

Confronting Theory with Experimental Data and vice versa

Lecture II: Social Preferences

**The Centre for Experimental Social Sciences (CESS)
Nuffield College, University of Oxford**

Jun 15-17, 2009

Background

- People often sacrifice their own payoffs in order to increase the payoffs of *anonymous* others.
- They do so even in circumstances that do not engage reciprocity motivations or strategic behavior.
- This has led economists to begin the systematic study of the *distributional preferences* that govern such behavior.

Social preferences theories

- *Social welfare*
 - persons pursue an aggregate of their own payoffs and those of others.
- *Inequality aversion*
 - persons care about differences between their own and others' payoffs.

Template for analysis

- The *dictator game* eliminates strategic behavior and reciprocity motivations and implicates only distributive preferences.
- Choices made by a person *self* that have consequences for her own payoff and the payoffs of an anonymous *other*.
- Throughout, we denote persons *self* and *other* by S and O , respectively, and the associated monetary payoffs by π_S and a π_O .

Given a *nondegenerate* utility function

$$U_S = u_S(\pi_S, \pi_O)$$

that captures the possibility of giving, person *self* is *selfish* when for any π and π'

$$u_S(\pi) \geq u_S(\pi') \text{ if and only if } \pi_S \geq \pi'_S$$

and otherwise displays some form of *altruism*.

Prototypical social preferences

Charness and Rabin (*QJE*, 2002) propose the following simple formulation

$$U_S(\pi_S, \pi_O) \equiv (\rho r + \sigma q)\pi_O + (1 - \rho r - \sigma q)\pi_S$$

where

$$r = 1 \text{ (} s = 1 \text{) if } \pi_s > \pi_o \text{ (} \pi_s < \pi_o \text{) and zero otherwise.}$$

Increasing the ratio ρ/σ indicates an increase in concerns for increasing aggregate payoffs rather than reducing differences in payoffs.

- (i) *competitive* preferences ($\sigma \leq \rho < 0$) – utility increases in the difference $\pi_S - \pi_O$
- (ii) *narrow self-interest* or *selfish* preferences ($\sigma = \rho = 0$) – utility depends only on π_S
- (iii) *difference aversion* preferences ($\sigma < 0 < \rho < 1$) – utility is increasing in π_S and decreasing in the difference $\pi_S - \pi_O$
- (iv) *social welfare* preferences ($0 < \sigma \leq \rho \leq 1$) – utility is increasing in both π_S and π_O .

Objections and replies

An unpublished working paper concludes

This puts the basis of our modeling on unobservable preferences, and raises the specter of extensive ad hoc modeling with a basis primarily in psycho babble.

Camerer (2003) replies

The goal is not to explain every different finding by adjusting the utility function just so; the goal is to find parsimonious utility functions, supported by psychological intuition...

Experimental design

In a typical dictator game, the problem faced by *self* is simply allocating a fixed total income between *self* and *other*.

Person *self* divides some *endowment* m between *self* and *other* in any way he wishes such that

$$\pi_S + \pi_O = m.$$

The dictator game, developed by Andreoni and Miller (*Econometrica*, 2002), allows for m to be spent on π_S and π_O at *price* levels p_S and p_O such that

$$p_S\pi_S + p_O\pi_O = m.$$

This configuration creates *budget sets* over π_S and π_O that allow for the thorough testing for consistency with utility maximization.

Experimental procedures

- A graphical computer interface that allows for the efficient collection of many observations per subject.
- The graphical representation does not force subjects into discrete choices that suggest extreme prototypical preference types.
- It generates a very rich data set well-suited to studying behavior at the level of the individual subject.

Econometric specification

- Our subjects' CCEI scores are sufficiently near one to justify treating the data as utility-generated.
- If choice data satisfy GARP we would ideally like to extract a rationalizing utility function.
- Afriat's theorem tells us that if a rationalizable utility function exists, it can be chosen to be increasing, continuous, and concave.

- The constant elasticity of substitution (CES) utility function is commonly employed in demand analysis.
- The patterns observed in the nonparametric approach suggest that it is appropriate to estimate a CES demand function.
- The CES is useful because attitudes towards giving can be adjusted by means of a single parameter.

The CES utility function is given by

$$U_S = [\alpha(\pi_S)^\rho + (1 - \alpha)(\pi_O)^\rho]^{1/\rho}$$

α - the relative weight on *self* versus *other*.

ρ - the curvature of the altruistic indifference curves.

$\rho > 0$ ($\rho < 0$) indicate preference weighted towards increasing total (reducing differences in) payoffs.

The CES demand function is given by

$$\pi_s(p, m') = \frac{A}{p^r + A} m'$$

where

$$r = -\rho / (\rho - 1)$$

and

$$A = [\alpha / (1 - \alpha)]^{1/(1-\rho)} .$$


This generates the following individual-level econometric specification for each subject n :

$$\frac{\pi_{sn}^i}{m_n^i} = \frac{A_n}{(p_n^i)^{r_n} + A_n} + \epsilon_n^i$$

where ϵ_n^i is assumed to be distributed normally with mean zero and variance σ_n^2 .

Estimate \hat{A}_n and \hat{r}_n using non-linear tobit maximum likelihood, and use this to infer the values of the CES parameters $\hat{\alpha}_n$ and $\hat{\rho}_n$.

Scatterplot of the CES estimates



Distinguishing social preferences from preferences for altruism

- Distributional preferences may be divided into two qualitatively different types which we call *preferences for altruism* and *social preferences*.
- Social preferences and distributional preferences are used interchangeably in the literature and our usage is not quite standard.
- Nevertheless, the distinctions that we draw are straightforward and capture important differences.

- *Preferences for altruism*

- tradeoffs between the payoffs to *self* and the payoffs to *others*.

- *Social preferences*

- tradeoffs between the payoffs to *others* (i.e. all persons except *self*).

A common assumption used in demand analysis allows for a clear demarcation between social preferences and preferences for altruism:

Independence For any π_S, π'_S , and profiles $\pi_O = (\pi_A, \pi_B)$ and π'_O
 $u_S(\pi_S, \pi_O) > u_S(\pi_S, \pi'_O)$ if and only if $u_S(\pi'_S, \pi_O) > u_S(\pi'_S, \pi'_O)$.


If the independence property is satisfied, then the utility function $u_S(\pi_S, \pi_O)$ is (weakly) *separable*.

There exists a *subutility* function $w_S(\pi_O)$ and a *macro* function $v_S(\pi_S, w_S)$ with v_S strictly increasing in w_S such that


$$u_S(\pi_S, \pi_O) \equiv v_S(\pi_S, w_S(\pi_O)).$$

- This formulation makes it possible to represent distributional preferences in a particularly convenient manner.
- The macro function v_S represents preferences for altruism, whereas the subutility function w_S represents social preferences.
- Separability imposes convenient (but specific and quite restrictive) patterns on demand behavior (Karni and Safra 2002).


Decision-level distribution of tokens given *toothers* as a fraction of total tokens kept and given



Decision-level distribution of expenditure on tokens given toothers
as a fraction of total expenditure on tokens



Decision-level distribution of expenditure on tokens given to personA
as a fraction of total expenditure on tokens given to others



Econometric specification

Suppose that w_S and v_S are members of the CES family:

$$w_S(\pi_O) = [\alpha' (\pi_A)^{\rho'} + (1 - \alpha')(\pi_B)^{\rho'}]^{1/\rho'}$$

and

$$v_S(\pi_S, w_S) = [\alpha (\pi_S)^\rho + (1 - \alpha) [w_S(\pi_O)]^\rho]^{1/\rho}$$

A family of CES functions that embed preferences for altruism and social preferences in a particularly convenient manner

$$U_S = [\alpha(\pi_S)^\rho + (1 - \alpha)[\alpha'(\pi_A)^{\rho'} + (1 - \alpha')(\pi_B)^{\rho'}]^{1/\rho}$$

The solution to the subutility maximization problem is given by

$$\pi_A(p_O, m_O) = \left[\frac{g'}{(p_B/p_A)^{r'} + g'} \right] \frac{m_O}{p_A}$$

where

$$r' = -\rho' / (1 - \rho'),$$

$$g' = [\alpha' / (1 - \alpha')]^{1/(1-\rho')}$$

and $m_O = p_O \pi_O$ is the total expenditure on tokens given to *others*.

The solution to the macro utility maximization problem is then given by

$$\pi_S(p, m) = \left[\frac{g}{q^r + g} \right] \frac{m}{p_S}$$


where

$$r = -\rho / (1 - \rho),$$


$$g = [\alpha / (1 - \alpha)]^{1/(1-\rho)}$$

and q is a *weighted relative price of giving*.


The distribution of the subutility CES parameter ρ'




Scatterplot of the CES estimates ρ and α in the three- and two-person experiments



The distribution of the CES parameter ρ in the three- and two-person experiments



Scatterplot of the CES estimates ρ and ρ'



Moral preferences

Harsanyi and Rawls argue for theories of *social justice* (equivalently, *fairness*) based on the choices that agents would make for society in the *original position*, behind a *veil of ignorance*.

... without knowing their own social and economic positions, their own special interests in the society, or even their own personal talents and abilities (or their lack of them). – Harsanyi (1975) –

Harsanyi and Rawls come to quite different conclusions, not because they view the original position differently, but because they treat uncertainty quite differently (Rawls denies orthodox decision theory).

- Both social preferences and preferences for altruism are highly heterogeneous, ranging from utilitarian to Rawlsian.
- In spite of this heterogeneity across subjects, there exists a strong positive within-subject correlation.
- A strong correlation between the equality-efficiency tradeoffs subjects make in their altruistic and social preferences.

Harsanyi's (1953, 1955) model for moral value judgments

Suppose an agent wants to make a moral value judgment about the relative merits of two alternative social systems.

... act in such a way as if he assigned the same probability to his occupying each social position under either system...

... then, he would clearly satisfy the impartiality and impersonality requirements to the fullest possible degree. – Harsanyi (1978) –

The agent has two different sets of preferences: *personal preferences* and *moral preferences* (preferences in the original position).

Two observations

- [1] Both Harsanyi and Rawls insist that moral preferences must conform to certain rationality requirements, and hence must have a special form – as opposed to personal preferences, which merely reflect taste.

- [2] Harsanyi and Rawls – and many other writers – view the original position as a purely hypothetical environment, and hence view moral preferences as a purely intellectual construct.

Our point of departure from the work of Harsanyi and Rawls – and the enormous literature they spawned – comes from two observations:

- [1] Choice behavior/preferences *behind* the veil of ignorance can be decomposed into choice behavior/preferences *in front of* the veil of ignorance:
 - choices that involve only personal consumption under uncertainty
 - and choices that involve social consumption – but no uncertainty.
- [2] Choices behind the veil of ignorance *can* be presented – and choices in the other two environments as well – in a controlled *laboratory* setting.

⇒ The linkage between preferences behind and in front of the veil of ignorance provides new ways of interpreting the theory of justice:

not just as a *normative* theory, but also as a *descriptive* theory and even as a *prescriptive* theory.

⇐ This linkage means that moral preferences *cannot* occupy such a privileged position – modulo certain assumptions, they are completely determined by risk preferences and social preferences.

Template for analysis

- Consider choice behavior by a single agent in each of three environments.
- Each choice has consequences for *self* (the agent) and for an (unknown) *other*.
- We consider only environments that involve binary choices and equiprobable lotteries.
- The results extend to more general choices and lotteries, and to unknown probabilities as well.

Consider lotteries over outcomes $[a, b]$, where a is consumption for *self* and b is consumption for *other*.

For our purposes, it suffices to consider binary lotteries with equal probabilities:

$$(.5)[a, b] + (.5)[c, d]$$

where $a, b, c, d \geq 0$. Write \mathcal{L} for the space of all such lotteries, and identify \mathcal{L} with the convex cone \mathbb{R}_+^4 .

Define closed convex subcones of \mathcal{L} :

$$\mathcal{R} = \{(.5)[a, 0] + (.5)[c, 0]\},$$

$$\mathcal{S} = \{(.5)[a, b] + (.5)[a, b]\},$$

$$\mathcal{M} = \{(.5)[a, b] + (.5)[b, a]\}.$$

We can interpret choice in each of the environments as choice in one of the corresponding cones by making an obvious identification:

– Risk: identify \mathbb{R}_+^2 with \mathcal{R} by

$$(x, y) \mapsto (.5)[x, 0] + (.5)[y, 0].$$

– Social: identify \mathbb{R}_+^2 with \mathcal{S} by

$$(x, y) \mapsto (.5)[x, y] + (.5)[x, y].$$

– Moral: identify \mathbb{R}_+^2 with \mathcal{M} by

$$(x, y) \mapsto (.5)[x, y] + (.5)[y, x],$$

which coincides *exactly* with Harsanyi's (1953, 1955) formalization of the original position.

Research questions

- [1] What is the relationship between moral preferences and personal/social (altruistic) preferences?
- [2] How can behavior behind [Harsanyi's] veil of ignorance be characterized experimentally?
- [3] Is behavior behind a veil of ignorance consistent with the utility maximization model?
- [4] Can the underlying moral preferences be recovered from observed choices?

Assumptions

Given a preference relation \succeq on \mathcal{L} , write $\succeq_{\mathcal{R}}$, $\succeq_{\mathcal{S}}$, $\succeq_{\mathcal{M}}$ for its restrictions to \mathcal{R} , \mathcal{S} , \mathcal{M} , respectively.

[i] \succeq satisfies the usual requirements: completeness, transitivity, reflexivity, continuity, and the Sure Thing Principle.

[ii] \succeq satisfies (weak) *independence*:

$$\begin{aligned} [a, b] \succeq_{\mathcal{S}} [a', b'] \text{ and } [c, d] \succeq_{\mathcal{S}} [c', d'] \\ \Rightarrow (.5)[a, b] + (.5)[c, d] \succeq (.5)[a', b'] + (.5)[c', d'] \end{aligned}$$

(*not* the usual independence axiom and does not have the usual consequences).

Next, we make two assumptions about *social* preferences:

[iii] **Worst outcome:** $[a, b] \succeq_{\mathcal{S}} [0, 0]$ for every $[a, b] \in \mathcal{S}$.

[iv] **Self-regarding:** for each outcome $[a, b]$ there is an outcome $[s, 0]$ such that $[s, 0] \succeq_{\mathcal{S}} [a, b]$.

[i] and [ii] are rationality requirements (should not necessarily be given any philosophical interpretation).

[iii] and [iv] limit the extent to which the *self* is (respectively) spiteful or altruistic toward *other*; they seem very natural requirements but they are not entirely innocuous.

Result: Every preference relation \succeq on \mathcal{L} that satisfies $[i]$ - $[iv]$ is determined by its restrictions $\succeq_{\mathcal{R}}$ and $\succeq_{\mathcal{S}}$.

Proof: Fix an outcome $[x, y]$. Because $\succeq_{\mathcal{S}}$ is self-regarding, there is some s such that $[s, 0] \succeq_{\mathcal{S}} [x, y]$.

Define the *selfish equivalent* of $[x, y]$ by

$$\sigma[x, y] = \inf\{s : [s, 0] \succeq_{\mathcal{S}} [x, y]\}.$$

Continuity and worse outcome guarantee that $[\sigma[x, y], 0] \sim_{\mathcal{S}} [x, y]$, and by construction,

$$[a, b] \sim_{\mathcal{S}} [\sigma[a, b], 0] \text{ and } [c, d] \sim_{\mathcal{S}} [\sigma[c, d], 0].$$

independence guarantees that

$$(.5)[a, b] + (.5)[c, d] \sim (.5)[\sigma[a, b], 0] + (.5)[\sigma[c, d], 0].$$

Hence

$$\begin{aligned} (.5)[a, b] + (.5)[c, d] &\succeq (.5)[a', b'] + (.5)[c', d'] \\ &\iff \\ (.5)[\sigma[a, b], 0] + (.5)[\sigma[c, d], 0] &\succeq \mathcal{R}(.5)[\sigma[a', b'], 0] + (.5)[\sigma[c', d'], 0] \end{aligned}$$

which decomposes preferences over \mathcal{L} into preferences over \mathcal{S} (selfish equivalents) and preferences over \mathcal{R} , as desired.

Given a linear budget constraint, we identify choice behavior in the Social Choice environment as

- selfish if the choice subject to every budget constraint is of the form $[y, 0]$ – giving nothing to *other*.
- symmetric if (a, b) is chosen subject to $px + qy \leq w$ iff (b, a) is chosen subject to the mirror-image budget constraint $qx + py \leq w$.

Corollary I: If the preference relation \succeq satisfies [i] and [ii] and choice behavior in the \mathcal{S} is selfish then choice behavior in \mathcal{R} coincides with choice behavior in \mathcal{M} .

Proof: Monotonicity and continuity guarantee that purely selfish behavior implies that $[x, 0] \sim_{\mathcal{S}} [x, y]$ for every x, y . independence implies that

$$(.5)[y, 0] + (.5)[x, 0] \sim (.5)[x, y] + (.5)[y, x].$$

It follows immediately that $\succeq_{\mathcal{R}}$ and $\succeq_{\mathcal{M}}$ coincide from whence choices in the Risk and Veil of Ignorance environments coincide, as asserted.

Corollary II: If the preference relation \succeq satisfies $[i]$ and $[ii]$ and choice behavior in \mathcal{S} is symmetric, then choice behavior in \mathcal{S} coincides with choice behavior in \mathcal{M} .

Proof: Suppose that (a, b) is chosen from some budget set B for the Social Choice environment, so that (b, a) is chosen in the mirror image budget set B' .

Say that (c, d) is chosen from the budget set B for the Veil of Ignorance environment, and that $(c, d) \neq (a, b)$.

Because $(c, d) \in B$, it follows that

$$(.5)[c, d] + (.5)[d, c] \succ_{\mathcal{M}} (.5)[a, b] + (.5)[b, a].$$

independence implies that

$$[c, d] \succ_{\mathcal{S}} [a, b] \text{ or } [d, c] \succ_{\mathcal{S}} [b, a],$$

which is inconsistent with the fact that (a, b) (resp. (b, a)) is chosen from the budget set B (resp. B').

It follows that risk attitude is irrelevant in the Veil of Ignorance environment, as asserted.


Experimental analysis

- Subjects in the experiments were recruited from all classes at UCLA and Yale Law School.
- Each decision problem is presented as a choice from a two-dimensional budget line.
- A choice (x, y) from the budget line represents an allocation between accounts x, y (corresponding to the horizontal and vertical axes).
- Choices are made through a simple point-and-click design using a graphical computer interface.

The actual payoffs of a particular choice in a particular environment/treatment are determined by the allocation to the x and y accounts:


- Risk: involves only pure risk; it is identical to the (symmetric) risk experiment of Choi, Fisman, Gale & Kariv (*AER*, 2007).
- Social Choice: involves only altruism; it is identical to the (linear) two-person dictator experiment of Fisman, Kariv & Markovits (*AER*, 2007).
- Veil of Ignorance: involves equiprobable binary lotteries over symmetric pairs of consumption for *self* and for *other*.

The distributions of token shares aggregated across subjects Social Choice




The tokens kept as a fraction of the sum of the tokens kept and given to other.

The distributions of token shares aggregated across subjects Risk




The fraction of tokens allocated to the cheaper account.

The distributions of token shares aggregated across subjects Veil of Ignorance




The fraction of tokens allocated to the cheaper account.

The distributions of token shares aggregated across subjects UCLA



Social: fraction of tokens kept by self. Risk and Veil: fraction of tokens allocated to the cheaper account.

The distributions of token shares aggregated across subjects Yale




Social: fraction of tokens kept by self. Risk and Veil: fraction of tokens allocated to the cheaper account.

Individual behavior


- The aggregate data tell us little about the choice behavior of individual subjects.
- Scatterplots of all choices of illustrative subjects – each entry plots $y/(x + y)$ as a function of $\log(p_x/p_y)$ in a particular treatment.
- There is no taxonomy that allows us to classify all subjects unambiguously.
- The characteristic of all our data is striking regularity *within* subjects and heterogeneity *across* subjects.

The relationship between the log-price ratio and the token share




X – Risk / X – Social Choice / X – Veil of Ignorance

ID62




X – Risk / X – Social Choice / X – Veil of Ignorance

ID102

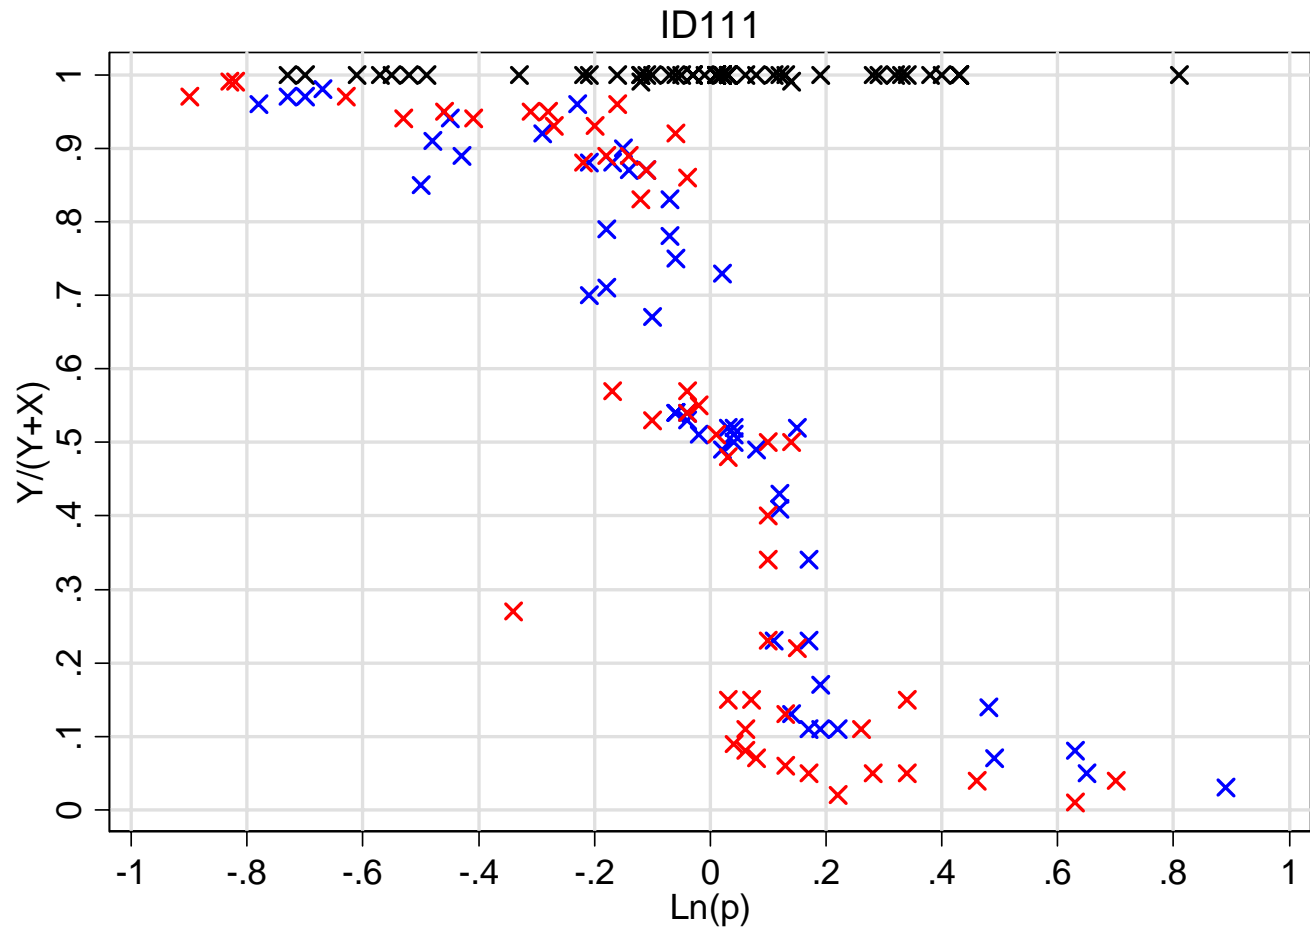


X – Risk / X – Social Choice / X – Veil of Ignorance

ID68




X – Risk / X – Social Choice / X – Veil of Ignorance




x – Risk / x – Social Choice / x – Veil of Ignorance

ID116




X – Risk / X – Social Choice / X – Veil of Ignorance

ID80




x – Risk / x – Social Choice / x – Veil of Ignorance

ID78




X – Risk / X – Social Choice / X – Veil of Ignorance

ID104




X – Risk / X – Social Choice / X – Veil of Ignorance

ID46



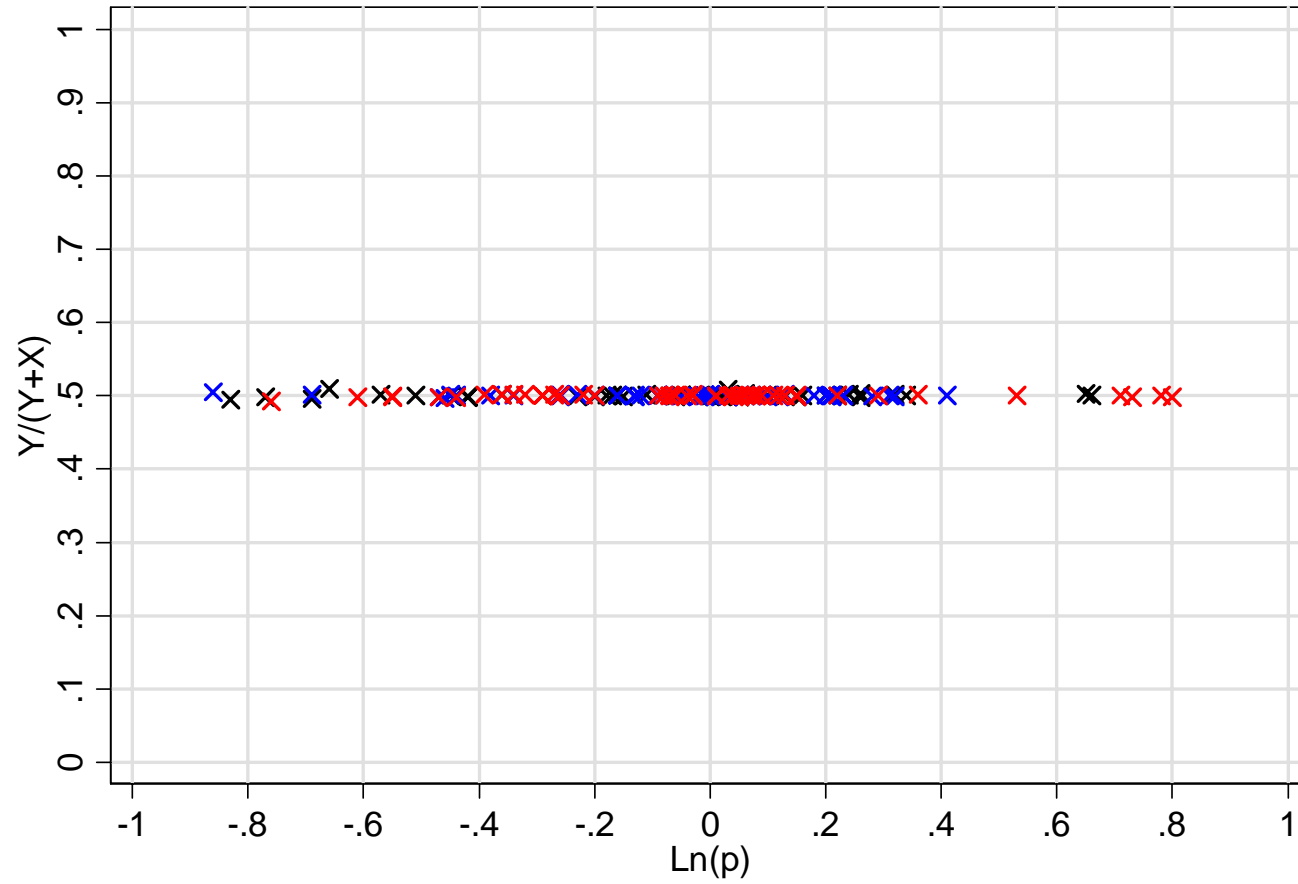
X – Risk / X – Social Choice / X – Veil of Ignorance

ID123




X – Risk / X – Social Choice / X – Veil of Ignorance

ID3




X – Risk / X – Social Choice / X – Veil of Ignorance

ID167




X – Risk / X – Social Choice / X – Veil of Ignorance




X – Risk / X – Social Choice / X – Veil of Ignorance

ID147



X – Risk / X – Social Choice / X – Veil of Ignorance



X – Risk / X – Social Choice / X – Veil of Ignorance

Testing the theory


- Many selfish subjects seem to display the same choice behaviors in the Risk and Veil of Ignorance environments, but a substantial number do not.
- Because of the nature of the data, “flexible” functional forms do not provide a plausible fit for the data.
- No satisfactory formulation to explain the “switching” between stylized behavior patterns exhibited by many subjects.
- Parametric approaches may be possible – keeping in mind that individual behaviors are extremely heterogeneous.

Non-parametric econometric approaches

Revealed preference

- The ratio of the CCEI score for the combined data set to the *minimum* of the CCEI scores for the separate data sets.
- A measure of the extent to which choice behaviors in any two environments coincide.
- Unfortunately, this test is weak – cannot discriminate between Risk and Veil of Ignorance behavior of selfish and non-selfish subjects.

The distributions of CCEI scores for the combined data set



Kolmogorov-Smirnov type tests

- A two-sample Kolmogorov-Smirnov tests of the equality of distributions of token and budget shares.
- The test is sensitive to differences in both location and shape of the empirical cumulative distribution functions of the two samples.
- Generalize the univariate Kolmogorov-Smirnov statistics for bivariate samples (Adler and Brown, 1986).

- There are subjects who fail Corollary I (selfish but display different behaviors in the \mathcal{R} and \mathcal{M}) and others who fail Corollary II.
- These subjects might have preferences over \mathcal{L} that do not obey independence (or might not be consistent with utility maximization).
- Individual preferences are very heterogeneous, ranging from utilitarian to Rawlsian.
- Actual preferences “mix-and-match” behavior in ways that no extant theory would regard as justified.

Takeaways

- A positive account of preferences for both personal and social consumption in rich choice environments.
- Two methodological contributions:
 - The establishment of theoretical links between preferences in various environments.
 - An experimental technique that allows for the collection of richer data about preferences.
- The experimental platform and analytical techniques are applicable to many other types of individual choice problems.