Behavioral Economics and the Study of Social Preferences – James Konow

International Conference on Behavioral Economics

– ISER and Osaka University

Subject Matter of Behavioral Economics

- Traditional economic model is of *homo economicus*: rational and self-interested agent
- **Behavioral economics** seeks a richer description of economic agents by relaxing these two assumptions:

1. **Bounded rationality**: limited cognitive abilities, viz., people are not always expected discounted utility maximizers
2. **Social preferences**: assumption of self-interest must sometimes be appended to account for interdependent preferences such as fairness and envy

- Focus of this presentation: social preferences branch of behavioral economics

Types of Social Preferences

Emerging consensus among behavioral economists:

- **Distributive preferences**: preferences over the final distribution, e.g., equity, efficiency, altruism; related to consequences or outcomes
- **Reciprocal preferences**: desire to reward or punish others beyond mere consequences, e.g., being “more than fair” to someone who has been fair to you; related to intentions or types of agents

Field Evidence of Social Preferences

Kahneman, Knetsch and Thaler (*AER* 1986)

- People care about fairness
- Ex: It is sometimes unfair to raise prices to eliminate a shortage, e.g., on snow shovels after a snowstorm
- This concern for fairness constrains profit seeking

Blinder and Choi (*QJE* 1990)

- Wage setters at firms confirmed the importance of fairness considerations in contributing to wage rigidity

Babcock, Wang and Loewenstein (QJE 1996)

- Survey teachers unions and school boards
- Strike activity is positively related to differences between teachers unions and school boards over the salaries they consider fair and comparable

Andreoni, Erard and Feinstein (*JEL* 1998)

- Amount of tax evasion is affected by the perceived fairness of the tax system

Zajac (*Political Economy of Fairness* 1995)

- Public support for the regulation of private industries depends on perceived fairness of firms' policies

Experimental Evidence of Social Preferences

Güth, Schmittberger and Schwarze (*JEBO* 1982)

- First experimental test of the "ultimatum game," a simple two player, two stage bargaining game
- Proposer offers a division of a fixed sum to an anonymous Responder, who then accepts or rejects the offer
  - If the Responder accepts, the sum is divided as agreed
  - If the Responder rejects, both receive nothing
- The subgame perfect equilibrium: the Proposer offers the minimum amount and the Responder accepts

Camerer (*Behavioral Game Theory* 2003)

- Hundreds of replications of the ultimatum game produce similar results
- Average Proposer offers are about 40%
- Responders reject offers below 20% about one-half of the time

Forsythe, Horowitz, Savin and Sefton (*GEB* 1994)

- Ultimatum game contains a possible confound: even self-interested Proposers might make generous offers if enough Responders are expected to reject unfair offers
- "Dictator game" removes this confound:

  Proposers can transfer any amount of a fixed sum, or nothing, and Responders have no recourse
- Results: Proposers offer less in the Dictator game than in the Ultimatum game but still give, on average, about 20-25% of the stakes

Fehr, Kirchsteiger and Riedl (*QJE* 1993)

- First experimental "gift-exchange" game:
- Experimental labor market in two stages:
  1. Employers first offer a wage to workers
  2. Workers then choose a costly effort level that increases the employers' earnings
- Results:

  1) worker effort is directly related to employer wage offers, and

  2) average wages exceed the market clearing wage

Charness (*JLaborE* 2004)

- Gift exchange experiment found that:

  1) worker effort varies more when the wage offered by an employer is intentional than when it is random, and

  2) both distributive and reciprocal preferences matter

Models of Social Preferences

- Social preferences exist if:

  $U_i(x_i, x_j), i \neq j$

  x_i, x_j: allocations of persons *i* and *j*, respectively
- Two person versions
- This formulation explains everything, but then it explains nothing
- **Becker (*JPE* 1974)**: Under altruism, utility of *i* is directly related to x_j, whereas under envy, they are inversely related
- Justifies in a general way deviations from self-interest but still need more structure to account for many specific observations
- Here, first treat theories and evidence of distributive preferences, and then reciprocal preferences

Distributive Preferences

Three models of distributive preferences:

1. Fehr and Schmidt (1999), Bolton and Ockenfels (2000):

   Inequality aversion

2. Charness and Rabin (2002):

   Rawlsian maximin and surplus maximization

3. Konow (2000):

   Inequity aversion

Fehr and Schmidt (*QJE* 1999)

- Inequality aversion: disutility from unequal allocations
- Frequently cited and frequently used model
- Linear preferences, viz., inequality aversion is piecewise linear
- Disutility from inequality greater if the other person is better off than if the other person is worse off
- Agents may differ with respect to the disutility from inequality
- Utility function of player *i* for two person case:

$$U_i(x_i, x_j) = x_i - \alpha_i \max\{x_j - x_i, 0\} - \beta_i \max\{x_i - x_j, 0\},$$

$$i \neq j, \beta_i \leq \alpha_i, 0 \leq \beta_i < 1$$

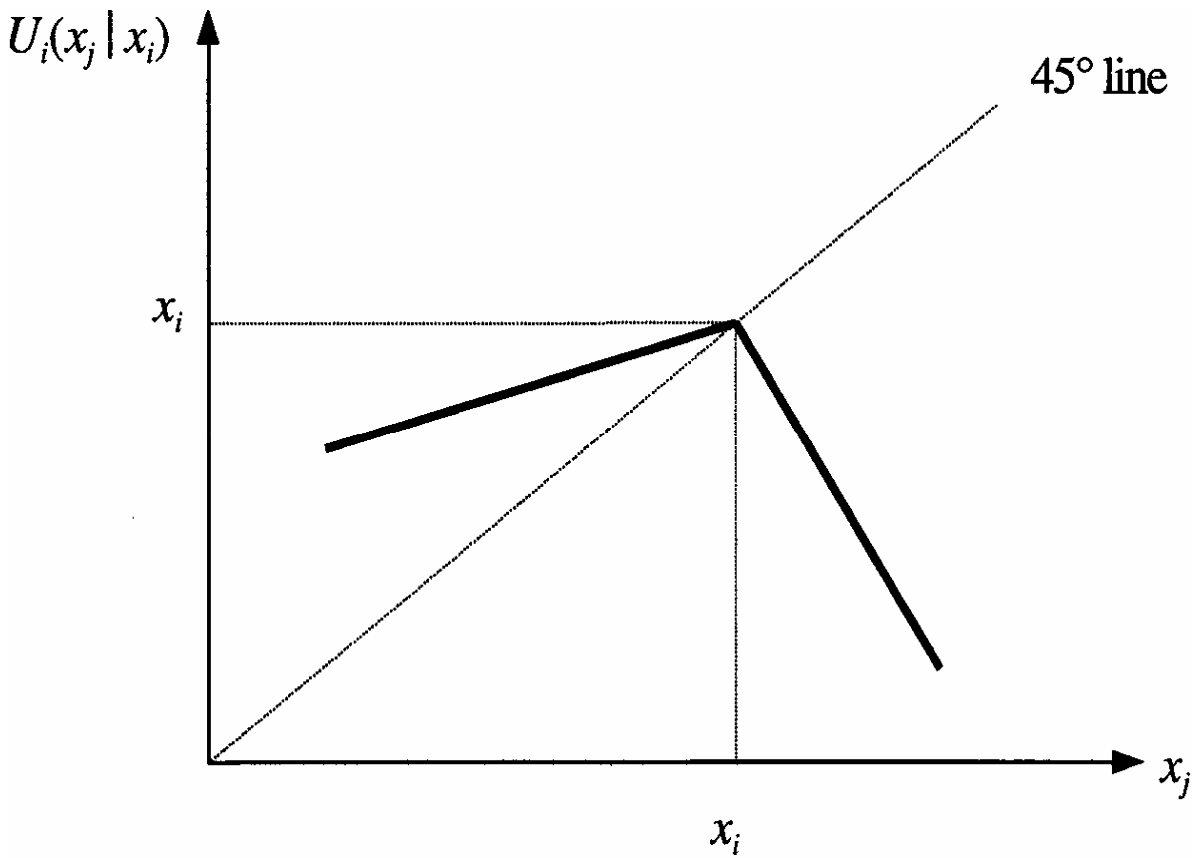

Advantage of Fehr and Schmidt model is it reconciles ostensibly contradictory experimental results:

1. Purely self-interested behavior
2. Deviations from self-interest, including both positive and negative actions toward others
3. A minority of selfish people can force a majority of fair ones to behave selfishly in market games or in simultaneous public good games
4. A minority of fair players can force a majority of selfish ones to give in gift exchange games and in public good games with punishment

Disadvantages of Fehr and Schmidt:

1. Linear functional form implies corner solutions in the dictator and other games, but many decisions lie between extremes
2. Absence of explicit reciprocity motive

Bolton and Ockenfels (*AER* 2000)

- Similar to Fehr and Schmidt with inequality aversion

Advantage:

1. More general functional form

Disadvantages:

1. More difficult to obtain closed form solutions
2. Makes some implausible predictions about the dictator and ultimatum games

Charness and Rabin (*QJE* 2002)

- Present series of distribution experiments and a model to explain the results
- Model incorporates both distributive and reciprocal preferences in addition to self-interest
  - focus in this section on distributive preferences component
- Distributive preferences incorporate two motives corresponding to two social welfare functions:
  1. Rawlsian maximin criterion
  2. Utilitarian welfare function
- Utility is a function of these distributive preferences and one's own payoff

- Quasi-maximin welfare function:

$$W(x_i, x_j) = \delta \cdot \min\{x_i, x_j\} + (1-\delta) \cdot (x_i + x_j),$$

$$\delta \in (0,1)$$

- Utility function with quasi-maximin preferences:

$$U_i(x_i, x_j) = (1-\gamma)x_i + \gamma \cdot W(x_i, x_j)$$

$$\gamma \in [0,1]$$

- Similar to Fehr and Schmidt, this implies a constant marginal rate of substitution, specifically, person *i* cares less about person *j* if *j* is better off than *i*

Advantage:

- Charness and Rabin experiments support a willingness to sacrifice personally in order to equalize and to increase the total

Disadvantage:

- Quasi-maximin preferences do not explain attempts to hurt others, e.g., Responder rejections in the ultimatum game or punishments in certain public good games

Charness and Rabin augment these distributive preferences with reciprocity motive

Drawbacks of Fehr and Schmidt and Quasi-maximin Preferences

- **Linear functional form**: affords analytical simplicity, but predicts corner solutions (e.g., one-half or all or nothing in dictator experiments) that conflict with results of many experiments (e.g., Andreoni and Miller, *Econometrica* 2002)
- **Inequality aversion**: important in many contextually lean experiments, but field studies and experiments with more real world elements routinely find distributive preferences deviate from equality (e.g., Blinder and Choi *QJE* 1990, Babcock et al *AER* 1995, Eckel and Grossman *Ec Journal* 1998, Konow *AER* 2000)

Konow (*AER* 2000, Mixed Feelings 2005)

- Equity does not always correspond to equality
- Model of inequity aversion, i.e., disutility from allocations that are not equitable (as opposed to merely unequal)
- Distributive preferences depend on an "entitlement," Φ_j, that person *i* considers the "right" allocation for person *j*
- Deviation of person *j*'s allocation from the entitlement generates disutility for *i*

Simplified Model of Inequity Aversion

- Utility of person *i*:
$$U_i(x_i, x_j, \phi_j) = u_i(x_i) - f_i(x_j - \phi_j)$$
- $u_i(x_i)$: material utility; positive and diminishing marginal material utility
- f : inequity aversion term; - f is strictly concave function of x_j with a maximum at Φ_j

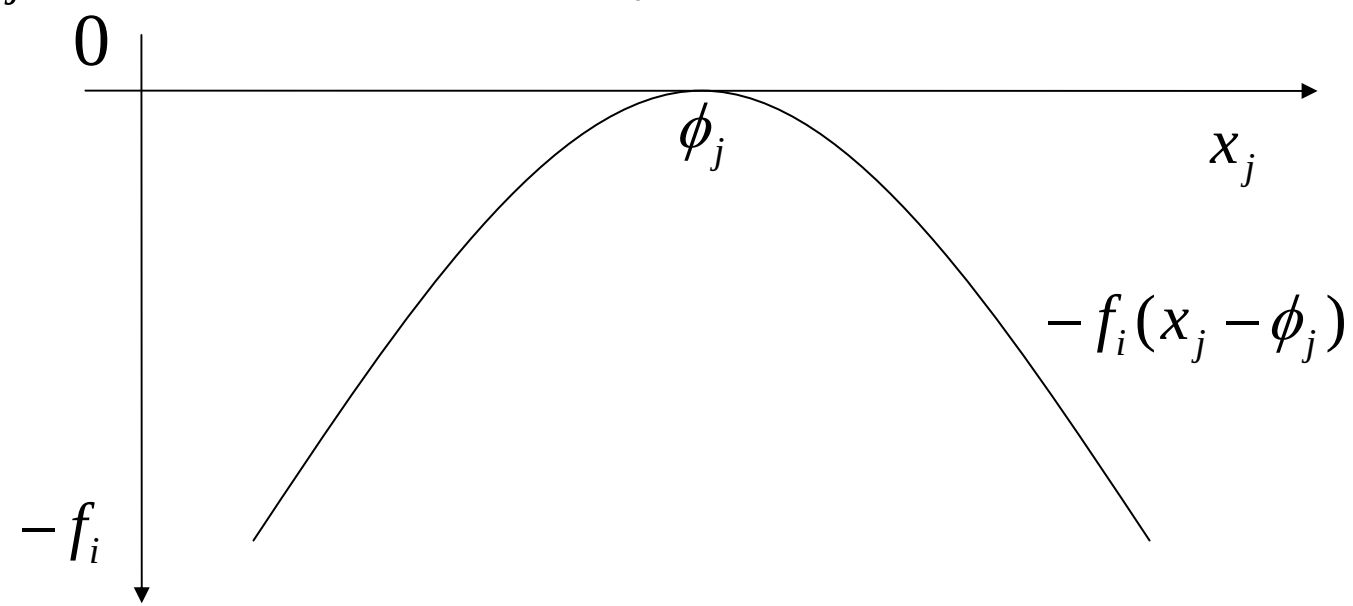

- In these papers, *i* is the decision maker, and *j*'s allocation depends on *i*'s choice, i.e., $x_j(x_i)$
- As with the previously discussed models of inequality aversion, the marginal disutility expressed in the f term can differ across agents
- The f term can also be parameterized to explain extreme as well as intermediate allocation decisions
- Most important to the predictive power of this model, however, is the specification of Φ_j

Specification of the Entitlement

- I argue (*JEBO* 2001, *JEL* 2003) the entitlement, Φ_j, is determined by three principles and context
- Context: relative importance of each principle depends on the context, i.e., information about the relevant actors and allocations
- Sometimes only one principle is salient
- Sometimes agents entertain multiple principles:
  - Trade-off between principles
  - Most important principle is the one that currently causes the greatest disutility on the margin
  - Inequity aversion can best be reduced by reallocating according to this principle
- The three distributive principles that (potentially) determine Φ_j are:
  1. Accountability
  2. Efficiency
  3. Need

Accountability Principle

- In the simplified version, fair allocations are proportional to contributions that agents control
- Example:
  - Suppose person *i* is twice as productive as *j*
  - Person *i*'s allocation should be twice *j*'s if *i*'s greater productivity is due entirely to factors he controls, called *discretionary variables* (e.g., effort, decisions, etc.)
  - Fair allocations are equal, however, if *i* and *j* have equal discretionary variables, and *i*'s greater productivity is due entirely to things he cannot control, called *exogenous variables* (e.g., endowed resources, innate skills, etc.)

Experimental test of accountability principle

Dictator experiment in two phases:

- Phase 1: All subjects generate earnings by preparing letters for mailing; these earnings are credited to a joint account assigned to each pair of subjects
- Phase 2: One subject is arbitrarily chosen to be Dictator and to allocate the earnings between Dictator and Recipient

One treatment variable concerns the first phase and how earnings are generated:

- *Discretionary* differences treatments: The same credit is given for all letters, but earnings differ across subjects because of differences in the number of letters prepared
- *Exogenous* differences treatments: Subjects are given the same materials and enough time to prepare all letters, but earnings differ because of arbitrary differences in the per letter credits; these differences always favor the Dictator by design

Another treatment variable concerns the identity of the Dictator in phase two of the experiment:

- Standard: The case just discussed where one of the two subjects in each pair is arbitrarily chosen to be Dictator
- Benevolent: A third subject, who does not prepare letters, is paid a fixed fee to decide the allocation of earnings for a pair of subjects

This is crossed with the phase one treatment variable to generate four cells:

- Standard/Discretionary
- Benevolent/Discretionary
- Standard/Exogenous
- Benevolent/Exogenous

Benevolent/Discretionary treatment results

- For each pair, one subject's contributions are measured on the x axis and the same subject's allocations on the y axis, both as fractions of total
- Benevolent Dictators have no stake in the allocation they decide and, therefore, are predicted to allocate fairly on the 45° line
- In fact, they allocate fairly, on average
- After correcting for an experimental artifact, about 4/5 allocate in exact proportion to discretionary contributions, and ar OLS regression of allocations on contributions generates an R^2 of .98

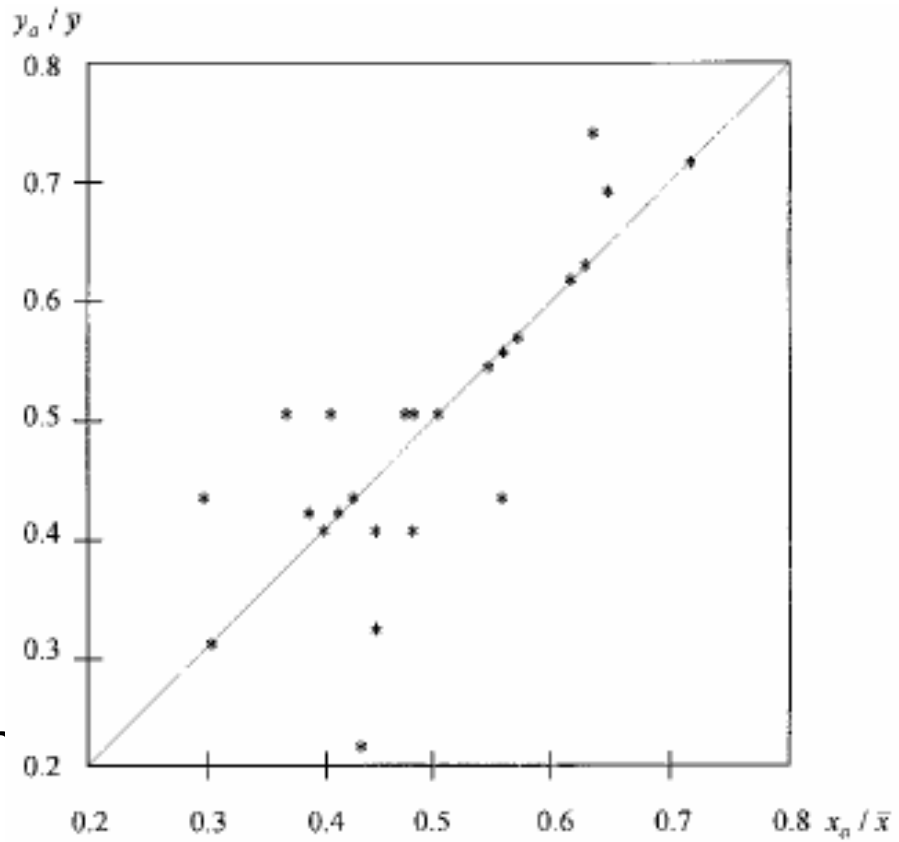

Benevolent/Exogenous treatment results

- Arbitrary per letter credits appear on the x axis, and the shares of earnings for the same subject appear on the y axis
- Benevolent Dictator allocations corroborate the theory: third parties with no stakes judge allocations impartially and ignore the irrelevant differences in per letter credits
- 87% of Benevolent Dictator allocations are equal shares, and the average allocation does not differ significantly from .5

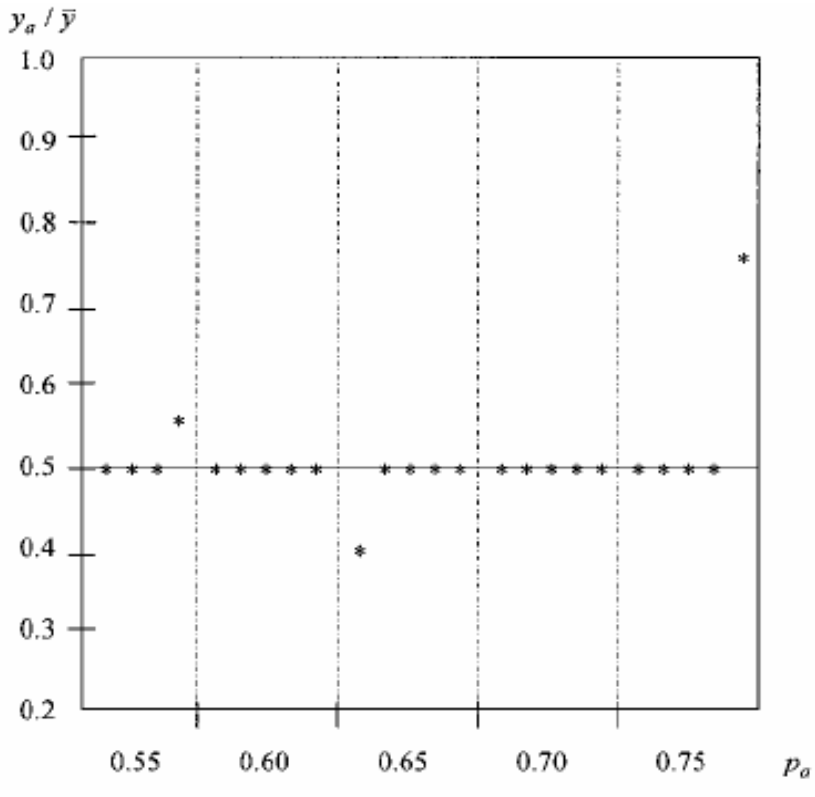

Standard/Discretionary treatment results

- Most Standard Dictators allocate the earnings fairly, i.e., close to the 45° line, or in the shaded area above it, which indicates the effect of self interest
- Transfers are significantly related to these "discretionary" contributions: for every $1 the Recipient contributes to joint earnings, the Dictator transfers about 30 cents more to the Recipient

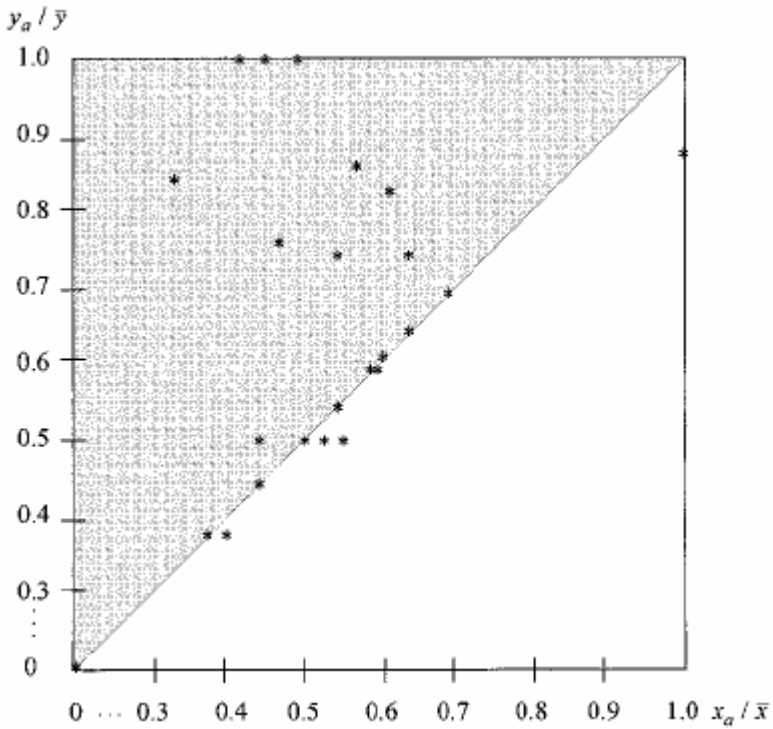

Standard/Exogenous treatment results

- Most Standard Dictators allocate equal shares of .5 along the horizontal line or greater than equal shares in the shaded area above the horizontal line, indicating the effect of self-interest
- A number of Standard Dictators allocate in proportion to the arbitrary per letter credits, which favor them

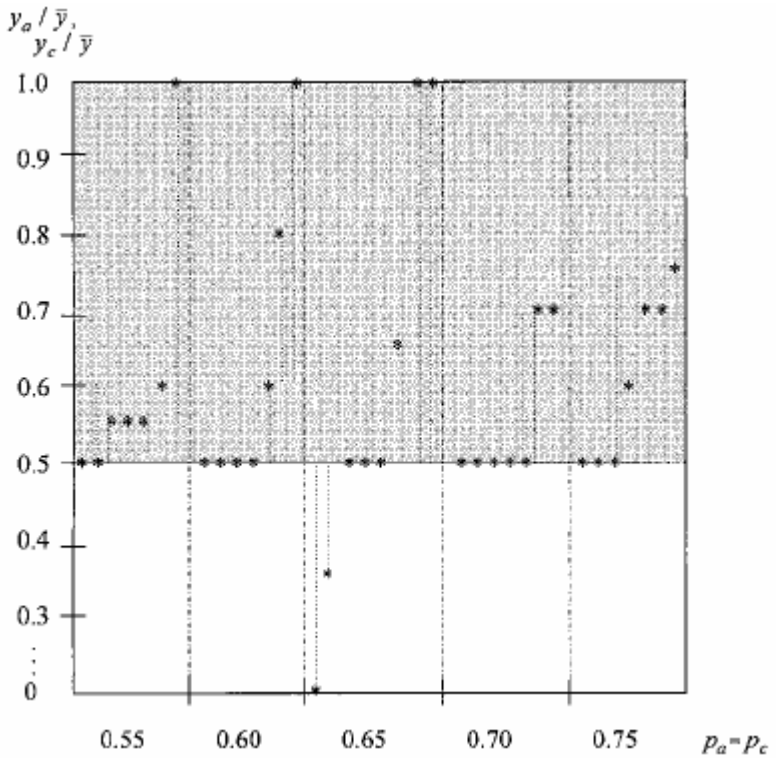

Social Preferences and Culture

Experimental studies of social preferences across cultures and national boundaries have come to contradictory conclusions

Significant differences:

- Roth, Prasnikar, Okuno-Fujiwara and Zamir (*AER* 1991): Small but significant differences in ultimatum game offers between US and Slovenian subjects versus Japanese and Israeli ones
- Henrich (*AER* 2000): A small scale society in the Amazon played the ultimatum game yielding smaller offers and fewer rejections than among US subjects

No significant differences:

- Brandts, Saijo and Schram (*Public Choice* 2004): Public goods experiment in the US, Japan, the Netherlands and Spain
- Okada and Riedl (1999, *GEB* 2005): Variation of the ultimatum game with Austrian and Japanese subjects

Konow, Saijo and Akai (2005)

- Konow, Saijo and Akai recently substantively replicated the results of Konow (2000) for the discretionary treatments with subjects in the US and Japan
- Benevolent dictators mostly allocated proportionately

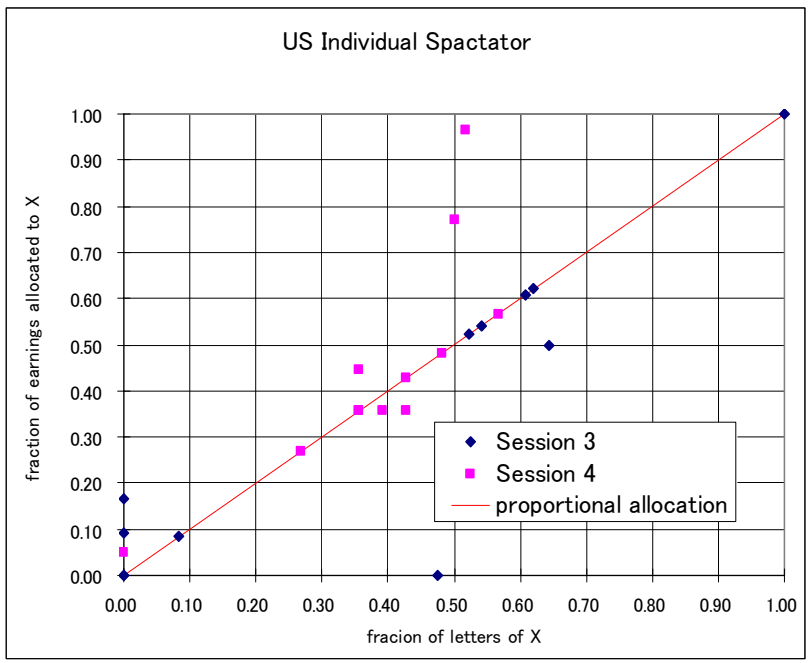

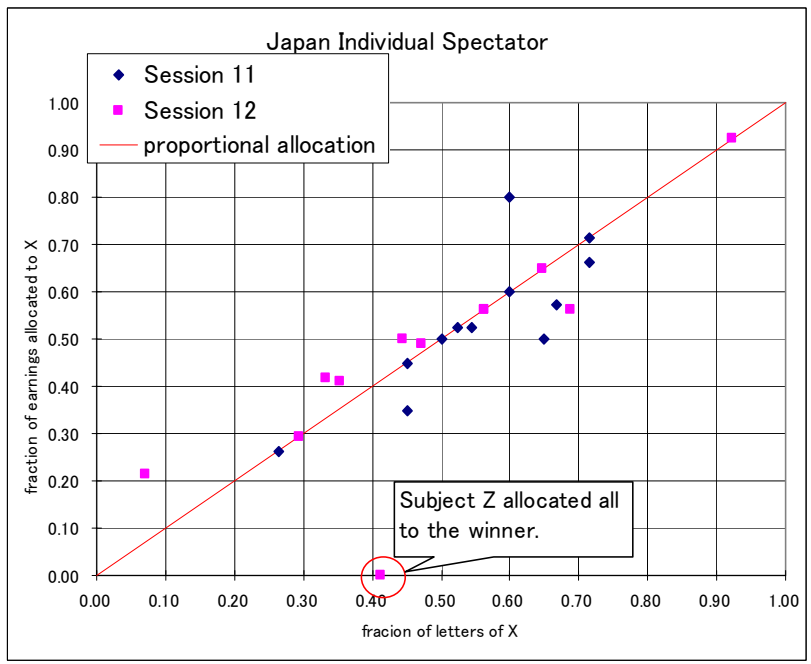

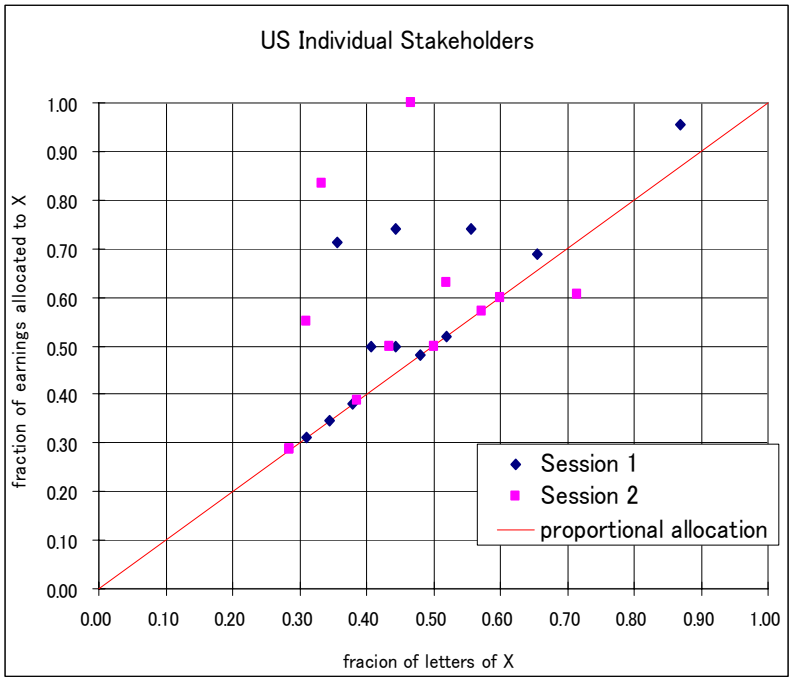

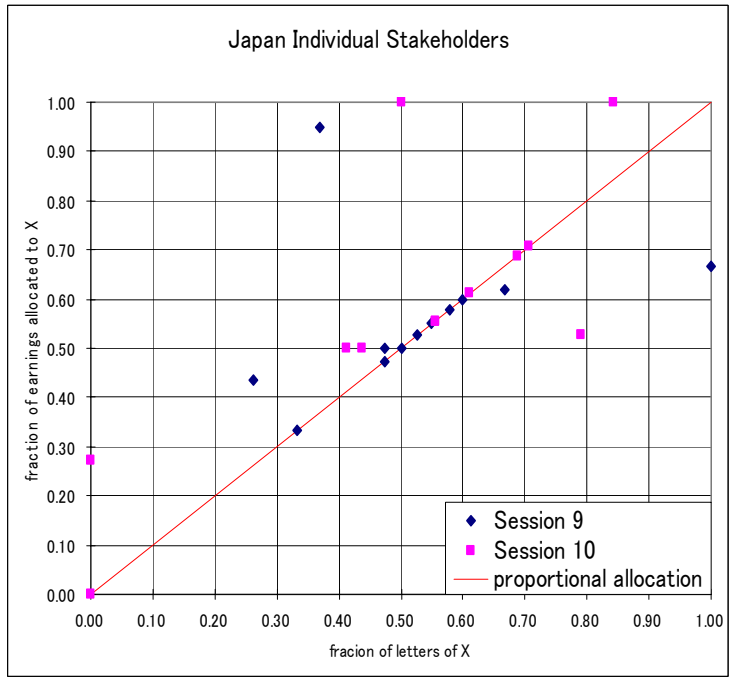

- US standard dictators took the proportional amount or more
- Japanese standard dictators, on the other hand, behaved almost like benevolent dictators, rarely taking more than the proportionate amount
- Preliminary conclusion: Similarities across nations can be explained by universal social preferences and differences by cultural differences in the trade-offs, e.g., between self-interest and fairness, or between competing distributive principles

Efficiency Principle

- Efficiency principle simply calls for the maximization of aggregate surplus
- Charness and Rabin incorporate this in their quasi-maximin preferences and experimentally test for it using a series of binary choices, including dictator like games
- They find a significant fraction of subjects in the US and Spain value surplus maximization, even when it comes at a personal cost
- Support diminishes, however, as efficiency increasingly conflicts with self-interest and/or equality

Results of Charness and Rabin Two-Person Dictator Games

- The dictator's payoff is the second entry
- The efficient allocation is always Distribution 2 on the right

Distribution 1	Distribution 2	% who Chose Dist.2
(0, 0)	(800, 200)	100
(200, 700)	(600, 600)	73
(400, 400)	(750, 400)	69
(400, 400)	(750, 375)	49
(300, 600)	(700, 500)	33

- The results indicate a desire to maximize surplus, even at a cost, but suggest a trade-off with self-interest and equality

Need Principle

- The need principle simply calls for allocations that meet the basic needs of people for food, clothing and housing
- Eckel and Grossman (*GEB* 1996): Dictators allocate 11% of stakes to anonymous student counterparts but a significantly greater 30% to a known charity
- Konow (Mixed Feelings 2005): This dictator study uses 1) known student counterparts and unfamiliar charities and 2) explicit statements of need in the case of charities
- Accountability/equity applies to students, and need to charities, whereby the latter implies a greater level of giving
- In comparison to Eckel and Grossman, this study finds even greater generosity toward charities and a larger difference in giving to students versus charities

Results of Standard Dictator and Charity Sessions in Mixed Feelings

- With $10 stakes, the modal gift of Standard dictators (i.e., those with student counterparts) is $0, or $5 among those who give any positive amount at all, whereas the modal gift among all Charity dictators (i.e., those with charity counterparts) is $10

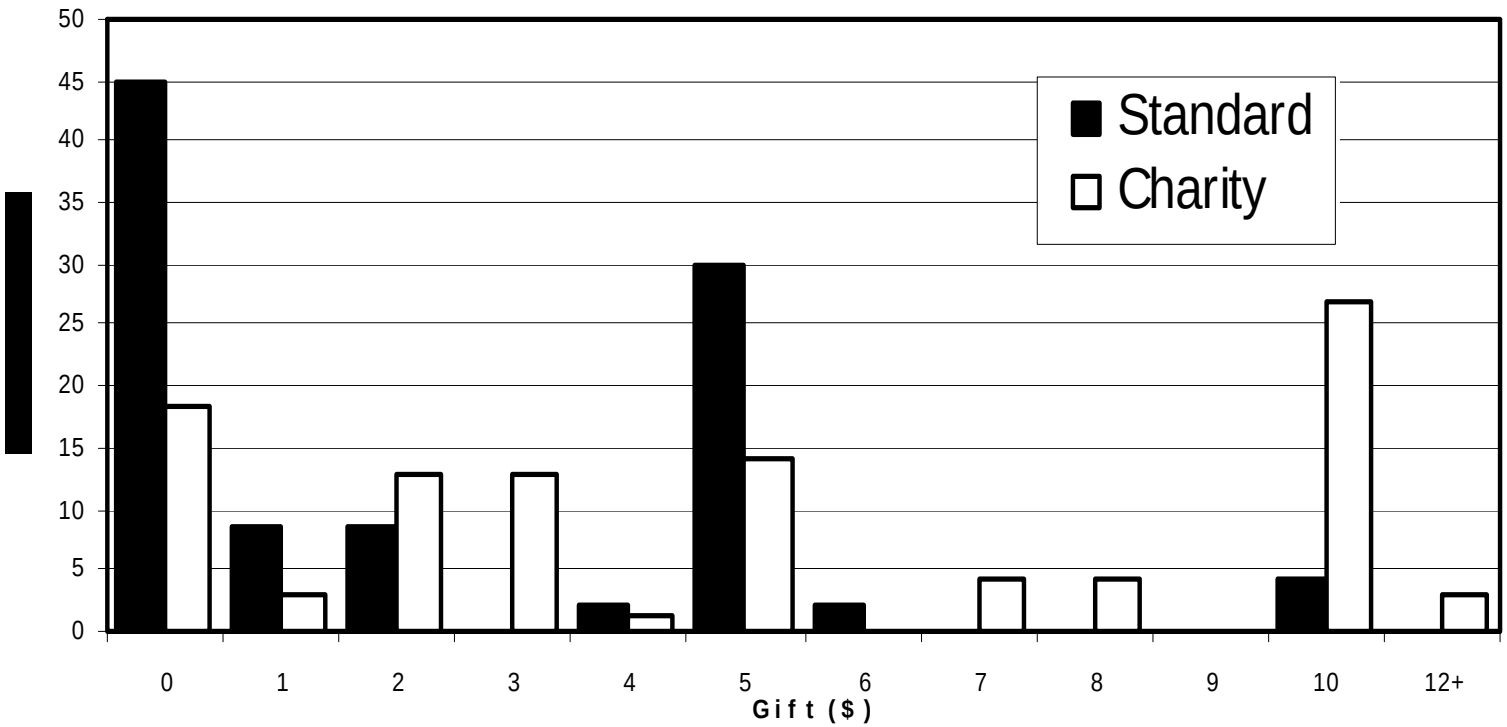

Reciprocal Preferences

- The chief limitation of models of distributive preferences is their inability to explain reciprocal behavior
- *Reciprocity* is an important behavioral phenomenon, especially in sustaining mutually beneficial relations when complete contracts are impossible or prohibitively expensive, e.g., in principal-agent relationships
- *Positive reciprocity*: being kind in response to an actual, perceived or expected kindness
- *Negative reciprocity*: being unkind in response to an actual, perceived or expected unkindness
- Discuss first experimental evidence on negative and positive reciprocity and then review theoretical attempts at formalizing this social preference

Negative reciprocity: Mini-ultimatum Games

- Falk, Fehr and Fischbacher (*Ec Inquiry* 2003), or FFF: These ultimatum games involved stakes of 10 points (8 Swiss Francs or about $6)
- The Proposer could always choose between just two Proposer/Responder distributions, x and y, e.g., x of 8/2 versus y of 5/5
- The Responder was then informed of the Proposer's alternatives and his choice and decided whether to accept or reject
- Distribution x always equaled 8 to Proposer and 2 to Responder (8/2), but the y distribution was varied across four treatments: 5/5, 2/8, 8/2, 10/0.

Two mini-ultimatum games

- Ex.: (a) 5/5 game vs. (b) 2/8 game

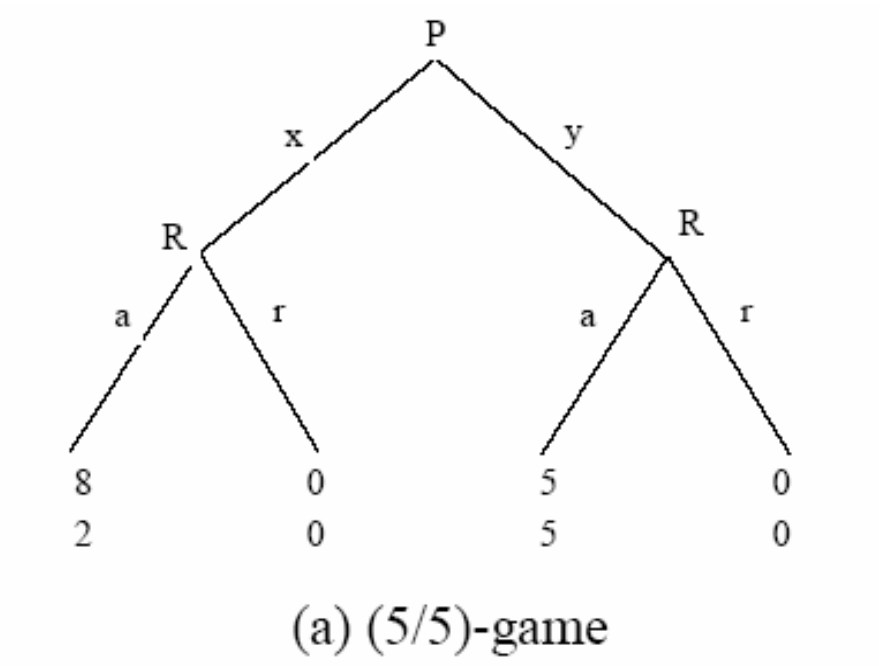

(a) (5/5)-game

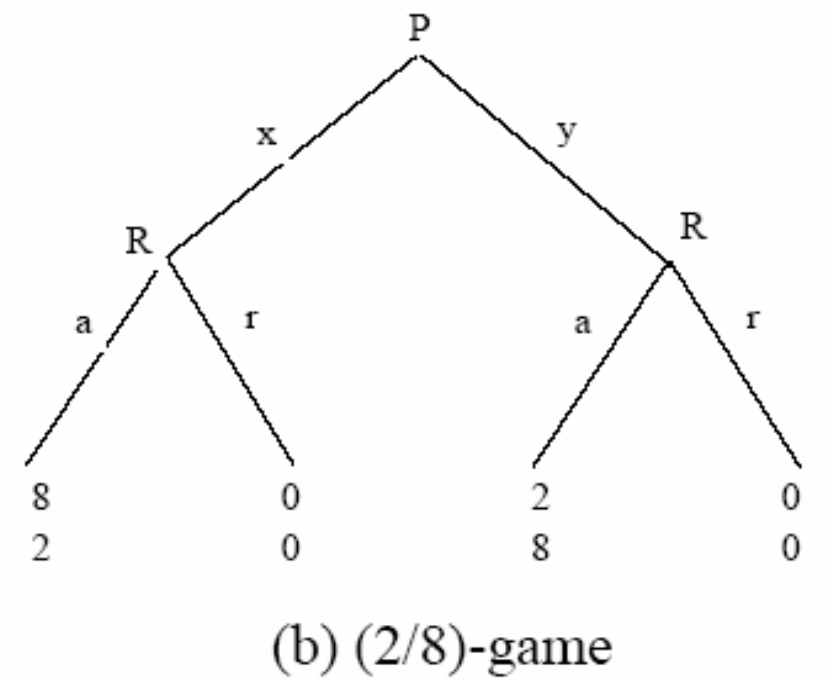

(b) (2/8)-game

Test of Negative Reciprocity

- If social preferences are purely distributive, the acceptance rate of Responders to the 8/2 offer should not depend on what alternatives the Proposer faced
- FFF, however, say that not only outcomes but also intentions matter: the Proposer who chose x in Game (b) turned down a similarly unequal distribution of 2/8, whereas the Proposer who chose x in Game (a) could have chosen an equal 5/5 distribution
- FFF predict that a larger fraction of Responders will reject the same 8/2 split in Game (a) than in Game (b) because they are motivated by negative reciprocity: they want to punish Proposers who have revealed their greater selfishness
- Predictions: the rejection rate is highest in the 5/5 game, and falls progressively in the 2/8, 8/2 and 10/0 games

Results of FFF mini-ultimatum games

- Predictions of negative reciprocity are confirmed:

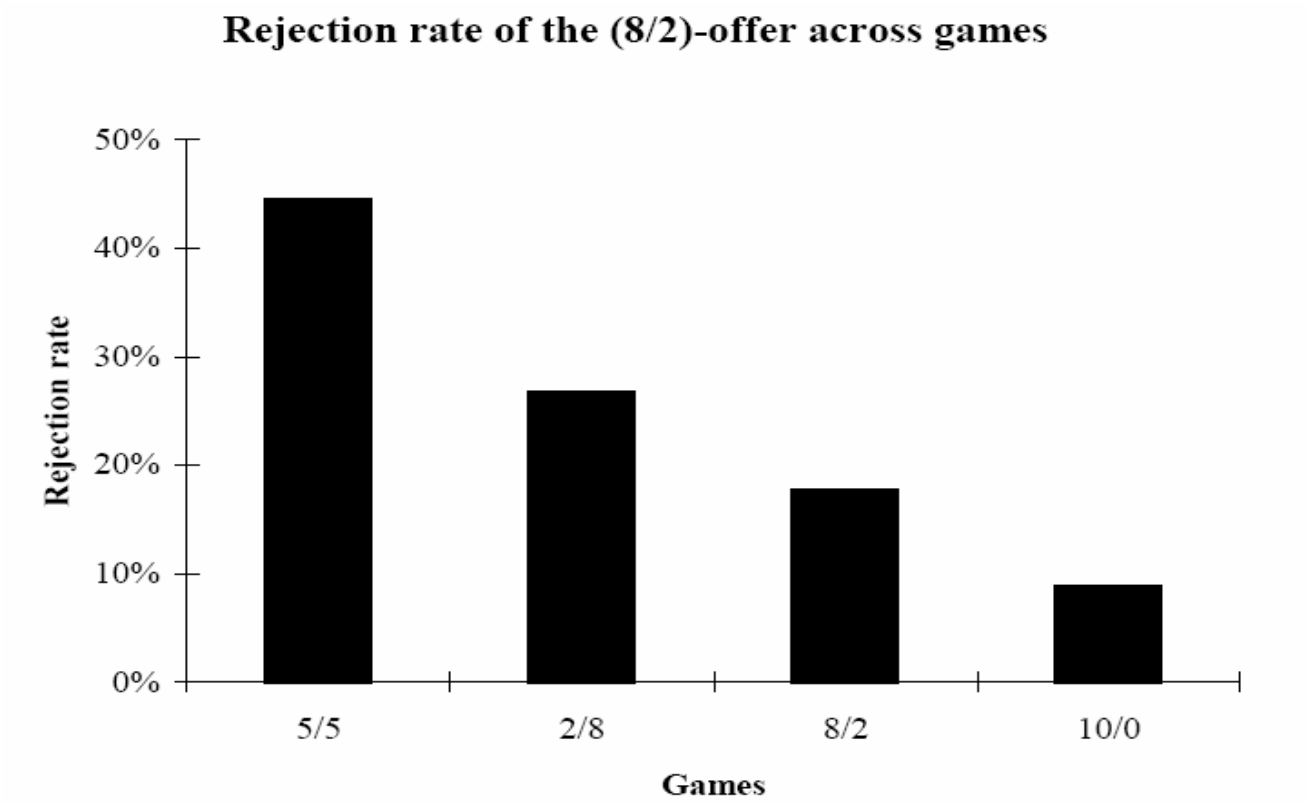

Positive Reciprocity: Trust Game

- Berg, Dickhaut and McCabe (*GEB* 1995), or BDM, introduced the "trust game"
- Frequently cited and replicated
- Simple, single shot gift exchange game:
- Anonymous subjects in separate rooms (call them subject A and subject B) receive $10 each; then there are two decision phases:
  1. Subject A (but not B) may send any fraction of his $10 fee to subject B, whereby any amount sent is tripled; Ex: if A sends $5, B receives $15
  2. Subject B may then send any amount of her gain back to subject A; Ex: B pays back $7 out of the $15 and keeps $8

Test of Trust and Positive Reciprocity

- Self-interest: The subgame perfect equilibrium with self-interested agents is, of course, for B to pay back nothing to A and, therefore, for A to send nothing to B
- Trust: BDM characterize transfers by A as signs of “trust”
- Reciprocity: BDM consider amounts paid back by B as indications of “reciprocity”
- Social history: BDM conducted this under two conditions:
  - *Social history* treatment: subjects were first informed of subject behavior in trust games
  - *No history* treatment: subjects were not so informed

  BDM conjecture that trust and reciprocity will be reinforced by information about prior trusting and reciprocal behavior

Results of BDM Trust Game

- Most A subjects sent something, on average $5.36
- The average total return, then, to B subjects was $16.08
- B subjects paid back, on average $6.46, rewarding A subjects by giving them an average net return of $1.10
- These results are for the *social history* treatment; the *no history* treatment revealed weaker trust and reciprocity
- Conclusions: people exhibit positive reciprocity, and social history reinforces these behaviors

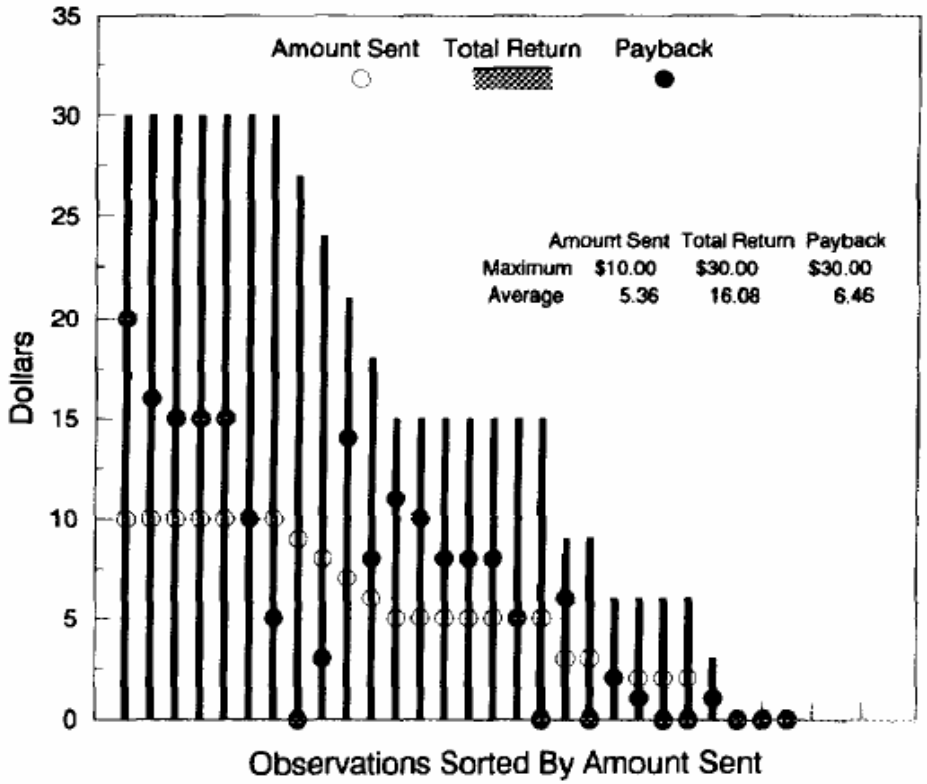

Positive and Negative Reciprocity: Moonlighting Game

- Abbink, Irlenbusch and Renner (*JEBO* 2000) introduce the "moonlighting game"
- Two players, A and B, are each endowed with 12 experimental "talers," equivalent to about 25 US cents
- Two decision stages:

1. First stage: player A chooses a transfer, $a \in \{-6,-5,\ldots,5,6\}$, to/from B; A can either take an amount from player B ($a<0$) or give B something ($a>0$)
   - If $a \geq 0$, A loses a talers and B receives $3a$ talers
   - If $a<0$, A gains $|a|$, and B loses $|a|$
2. Second stage: B observes a, and chooses a transfer to/from A, $b \in \{-6,-5,\ldots,17,18\}$
   - If $b \geq 0$, B loses b talers and A gains b
   - If $b<0$, B pays $|b|$ talers and reduces A's payoff by $|3b|$

- **Positive and negative reciprocity:**

  Moonlighting game makes explicit kindness/kindness and reward/punishment
- **Subgame perfect equilibrium:**
  - Since any transfer is costly, B neither rewards nor punishes
  - A, therefore, takes the maximum from B
- **Communication:**
  - In one treatment, B could propose a non-binding contract to A, which A could accept or reject (cheap talk)
  - In the other treatment, no such communication was allowed

Results of Moonlighting Game

- Results for both communication treatments were similar:

  B players were more inclined to punish $a<0$ than to reward $a>0$

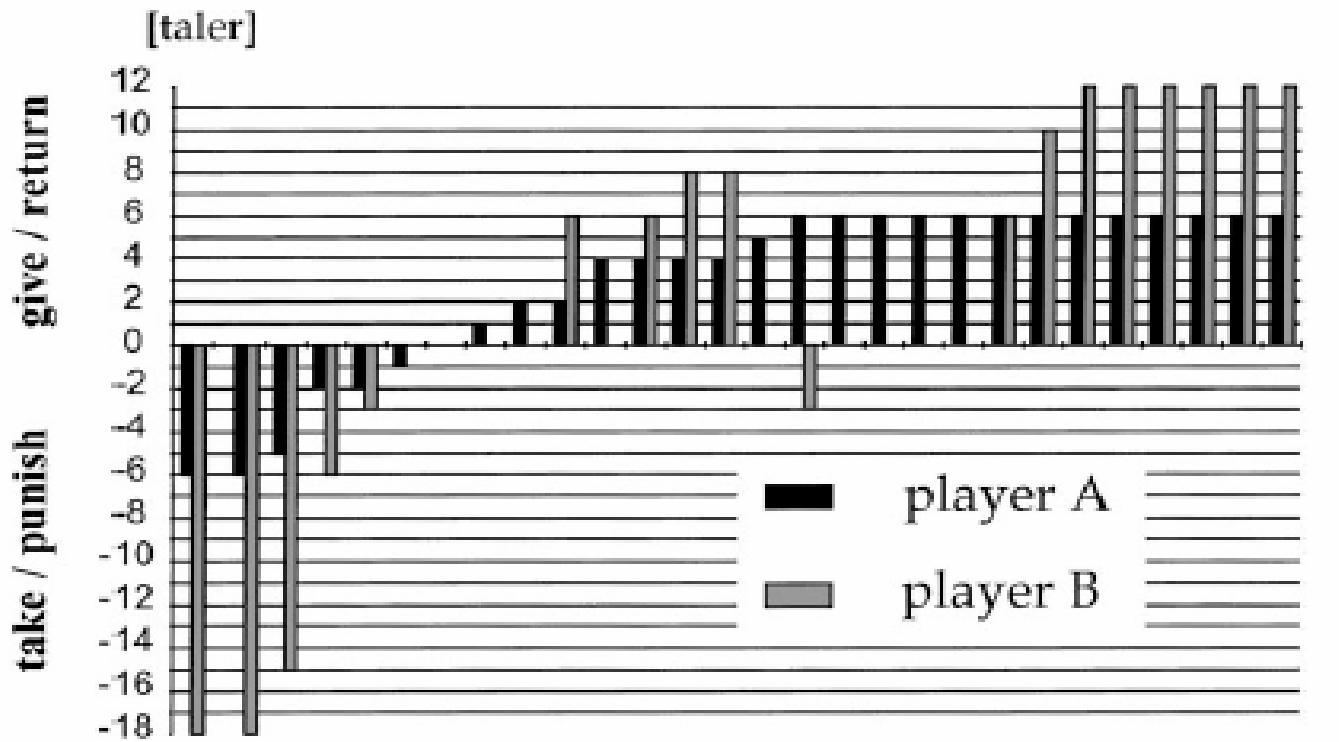

Models of Reciprocal Preferences

- Numerous models of reciprocal preferences have emerged, mostly aiming to capture the results of experiments
- These are more complex than models of distributive preferences
- There are fundamental differences in approach to reciprocal preferences
- Some incorporate both distributive and reciprocal preferences
- The seminal paper is Rabin's model of reciprocity

Rabin (*AER* 1993)

- In a departure from traditional game theory, Rabin employs "psychological game theory"
- Utility depends, not on allocations, but on players' beliefs:

$$U_i = \pi_i + \tilde{f}_j[1 + f_i]$$

where

π_i is the monetary payoff to player *i*,

$\tilde{f}_j$ is player *i*'s belief about how kind player *j* is being to him,

f_i is how kind player *i* is being to player *j*,

and where all terms depend on beliefs of players about strategies

- Negative reciprocity (f_i <0) increases utility if *i* believes *j* is being unkind ($\tilde{f}_j$ <0)
- Positive reciprocity (f_i >0) increases utility if *i* believes *j* is being kind ($\tilde{f}_j$ >0)

Pros:

- Seminal paper to formalize reciprocal preferences
- With sufficiently strong reciprocal preferences, cooperation is an equilibrium strategy in the prisoners dilemma, consistent with the evidence of many experiments

Cons:

- Does not explain third party reward and punishment that some experiments have revealed
- No role for distributive preferences
- There is a multiplicity of equilibria due to self-fulfilling beliefs about intentions
- Rabin's model is restricted to two player, normal form games

Subsequent authors have formulated more general models:

- Dufwenberg and Kirchsteiger (*GEB* 2004), or DK, generalize Rabin's reciprocity model to N-person extensive form games
- Falk and Fischbacher (*GEB* 2005), or FF, generalize Rabin to N-person extensive form games and also include distributive preferences in the form of inequality aversion
- Charness and Rabin (*QJE* 2002), or CR, model psychological games combining reciprocal preferences with distributive preferences, specifically, the previously discussed quasi-maximin preferences

Pros and Cons of These Models

Advantages:

- All generalize Rabin's model in some way
- FF and CR add distributive preferences, specifically, inequality aversion
- CR also incorporate the efficiency motive
- Their models are consistent with the stylized facts of many experiments

Disadvantages:

- As with Rabin, there are multiple equilibria because of self-fulfilling beliefs
- The complexity of these models makes it difficult to generate clear predictions in many cases
- Their complexity also makes it difficult to test them empirically

Levine (*Rev Ec Dynamics* 1998)

- Levine improves tractability by modeling preferences for others' types, rather than beliefs about their intentions
- The utility of person *i* is:

$$U_i(x_i, x_j) = x_i + \frac{a_i + \lambda a_j}{1+\lambda} x_j$$

where

$a_i \in$ (-1,1) is *i*'s coefficient of altruism

$a_j \in$ (-1,1) is *i*'s current estimate of *j*'s type, and

$\lambda \in$ [0,1] is a weight parameter

- Supposing first λ=0, the model reduces to $U_i = x_i + a_i x_j$; if a_i>0, person *i* is an altruist, whereas if a_i<0, *i* is spiteful. This does not explain, however, why a given person is sometimes altruistic and sometimes spiteful
- Allowing λ>0, an altruistic *i* feels more altruistic towards another altruist, etc.

- Sequential games become signaling games about player types

Advantages:

- Tractable model
- Can explain behavior in ultimatum games, public good games, and competitive market games

Disadvantages:

- Multiple equilibria are possible
- Does not incorporate distributive preferences
- Does not explain positive giving in the dictator game or player B paybacks in the trust game

Cox, Friedman and Gjerstad (2004)

- Model of both distributive and reciprocal preferences that aims to be tractable
- Preferences are formulated in terms of objectively defined variables
- Cox, et al use a modified CES utility function:

$$U_i(x_i, x_j) = \begin{cases} 1/\alpha \cdot (x_i^{\alpha} + \theta x_j^{\alpha}), & \alpha \neq 0; \\ x_i x_j^{\theta}, & \alpha = 0. \end{cases}$$

$$\alpha \in (-\infty, 1]$$

- Indifference curves for x_i and x_j vary with the value of the convexity parameter α from linear to Leontief consistent with allocations observed in many experiments
- The emotional state, θ, can be positive (gratitude) or negative (resentment), and is a function of the reciprocity motive and "status" motives (e.g., age, gender)
- The reciprocity motive is a function of the feasible payoffs and of a value, m_0, that is a reference value for determining whether an action should be reciprocated

Pros and Cons of Cox, et al.

Advantages:

- Tractable but still fairly general in its applications
- Preferences are formulated in terms of objectively defined variables rather than beliefs, avoiding some of the problems of psychological games

Disadvantages:

- Parameterized model rather than general functional form
- Model is formulated for only two players
- Theory has been tested on only two data sets

Conclusions

- We now have ample evidence that social preferences impact economic variables in a variety of important domains
- These preferences can be broadly assigned to one of at least two categories: distributive and reciprocal
- Incorporating these motives in models of individual preferences can improve the explanatory and predictive power of those models
- Existing models often prove complex, produce multiple equilibria or fail to account for important observations
- Additional work is needed to formulate better models and to collect relevant evidence on the models