BEHAVIORAL GAME THEORY

Heinrich H. Nax & Bary S. R. Pradelski

heinrich.nax@uzh.ch bpradelski@cnrs.fr

University of Zurich[UZH]

ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

cnrs Centre National de la Recherche Scientifique

Plan for this talk

1. What is behavioral game theory?
2. Modeling human behavior
3. Behavioral evidence
4. Two Fruitflies: UG and VCM

So, what is behavioral game theory?

Strategic interactions/ game theory

> "... the study of mathematical models of conflict and cooperation between *intelligent* *rational* decision-makers."

Myerson, *Game Theory*, 1991.

3 ingredients (J. v. Neumann 1928)

1. ***individuals*** $N = \{1, 2, ..., n\}$ agents/subjects/players
2. ***actions*** $s_i \in S_i$ strategies/decisions
3. ***utilities*** $u : \mathcal{S} \to \mathbb{R}^n$ payoffs/profit/outcomes

Game theory describes interactions

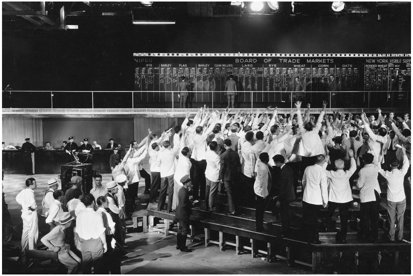

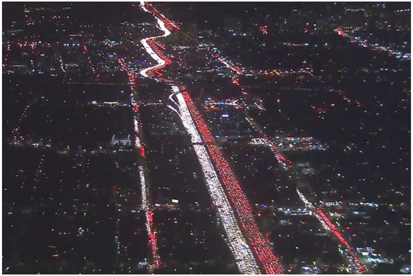

Markets:

- Individuals (traders)
- Strategies (buy/sell)
- Outcome (profit/loss)

wired.co.uk

Routing and congestion:

- Cars (drivers)
- Decisions (routes)
- Traffic (travel time)

valleyproirrigation.com

Game theory describes interactions

Games animals play:

- Individuals (honeybees)
- Strategies (foraging nectar)
- Outcome (survival)

beecare.bayer.com

Social dilemmas:

- Users (farmers)
- Actions (water usage)
- Result (depletion/profit)

thewrap.com

'The' solution concept

The 3 ingredients do not tell us *what* people do.

1. ***individuals*** N
2. ***actions*** s
3. ***utilities*** u

Nash Equilibrium (PhD, 1950)

A *Nash equilibrium* is a strategy profile s^* such that for every player i,

$$u_i(s_i^*, s_{-i}^*) \geq u_i(s_i, s_{-i}^*) \text{ for all } s_i.$$

- At s^*, no i can unilaterally improve by not playing s_i^*.
- Fixing all the other players at s_{-i}^*, then s_i^* is a best response for all i.

In other words,

Nash Equilibrium

is an outcome (= *strategy choice by each player*), where no individual can *unilaterally* improve his own position by changing his strategy.

- the best you can do against what your opponents do, when they also do their best against what [...]
- need not result in the best outcome for the individual
- nor indeed in the best collective outcome

Equilibrium predictions

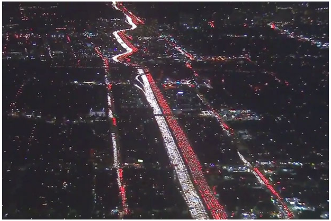

Simple congestion games:

- sharp pinpoint prediction
- same travel time
- not individual best (!taxi)

Social dilemmas:

- individual over-usage
- tragedy of the commons
- worst outcome

Equilibrium foundations from neoclassical economics

Rationality assumptions

- **A1. common knowledge:** game and payoffs
- **A2. correct beliefs:** about each other
- **A3. optimization:** maximization of expected utilities (satisfying Bayes)

Homo Oeconomicus

- **A4. narrow self-interest:** own material payoff only
  - **no concern** for others' payoffs
  - **no consideration** of one's effects on higher-order norms or similar

Clearly not a descriptive theory – in the real world,

~~A1. common knowledge~~ the exact game structure is typically **unknown**

~~A2. correct beliefs~~ humans usually do not know the exact motivations of others and are often **wrong** about each other

~~A3. optimization~~ the human brain is at most an **imperfect** constrained-optimization machine

~~A4. narrow self-interest~~ **social** motivations often include the welfare (positive and negative) of others

Outside of game theory (e.g. Kahnemann & Tversky 1979), *heuristics* have replaced these postulates, but most

- game theorists uphold A1-A4,
- *'behavioral'* game theorists uphold *A1-A3* *('subjective utility correction project')*.

What is more, most real-world interactions are dynamic and/or repeated

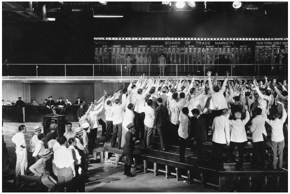

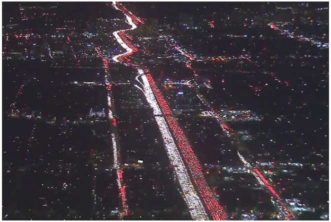

Toward a descriptive theory of learning in games

More than just incomplete information. *instead*
learning about the game, nature and the environment

No common knowledge. *instead*
learning about others, their actions and motives

Social norms and social motivations matter. *hence*
a theory of forming and adapting social motives as **norms are evolving** too

Depart fundamentally from utility maximization. *instead*
following **rules and heuristics**, as well as making reasoned choices

..from *prescriptive* to *descriptive*..

Learning in games: theory and *evidence*

An experimental agenda for behavioral game theory

Imagine the Voluntary Contributions Game – You!

1. **Voluntary contributions:** Decide on a contribution c from $0, 1, 2, 3, 4$ – whatever you do not contribute is yours automatically
2. **Public good:** Let C be the total contributions: $C = \sum c$. A *public good* is created from C worth $3 \cdot C$
   - Each of you will enjoy an equal share of the public good: $\frac{3}{n}C$
3. **Payoffs:** So you will earn a total of $4 - c + \frac{3}{n}C$

Evidence from the lab

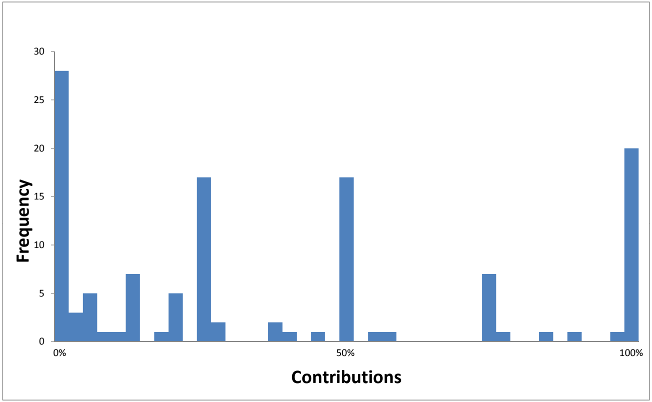

from Nax, Burton-Chellew, West and Young (JEBO 2016)

Key characteristics

- **Social optimum:** all everything (total payoffs of $12n$)
- **Social worst-case:** all nothing (total payoffs of $4n$)
- **Individual best:** others everything, you nothing
- **Individual worst:** others nothing, you everything
- **Dominant strategy (if selfish):** nothing

Nash Equilibrium: *all nothing* $\leftarrow$ **Social dilemma**

Behavioral evidence

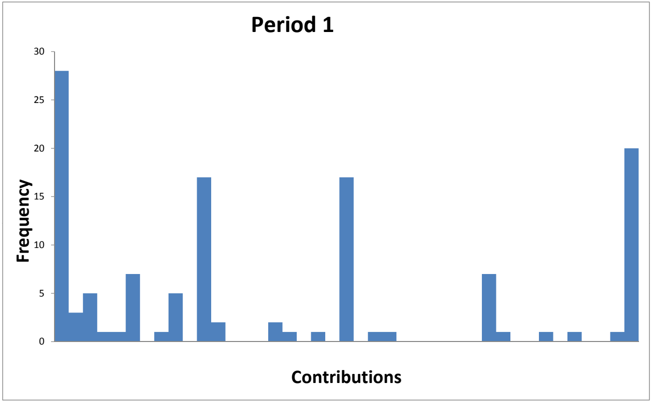

Contribution dynamics

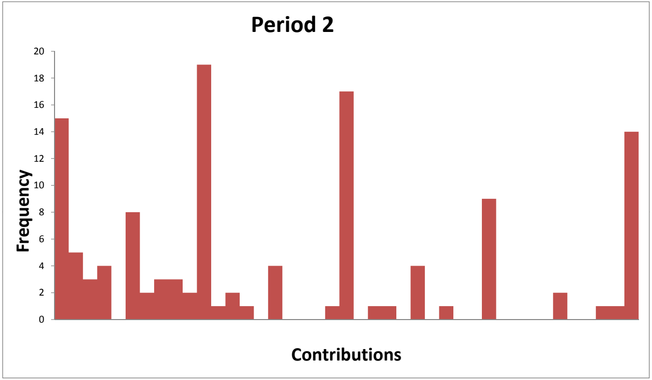

Decline

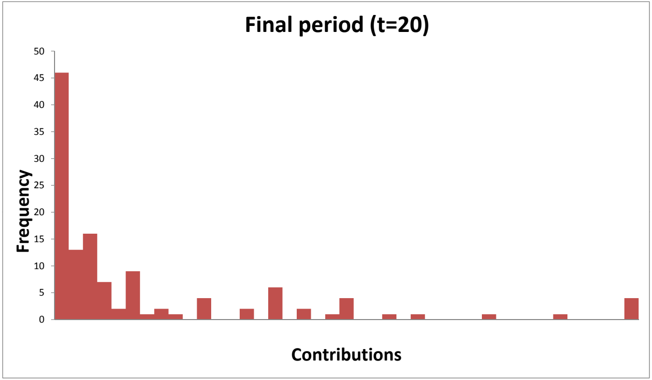
Final period (t=20)
Frequency
50
45
40
35
30
25
20
15
10
5
0
Contributions

The decline of cooperation

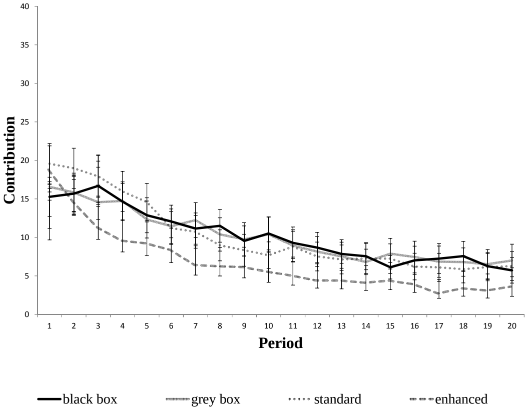

from Burton-Chellew, Nax and West (Proceedings B 2015)

Mixed real-world evidence

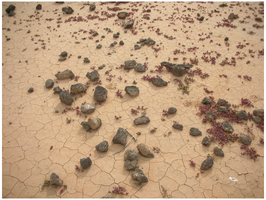

- **Bad:** groundwater depletion, over-fishing, over-forestry, climate change, littering, etc.
- large interactions, institutions

+ **Good:** sustainable resource usage, crowd-sourcing, crowd-funding, club goods, etc.
- local interactions, institutions

Interpreting behavioral evidence: 2 approaches

Digging into the micro-motives of the cooperation decline

Approach 1: (~~*A4. narrow self-interest*~~)

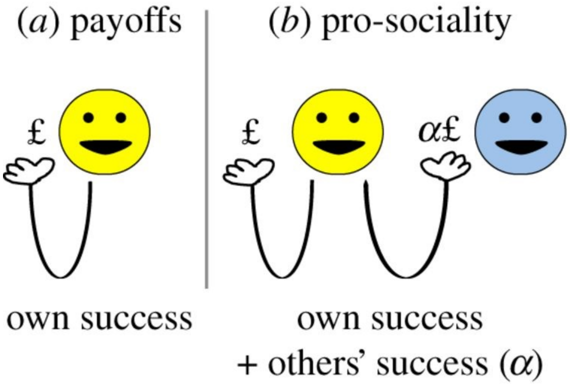

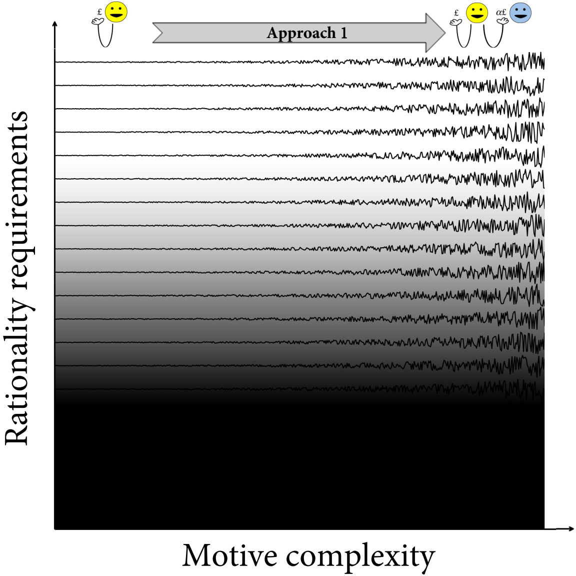
£
Approach 1
£
α£
Rationality requirements
Motive complexity

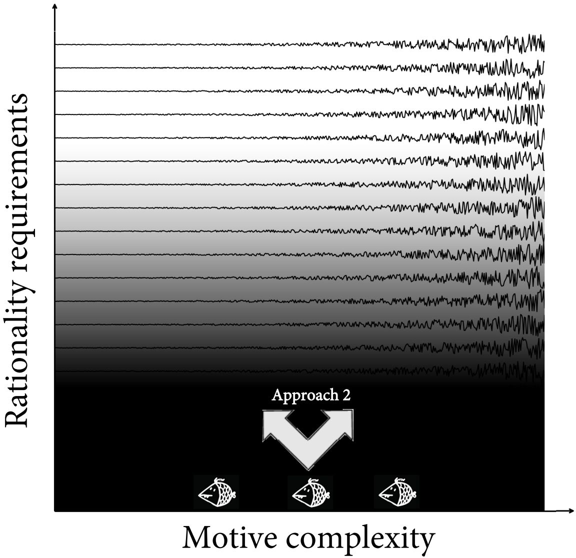

Rationality requirements
Approach 2
Motive complexity

Rule rationale

Approach 2: (~~*A1-A3. neoclassical rationality*~~)

- **near** with success
- **far** with failure
- **directed** with success

Black Box learning. Based on the *Law of Effect* (*à* la Thorndike, Pavlov, Skinner), this form of *trial-and-error* is a well-known description of (animal) behavior, implementing welfare-maximal Nash equilibria in games.

2 different approaches (summary)

- **Approach 1:** presence of *reasoned motives* and **reciprocity**
  - conditional cooperation
  - interactive preferences
  - **heterogeneity** in motives and reciprocity

- **Approach 2:** *learning* dependent on information
  - following the **law of effect**
  - agent homogeneity
  - present also in complex and high-information settings

“zoon logicon"

Aristotle, *Metaphysics*, ca. 330 B.C.

Today, at hand of a 2 experimental game examples (2 fruit-flies of experimental game theory), I want to motivate a view that *rules* and *heuristics* are crucial elements of individual decision-making, that they are used intuitively and deliberately, and that one can formulate a behavioral theory of game play.

rational animal & rational **animal**

Some historical background: Experiments in Economics

Experiments on decision problems/risk/1-player games:

Allais 1953, Ellsberg 1961, Ainslie 1975, Kahneman and Tversky 1979: experiments that challenge the axioms of standard decision theory and with it the notion of man as a "perfectly rational" expected utility maximizer (Ramsey 1931, von Neumann and Morgenstern 1944, Savage 1954)

Recall Lecture 3 on utilities: The clean "theory of expected utility maximization" (Ramsey-Savage-von Neumann) contradicted by simple experiments such as those by Allais/ Ellsberg/ Kahneman-Tversky lead to **Behavioral Economics!**

Background: Experiments in Biology

Experiments on animal behavior:

Thorndike 1898, Morgan 1903, Pavlov 1927, Thorpe 1956: classic experiments that reveal that the "law of effect", i.e. a consequentialist view of trial and error, explains animal behavior (later formalized as "radical behaviorism"/"reinforcement learning" Skinner 1974, Hoppe 1931, Estes 1950, Bush and Mosteller 1955, Heckhausen 1955, Herrnstein 1961, Roth and Erev 1995, Erev and Roth 1998)

Follow the path of success/ avoid the path of failure.

Homo oeconomicus: “perfect rationality strawman”

Perfect rationality

- **common knowledge:** about the structure of the game, about the structure of payoffs
- **common beliefs:** players have beliefs about each others’ behavior, and these beliefs are correct
- **optimization:** individual behavior is governed/ described by optimization/ maximization in terms of expected utilities

Pure self-interest

- **narrow self interest:** agent cares about own material payoff only
- no concern for other players’ payoffs
- no consideration of the effects of his actions on upholding higher-order norms or similar
- decisions are not subject to social influence

Perhaps more **realistic** environments

Knowledge and information

- the game structure is often unknown, or at least large parts of it
- players may not be able to observe information about relevant players' in the game, may sometimes not know they even exist
- a player may know little about others' utility functions, about how he affects them and how they affect him
- i.e. info too low for neoclassical assumptions to make sense

Behavior and motivations

- instead of optimizing behavior, players may follow behavioral heuristics
- players may learn about the game and which strategies to play as the game goes on
- instead of narrow self interest, an agent may also care about others' payoffs and/ or the distribution of payoffs
- agents may follow social norms, and may be subject to social influence

Today's focus is on human behavior in 2 games: Drosophila

Ultimatum game

- one side **proposer** moves first: makes a proposal as to how to split a cake
- the other side **recipient** responds: either accepts the offer so that it will be realized, or destroys the cake (both get zero)
- **Nash equilibria:** any proposal made, responder accepts
- **Subgame perfection:** proposer takes all, accept nevertheless

Public goods game

- the game we just played
- contributions are socially valuable (increase total payoffs as $R > 1$)
- but each individual has an incentive to withhold his own contribution (free-ride as $R/n < 1$)
- **Nash equilibrium:** universal non-contribution

What is studied with these games?

Ultimatum game

- introduced to model negotiations by Gueth et al. (1982), Binmore et al. (1985) and Gueth and Tietz (1987), Ochs and Roth (1989)

A: Nash equilibrium (responder should always accept)

B: Subgame perfection (proposer gives nothing)

C: Reputation models (Kreps and Wilson 1982) in case of repetition

D: Social preferences such as fairness, pro-sociality, spitefulness

Public goods game

- introduced to model social dilemma situations by Bohm (1972, 1983), Dawes (1980), Isaac et al. (1985), Isaac and Walker (1988), Andreoni (1988)

A: Nash equilibrium

D: Social preferences such as fairness, pro-sociality, conditional cooperation, reciprocity

E: Mechanisms such as punishment, rewards, etc.

We can think of different **information settings** for these experiments

Ultimatum game

- **high information:** players know the structure of the game, know their own position in the game, know the payoff structure, the game is anonymous
  - proposer: moves first, knows who the responder is / how he is selected
  - responder: moves second, observes the offer

Public goods game

- **high information:** players know the structure of the game, know their own position in the game, know the payoff structure, the game is anonymous
  - players decide how much to contribute
  - learn about others' decisions of past rounds as the game goes on

or less information

Ultimatum game

- **low information:** players do not know the payoff structure of the game, do not observe others' actions, learn only about payoffs as they realize
  - proposer: moves first picking a number between zero and everything, knows nothing about the nature of his "proposal"
  - responder: selects either option A ("accept") or option B ("reject") without knowing their significance

Public goods game

- **low information:** players do not know the payoff structure of the game, do not observe others' actions, learn only about payoffs as they realize
  - players decide how much to enter into a "black box"
  - players learn about the payoff consequences of their own actions only, receive no information about others

Motivation for experimental game theory:

A large body of economic theory presumes rather extreme behaviors in terms of

- rationality
- optimization
- strategizing

What do real humans do?

Experiments

The “clean” equilibrium predictions based on the theories of von Neumann-Nash contradicted by simple experiments such as the ones we will talk about today (Ultimatum Games/ Voluntary Contribution Games).
These experiments lead to **Experimental/Behavioral Game Theory.** (Zurich being one of ‘the’ places in the world where this line of research is pursued.)

Recall our two games

Ultimatum Game:

- One player offers a share of a pie, then the other accepts or rejects.

Voluntary Contributions Game:

- Players simultaneously decide how much to contribute to a joint effort that creates a public good.

Game 1: Ultimatum Game

- Gueth et al. (1982), one-shot
- Rubinstein (1985), multiple rounds
- Review: "Thirty years of UG" (Gueth and Kocher 2013)

THE GAME

1. the proposer (player 1) suggests a split between him and the receiver (player 2)
2. Player 2 can either accept or reject:
    1. If he accepts, the shares proposed by player 1 realize
    2. If he rejects, both players receive nothing.

- **Nash equilibria:** any split supportable as a Nash equilibrium
- Unique **subgame-perfect Nash equilibrium**: (1 all, 2 nothing)

Testing the extreme SPNE prediction

- The unique subgame-perfect Nash equilibrium is an extreme allocation
- Any rejection by the responder kills own and other's payoff
- Any positive proposal, presuming (rational) acceptance, seems like a gift;
- however, presuming (off the equilibrium-path) rejection of low offers, a substantial proposal may be strategically rational
- hence, it may be rational to have a **rejection reputation**
- Meta-analysis suggests
  - proposals of roughly 40%;
  - high rejection rates for proposals under 20%, intermediate rejection rates for proposals of 20%-40%, and almost zero rejection rates for proposals >40%
  - Over time, decline or no decline of proposals depending on experimental/matching protocol

Acceptance rates

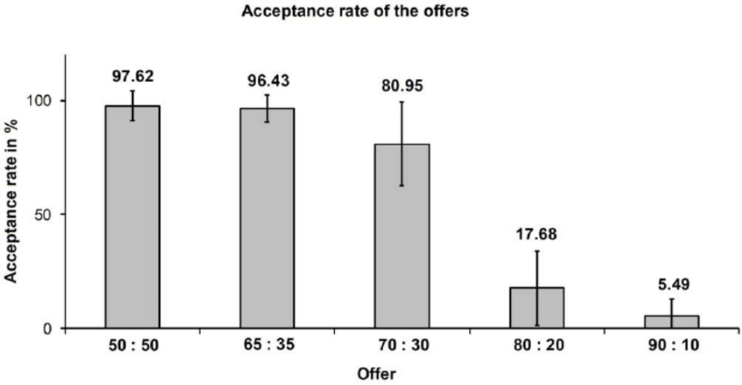

from Hollmann et al., PLoS ONE 2011

Offers

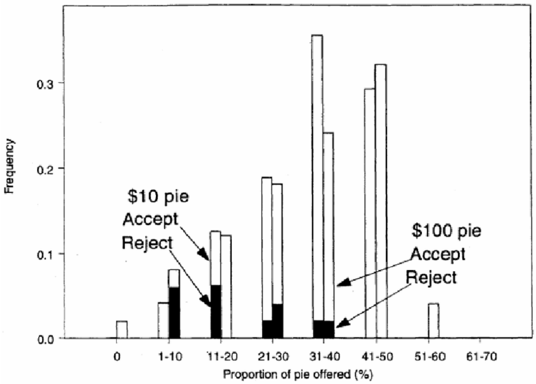

from Hoffman et al., IJGT 1996

Game 2: Voluntary Contributions Game

- Marwell and Ames (1979), one-shot
- Andreoni (1988): random (re-)matching
- **Review:** "Sustaining cooperation in laboratory public goods experiments" (Chaudhuri 2011) - older review by Ledyard (1995)

THE GAME

1. the game we played
   $\phi_i(c) = (B - c_i) + \sum_{j \in N} mpcr * c_j$
2. **Unique Nash equilibrium** if agents are selfish: all give nothing.

Characteristics of the NE

- Again, the Nash equilibrium is an extreme allocation
- Lowest social welfare
- Pareto-dominated by social optimum
- Any positive contribution decreases own payoff but increases those of others and increases total welfare
- Meta-analysis suggests
    - average contributions of roughly 40%-50% when game is played once or in the first round when repeated;
    - when repeated (with random re-matching w/o any mechanism): over time, contributions roughly halve every 10-20 periods depending on matching protocol

Contributions 1

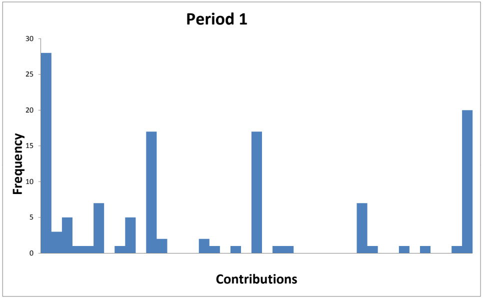

from Nax et al., JEBO 2016

Contributions 2

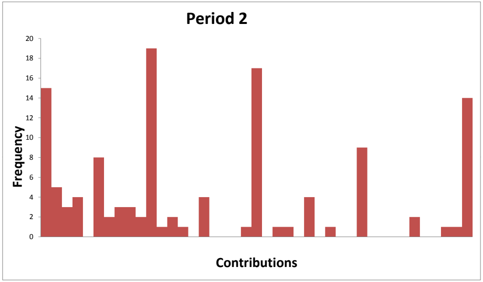

from Nax et al., JEBO 2016

Contributions final

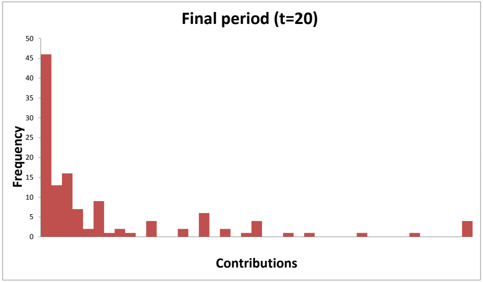

from Nax et al., JEBO 2016

Contributions trend

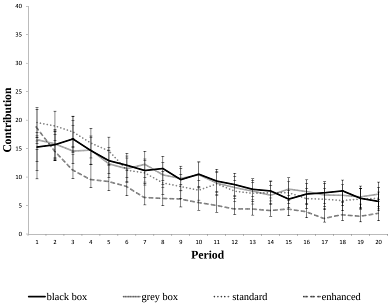

from Nax et al., JEBO 2016

Interpretation 1: The ‘subjective utility correction project’

- The failure to play according to Nash equilibrium as predicted by pure self-interest is explained using alternative payoff functions that include social preferences and concerns for other players’ payoffs such as
  - Fairness considerations (Fehr-Schmidt)
  - Inequality/inequity aversion (Bolton-Ockenfels)
  - Altruism (Fehr-Gachter, Gintis-Bowles-Boyd-Fehr, Fehr-Fischbacher)
  - Reciprocity (Fischbacher-Gachter-Fehr)
- Note: This approach (by the Zurich school) mirrors the various “corrections” to utility functions motivated by ambiguity aversion, etc.

Homo Oeconomicus and Friends

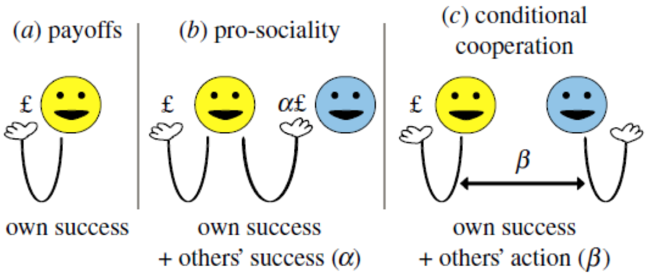

Rational choice theory assumes individuals to be fully rational and thus capable of expressing their preferences perfectly through the consequences of their actions (Becker 1976).

What would someone according to the ‘subjective utility correction project’ do in the voluntary contributions game?

- In the one-shot game and in the final period of a repeated game, he would contribute zero.
- However, if his utility contains a concern for the other player, and is, for example, Cobb-Douglas of the form $u_i(c) = (\phi_i^{1-\alpha_i} * \phi_{-i}^{\alpha_i})$,
- where ϕ^{α}_{-i} is the average payoff to players $j \neq i$, then...

...we have a range of social personas...

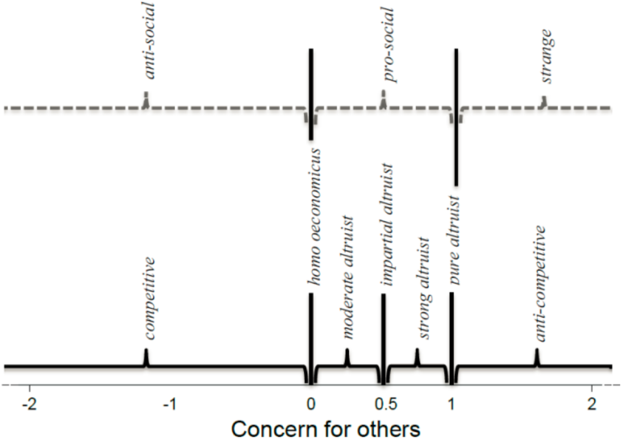

And positive contributions are evidence of concerns for others in this range:

- (0,0.5) $\longrightarrow$ moderate altruist
- 0.5 $\longrightarrow$ impartial altruist
- (0.5,1) $\longrightarrow$ strong altruist
- 1 $\longrightarrow$ pure altruist

...and in the final period we have...

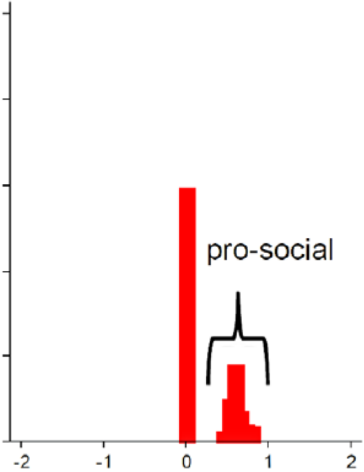

Interpretation 2: Mistakes equilibrium

The failure to play according to Nash equilibrium as predicted by pure self-interest is explained by relaxing the rationality assumption. Examples of such models include

- "Noise"/ QRE (Palfrey-Prisbey)
- Intuitive versus contemplative players (Rubinstein)

According to such a model, positive contributions are evidence of "less" or bounded rationality.

Deviations

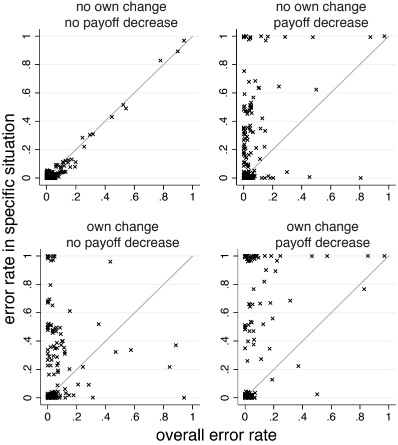

Players best respond but deviate (Maes and Nax, JET 2016)

Interpretation 3: Learning

The failure to play according to Nash equilibrium as predicted by pure self-interest is explained by adaptive processes of learning to play the game. Examples of such models include

- Reinforcement learning (Roth-Erev)
- Directional learning (Selten)
- Perturbed best reply (Young)
- Belief-based learning (Fudenberg-Levine)
- EWA (Camerer-Ho)

Is there a way to tell what is what?

- Can we distinguish between motivations?
- How much can we attribute to which explanation?

Experiments: Set-Up

- Experiments were conducted @ CESS Nuffield of University of Oxford (involving 236 subjects in 16 sessions)
- In each session, 16 players played four of our games
- The *mpcr* was 0.4 or 1.6
- The budget was 40 coins each round
- Each game was repeated for 20 rounds
- Players received instructions containing different amounts of information about the game and sometimes (anonymous) feedback about previous-period play
- Play was incentivized with real money (e.g. one coin=0.01 CHF)

Consistent deviations from homo oeconomicus?

- By design of the experiment, games differed with respect to whether contributing zero was a strictly dominant strategy
- In half of the games, contributing everything was a strictly dominant HOE strategy (e.g. by setting the $mpcr = 1.6 > 1$)
- In half of the games, contributing nothing was a strictly dominant HOE strategy (e.g. by setting the $mpcr = 0.4 < 1$)

Contributions (final round)

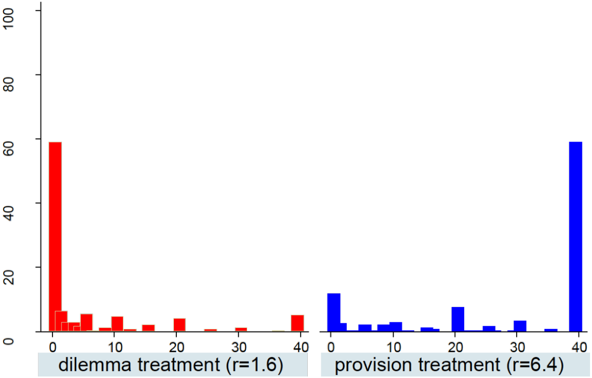

Implied Preferences (final round)

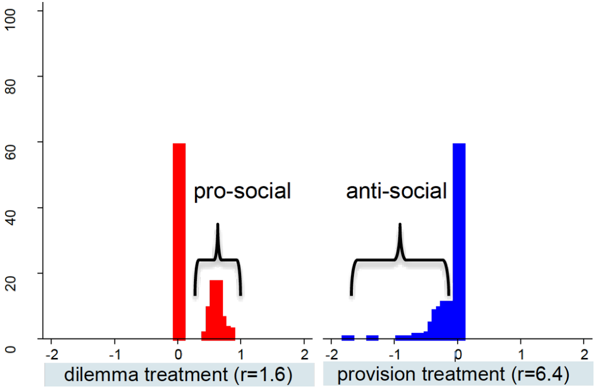

Combined Preferences

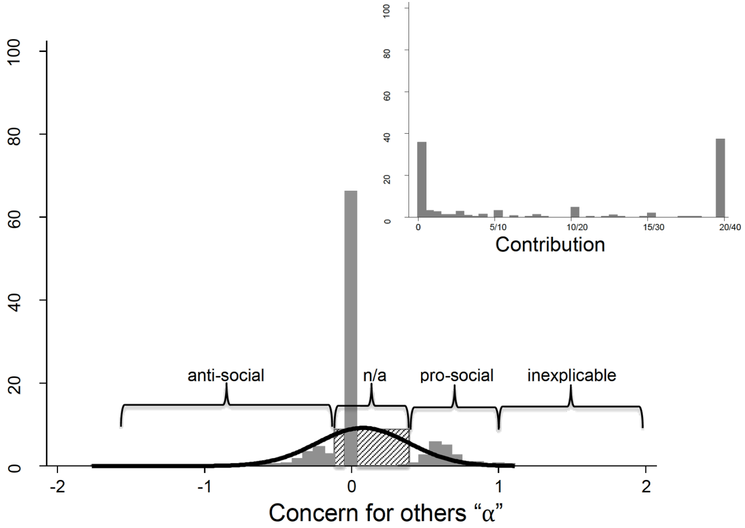

Summary

In total, there therefore are

- 46.7% players consistent with homo oeconomicus.
- 15.4% are consistent and anti-social.
- 21.4% are consistent and pro-social.
- 16.5% are inconsistent, meaning pro-social in one and anti-social in the other — mistakes
- The median is neutral, the mean close to neutral.
- Note that inconsistent players in terms of social preferences may by consistent in terms of ‘erroneous play’

Dynamics: the role of learning and conditional cooperation

- Games differed with respect to the amount of information about the structure of the game, and about other players' past actions and payoffs,
- allowing us to look into the question whether and how players react to what others do and how they learn from experience.

Two types of information

Black box

- Players do not know the structure of the game
- Players learn nothing about other players' actions or payoffs
- Players know their own history of actions and payoffs only

Standard (enhanced)

- Players know the structure of the game
- Players learn what others did in the past as the game is repeated
- (Players are explicitly told what payoffs others got)

Patterns in different treatments

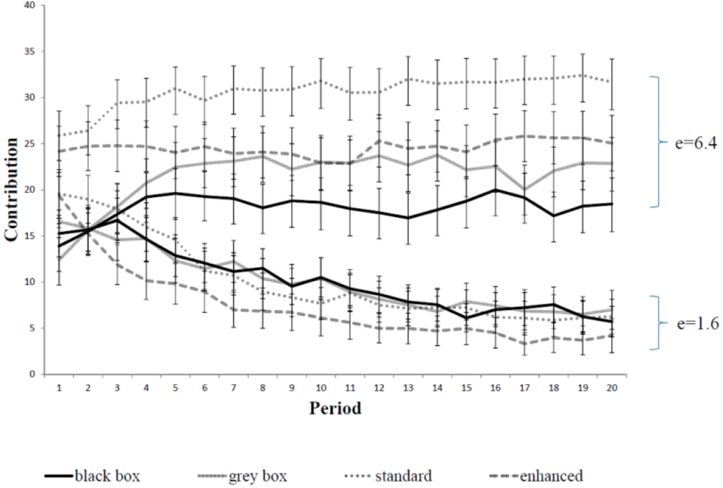

Learning: a simple model

Suppose players initially make random contributions. Thereafter,

- they follow the direction of payoff increases
- they avoid the direction of payoff decreases

Notice such a learning rule is completely uncoupled (Foster and Young 2006) from information about others' actions and payoffs, relying only on own realized payoffs.

Conditional cooperation

Suppose players contribute/free-ride if others do too (Fischbacher et al, EL 2001).

- the increase their contributions if others increase their contributions
- they decrease their contributions if others decrease their contributions

Notice such a learning rule is uncoupled (Hart and Mas-Colell 2003) from information about others' payoffs, relying only on own realized payoffs and others' actions.

Evidence of conditional cooperation in standard treatment

	black box	standard	enhanced
payoff-based learning	✓ 0.30*	✓ 0.25*	✓ 0.14*
pro-social learning[a]	✗ −0.13*	✗ −0.23*	✗ −0.29*
conditional cooperation[a]	n.s. 0.05	✓ 0.21*	n.s. −0.001

*significance < 0.001.
[a]Controlling for payoff-based learning.

A richer Black box learning model

Suppose players initially make random contributions. Thereafter, adjustments follow four regularities:

1. **Asymmetric inertia:** stay with your current strategy more often after success than after failure
2. **Search volatility:** search for new strategies more randomly after failure than after success
3. **Search breadth:** search for new strategies further away after failure than after success
4. **Directional bias:** follow the direction of payoff increases, and avoid the direction of payoff decreases

Does this remind you of something?

Figure 7: SEARCH (all treatments, both rates of return).

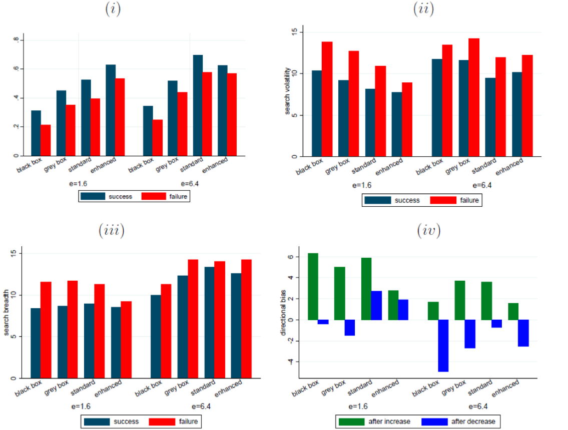

The bar charts summarize the four search components in all treatments for both rates of return separately. The respective panels are: (i) inertia, (ii) search volatility, (iii) search breadth, (iv) directional bias. For (ii) – (iv), the y-axis are units of contributions; for (i), the y-axis are probabilities.

Summary: theoretical game theory versus reality

Mainstream game theory relies on rather extreme assumptions such as

- complete information,
- common knowledge,
- unbounded rationality, and
- optimizing behavior.

In many real-world situations, these assumptions are untenable because

- the game may be too complex,
- behavior of others may be unobservable,
- players may not know others' utility functions, and
- the structure of the game may be unknown.

In addition, real-world humans care about others, and follow certain rules/norms.

The economic laboratory promises some answers

- Play often does not coincide with the Nash equilibrium predictions.
- There are robust deviations from predictions, and many experiments have made similar observations.
- To explain these deviations, we must
  - abandon the assumption of narrow self-interest in favor of social preferences

  and/or
  - abandon the assumption of strictly optimizing behavior in favor of behavior that allows for heuristics/learning

Learning

- Over time, play approaches equilibrium in most settings, including those where very limited information is available.
- There is a rich theoretical literature on these convergence properties, but relatively little of it has been tested in the laboratory.
- And there is a lack of acknowledgement in experimental research of the fact that simple heuristics may explain behavior not only in low-information but also in richer information environments.
- There is plenty of room for innovative experimental-theoretical work in this area.

Some concluding remarks

- Aristotle called man a “rational animal” (“zoon logikon” or “zoon logon echon”)
- There is a side to human nature which is rational, describable by (corrected) utility maximization
- Utility may include components concerning others’ material payoffs too
- There is also a side not describable that way – but instead by heuristics and by learning models
- It is my belief that such ‘rules’ may themselves be more rational than is usually considered

References (some own work)

- M Burton-Chellew, HH Nax & S West, “Payoff-based learning explains the decline in cooperation in public goods games”, Proceedings of the Royal Society of London B, 2015
- HH Nax, M Burton-Chellew, S West & HP Young, “Learning in a black box”, JEBO 2016
- M Maes & HH Nax, “A behavioral study of ‘noise’ in coordination games”, JET 2016

THANKS EVERYBODY

Keep checking the website for new materials as we progress:

http://gametheory.online/project_show/9