The Limits to Altruism

Richard Povey

Hertford College and St Hilda's College, University of Oxford

May 18^{th} 2018

Presentation given at the Global Priorities Institute, University of Oxford. Slides are available at: `users.ox.ac.uk/~sedm1375/GPIpresentation.pdf`

Positive and Normative Limits to Altruism

- **Positive Limits** - *Evidence for Imperfect Altruism*
  - Experimental / Behavioural Game Theory
  - Economics of Happiness
  - Political Economy
  - Biased Altruism Within the Family
- **Normative Limits** - *Social Optimality of Imperfect Altruism*
  - Fundamental Theorems of Welfare Economics
  - Rotten Kid Theorem and Samaritan's Dilemma
  - Punishment in Dynamic Games
  - Evolutionary Theory
- **Normative Limits** - *Policy Consequences of Imperfect Altruism*
  - Market failures and corrective policies
  - Cost-benefit analysis

Rings of Altruism

The following diagram provides a visualisation of an individual's altruistic linkages to other individuals. The inner right might represent a family or proximate group.

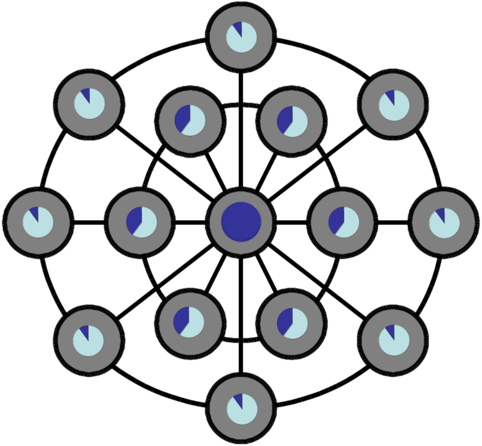

Strong and Weak Reciprocity

- **Weak reciprocity** is the form of altruism that can be seen as enlightened self-interest – individuals do each other a good turn because they rationally expect to be "paid back". However, this is insufficient to explain the complex functional integration of human societies [Fehr & Gächter, 2000].
- **Strongly reciprocal altruism** takes a **positive** and **negative** form, where individuals either help or harm others at material cost to themselves. This acts as "glue" holding institutions together, because the willingness of strong reciprocators to punish "cheats" forces selfish individuals to also behave well.
- The importance of the willingness of individuals to engage in altruistic punishment has also been reflected in **cultural selection theory**. Altruistic punishment is a key mechanism which acts as an "altruism amplification device", because it is usually less costly to punish another individual (e.g. by ostracising them) than it is to make an altruistic sacrifice for their benefit [Sober & Wilson, 1999].

Experimental / Behavioural Game Theory

- Some of the most important games that have been extensively tested in laboratory environments are:
  - Finitely-repeated prisoners' dilemma
  - Public goods games
  - Dictator or ultimatum games
  - Centipede game
- There is a vast literature which it would be foolish to attempt to summarise here, but a strong consensus that the predictions of classical game theory (based upon self interest and perfect common knowledge of rationality) are systematically violated. There is also a consensus that in order to explain observed behaviour it is necessary to introduce both limitations upon perfect common knowledge of rationality *and* other-regarding preferences.

Prisoners' Dilemma

	H	**L**
H	$b-c$, $b-c$	$-c$, $\underline{b}$
L	$\underline{b}$, $-c$	$\underline{0}$, $\underline{0}$

- If agents play a high altruism strategy (H) then convey a benefit b upon the other player but themselves incur a cost c.
- Since L is a strictly dominant strategy, backwards induction can be used to show that any finitely-repeated prisoners' dilemma results in a unique subgame-perfect Nash equilibrium with L played by both players in every period.
- However, significant co-operation occurs in finitely-repeated experimental prisoners' dilemma games. Andreoni and Miller conclude on the basis of experimental evidence that rational reputation-building on the part of most agents plus true altruistic behaviour on the part of a minority offers the best explanation for this phenomenon [Andreoni & Miller, 1993].

Public Goods Games

- Public Goods games are similar to N-player prisoners' dilemma but where each player can choose how much to contribute, with each unit of contribution creating a benefit b which is shared over the group but at a cost $b > c > \frac{b}{N}$.
- Evidence [Dawes & Thaler, 1988] shows that for small groups average contributions are usually in the region of 40%-60% of the optimal level. When the game is repeated with the same individuals playing, the average level of contributions tends to drop over time. However, the ability to altruistically punish non-co-operators and non-punishers greatly increases the ability to sustain co-operation [Fehr & Gächter, 2000] [Fehr & Fischbacher, 2003].

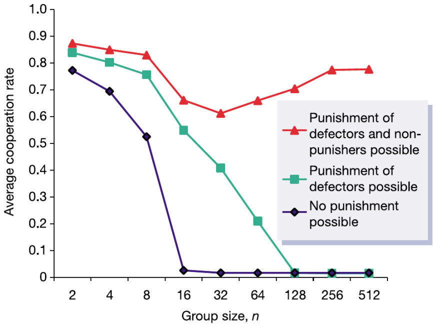

The Ultimatum Game

- The "ultimatum game" is played between two individuals. The first individual proposes the division of £x between the two individuals and the second individual can either accept the offer or refuse, in which case both get a payoff of 0.
- If both individuals are selfish and there is full common knowledge of rationality, classical game theory predicts that the first individual offers the smallest amount they can that is higher than 0 (i.e. 1p) and that the second individual accepts.
- However, when the game is played in experimental situations, the predicted outcome occurs extremely rarely, and there is significant variation between cultures regarding the amount that the first individual offers to the second. The empirical evidence has been summarised as showing that offers are usually between 30% and 40%, with the mode often being 50%. Very few offers are below 20%, and those which are this low are often rejected [Camerer & Thaler, 1995].

The Ultimatum Game

- Andreoni et al. have extended the ultimatum game to convexify the strategy space of the second individual by allowing them to continuously shrink the "pie" after the allocation is chosen by the proposer. Around 40% of subjects were found to have convex preferences for equity as illustrated by the diagrams below, whilst around 50% were found to have selfish preferences [Andreoni et al., 2003]:

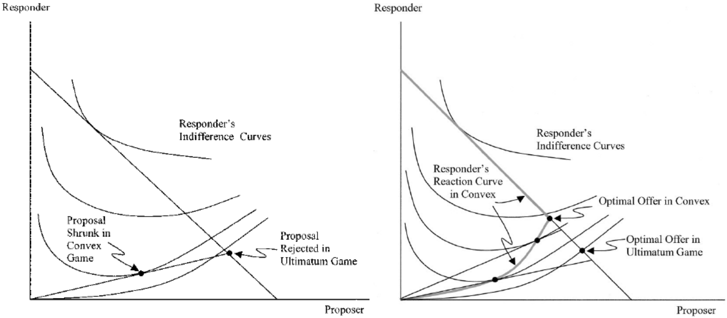

Centipede Game

- McKelvey and Palfrey conduct experimental centipede games and find that typically players pass for a number of periods before somebody takes the larger pile [McKelvey & Palfrey, 1992]. They explain this using the idea that a proportion of the population is altruistic, and that selfish individuals can pretend to be altruistic in order to get their opponent to co-operate. By calibrating the model to their data, they estimate that 5% of the population is believed to be altruistic [McKelvey & Palfrey, 1992].

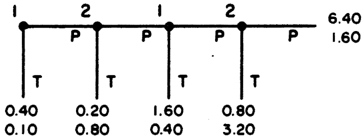

FIGURE 1.—The four move centipede game.

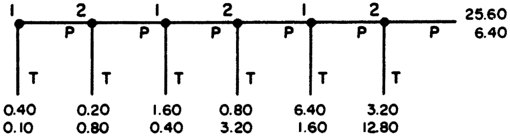

FIGURE 2.—The six move centipede game.

Economics of Happiness

A key result of this literature is that the positive relationship between happiness and income is greater within a society than it is over time as a society develops. This strongly suggests the presence of negative relative income effects ("keeping up with the Jones' ") which would seem to have a connection with negative strong reciprocity / preferences for fairness. [Easterlin, 1974] [Clark et al., 2008] [Layard, 2006].

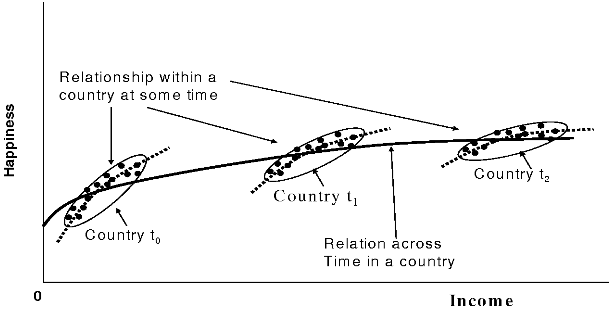

Biased Altruism within the Family

- Research into the discount rates parents are revealed to apply to their children's welfare in making costly decisions related to children's health, as measured by lead contamination, has shown them to be similar to market interest rates for wealthier parents and higher for poorer parents, but nowhere near as high as the discount rates of 20% - 50% found to be applied to consumer durables [Agee & Crocker, 1996]. This suggests a strong degree of parental altruism.

- Researchers into differentials in wages and human capital investment between males and females in the US [Behrman et al., 1986] have attempted to determine whether or not this is driven by greater weight upon the success of male children in parents' altruistic utility functions. The conclusion of this study was that existing wage differentials reinforce inequalities in human capital investment, but that parents do not, at root favour boys (in fact, if anything, the raw weighting on the welfare of girls is slightly higher). A similar study of the Phillipines [Davies & Zhang, 1995], however, found evidence for pure gender bias, underlining the fact that altruistic imperfections are culturally relative.

- Evidence on the treatment of step children in the US [Case et al., 1999], on the other hand, suggests that they do receive a smaller proportion of family income on food if they live with a stepmother, after controlling for income.

Political Economy

- Research on the relationship between environmental valuations reported in survey data and life expectancy [Popp, 2001] has also found evidence of a role for partial altruism. If people are fully altruistic, their life expectancy should not affect their contingent valuation of environmental goods. On the other hand, if they are fully selfish, the valuation should be, on average, zero as life expectancy goes to zero. The evidence, however, rejects both these hypotheses, suggesting the presence of partial altruism. The central estimate is of an equal weighting between individual welfare and the average welfare of future generations, but this estimate is very sensitive to the assumed discount rate.

- Some of the most striking evidence for imperfections to altruism comes from the realm of international political economy. A 1998 study which sought to estimate the marginal cost of additional life expectancy in different countries found that the implicit valuation of a life year in the richest countries was 300 times that in the poorest countries. Once difference in average life expectancy are taken into account, the cost of saving an entire lifespan in the richest countries came to 1000 times that of the poorest countries [Dowrick et al., 1998].

Normative Limits

There is much confusion of the ideal that a person ought to be allowed to pursue his own aims with the belief that, if left free, he will or ought to pursue solely his selfish aims.

[Hayek, 1960]

In my view the ideal society would be one in which each citizen developed a real split personality, acting selfishly in the market place and altruistically at the ballot box.

[Meade, 1973]

The Fundamental Theorems of Welfare Economics

Theorem

First Fundamental Theorem of Welfare Economics

If markets exist for all goods and all agents are perfectly competitive price takers then the general equilibrium of a system of markets is a Pareto-efficient allocation of resources.

Theorem

Second Fundamental Theorem of Welfare Economics

If the assumptions for the First Fundamental Theorem are satisfied, all preferences and technologies are convex and appropriate lump sum redistribution of endowments can be achieved, then any Pareto-efficient allocation can be achieved as a general equilibrium of a system of perfectly competitive markets.

- Comment: The assumptions required are very strong, plus Pareto-efficiency is only a bare minimum requirement for a socially desirable allocation. However, it is noteworthy that the FFTWE and SFTWE can hold for both selfish and altruistic preferences (provided, as we shall see, altruistic preferences take a non-paternalistic form so that no externalities/missing markets are created).

The Edgeworth Box

- The Edgeworth Box contains all the feasible allocations of goods in the economy. For example, take an economy with just two goods, where the total amount of the first good is represented by X and the total amount of the second good by Y. Suppose there is also only two people, person 1 and person 2.
- The width of the Edgeworth box is X and the height is Y. The amount of good X consumed by persons 1 and 2 is X_1 and X_2 respectively. Similarly, Y_1 and Y_2 are the amounts of good Y consumed by each individual.
- A feasible allocation is one where $X_1+X_2 \leq X$ and $Y_1+Y_2 \leq Y$.
- At point a, person 1 gets nothing and person 2 gets everything. At point b, person 2 gets nothing.

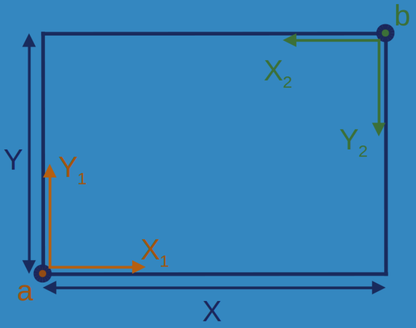

Although Edgeworth boxes are usually drawn for only two individuals and two goods, a similar framework would apply for any number of individuals and goods.

The Edgeworth Box

- Suppose that person 1 starts off owning a certain amount of goods X and Y, and so does person 2. This initial allocation is known as the endowment, and is represented by a point (point e) within the Edgeworth box. Here person 1 owns more of good X and person 2 owns more of good Y.

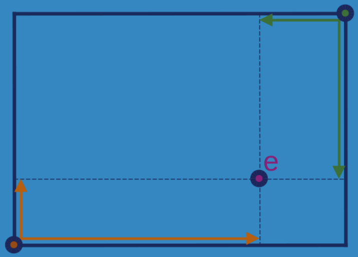

- Suppose that person 1 starts off owning a certain amount of goods X and Y, and so does person 2. This initial allocation is known as the endowment, and is represented by a point (point e) within the Edgeworth box. Here person 1 owns more of good X and person 2 owns more of good Y.

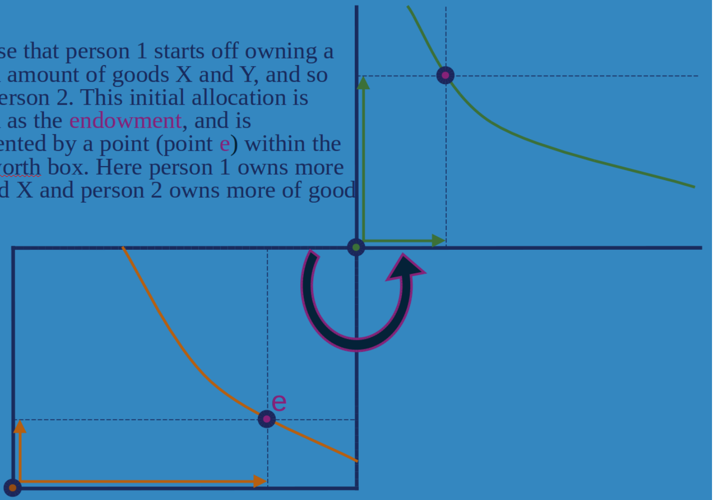

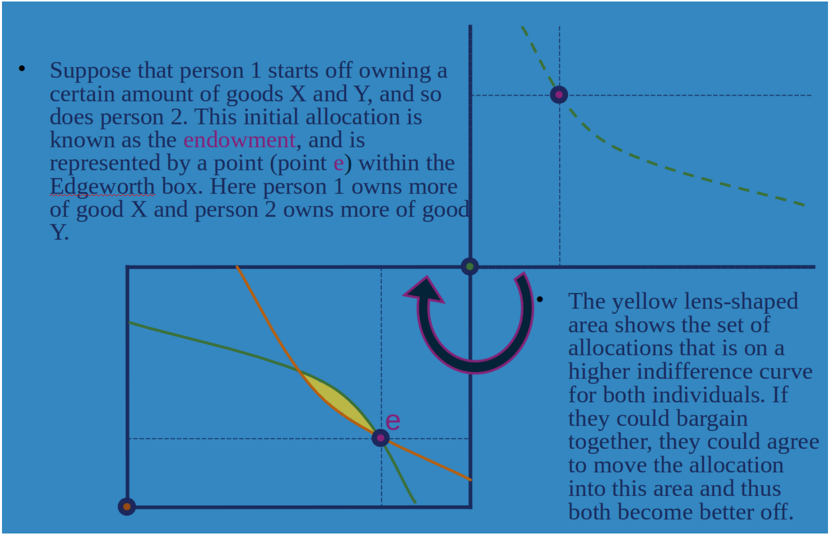
• Suppose that person 1 starts off owning a certain amount of goods X and Y, and so does person 2. This initial allocation is known as the endowment, and is represented by a point (point e) within the Edgeworth box. Here person 1 owns more of good X and person 2 owns more of good Y.
e
• The yellow lens-shaped area shows the set of allocations that is on a higher indifference curve for both individuals. If they could bargain together, they could agree to move the allocation into this area and thus both become better off.

The Edgeworth Box

- At a Pareto efficient allocation, neither individual can be made better off without making the other one worse off. For this to be the case, both individuals' indifference curve must be tangent. This means they must just touch but not cross. The set of Pareto efficient allocations is known as the contract curve.

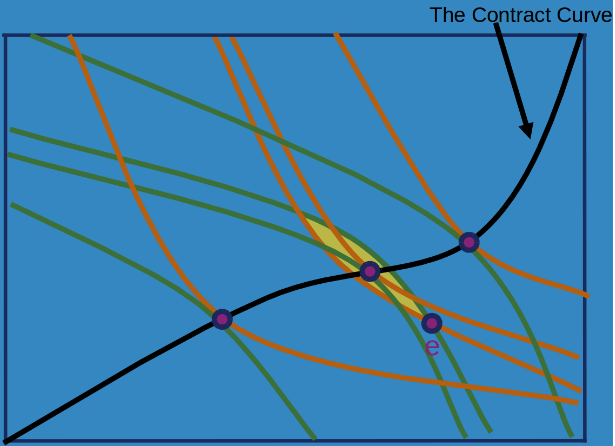

The Fundamental Theorems of Welfare Economics

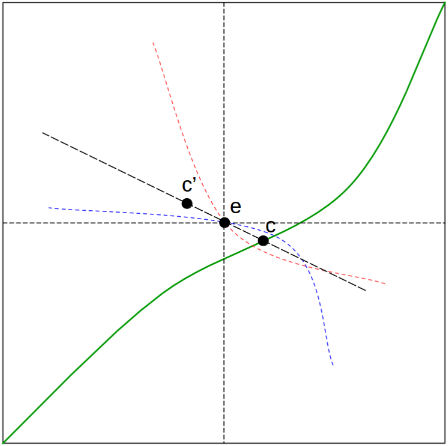

Person 1 and Person 2 trade from the endowment point *e* to the competitive general equilbrium point *c* on the contract curve. The slope of the dotted line is the ratio of the prices of the two goods in general equilibrium. The general equilibrium point c necessarily has the property that it is *equitable* – both individuals would be worse off if they were to exchange their equilibrium bundles.

The Fundamental Theorems of Welfare Economics

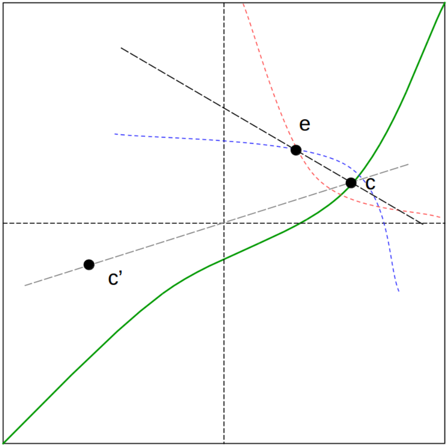

When initial endowments are unequal, the equity property of the allocation no longer holds. Here, person 1 starts off with more of both goods than person 2. Although they trade to a Pareto-efficient equilibrium point c, person 2 would prefer to exchange bundles and move to point c'. The allocation from the perfectly competitive market is therefore not equitable.

The Fundamental Theorems of Welfare Economics

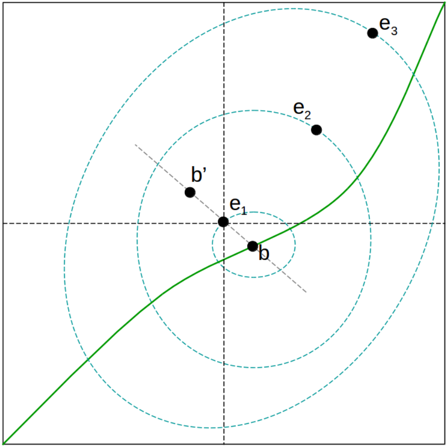

Whatever the initial endowment (i.e. however equal or unequal it is), since both players share the same indifference curves, they will always bargain to end up at their common bliss point b. This is clearly preferred by both players to b'. So, perfect altruism would solve the problem of unequal endowments and always ensure an equitable resource allocation.

The Non-Twisting Theorem

$$V_1 = \sqrt{X_1} + \sqrt{Y_1} \qquad V_2 = \sqrt{X_2} + \sqrt{Y_2}$$

$$U_1 = V_1 + \theta_1 V_2 = \sqrt{X_1} + \sqrt{Y_1} + \theta_1\left(\sqrt{X_2} + \sqrt{Y_2}\right)$$

$$U_2 = V_2 + \theta_2 V_1 = \sqrt{X_2} + \sqrt{Y_2} + \theta_2\left(\sqrt{X_1} + \sqrt{Y_1}\right)$$

"Shrunken contract curve" lies between points d and c. Point a is a Pareto-efficient CGE. Point b represents a possible Pareto-inefficient CGE if redistributive bargaining is not possible (e.g. due to a large number of each type of consumer).

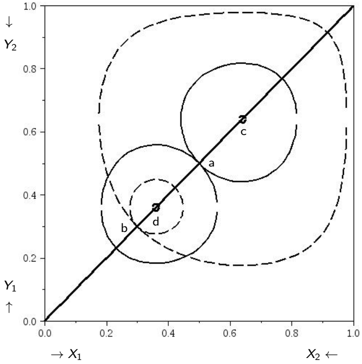

The Rotten Kid Theorem and The Samaritan's Dilemma

Theorem

Rotten Kid Theorem: *All members of a family will behave efficiently, even if they are completely selfish (or imperfectly altruistic) provided that the head of the family is sufficiently altruistic to make an operative transfer [Becker, 1974] [Bergstrom, 1989].*

- Consider the decision of a "rotten kid" over whether to take an action that increases or decreases their pre-transfer income at the expense or gain of another family member. The head of household will take into account the decision made by the rotten kid when deciding how big a transfer of resources to give him. Any action which increases the overall collective family wealth therefore makes the rotten kid better off. The rotten kid thus behaves fully efficiently.
- The key assumption is that the actions taken by the rotten kids must be of such a nature that they cause a shift rather than a change in the slope of the family's utility possibilities frontier.
- When "rotten kids" can take actions that distort the slope of the utility possibilities frontier, they may take inefficient actions from the perspective of the family. One of the main situations in which this condition does not apply is when there is a moral hazard problem between the head of the family and the rotten kid due to certain goods in the rotten kid's utility function not being under the direct control of the head of household via the transfer process. This is also known as the "Samaritan's Dilemma" [Bruce & Waldman, 1990] [Andreoni, 1989].

Intrinsic and Extrinsic Incentives

- **Intrinsic Incentives** - Altruistic preferences provide an *intrinsic* motivation for individuals to exhibit altruistic behaviour.
- **Extrinsic Incentives** - Punishment systems provide an *extrinsic* motivation.
- Often it is empirically difficult to distinguish between the two (e.g. enlightened self-interest in the repeated prisoners' dilemma) [Hammond, 1975].
- These two forms of incentives represent alternative "social technologies" that can potentially be used to achieve socially beneficial outcomes, but which can interfere with one another in a perverse manner. The moral preferences and institutions which have evolved in human society represent a particular "policy mix" which may (or may not) be socially optimal.

Punishment in Dynamic Games

[W]hen altruism improves static non-cooperative outcomes, it lessens the severity of credible punishments. An altruist may well be perceived as a "softy" and his threats may not be taken seriously.

[Bernheim & Stark, 1988]

[T]he most efficient way to provide low payoffs, in terms of incentives to cheat, is to combine a grim present with a credibly rosy future.

[Abreu, 1986]

Punishment in Dynamic Games

> *[I]n comparison with a situation wherein altruism is absent altogether, the prevalence of just some altruism could result in Pareto inferior outcomes. Hence, if the formation of altruism may not only fail to do any good but may actually make things worse whereas the formation of sufficiently high levels of altruism is almost always beneficial,...a troubling discontinuity arises: to the extent that the formation of altruism is like the rising of bread dough (i.e. it has to be gradual) groups yearning to build up their social stock of altruism may have to endure Paretial deterioration before experiencing Paretial gains. Perhaps one reason why a great many societies consist of self-interested economic men and women rather than altruistic economic men and women has to do with this nonmonotonicity.*
>
> *[Stark, 1989]*

Theorem

As altruism becomes perfect, the Nash equilibrium outcome in a repeated game becomes arbitrarily close to the socially efficient outcome [Bernheim & Stark, 1988].

The Sequential Punishment Model - A Word on Idiom

The sequential punishment model presented in this paper is intended as a highly abstract and stylized representation of social interaction, rather than as a realistic model of a specific situation. A simple "parable" can often help with the intuition. Models with a similar idiom include:

- Robinson Crusoe economy [Ruffin, 1972].
- Samuelson's "chocolate pension game" [Samuelson, 1958].
- Diamond's model of fiat money in a "coconut economy" [Diamond, 1984].

So, in that spirit, a desert island parable seems appropriate...

The Sequential Punishment Model - An Island Parable

Individuals (who have been on the island long enough to set up a "back garden") finish off a cold beer one at a time and must decide whether to walk to the bin or just throw their bottle into another individual's garden:

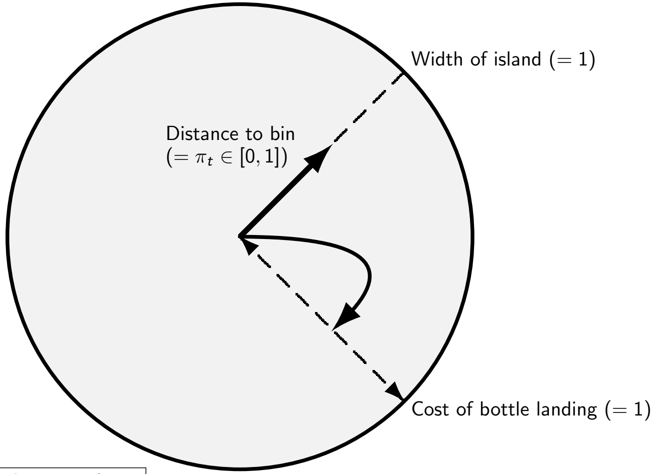

Coefficient of altruism - $\theta < 1$

Discount factor - $0 \leq \delta < 1$

The Sequential Punishment Model - Players' Preferences

- Felicity represents "private utility" from "economically fundamental" goods.
- In period t, player t moves so as to maximize his **expected social utility** u_t, discounted looking forward:
- Utility thus includes broader "moral preferences". This is of course only one among a number of alternative ways to define altruism. The advantage is that it enables us to simplify away from any "multiplier effects". Not to say that these do not exist and are not important in the real world, but in the sequential punishment model we wish to focus on the role of punishment and its interaction with altruistic preferences in as clean and simple an environment as possible.
- Social welfare function is utilitarian in felicities, or we can argue that Pareto efficiency requires an equilibrium where no bottles are thrown.

The Sequential Punishment Model - Three Effects

- **Temptation Effect** - Individuals with higher altruism are less *tempted* to inflict harm upon another individual for their own gain. (This is the main benefit from higher altruism.)
- **Willingness Effect** - Individuals with higher altruism are less *willing* to punish another individual for a previous misdemeanour by inflicting harm upon them. (This is a cost to higher altruism.)
- **Severity Effect** - Individuals with higher altruism also find some kinds of punishment less severe. In particular, if a fine was imposed, and some or all of the revenue is redistributed to another individual whose felicity has some weight in the utility function of the individual we are trying to punish, then any given size of fine is less severe for the punishee. (Another cost to higher altruism.)

The Sequential Punishment Model - Notation

- δ - Discount factor.
- θ - Coefficient of altruism.
- $\pi_t \in [0,1]$ - Benefit from harming / punishing in period t (randomly distributed between $\hat{\pi}$ and 1).
- θ^* - Socially optimal level of altruism - Enables efficient equilibrium to be sustained for largest possible range of δ.
- δ^* - Lowest possible value of δ for which the socially efficient outcome can be sustained. (Corresponds to θ^*.)
- $\kappa(\theta)$ - Net loss in utility when deviating from socially efficient equilibrium when optimal punishment is applied.

Overview - Socially Efficient Equilibria

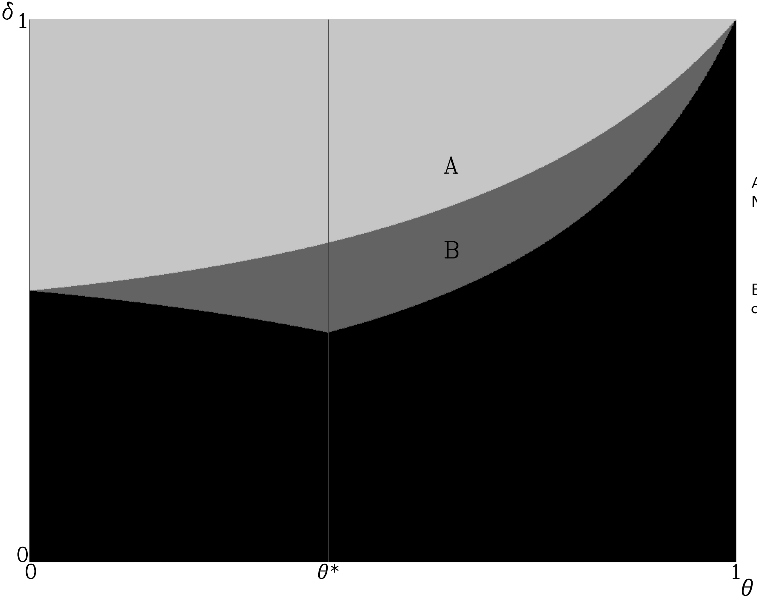

A - Supportable using Nash-reversion path

B - Supportable using optimal path

The Sequential Punishment Model - Solution Concepts

- **Folk Theorem** - [Aumann & Shapley, 1992] [Rubinstein, 1979] [Fudenberg & Maskin, 1986] For any given θ, as $\delta \longrightarrow 1$, the socially efficient outcome becomes supportable. We are interested here, however, in what happens as $\theta \longrightarrow 1$ for any given $\delta < 1$.
- **Optimal Penal Codes** - [Abreu, 1988] Abreu's framework of optimal penal codes in the form of punishment paths provides a natural framework that can be adapted to analyse socially efficient equilibria in the sequential punishment model.
- **Renegotiation Proofness** - [Farrell & Maskin, 1989] [Benoit & Krishna, 1993] We assume that society is able to avoid the temptation to let malefactors "off the hook". Thus we stick with subgame perfection rather than further refining the equilibrium criterion.

Results - Illustrated Using Uniform Distribution of Benefit

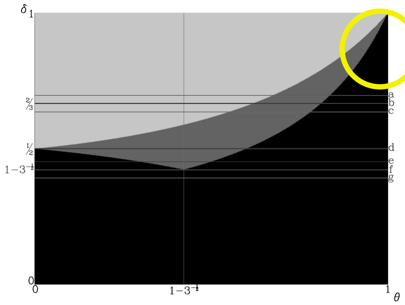

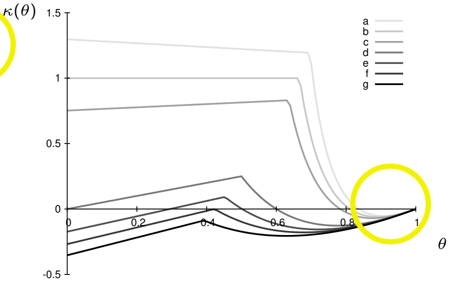

Theorem

As $\theta \longrightarrow 1^{-}$, $\kappa(\theta) \longrightarrow 0^{-}$.

Proof.

Intuition: As $\theta \longrightarrow 1^{-}$, the willingness and severity effect become negligible, and the temptation effect ensures that $\frac{d\kappa}{d\theta} > 0$. □

Results - Illustrated Using Uniform Distribution of Benefit

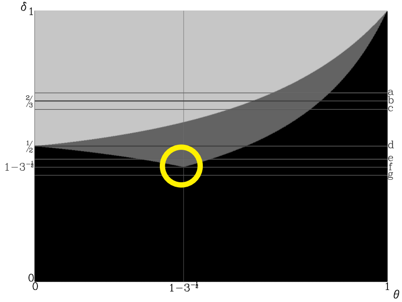

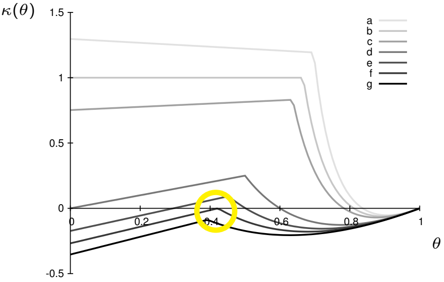

Theorem

$\theta^* \in (0,1)$ *(* $\theta^* = 1 - \frac{1}{\sqrt{3}} \approx 42\%$ *for uniform benefit distribution)*

Proof.

Intuition: Let $\delta = \delta^* = \theta^*$. If $\theta = \theta^* + \epsilon$ then willingness and severity effect dominate temptation effect, so $\frac{d\kappa}{d\theta} < 0$. If $\theta < \theta^*$ then (because punishment is maximal) willingness effect is 0, temptation effect dominates severity effect, so $\frac{d\kappa}{d\theta} > 0$. □

Other Benefit Distributions

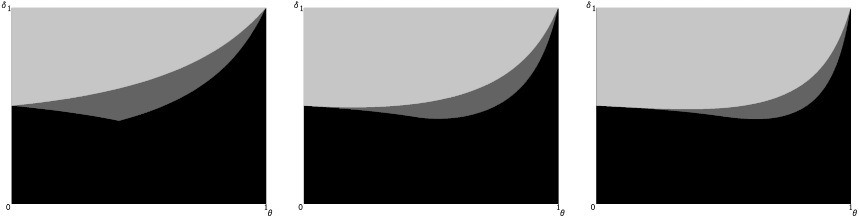

Figure: Socially efficient equilibria for $g(\pi) = 1$, $g(\pi) = 2\pi$ and $g(\pi) = 3\pi^2$.

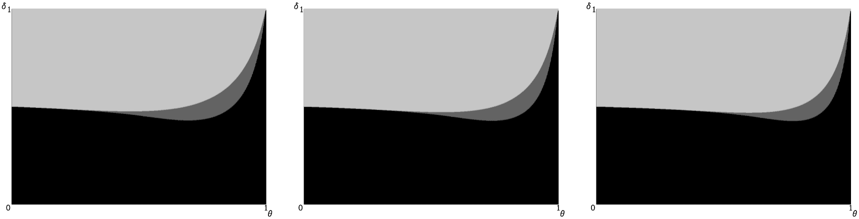

Figure: Socially efficient equilibria for $g(\pi) = 4\pi^3$, $g(\pi) = 5\pi^4$ and $g(\pi) = 6\pi^5$.

Multilevel Selection Theory

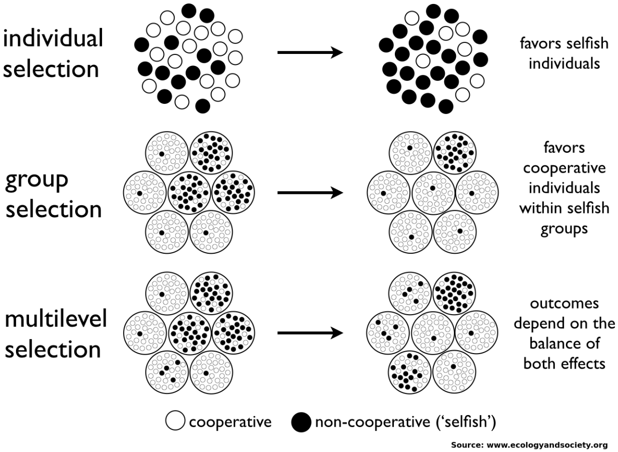

Haystacks Model

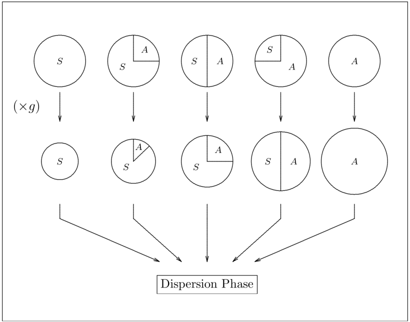

Note: Although the Haystacks model provides a neat mathematical framework for analysing group selection, there are other mechanisms in play that have a similar effect and make group selection a lot more plausible for cultural evolution:

- Ostracism
- Inter-group conflict

Haystacks Model

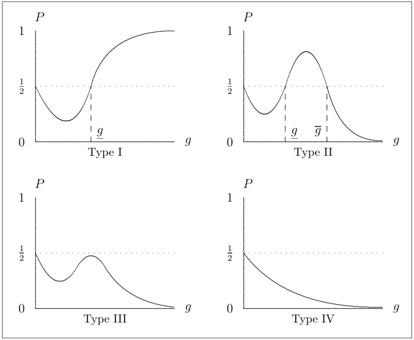

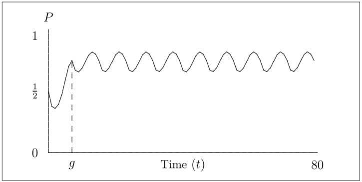

Source: [Cooper & Wallace, 2004].
A constant positive probability of mutations rules out Type I cycles, and we then need the number of periods of isolation to be in a "Goldilocks zone" to ensure Type II cycles and enable altruism to survive in the population.

Evolutionary Sequential Punishment Model

A's Phenotype / **B's Phenotype**		**H** $\theta_A \geq \hat{\pi}$	**L** $\theta_A < \hat{\pi}$
H	$\theta_B \geq \hat{\pi}$	0 $\boxed{0}$ 0	$\hat{\pi}$ $\boxed{\hat{\pi} - 1}$ $-\frac{1}{2}$
L	$\theta_B < \hat{\pi}$	$-\frac{1}{2}$ $\boxed{\hat{\pi} - 1}$ $\hat{\pi}$	$\frac{\hat{\pi}}{2}$ $\boxed{\hat{\pi} - 1}$ $\frac{\hat{\pi}}{2}$

Middle payoff is social welfare (sum of payoffs of players 1 to 3). Player A and B face a 50-50 lottery on whether they get to move first as player 1 or second as player 2. With universal high altruism, neither player ever inflicts harm. With universal low altruism, player 2 inflicts harm on player 3 but the credible threat of "switching" this into player 1 deters player 1 from inflicting harm. If player 1 is altruistic and player 2 is selfish then player 1 does not inflict harm but player 2 does (upon player 3). If player 1 is selfish and player 2 is altruistic then player 2 does not inflict harm so player 1 does (upon player 2).

Evolutionary Sequential Punishment Model

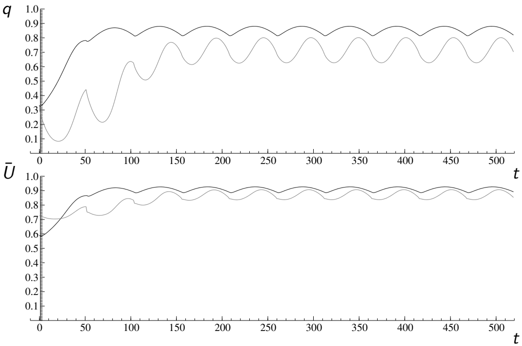

Grey line shows altruism in sequential punishment model, black line in analogous prisoners' dilemma model.

Figure: 52 periods of isolation, $\pi = 0.075$

Theorem

Altruism always survives more easily in the prisoners' dilemma than in the sequential punishment model [Povey, 2014].

Policy Consequences

- The impact of the presence of altruism has been analysed in the context of the theory of cost-benefit analysis. It has been shown that in the presence of non-paternalistic altruism, household willingness to pay for a public good exceeds the sum of individuals' willingness to pay [Quiggin, 1998]. An attempt to estimate the results of the presence of paternalistic altruism on the value of statistical life by calibrating to 2.5 individuals per altruistic family predicted that this value is 10%-40% higher than the individual value [Jones-Lee, 1992].
- A number of overlapping generations models of environmental degradation have shown that the presence of partial altruism does not guarantee an efficient internalization of these externalities [Jouvet et al., 2000] [Turner, 1997]. It has also been found that co-operation between nations to internalise current environmental externalities may lead to a deterioration of future environment relative to non-co-operation because improved environmental technology frees up more resources for consumption [John & Pecchenino, 1997].
- It has been shown that altruism can either increase or decrease optimal Pigovian taxation depending on its precise form [Johansson, 1997]. The optimal subsidy for voluntary giving has also been found to depend in a dramatic way upon the nature of the altruistic motivation for giving [Kaplow, 1998].
- It has been shown that infectious happiness (similar to altruism if agents are conscious of its effects) reduces market failures such as over-supply of labour (due to relative income effects) and under-supply of public goods [Povey, 2015].

References I

Abreu, Dilip (1986).
"Extremal Equilibria of Oligopolistic Supergames".
Journal of Economic Theory, 39, 191–225.

Abreu, Dilip (1988).
"On the Theory of Infinitely Repeated Games with Discounting".
Econometrica, 56(2), 383–396.

Agee, Mark D. and Crocker, Thomas D. (1996).
"Parents' Discount Rates for Child Quality".
Southern Economic Journal, 63(1), 36–50.

Andreoni, James (1989).
"Giving with Impure Altruism: Applications to Charity and Ricardian Equivalence".
The Journal of Political Economy, 97(6), 1447–1458.

Andreoni, James and Castillo, Marco and Petrie, Ragan (2003).
"What Do Bargainers' Preferences Look Like? Experiments with a Convex Ultimatum Game".
The American Economic Review, 93(3), 672–685.

Andreoni, James and Miller, John H. (1993).
"Rational Cooperation in the Finitely Repeated Prisoner's Dilemma: Experimental Evidence".
The Economic Journal, 103(418), 570–585.

Aumann, Robert J. and Shapley, Lloyd S. (1992).
"Long Term Competition - A Game Theoretic Analysis".
UCLA Economics Working Papers 676, UCLA Department of Economics.

References II

Becker, Gary S. (1974).
"A Theory of Social Interactions".
Journal of Political Economy, 82, 1063–1093.

Behrman, Jere R. and Pollak, Robert A. and Taubman, Paul (1986).
"Do Parents Favor Boys?".
International Economic Review, 27(1), 33–54.

Benoit, Jean-Pierre and Krishna, Vijay (1993).
"Renegotiation in Finitely Repeated Games".
Econometrica, 61(2), 303–23.

Bergstrom, Theodore C. (1989).
"A Fresh Look at the Rotten Kid Theorem–and Other Household Mysteries".
The Journal of Political Economy, 97(5), 1138–1159.

Bernheim, Douglas B. and Stark, Oded (1988).
"Altruism within the Family Reconsidered: Do Nice Guys Finish Last?".
The American Economic Review, 78(5), 1034–1045.

Bruce, Neil and Waldman, Michael (1990).
"The Rotten-Kid Theorem Meets the Samaritan's Dilemma".
The Quarterly Journal of Economics, 105(1), 155–165.

Camerer, Colin and Thaler, Richard H. (1995).
"Anomalies: Ultimatums, Dictators and Manners".
The Journal of Economic Perspectives, 9(2), 209–219.

References III

Case, Anne and Lin, I. Fen and McLanahan, Sara (1999).
"Household Resource Allocation in Stepfamilies: Darwin Reflects on the Plight of Cinderella".
The American Economic Review, 89(2), 234–238.

Clark, Andrew E. and Frijters, Paul and Shields, Michael A. (2008).
"Relative Income, Happiness, and Utility: An Explanation for the Easterlin Paradox and Other Puzzles".
Journal of Economic Literature, 46(1), 95–144.

Cooper, Ben and Wallace, Chris (2004).
"Group Selection and the Evolution of Altruism".
Oxford Economic Papers, 56, 307–330.

Davies, James B. and Zhang, Junsen (1995).
"Gender Bias, Investments in Children, and Bequests".
International Economic Review, 36(3), 795–818.

Dawes, Robyn M. and Thaler, Richard H. (1988).
"Anomalies: Cooperation".
The Journal of Economic Perspectives, 2(3), 187–197.

Diamond, Peter (1984).
"Money in Search Equilibrium".
Econometrica, 52(1), 1–20.

Dowrick, Steve and Dunlop, Yvonne and Quiggin, John (1998).
"The Cost of Life Expectancy and the Implicit Social Valuation of Life".
The Scandinavian Journal of Economics, 100(4), 673–691.

References IV

Easterlin, Richard A. (1974).
"Does economic growth improve the human lot? : some empirical evidence".
Nations and households in economic growth : essays in honor of Moses Abramovitz, (pp. 89–125).

Farrell, Joseph T. and Maskin, Eric S. (1989).
"Renegotiation in Repeated Games".
Games and Economic Behaviour, 1(1), 327–360.

Fehr, Ernst and Fischbacher, Urs (2003).
"The Nature of Human Altruism".
Nature, 425, 785–791.

Fehr, Ernst and Gächter, Simon (2000).
"Cooperation and Punishment in Public Goods Experiments".
The American Economic Review, 90(4), 980–994.

Fudenberg, Drew and Maskin, Eric (1986).
"The Folk Theorem in Repeated Games with Discounting or with Incomplete Information".
Econometrica, 54(3), 533–54.

Hammond, Peter (1975).
Charity: Altruism or Cooperative Egoism?
In E. S. Phelps (Ed.), *Altruism, Morality and Economic Theory* (pp. 115–131). Russell Sage Foundation, New York.

Hayek, Friedrich A. (1960).
The Constitution of Liberty.
Routledge and Kegan Paul Ltd, London.

Johansson, Olof (1997).
"Optimal Pigovian Taxes under Altruism".
Land Economics, 73(3), 297–308.

John, A. Andrew and Pecchenino, Rowena A. (1997).
"International and Intergenerational Environmental Externalities".
The Scandinavian Journal of Economics, 99(3), 371–387.

Jones-Lee, Michael. W. (1992).
"Paternalistic Altruism and the Value of Statistical Life".
The Economic Journal, 102(410), 80–90.

Jouvet, Pierre-Andre and Michel, Philippe and Vidal, Jean-Pierre (2000).
"Intergenerational Altruism and the Environment".
The Scandinavian Journal of Economics, 102(1), 135–150.

Kaplow, Louis (1998).
"Tax Policy and Gifts".
The American Economic Review, 88(2), 283–288.

Layard, Richard (2006).
"Happiness and Public Policy: A Challenge to the Profession".
The Economic Journal, 116(510), pp. C24–C33.

McKelvey, Richard D. and Palfrey, Thomas R. (1992).
"An Experimental Study of the Centipede Game".
Econometrica, 60(4), 803–836.

References VI

Meade, James E. (1973).
Theory of Economic Externalities: The Control of Environmental Pollution and Similar Social Costs.
Sijthoff, Leiden.

Popp, David (2001).
"Altruism and the Demand for Environmental Quality".
Land Economics, 77(3), 339–349.

Povey, Richard (2014).
"Punishment and the potency of group selection".
Journal of Evolutionary Economics, 24(4), 799–816.

Povey, Richard (2015).
"The welfare economics of infectious happiness".
Economics Letters, 133, 1 – 3.

Quiggin, John (1998).
"Individual and Household Willingness to Pay for Public Goods".
American Journal of Agricultural Economics, 80(1), 58–63.

Rubinstein, Ariel (1979).
"Equilibrium in Supergames with the Overtaking Criterion".
Journal of Economic Theory, 21(1), 1–9.

Ruffin, Roy J. (1972).
"Pollution in a Crusoe Economy".
The Canadian Journal of Economics, 5(1), 110–118.

Samuelson, Paul A. (1958).
"An Exact Consumption-Loan Model of Interest with or without the Social Contrivance of Money".
Journal of Political Economy, 66, 467.

Sober, Elliot and Wilson, David S. (1999).
Unto Others.
Harvard University Press, Cambridge, Massachussets.

Stark, Oded (1989).
"Altruism and the Quality of Life".
The American Economic Review, 79(2), 86–90.

Turner, Matthew A. (1997).
"Parental Altruism and Common Property Regulation".
The Canadian Journal of Economics, 30(4a), 809–821.