allowed to pass on the entire loss to transactors.

This can occur even when the transactor will encounter substantial inconvenience, as in the displaced tenant.

Profit

Q9A) A small company employs several workers and has been paying them average wages. There is severe unemployment in the area and the company could easily replace its current employees at a lower wage. The company has been making money. The owners reduce the current workers' wages by 5 percent. **Acceptable 23%** **Unfair 77%**

Q9B) … The company has been losing money. The owners reduce the current workers' wages by 5 percent.
Acceptable 68% **Unfair 32%**

Firms are only allowed to protect themselves against losses.

Q 11A) A small factory produces tables and sells all that it can make at $200 each. Because of changes in the price of materials, the cost of making each table has recently decreased by $40. The factory reduces its price for the tables by $20. **Acceptable 79%** **Unfair 21%**

Q11B) The cost of making each table has recently decreased by $20. The factory does not change its price for the tables.

Acceptable 53% **Unfair 47%**

Firms are not just expected to be brokers in passing on costs. In accordance with dual-entitlement fairness rules, firms are permitted to pass on losses to customers, but are not held to the same standard for passing on gains.

Firms are allowed higher profits through cost reductions, but NOT from imposing losses on customers.

Market Power

Thus, consumers seem to have a good understanding that monopolies are gaining partly by consumers' expense:

Q12) A severe shortage of Red Delicious apples has developed; none of the grocery stores have any. One grocer receives a single shipment of Red Delicious apples at the regular wholesale cost and raises the retail price by 25% over the regular price. **Acceptable 37%** **Unfair 63%**

Q13) A grocery chain has stores in many communities. Most of them face competition from other groceries. In one community the chain has no competition. Although its costs and volume of sales are the same there as elsewhere, the chain sets prices 5% higher than in other communities.

Acceptable 24% **Unfair 76%**

Market Power

Likewise, perfect price discrimination is particularly outrageous:

Q14) A landlord rents out a small house. When the lease is due for renewal, the landlord learns that the tenant has taken a job very close to the house and is therefore unlikely to move. The landlord raises the rent $40 per month more.

Acceptable 9% Unfair 91%

An action that deliberately exploits the special dependence of a particular individuals is especially offensive.

Other examples include price gouging for supplies after natural disasters.

Market Power

Q15) A store has been sold out of the popular Cabbage patch dolls for a month. A week before Christmas a single doll is discovered in the storeroom. The managers announce that they will sell the doll by auction to the customer who offers to pay the most. **Acceptable 26%** **Unfair 74%**

Add: "The proceeds from the auction will go to UNICEF"

Acceptable 79% **Unfair 21%**

Q16) A business in a community with high unemployment needs to hire a new computer operator. Four candidates are equally qualified. The manager asks the candidates to state the lowest salary they would be willing to accept, and then hires the one who demands the lowest salary. Acceptable 36% **Unfair 64%**

Market Power

The aversion to the auction format here and at the beginning of class is because the firm is gaining from the competition among buyers, which imposes a loss on some buyers. People are fine when the additional profits go to another source, such as charity.

Note that this aversion does not apply to all markets – the individual who sells securities at twice the price paid for them is admired, not reviled

Goods for which there is an active resale market, and non-consumable stores of value, can be sold freely at auction.

Implications

What are the implications of fairness considerations for standard economics? (Besides the fact that 20% of people think any change whatsoever is unfair and 20% of people will be fine with anything)

1) **When excess demand is unaccompanied by increases in suppliers' costs, the market will fail to clear in the short run.**

People will view the increase in price as unfair profit-seeking. Only after time, and the reference transaction has adjusted, will the high price become fair and appropriate.

Empirical studies have found that short-run shifts in demand have small and often insignificant effects

Implications

2) **When a single supplier provides a family of goods for which there is differential demand, without differences in input costs, there will be shortages of the most valued item.**

This is why there is excess demand at vacation spots during popular season times, and long lines at restaurants on Saturday nights.

Although it would be efficient to charge a higher price during peak travel times or a $5 surcharge on Saturdays, consumers would view this as exploitation.

3) **Price changes will be more sensitive to variations of costs than to variations in demand. In particular, price changes will be more sensitive to cost increases than to cost decreases.**

Asymmetric price rigidity – Cost increases tend to be passed on quickly and completely, whereas cost decreases can be retained in part if not in full.

4) **Price decreases will often take the form of discounts rather than reductions in the list price.**

This is for two reasons: Temporary discounts are viewed more favorably than temporary surcharges (due to loss aversion)

Further, temporary discounts can be recalled with less resistance than temporary increases

Likewise for wages – firms that frame part of their compensation package as bonuses should have more flexibility

Game Theory Primer

For those of you with imperfect memory:

Game Theory- the study of how agents behave in strategic situations

By strategic, we mean a situation in which each person, in deciding what actions to take, must consider how others might respond to that action

"Game" is not meant to imply a frivolous competition. A game is simply defined by players, their available actions, resulting payoffs, and structure (timing of moves, whether Nature moves)

Game Theory Primer

Prisoner's Dilemma: Two partners in crime have been captured. The police have enough evidence to convict them on a lesser charge (weapons charge = 1 year) but lack the evidence to convict them on the more serious charge of bank robbery. Because the prosecutor wants conviction on the more serious charge but lacks hard evidence, he needs a confession from one of the criminals

The police look the two in separate rooms and offer each of them a deal:

"We can lock you up for 1 year. However, if you confess to the bank robbery and implicate your partner, we will give you immunity. You will go free and your partner will get 20 years in jail. If you both confess, we won't need your testimony and avoid the cost of a trial so you will both get 8 years."

Game Theory Primer

The decision can thus be modeled using a matrix:

		Clyde: Confess	Clyde: Don't
Bonnie	Confess	-8, -8	0, -20
	Don't	-20, 0	-1, -1

<u>Best response</u>- the strategy in a single period that creates the most favorable immediate outcome for the current player, taking other players' strategies as given

Game Theory Primer

Nash Equilibrium – an outcome from which neither player would want to deviate, taking the other player's behavior as given

Note that a Nash equilibrium occurs whenever each player is best responding to the other player's best response

A game may have zero, one, or more pure strategy Nash equilibrium

Dominant strategy- A strategy that is best for a player in a game regardless of the strategies chosen by the other players

Here, both players have a dominant strategy: to confess

Game Theory Primer

A game without a dominant strategy: Consider two teenagers playing the game of Chicken to look cool

		Dean	
		Swerve	**Don't**
	Swerve	**0, 0**	**-50, 50**
James			
	Don't	**50, -50**	**-100, -100**

If Dean is going to swerve, then James's best response is to not swerve.

But if Dean is not going to swerve, then James's best response is to swerve.

Thus this game has two pure-strategy Nash equilibria: (swerve, don't swerve) and (don't swerve, swerve)

Game Theory Primer

In some games, one player moves first, Player 2 observes the action, and then Player 2 decides.

How do we solve these games? Backward Induction.

For example, consider sequential Paper, Rock, Scissors.

More seriously, consider an incumbent monopoly that knows a 2nd firm is considering entry. The monopolist can either accommodate (maintain its high price) or start a price war (lower its price). The payoffs are as follows

	Enter	Don't
High	250, 250	500 , 0
Low	-100 , -100	300 , 0

(High, Enter) is the simultaneous NE. What is the sequential equilibrium?

Games

Prisoner's Dilemma Games:

	C	D
C	3, 3	0, 5
D	5, 0	1, 1

(C, C) maximizes social surplus, but D is the individually rational dominant strategy

So, people should choose D if they are purely self-interested

However, in experiments, around 30% of people will choose C

Games

Public Good Game: - Basically a multi-person Prisoner's Dilemma

N players, each given initial endowment e_i

Decides how much to contribute c_i to a public good, which is multiplied and shared by everyone

Payoff is thus: $e_i - c_i + m*(\Sigma c)/N$

m denotes how much contributions are multiplied by

m < N, so a player does not benefit enough personally to contribute for private gain

m > 1, so everyone contributing would be Pareto improving

Example: 4 players, each given $10

Can contribute any amount to the public pool, which will be tripled and then shared equally among all players

Everyone gave all $10: Everyone would end up with $30

But suppose the other 3 are all contributing. If you didn't contribute, you'd get $10 + 90/4 = $32.50

In reality, mean contributions are about ½ of endowments. (Although there is wide dispersion – most contribute either all or nothing)

Further, there is positive correlation between amount contributed and how much a subject expects others to contribute.

Players who contribute say they expect others to contribute

Games

Ultimatum Games: Take it or leave it bargaining

The Ultimatum game is a sequential game.

There is a pie of size P.

The 1st player, the Proposer, decides what share s of the pie to offer to player 2, and what share to keep for himself.

The 2nd player, the Responder, either accepts or rejects this offer.

If Player 2 accepts, her payoff is the share s*P and player 1 keeps the (1-s)*P

If Player 2 rejects, both players get 0.

What should self-interested players do?

Games

Pure self interest implies that the Respondent should accept ANY s > 0. Anticipating this, the Proposer should offer s = .01

How does this game actually play out?

-Proposers virtually never choose s > 50%

-Modal and median offers are 40-50%, and means are 30-40%

-There are hardly any offers of 0-10%

-Offers of 40% are rarely rejected

-Offers below 20% are rejected about half the time

-The probability of rejection increases as s gets smaller.

Note that the truly weird behavior isn't the Proposing Player 1, it's the Responding Player 2.

Games

Given the fact that Respondents may reject some unfair offers, it might not be a bad idea then to offer larger shares. This is especially true if Proposers are risk-averse.

Note that if Player 2 rejects all offers $s < s^*$, Player 1 offering s^* is a Nash equilibrium (it's just not subgame perfect)

Given the distribution of rejection rates, Proposers are generally playing pretty close to optimally (even controlling for risk aversion), although often slightly more generous than income-maximizing.

So, observing Proposer offers $s > 0$ could be due to either social preferences or fear of rejection / risk aversion (or both)

Games

Dictator game – “Control” for the Ultimatum game

(Although technically not a game, just an optimization problem)

Eliminating the possibility of rejection allows us to see how much of the offer was due to fear of rejection and how much due to social preference.

There is a pie of size P.

Player 1, the Proposer, decides what share s of the pie to offer to player 2, and what share to keep for herself.

That’s it – Player 2 doesn’t do anything but take what Player 1 gives him.

The mean allocation is about 20%

Games

The fact that dictator game offers are lower than ultimatum game offers shows that Proposers are being strategic (offering more to avoid being rejected)

However, the fact that dictator game offers are still positive shows that people do have some social preferences

Note that because people do not behave according to the predictions of game theory does NOT mean that game theory is wrong, any more than cashiers giving the wrong change means that the laws of arithmetic are wrong

It simply means that people are not perfectly self-interested as standard game theory assumes

Altruism

Let $x = (x_1, x_2)$ be a vector of material payoffs.

x_1 denotes player 1's material payoff.

X_2 denotes player 2's material payoff.

Pure self-interest: Player 1's utility/preferences are $u_1(x_1)$

Social preferences: Player 1's utility/preferences are $u_1(x_1 , x_2)$

A person exhibits simple altruism if her utility is increasing in other people's material payoffs.

Simple formulation:

$u_1(x_1 , x_2) = x_1 + \varphi x_2$ for $\varphi > 0$

Altruism

Can simple altruism explain the experimental evidence? Suppose Player 1 is altruistic:

Prisoner's Dilemma:

	C	D
C	**3 +φ3, 3**	**0+φ5, 5**
D	**5, 0**	**1+φ1, 1**

Choosing C hurts oneself but helps the other player. Because an altruist get utility from own consumption as well as the other player's, if the gain to the other from choosing C outweighs the diminished own consumption, the altruist will cooperate.

Here, for $\varphi > 2/3$, it is optimal to choose C.

Altruism

Altruism in Public Goods games:

Because contributing to the public good yields additional payoffs to others, altruistic preferences should increase the likelihood of contributing.

If φ is high enough, it becomes optimal to always contribute.

However, note that for linear public goods games, if players have concave utility, if others are contributing a lot, the marginal return from the altruism component is low, so players should *reduce* their contributions if everyone else is contributing.

This is exactly opposite of how people actually behave.

Altruism

Altruism in the dictator game:

Because choosing $s > 0$ increases the Respondent's payoff, simple altruism works in favor of choosing $s > 0$

If φ is large enough, it will be optimal for the Dictator to choose $s > 0$.

Altruism

Altruism in the Ultimatum game:

Proposer: Similar to the Dictator game, choosing s > 0 increases the Respondent's payoff, simple altruism works in favor of choosing s > 0

However, for the Respondent, rejecting an offer s > 0 decreases both the Respondent's own payoff AND the Proposer's payoff. Altruism implies that the Respondent should be even more likely to accept offers than the pure self-interest case

Altruism cannot explain Responder rejections.

Main Conclusion – Simple altruism can explain some experimental results, but it can't explain everything

Inequity Aversion

Fehr and Schmidt – Inequality aversion – people don't just care about their own payoff, they care about how much they make relative to others

People may even be willing to pay (reduce their own payoffs) to reduce inequality

Simple model:

$$u_1 (x_1, x_2) = x1 - g (x_1, x_2)$$

$$\text{where } g (x_1, x_2) = \begin{cases} 0 & \text{if } x_1 = x_2 \\ \alpha (x_2 - x_1) & \text{if } x_1 < x_2 \\ \beta (x_1 - x_2) & \text{if } x_1 > x_2 \end{cases}$$

With $0 \leq \beta < 1$ and $\alpha \geq \beta$

Inequity Aversion

β captures how much aversion a person has to being better off than other people

If β = 0 corresponds to the self-interest case – a person doesn't care if he is better off than others

Note that we assumed β ≥ 0. β < 0 would imply that a person enjoys being better off than others. Although this is certainly true for some people in the real world, for the games we're considering here, they would have no impact on equilibrium behavior.

β > 0 implies that the player is willing to give some money to player 2 to help reduce the inequality between them.

Inequity Aversion

For example, β = 0.25 implies that Player 1 is willing is willing to give up a dollar if it means Player 2 would get 3 dollars.

Ex) Player 1 has $10 and player 2 has $2

$u_1 (x_1, x_2) = 10 – 0.25*(10-2) = 8$

Lose $1, player 2 gets $3:

$u_1 (x_1, x_2) = 9 – 0.25*(9-5) = 8$

β = 0.50 implies that Player 1 is just indifferent between keeping a dollar for himself and giving it to player 2.

Player 1 has $10, player 2 has $2

$u_1 (x_1, x_2) = 10 – 0.50*(10-2) = 6$

Inequity Aversion

Give $1 to player 2:

$u_1 (x_1, x_2) = 9 – 0.50*(9-3) = 6$

$0.50 < \beta < 1$ means that a person is willing to support "inefficient transfers" or "leaky buckets" to reduce inequity.

$\beta = .75$: Player 1 has $10, player 2 has $2

$u_1 (x_1, x_2) = 10 – 0.75*(10-2) = 4$

Lose $1, player 2 gets 33 cents:

$u_1 (x_1, x_2) = 9 – 0.75*(9-2.33) = 4$

Inequity Aversion

Why the upper bound $\beta < 1$. Note that at B = 1, a person is indifferent between keeping a $1 and just throwing it away for the purpose of reducing inequality. That seems a bit much.

α measures how averse someone is to unfavorable inequity – inequity when they've got the short end of the stick

Why $\alpha \geq \beta$? If there's going to be inequity, someone would rather it be in favor of themselves. Loss aversion in social comparisons.

Unlike β, there is no need to place an upper bound on α. People may be willing to forgo large amounts to reduce unfavorable inequality.

Inequity Aversion

Example:

α = .75: Suppose player 1 has $6 and player 2 has $10:

$u_1 (x_1, x_2) = 6 – 0.75*(10-6) =3$

Player 1 would be willing to sacrifice $1 if it would reduce player 2's payoff by $2.33

$u_1 (x_1, x_2) = 5 – 0.75*(7.67-5) =3$

Likewise, α = 4 implies that the player would be willing to spend $1 to reduce the other player's payoff by $1.25:

$u_1 (x_1, x_2) = 6 – 4*(10-6) = -10$

$u_1 (x_1, x_2) = 5 – 4*(8.75-5) = -10$

Inequity Aversion

This model depicts inequality aversion based on the *difference* between payoffs, $(x_2 - x_1)$ or $(x_1 - x_2)$. Other models depict it based upon the *relative* payoffs, x_2 / x_1 or x_1 / x_2

Note: The previous formulation assumes utility is linear in payoffs. It is certainly possible to relax this assumption.

Likewise, this formulation depicts the utility function as linear in inequality aversion. This implies that the MRS between monetary income and inequality is constant. Again, this can be easily relaxed, although even this simple formulation explains much of the experimental evidence.

Inequity Aversion

Inequity aversion in the prisoner's dilemma:

Prisoner's Dilemma:

	C	D
C	**3 , 3**	**0 –α5, 5**
D	**5-β5, 0**	**1, 1**

Suppose Player 1 knows that Player 2 will cooperate. If Player 1 defects, she will experience guilt over the advantageous inequality. For β large enough, it becomes optimal to cooperate.

If both players have such inequity aversion, this game will now have TWO equilibria: (C, C) and (D, D)

Inequity Aversion

Inequity Aversion in Public Good:

If know one else is contributing, and you contribute, disadvantageous inequality would kick in, so this is definitely not a good idea.

If everyone else is contributing, and you don't, then advantageous inequality would kick in. For large β, it thus becomes optimal to donate as well.

Note that inequality aversion does capture the finding that donations and beliefs and expectations about others' contributions are positively correlated.

Inequity Aversion

Inequity Aversion in Dictator Games:

Keeping all the money would create advantageous inequality, so a high enough β implies that a Dictator should give some positive s.

Note that the current formulation predicts that s = 0.5 if β > 0.5 and s = 0 if β < 0.5 That is, it predicts only very fair or very unfair outcomes.

This is where utility that is concave in the amount of advantageous inequality comes in handy, and allows for interior solutions in [0 , 0.5]

Inequity Aversion

Inequality Aversion in Ultimatum Games:

Again, keeping all the money would create advantageous inequality, so a high enough β implies that a Proposer should give some positive s.

Suppose the Responder is evaluating an unfair offer s < .5. Accepting it yields the material payoff but creates disadvantageous

Accept: $s - \alpha((1-s) - s)$

Reject: 0

This yields an acceptance threshold of: $s(\alpha) = \alpha / (1+2\alpha)$. Accept any offer higher than this.

Inequity Aversion

$\alpha = 0 \rightarrow$ Accept any $s > 0$

$\alpha = .5 \rightarrow$ Accept any $s > .25$

$\alpha = 1 \rightarrow$ Accept any $s > .33$

$\alpha = 4 \rightarrow$ Accept any $s > .44$

Now re-consider the decision of the Proposer. Suppose he does not know the Responder's α, but knows the distribution of α across the population. Now his optimal offer is both a function of his own inequality aversion β and the population distribution α.

For example, even a Proposer with $\beta = 0$ might maximize expected utility by offering $s > 0$ if he knows a substantial fraction of the population has high α and would thus reject unfair offers.

Intentions

Another aspect of social preferences is that intentions matter: Up until now, we have only considered "payoff based approaches" wherein all utility is based upon the final outcome

But people do not just care about the final outcome, but also how they got there

A boy is walking home from school with a friend when he finds two apples. He kept the larger one for himself and gave the smaller one to his friend. "That wasn't very nice of you to keep the larger one for yourself" the friend said. "Well, what would you have done?" the first boy asked. "I'd have given you the larger one and kept the small one for myself" said the friend. The first boy replied, "Then we each got what we wanted, so what are you complaining about?"

Intentions

There's more to it than just altruism – people do not seek to uniformly help others, but want to be nice to people who treat them nicely, and meanly towards those who harm them

Consider the following variant of the Ultimatum game:

Instead of choosing any allocation, the Proposer must choose between one of 2 choices (and the Responder knows this):

A: (8,2) or (10,0)

B: (8,2) or (2,8)

C: (8,2) or (5,5)

How do you think Responders behave towards an (8,2) offer in each case?

Intentions

Likewise, consider another Ultimatum variant:

Version 1: Player 1 offers a share of $10 from {0, .50, 1.00, 1.50 … 9.50, 10}

Player 2 accepts or rejects

Version 2: A computer randomly chooses a share of $10 from {0, .50, 1.00, 1.50 … 9.50, 10}

Player 2 accepts or rejects

Version 1: 29% accept {9.50, .50} ; 71% reject

Version 2: 80% accept {9.50, .50} ; 20% reject

Intentions

Rabin's "fairness equilibrium" – develop a formal game-theoretic model of conditional reciprocity

3 key features:

1) **People are willing to sacrifice their own material well-being to help those who are being kind**

2) **People are willing to sacrifice their own material well-being to punish those who are being unkind**

3) **Both motivations (1) and (2) will become stronger the smaller are the material payoffs**

Intentions

The central concept in Rabin's approach is a kindness function – payoffs do not just depend on the actions taken, but also on the beliefs about the other player's motives

As it turns out, it is impossible to capture this simply by modifying payoffs. As such, it is necessary to incorporate beliefs

How to assess kindness?

First, from a player's point of view, how kind is he being to the other player:

Intentions

Suppose player 1 has belief b_2 about what player 2 will do. Then from player 1's point of view, his choice is an allocation out of all possible payoffs in $\Pi(b_2)$

The highest and lowest of these payoffs, for player 2, is Π_2^{max} and Π_2^{min} . Define the equitable or fair payoff out of this as Π_2^{fair} . (For example, one candidate is the average of the highest and lowest payoff)

Then player 1's kindness towards player 2, which depends on her actual choice a_1 is:

$$f_1(a_1, b_2) = \frac{\Pi_2(b_2, a_1) - \Pi_2^{\text{fair}}(b_2)}{\Pi_2^{\max}(b_2) - \Pi_2^{\min}(b_2)}$$

Intentions

Thus kindness is the fraction above or below the fair point that player 2's payoff lies.

A positive f_1 is kind because it means player 2 gets a payoff higher than the fair one; a negative value is mean because player 2 got a lower payoff than the fair one.

Player 1 also cares about how kind player 2 is being to him. To assess this, he forms a perceived kindness, based on what he believes Player 2 believes he will do, c_1:

$$f_2(b_2, c_1) = \frac{\Pi_1(c_1, b_2) - \Pi_1^{\text{fair}}(c_1)}{\Pi_1^{\max}(c_1) - \Pi_1^{\min}(c_1)}$$

Intentions

Given all this, player 1's social preferences are:

$$U_i(a_i, b_j, c_i) = \Pi_i(a_i, b_j) + \alpha f_2(b_2, c_1) + \alpha f_2(b_2, c_1) * f_1(a_1, b_2)$$

Players care about their monetary payoffs (the first term), whether they are being treated kindly or not (the second term), and whether they are being kind to a kind person or vice versa (the third term). That is, whether they are coordinating on kindness

α simply converts fairness into utilities, and hence α = 0 is the case of pure self-interest

Consider two people shopping, and there are 2 cans of soup left. Each person can either try to grab both cans or share:

	Grab	Share
Grab	x , x	2x, 0
Share	0, 2x	x, x

The Nash equilibrium is (Grab, Grab)

If we added fairness concerns though, (people get utility from being nice to each other), (Share, Share) would become a "fairness equilibrium"

As material payoffs are smaller, fairness concerns become more dominant. For two people fighting over soup, the social grief and bad tempers are likely to be of greater importance than whether they actually get the cans

Intentions

Example: Prisoner's Dilemma with Fairness Payoffs

	C	D
C	4 +.75α, 4+.75α	0-.5α, 6
D	6 , 0-.5α	0,0

For player 1, if she expects player 2 to choose D, she feels even worse choosing C than usual, because player 1 is being kind. Oppositely, if player 1 expects player 2 to choose C, for high enough α she might want to choose C as well. Reciprocating cooperation is nice.

Intentions

Chicken with fairness adjusted payoffs:

	Dare	Chicken
Dare	**-2, -2**	**2 , 0 -.5α**
Chicken	**0 - .5α, 2**	**1+.75α , 1 + .75$^{\alpha}$**

Suppose Player 1 expects player 2 to Dare. If fairness effects are strong enough, Player 1 might prefer to Dare as well and harm himself, rather than let the mean player get an advantage.

Similarly, if the other player is being nice and playing Chicken, he might want to repay the kindness and choose Chicken as well.

Thus (D,D) and (C,C) are fairness equilibria – the exact opposite of the standard equilibria

For any specific outcome:

1) **Are the material payoff concerns in favor or against it?**

(iff it's a Nash equilibrium)

2) **Do fairness concerns work in favor or against it?**

If both forces are in favor of that outcome, then it's a fairness equilibrium.

If both forces work against that outcome, then it's never a fairness equilibrium.

If material payoffs are for an outcome while fairness concerns are against, the material concerns dominate as payoffs increase.

However, if payoffs are sufficiently small, behavior may be driven entirely by fairness concerns.

Intentions

In every game, there is a weakly negative fairness equilibrium.

Unfortunately, the opposite is not necessarily true – we can never guarantee that players will part with a warm glow.

Why the asymmetry? If a player is maximizing his own payoffs, he's either being mean or simply neutral.

Players may not view standard maximizing behavior as kind, because they realize the other player is simply being self-interested.

However, it's always possible to deviate from the optimal choice, and this may be perceived as being mean.