Bayesian Games

You are Player A in the following game. What should you do?

		Player B: S_1	Player B: S_2
Player A	S_1	3 ?	-2 ?
Player A	S_2	0 ?	6 ?

Question: When does this situation arise?

Recipe for Nash-Equilibrium-Based Analysis of Such Games

- Assume you've been given a problem where the *i*'th player chooses a real number x_i
- Guess the existence of a Nash equilibrium

$$(x_1^*, x_2^* \cdots x_n^*)$$

- Note that , $\forall i$,

$$x_i^* = \arg\max_{x_i} \begin{bmatrix} \text{Payoff to player } i \text{ if player } i \\ \text{plays "} x_i \text{" and the } j\text{'th player} \\ \text{plays } x_j^* \text{ for } j \neq i \end{bmatrix}$$

- Hack the algebra, often using "at x_i^* we have

Hockey lovers get 2 units for watching hockey, and 1 unit for watching football.

Football lovers get 2 units for watching football, and 1 unit for watching hockey.

Pat's a hockey lover.

Pat thinks Chris is probably a hockey lover also, but Pat is not sure.

Pat \ Chris	H	F
H	2 2	0 0
F	0 0	1 1

With 2/3 chance

Pat \ Chris	H	F
H	2 1	0 0
F	0 0	1 2

1/3 chance

In a Bayesian Game each player is given a type. All players know their own types but only a prob. dist. for their opponent's types

An n-player Bayesian Game has

a set of action spaces $A_1 \cdot\cdot\; A_n$
a set of type spaces $T_1 \cdot\cdot\; T_n$
a set of beliefs $P_1 \cdot\cdot\; P_n$
a set of payoff functions $u_1 \cdot\cdot\; u_n$

$P_{-i}(t_{-i}|t_i)$ is the prob dist of the types for the other players, given player i has type i .

$u_i(a_1 , a_2 \cdots a_n , t_i)$ is the payout to player i if player j chooses action a_j (with $a_j \in A_j$) (forall $j=1,2,\cdots n$) and if player i has type $t_i \in T_i$

Bayesian Games: Who Knows What?

We assume that all players enter knowing the full information about the A_i's , T_i's, P_i's and u_i's

The *i*'th player knows t_i, but not $t_1\ t_2\ t_3 \cdots t_{i-1}\ t_{i+1} \cdots t_n$

All players know that all other players know the above

And they know that they know that they know, *ad infinitum*

Definition: A <u>strategy</u> $S_i(t_i)$ in a Bayesian Game is a mapping from T_i? A_i : a specification of what action would be taken for each type

Example

$A_1 = \{H,F\}$ $A_2 = \{H,F\}$

$T_1 = \{\text{H-love}, \text{Flove}\}$ $T_2 = \{\text{Hlove}, \text{Flove}\}$

$P_1\ (t_2 = \text{Hlove} \mid t_1 = \text{Hlove}) = 2/3$
$P_1\ (t_2 = \text{Flove} \mid t_1 = \text{Hlove}) = 1/3$
$P_1\ (t_2 = \text{Hlove} \mid t_1 = \text{Flove}) = 2/3$
$P_1\ (t_2 = \text{Flove} \mid t_1 = \text{Hlove}) = 1/3$
$P_2\ (t_1 = \text{Hlove} \mid t_2 = \text{Hlove}) = 1$
$P_2\ (t_1 = \text{Flove} \mid t_2 = \text{Hlove}) = 0$
$P_2\ (t_1 = \text{Hlove} \mid t_2 = \text{Flove}) = 1$
$P_2\ (t_1 = \text{Flove} \mid t_2 = \text{Hlove}) = 0$

$u_1\ (H,H,\text{Hlove}) = 2$ $u_2\ (H,H,\text{Hlove}) = 2$
$u_1\ (H,H,\text{Flove}) = 1$ $u_2\ (H,H,\text{Flove}) = 1$
$u_1\ (H,F,\text{Hlove}) = 0$ $u_2\ (H,F,\text{Hlove}) = 0$
$u_1\ (H,F,\text{Flove}) = 0$ $u_2\ (H,F,\text{Flove}) = 0$
$u_1\ (F,H,\text{Hlove}) = 0$ $u_2\ (F,H,\text{Hlove}) = 0$
$u_1\ (F,H,\text{Flove}) = 0$ $u_2\ (F,H,\text{Flove}) = 0$
$u_1\ (F,F,\text{Hlove}) = 1$ $u_2\ (F,F,\text{Hlove}) = 1$
$u_1\ (F,F,\text{Flove}) = 2$ $u_2\ (F,F,\text{Flove}) = 2$

Bayesian Nash Equilibrium

The set of strategies (s_1^* ,s_2^* $\cdots$ s_n^*) are a

Pure Strategy Bayesian Nash Equilibrium

iff for each player *i*, and for each possible type of *i* : $t_i \in T_i$

$s_i^*(t_i)$ =

$$\arg\max_{a_i \in A_i} \sum_{t_{-i} \in T_{-i}} u_i\left(s_1^*(t_1),\ldots s_{i-1}^*(t_{i-1}), a_i, s_{i+1}^*(t_{i+1})\ldots s_n^*(t_n)\right) \times \mathrm{P}_i\left(t_{-i} \middle| t_i\right)$$

i.e. no player, in any of their types, wants
to change their strategy

NEGOTIATION: A Bayesian Game

Two players:	S, (seller) and B, (buyer)
$T_s = [0,1]$	the seller's type is a real number between 0 and 1 specifying the value (in dollars) to them of the object they are selling
$T_b = [0,1]$	the buyer's type is also a real number. The value to the buyer.

Assume that at the start

$V_s \in T_s$ is chosen uniformly at random

$V_b \in T_b$ is chosen uniformly at random

The "Double Auction" Negotiation

S writes down a price for the item (g_s)

B simultaneously writes down a price (g_b)

Prices are revealed

If $g_s = g_b$ no trade occurs, both players have payoff 0

If $g_s = g_b$ then buyer pays the midpoint price $\frac{(g_s+g_b)}{2}$ and receives the item

Payoff to S : $1/2(g_s+g_b)-V_s$

Payoff to B : $V_b-1/2(g_s+g_b)$

Negotiation in Bayesian Game Notation

$T_s = [0,1]$ write $V_s \in T_s$

$T_b = [0,1]$ write $V_b \in T_b$

$P_s(V_b|V_s) = P_s(V_b)$ = uniform distribution on [0,1]

$P_b(V_s|V_b) = P_b(V_s)$ = uniform distribution on [0,1]

$A_s = [0,1]$ write $g_s \in A_s$

$A_b = [0,1]$ write $g_b \in A_b$

$u_s(P_s, P_b, V_s) =$ What?

$u_b(P_s, P_b, V_b) =$ What?

Double Negotiation: When does trade occur?

…when

$g_b^*(V_b) = 1/12 + 2/3\, V_b > 1/4 + 2/3\, V_s = g_s^*(V_s)$

i.e. when $V_b > V_s + 1/4$

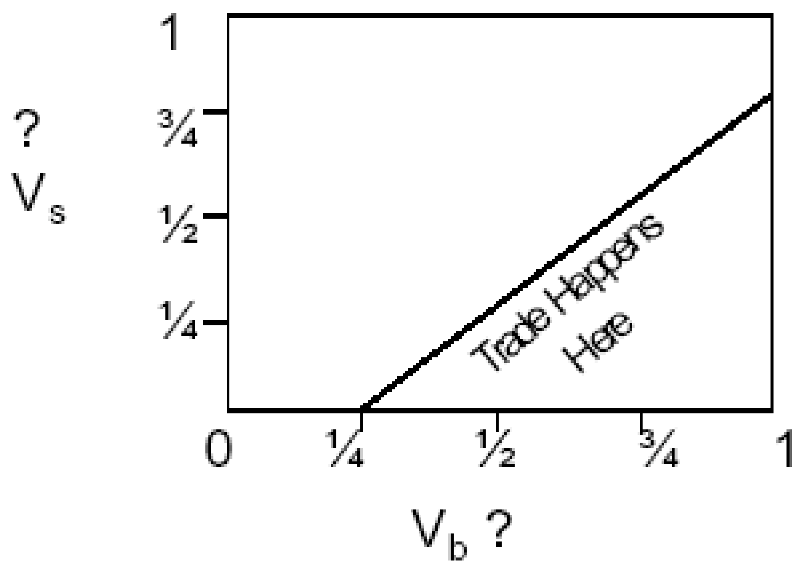

Prob(Trade Happens) = $1/2 \times (3/4)^2 = 9/32$

What You Should Know

Strict dominance

Nash Equilibria

Continuous games like Tragedy of the Commons

Rough, vague, appreciation of threats

Bayesian Game formulation

What You Shouldn't Know

- How many goats your lecturer has on his property
- What strategy Mephistopheles uses in his negotiations
- What strategy this University employs when setting tuition
- How to square a circle using only compass and straight edge
- How many of your friends and colleagues are active Santa informants, and how critical they've been of your obvious failings

Subgame-perfect Nash equilibrium

- A Nash equilibrium of a dynamic game is subgame-perfect if the strategies of the Nash equilibrium constitute or induce a Nash equilibrium in every subgame of the game.

- Subgame-perfect Nash equilibrium is a Nash equilibrium.

Games for Social Psych

Game theory is two different enterprises:

(1) Using games as a language or taxonomy to parse the social world; *(language for theory construction)*

(2) deriving precise predictions about how players will play in a game by assuming that players maximize expected "utility" (personal valuation) of consequences, plan ahead, and form beliefs about other players' likely actions. *(This is one theory expressed in the language)*

Changing the assumptions at (2) allows for modeling what people actually do, using a precise theoretical language.

Eliciting Social Preferences

But they are *MY* preferences

I'm a *team player*

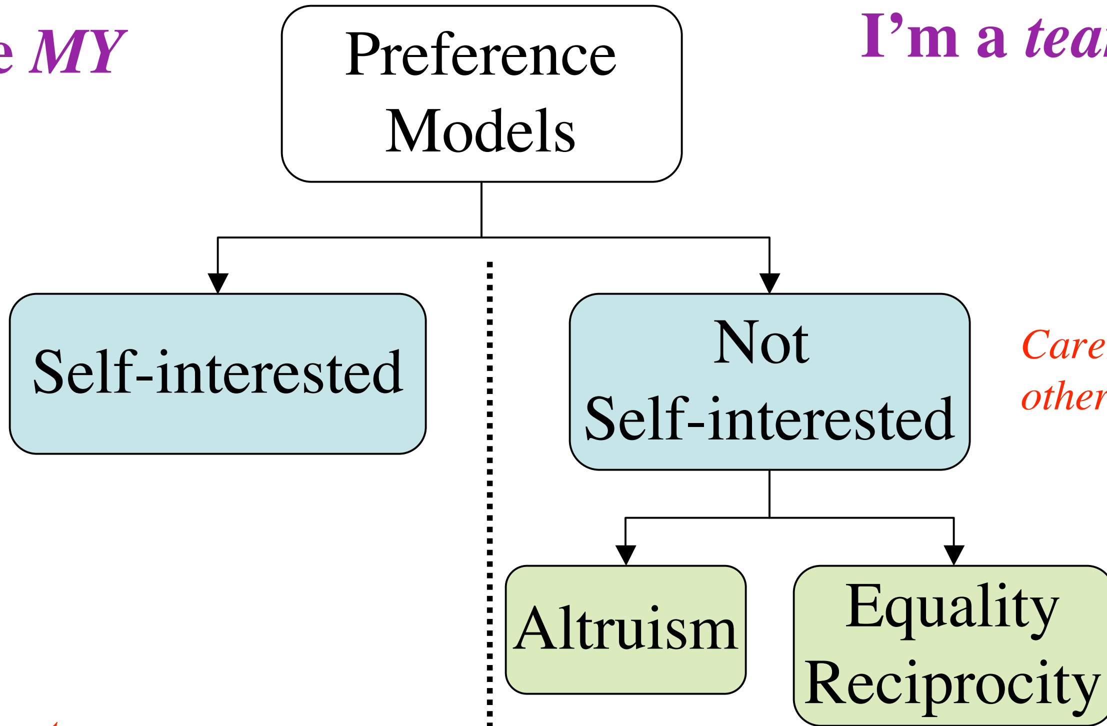

Don't care about the payoffs of others

Cares about the payoffs of others

Apparent interest in others really self-interest disguised. Any payment is in expectation of a larger self-benefit later.

Willing to pay for other's benefit

Willing to pay to punish violations of equality and reciprocity, and pay to reward obediance to norms

Games for eliciting social preferences

Table 1: Seven experimental games useful for measuring social preferences

Game	Definition of the Game	Real life Example	Predictions with rational and selfish players	Experimental regularities, References	Interpretation
Prisoners' dilemma Game	Two players, each of whom can either cooperate or defect. Payoffs are as follows: Cooperate / Defect Cooperate: H,H / S,T Defect: T,S / L,L H>L, T>H, L>S	Production of negative externalities (pollution, loud noise), exchange without binding contracts, status competition.	Defect	50% choose Cooperate. Communication increases frequency of cooperation Dawes (1980)**	Reciprocate expected cooperation
Public Goods Game	n players simultaneously decide about their contribution g_i. ($0 \leq g_i \leq y$) where y is players' endowment; each player *i* earns $\pi_i = y - g_i + mG$ where G is the sum of all contributions and $m<1<mn$.	Team compensation, cooperative production in simple societies, overuse of common resources (e.g., water, fishing grounds)	Each player contributes nothing, i.e. $g_i = 0$.	Players contribute 50% of y in the one-shot game. Contributions unravel over time. Majority chooses $g_i=0$ in final period. Communication strongly increases cooperation. Individual punishment opportunities greatly increase contributions. Ledyard (1995)**.	Reciprocate expected cooperation
Ultimatum Game	Division of a fixed sum of money S between a Proposer and a Responder. Proposer offers x. If Responder rejects x both earn zero, if x is accepted the Proposer earns S – x and the Responder earns x.	Monopoly pricing of a perishable good; "11th-hour" settlement offers before a time deadline	Offer $x=\varepsilon$; where ε is the smallest money unit. Any $x>0$ is accepted.	Most offers are between .3 and .5S. x <.2S rejected half the time. Competition among Proposers has a strong x-increasing effect; competition among Responders strongly decreases x. Güth et al (1982)*, Camerer (2003)**	Responders punish unfair offers; negative reciprocity
Dictator Game	Like the ultimatum game but the Responder cannot reject, i.e., the "Proposer" dictates (S-x, x).	Charitable sharing of a windfall gain (lottery winners giving anony-mously to strangers)	No sharing, i.e., x = 0	On average "Proposers" allocate x=.2S. Strong variations across experiments and across individuals Kahneman et al (1986)*, Camerer (2003)**	Pure altruism

More Games

Trust Game	Investor has endowment S and makes a transfer y between 0 and S to the Trustee. Trustee receives 3y and can send back any x between 0 and 3y. Investor earns S – y + x. Trustee earns 3y – x.	Sequential exchange without binding contracts (buying from sellers on Ebay)	Trustee repays nothing: x = 0. Investor invests nothing: y = 0.	On average y = .5S and trustees repay slightly less than .5S. x is increasing in y. Berg et al (1995)*, Camerer (2003)**	Trustees show positive reciprocity.
Gift Exchange Game	"Employer" offers a wage *w* to the "worker" and announces a desired effort level $\hat{e}$. If worker rejects $(w, \hat{e})$ both earn nothing. If worker accepts, he can choose *any e* between 1 and 10. Then employer earns 10e – w and worker earn $w - c(e)$. $c(e)$ is the effort cost which is strictly increasing in *e*.	Noncontractibility or nonenforceability of the performance (effort, quality of goods) of workers or sellers.	Worker chooses e = 1. Employer pays the minimum wage.	Effort increases with the wage *w*. Employers pay wages that are far above the minimum. Workers accept offers with low wages but respond with e = 1. In contrast to the ultimatum game competition among workers (i.e., Responders) has no impact on wage offers. Fehr et al (1993)*	Workers reciprocate generous wage offers. Employers appeal to workers' reciprocity by offering generous wages.
Third Party Punishment Game	A and B play a dictator game. C observes how much of amount S is allocated to B. C can punish A but the punishment is also costly for C.	Social disapproval of unacceptable treatment of others (scolding neighbors).	A allocates nothing to B. C never punishes A.	Punishment of A is the higher the less A allocates to B. Fehr and Fischbacher (2001a)*	C sanctions violation of a sharing norm.

Note: ** denotes survey papers, * denotes papers that introduced the respective games.

Public Good Games (Tragedy)

- **Public goods games**: *Every player is best off by contributing nothing to the public good, but contributions from everyone would make everyone better off.*
- Example: *n* subjects per group, each with an endowment of \$y. Each Each contributes \$0-\$y to a group project. Common payoff of \$*m* per \$1 in group project (share in the investment). In addition, *mn* > 1 (the group return for one more dollar > \$1). A dollar saved is a dollar earned, so:
- Payoff for player *i:*
- $p_i = y - g_i + mG, \quad G = \sum_{i=1}^{n} g_i$
- $\$g_i$ = *i*'s investment,
- Self-interested subjects should contribute nothing to the public good, regardless of how much the other subjects contribute.

Figure 1: Average contributions over time in public good games with a constant group composition (Source: Fehr and Gächter 2000)

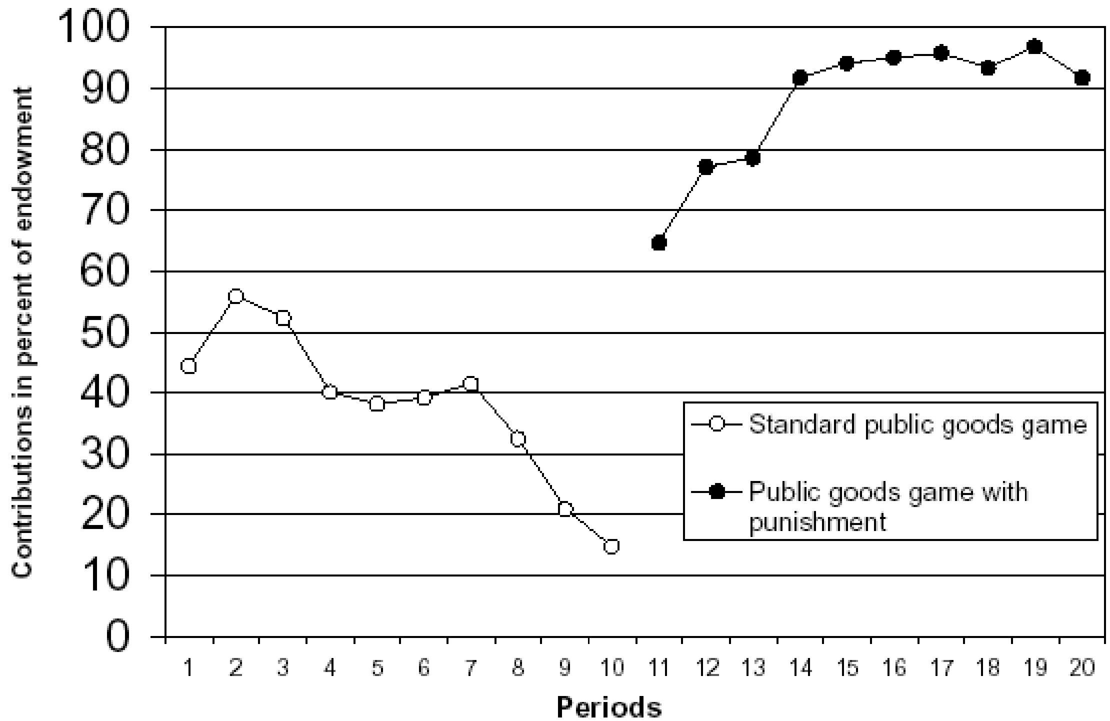

Ultimatum

extensive form

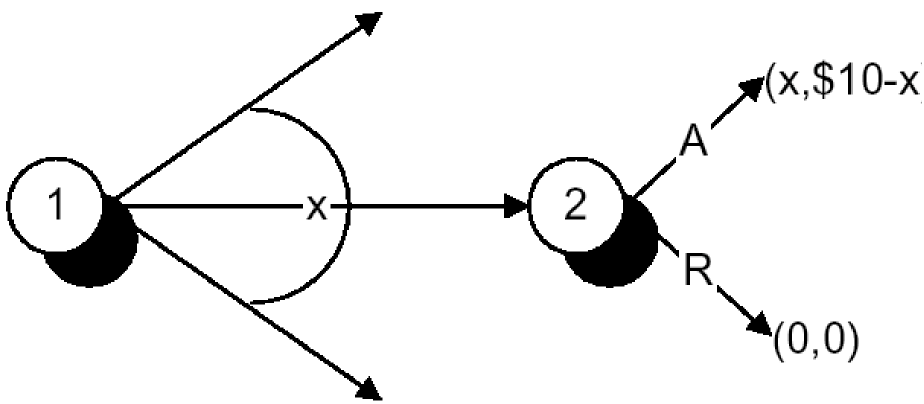

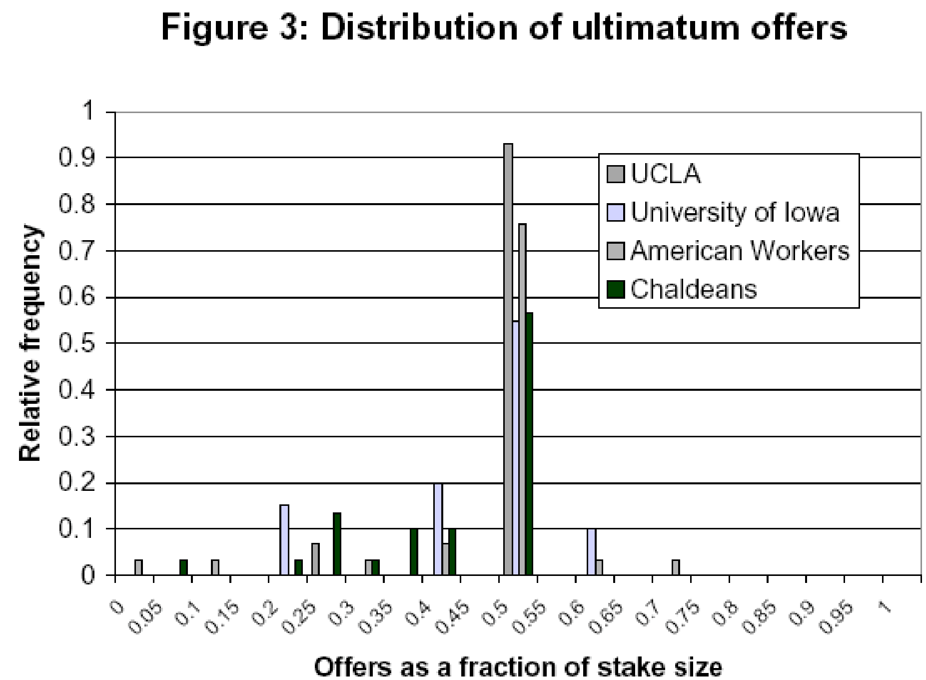

Observed offer: ~40%, relatively independent of stake size
Predicted offer: smallest increment
weak or unreplicated effects:
• gender, major (econ majors offer and accept less), physical attractiveness (women offers >50% to attractive men), age (young children accept lower offers), and autism (autistic adults offer very little; see Hill and Sally, 2002), sense of entitlement

Ultimatum with competition

- Competing receivers- lower offers - ~20%
- Competing proposers- higher offers ~75%

- Why?
  - Altruism (a preference for sharing equally)
    - *Non self-interested*
  - "Strategic fairness" (a fear that low offers will be rejected)
    - *self-interested*

Dictator

Dictator game:

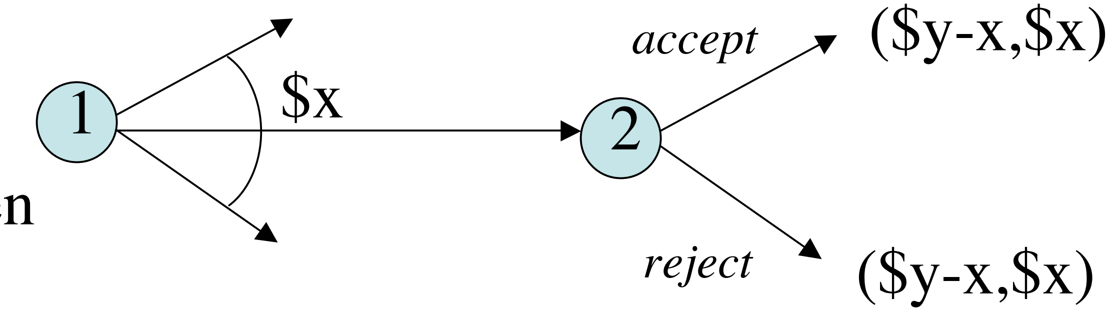

Proposer division of $*y* between self and other player

Self-interested prediction

Propose: $0

Students: ~10-25%,

Kansas workers/Chaldeans

~50% same as in Ultimatum

Figure 4: Dictator game allocations

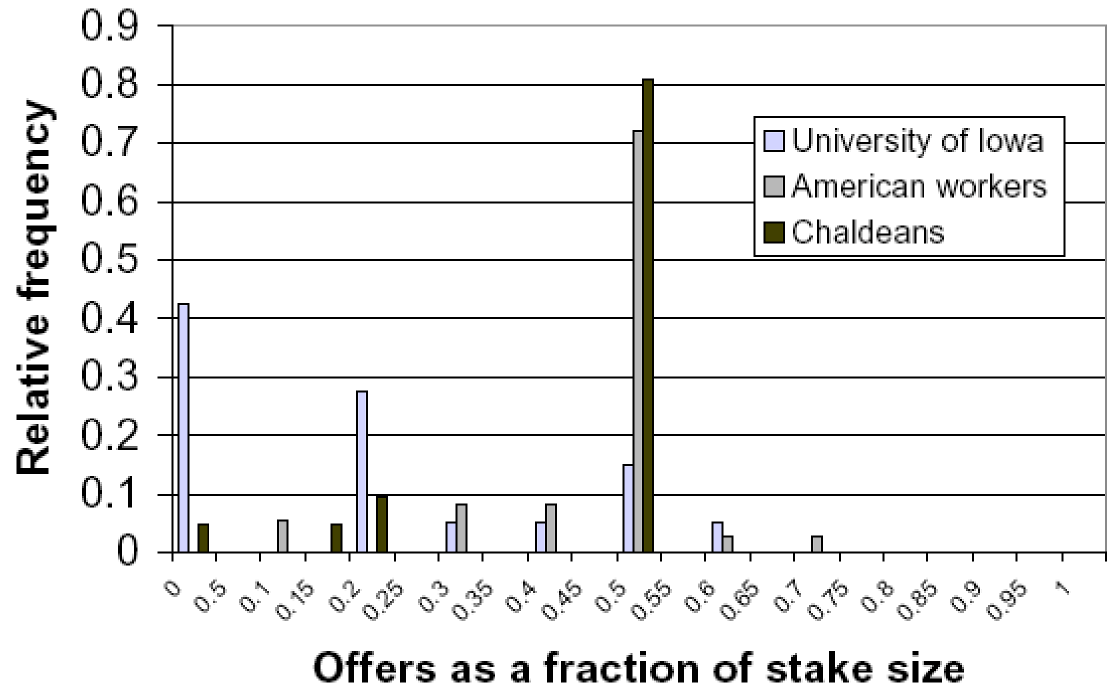

Modeling Social Preferences

- Two model flavors have been proposed—
- Inequality-aversion: *players prefer more money and also prefer that allocations be more equal.* Fehr and Schmidt (1999)

  x_i = payoff of player *i*

  $U_i(x) = x_i - a_i(x_j - x_i)$ if player *i* is worse off than player *j* ($x_j - x_i \le 0$), and

  $U_i(x) = x_i - b_i(x_j - x_i)$ if player i is better off than player *j* ($x_j - x_i \ge 0$).

  Envy: a_i measures player *i's* dislike of disadvantageous inequality

  Guilt: b_i measures player *i's* dislike of advantageous inequality

- Models of reciprocity. Rabin Utility model

$$U_i(\$, q=personality) = U_i(\$) + w\, U_{pi}(q_i) * U_{pk}(q_k)$$

U_{pi}('niceness')>0, U_{pi}('meanness')<0,

Thus, if the other player is nice (positive niceness) they want to be nice too, so the product of nicenesses will be positive. But if the other player is mean (negative niceness) they want to be negative too so the product of nicenesses will be positive.

Captures the fact that a single player may behave nicely or meanly depending on how they expect to be treated - it locates social preferences and emotions in the combination of a person, their partner, and a game, rather than as a fixed personal attribute.

Modeling social preferences via utilities on opponent's outcomes

- Set of players: {Prisoner 1, Prisoner 2}
- Sets of strategies: $S_1 = S_2 =$ {**M**um, **C**onfess}
- Utility functions are now on both players payoffs

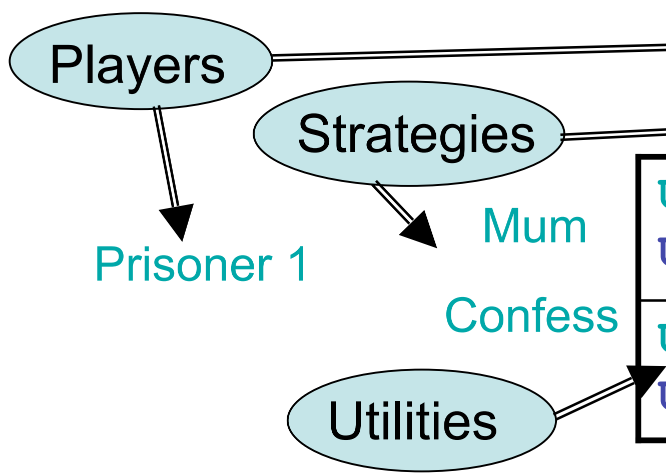

Prisoner 2

	Mum	Confess
Mum	$U_1(-1,-1)$, $U_2(-1,-1)$	$U_1(-9,0)$, $U_2(-9,0)$
Confess	$U_1(0,-9)$, $U_2(0,-9)$	$U_1(-6,-6)$, $U_2(-6,-6)$

Altruistic Preferences

- players $i = 1, \ldots, n$
- at terminal nodes *direct utility* of u_i
- coefficient of altruism $-1 < a_i < 1$
- *adjusted utility*

$$v_i = u_i + \sum_{j \neq i} a_i u_j$$

$$v_i = u_i + \sum_{j \neq i} \frac{a_i + \lambda a_j}{1 + \lambda} u_j \,.$$

Fairness seeking

- $0 \leq \lambda \leq 1$
- objective is to maximize adjusted utility
- since the stakes are small, ignore risk aversion, and identify direct utility with monetary payoffs
- prior to start of play, players drawn independently from population with a distribution of altruism coefficients represented by a common cumulative distribution function. $F(a_i)$
- each player's altruism coefficient a_i is privately known
- the distribution *F* is common knowledge

Feeling: This is your brain on unfairness (Sanfey et al. Sci 13 March '03)

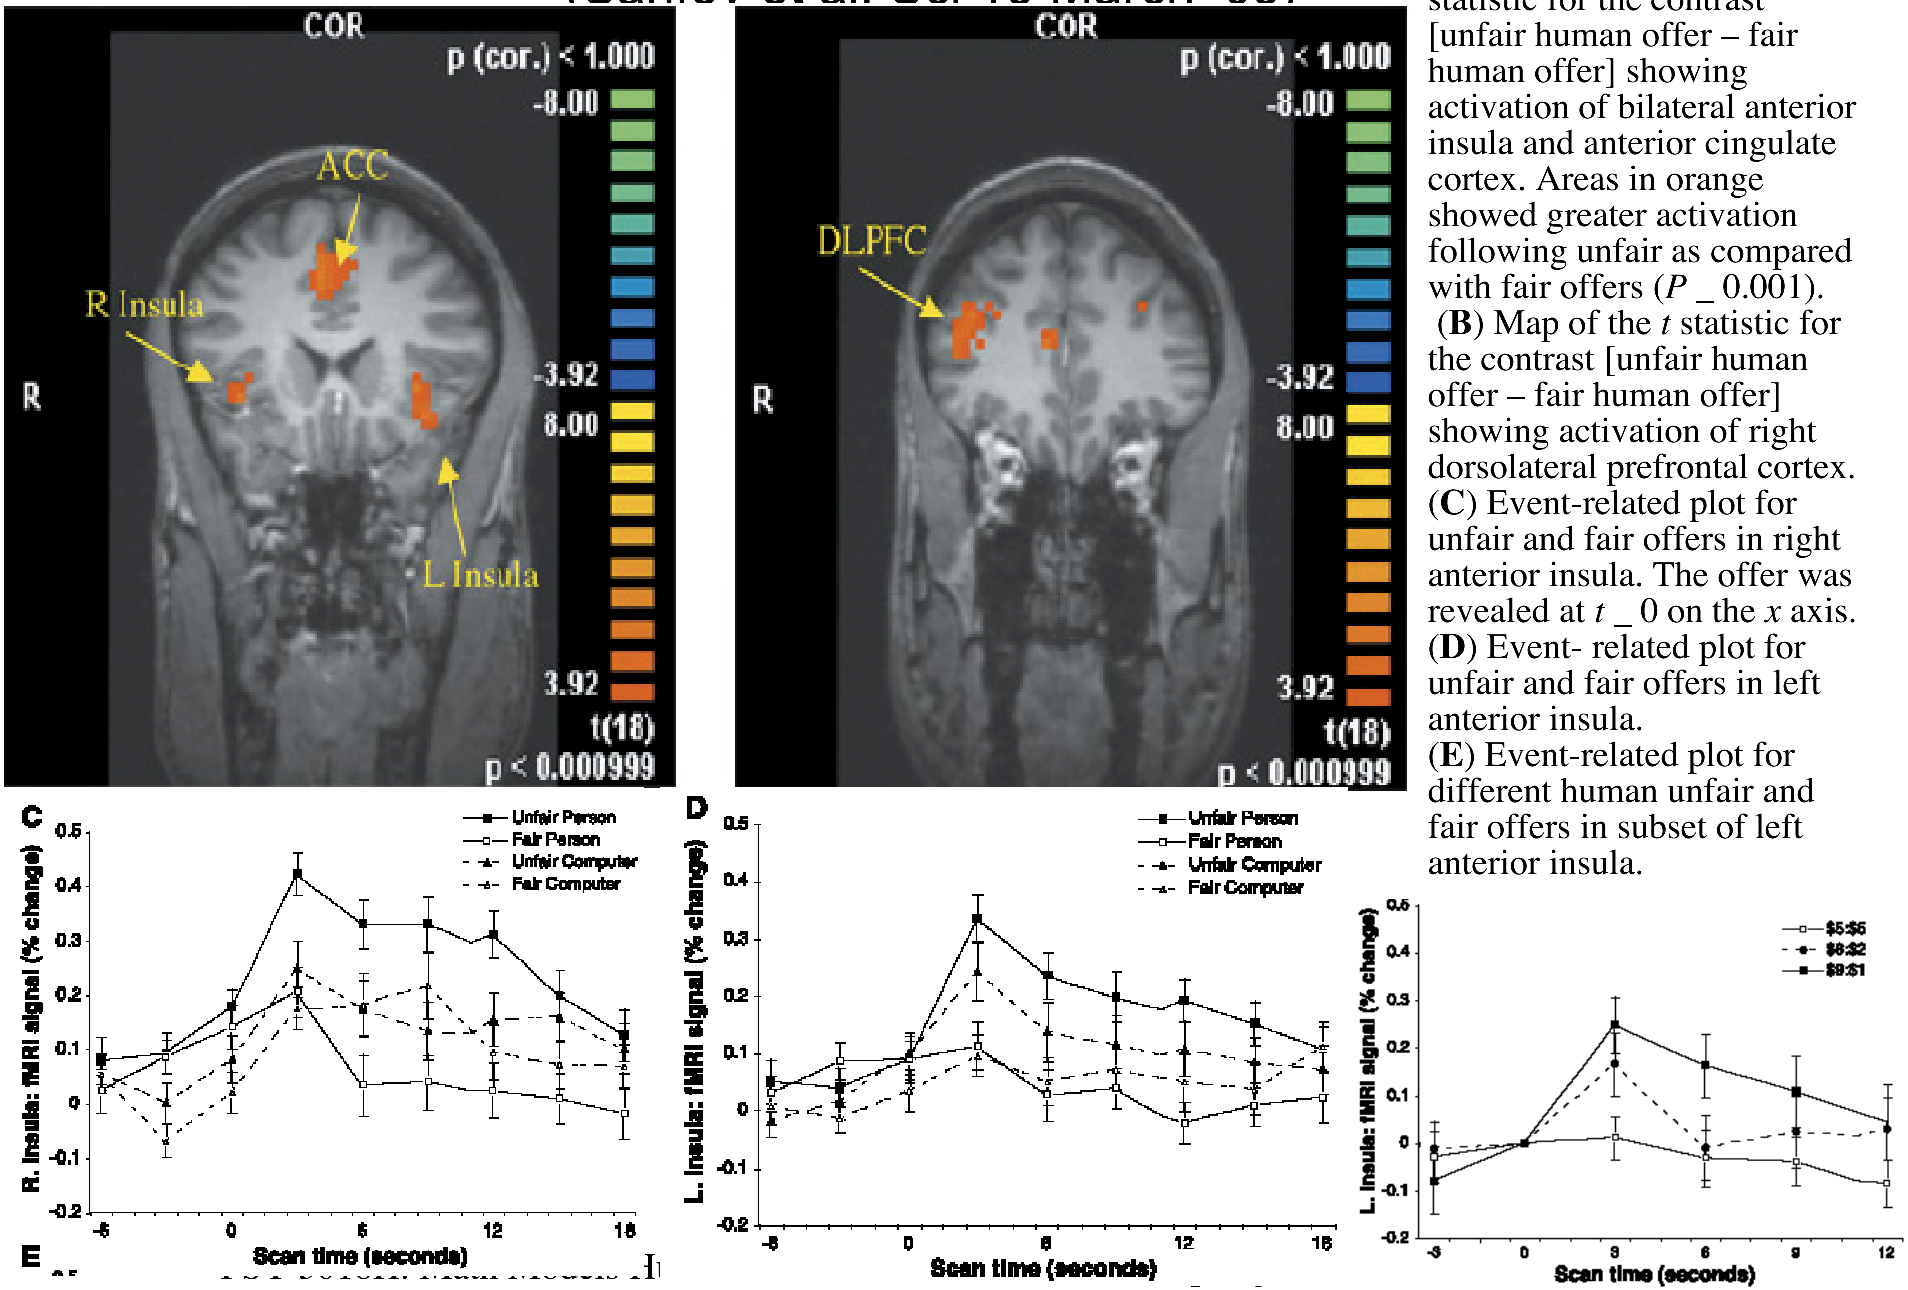

Fig.2. Activation related to the presentation of an unfair offer. (**A**) Map of the *t* statistic for the contrast [unfair human offer – fair human offer] showing activation of bilateral anterior insula and anterior cingulate cortex. Areas in orange showed greater activation following unfair as compared with fair offers (*P* _ 0.001). (**B**) Map of the *t* statistic for the contrast [unfair human offer – fair human offer] showing activation of right dorsolateral prefrontal cortex. (**C**) Event-related plot for unfair and fair offers in right anterior insula. The offer was revealed at *t* _ 0 on the *x* axis. (**D**) Event- related plot for unfair and fair offers in left anterior insula. (**E**) Event-related plot for different human unfair and fair offers in subset of left anterior insula.

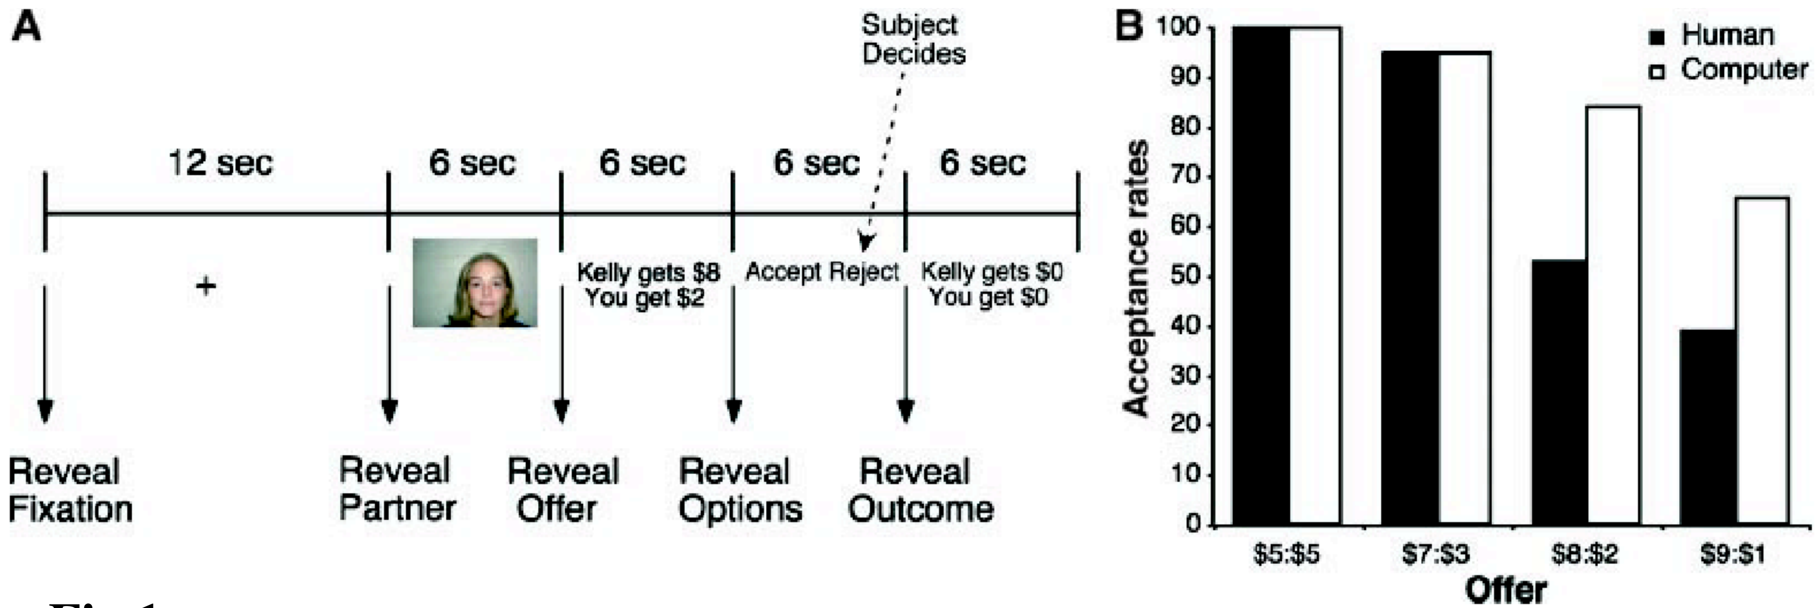

Fig.1.
(**A**) Time line for a single round of the Ultimatum Game, each lasting 36 s. Each round began with a 12-s preparation interval. The participant then saw the photograph and name of their partner in that trial for 6 seconds. A picture of a computer was shown if it was a computer trial, or a roulette wheel if it was a control trial. Next, participants saw the offer proposed by the partner for a further 6 s, after which they indicated whether they accepted or rejected the offer by pressing one of two buttons on a button box.

(**B**) Behavioral results from the Ultimatum Game. These are the offer acceptance rates averaged over all trials. Each of 19 participants saw five $5:$5 offers, one $7:$3 offer, two $8:$2 offers, and two $9:$1 offers from both human and computer partners (20 offers in total).

Ultimatum offer experimental sites

The Machiguenga

independent families
cash cropping

slash & burn
gathered foods
fishing
hunting

African pastoralists (Orma in Kenya)

Whale Hunters
of
Lamalera, Indonesia

High levels of cooperation among hunters of whales, sharks, dolphins and rays. Protein for carbs, trade with inlanders. Carefully regulated division of whale meat

Researcher: Mike Alvard

Ultimatum offers across societies
(mean shaded, mode is largest circle…)

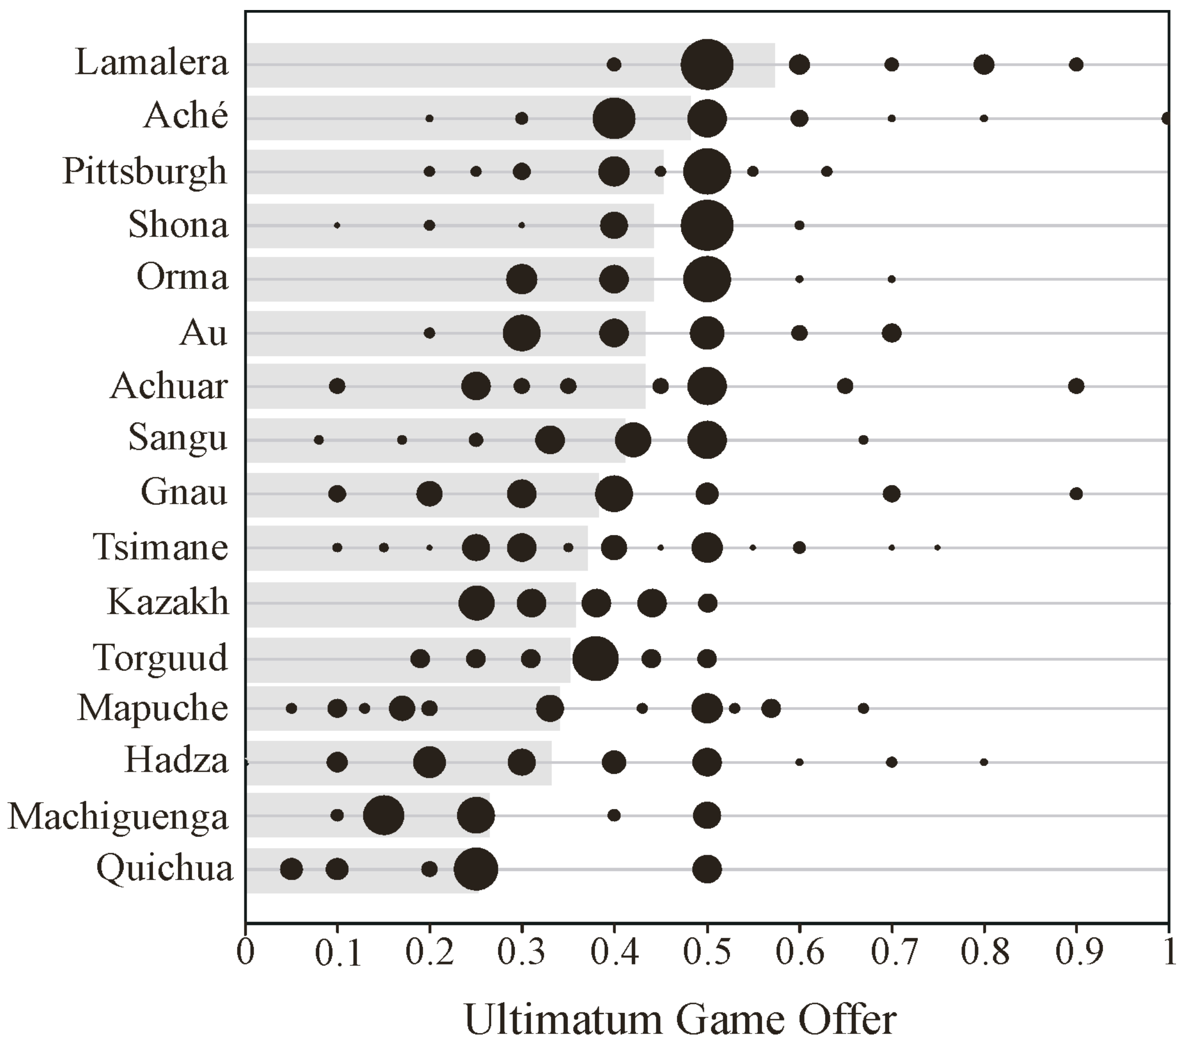

Israeli subject (autistic?) complaining post-experiment (Zamir, 2000)

‘I did not earn any money because all the other players are *stupid!* How can you reject a positive amount of money and prefer to get zero? They just did not understand the game! You should have stopped the experiment and explained it to them...’

Behavioral game theory

BGT: How people actually play games

Key extensions over traditional Game Theory

- **Framing:** Mental representation
- **Feeling**: Social preferences (Fehr et al)
- **Thinking**: Cognitive hierarchy (τ)
- **Learning**: Hybrid fEWA (Experience-weighted attraction) adaptive rule (λ)
- **Teaching**: Bounded rationality in repeated games (α, λ)

BGT Notes based on notes from Colin F. Camerer, Caltech http://www.hss.caltech.edu/~camerer/camerer.html

Behavioral Game Theory, Princeton Press 03 (550 pp); *Trends in Cog Sci*, May 03 (10 pp); *AmerEcRev*, May 03 (5 pp); *Science*, 13 June 03 (2 pp)

Thinking: A one-parameter cognitive hierarchy theory of one-shot games

(Camerer, Ho, Chong)

- Model of constrained strategic thinking
- Model does several things:
  - 1. Limited equilibration in some games (e.g., pBC)
  - 2. Surprisingly fast equilibration in some games (e.g. entry)
  - 3. De facto purification in mixed games
  - 4. Limited belief in noncredible threats
  - 5. Has economic value
  - 6. Can prove theorems
    - e.g. risk-dominance in 2x2 symmetric games
  - 7. Permits individual diff's & relation to cognitive measures

Different equilibrium notions

Principle	Nash	CH	QRE
Strategic Thinking	X	X	X
Best Response	X	X	
Mutual Consistency	X		X
	Nash: *Everyone's the same, ideal and make best self-interested response*	***CH:*** *Everyone's NOT the same, but makes best response given values*	***QRE:*** *Everyone's the same, but NOT best response given values*

QRE: quantal-response equilibrium. Players do not choose the best response with probability one (as in Nash equilibrium). Instead, they "better-respond", choosing responses with higher expected payoffs with higher probability.

CH: Camerer-Ho

The cognitive hierarchy (CH) model (I)

- Discrete steps of thinking:

  Step 0's choose randomly (nonstrategically)

  K-step thinkers know proportions $f(0),\ldots f(K\text{-}1)$

  (can't imagine what smarter people would do, but can for simpler)

  Calculate what $0, \ldots K\text{-}1$ step players will do

  Normalize beliefs $g_K(n)=f(n)/\ \Sigma_{h=0}^{K\text{-}1} f(h)$.

  Calculate expected payoffs and best-respond

- Exhibits "increasingly rational expectations":
  - Normalized $g_K(n)$ approximates f(n) more closely as K→ ∞
    - i.e., highest level types are "sophisticated"/"worldly and earn the most
    - Also: highest level type actions converge as K→ ∞
    - (→ marginal benefit of thinking harder →0)

The cognitive hierarchy (CH) model (II)

- Two separate features:
  - Not imagining k+1 types
  - Not believing there are many other k types

Models Overconfidence:
K-steps think others are all one step lower (K-1)
(Nagel-Stahl-CCGB)
"Increasingly irrational expectations" as K→ ∞
Has some odd properties (cycles in entry games…)

What if self-conscious?:
Then K-steps believe there are other K-step thinkers
Predictions "Too similar" to quantal response equilibrium/Nash
(& fits worse)

The cognitive hierarchy (CH) model (III)

- What is a reasonable simple f(K)?
  - A1*: f(k)/f(k-1) ~1/k
    → Poisson $f(k)=e^{-\tau}\tau^{\kappa}/k!$ mean, variance τ

  With additional assumptions, it is possible to pin down the parameter τ

  - A2: f(1) is modal → $1< \tau < 2$
  - A3: f(1) is a 'maximal' mode
    or f(0)=f(2) → $t=\sqrt{2}=1.414..$
  - A4: f(0)+f(1)=2f(2) → t=1.618 (golden ratio _)

*Amount of working memory (digit span) correlated with steps of iterated deletion of dominated strategies (Devetag & Warglien, 03 J Ec Psych)

Poisson distribution

- Discrete, one parameter
  - (→ “spikes” in data)
- Steps > 3 are rare (working memory bound)
- Steps can be linked to cognitive measures

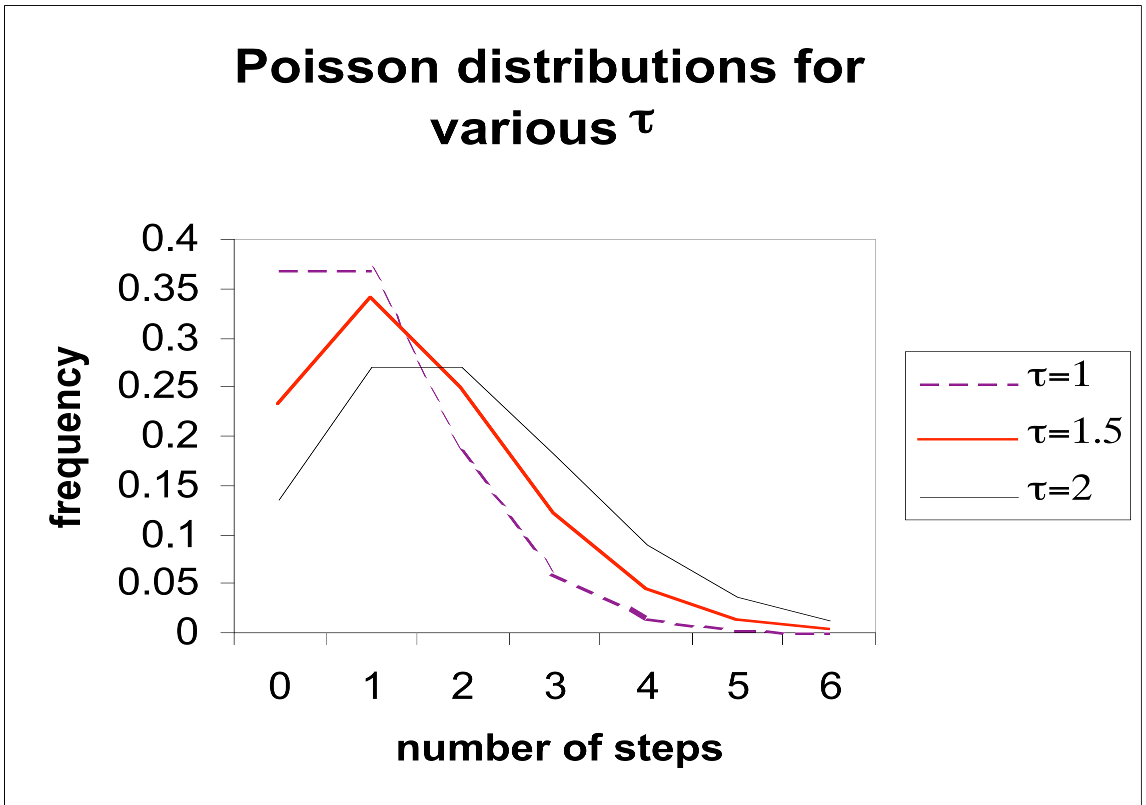

Beauty contest game

- N players choose *real* numbers x_i in [0,100]
- Compute target $(2/3)*(\Sigma x_i /N)$
- Closest to target wins $20

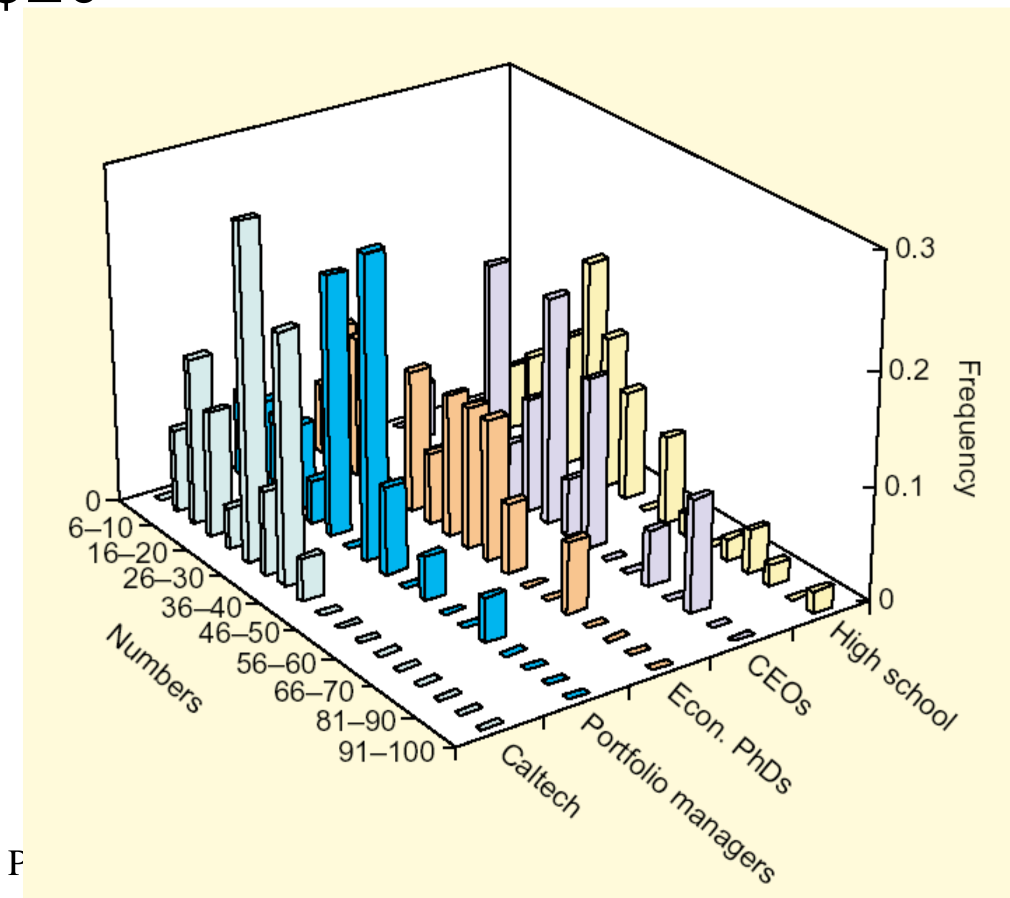

- Nash Eq?
  - Real?
    - $(2/3)^n * \text{mean}$, n → inf
  - Integers?

1. Limited equilibration in p-BC:
Pick [0,100]; closest to (2/3)*(average) wins

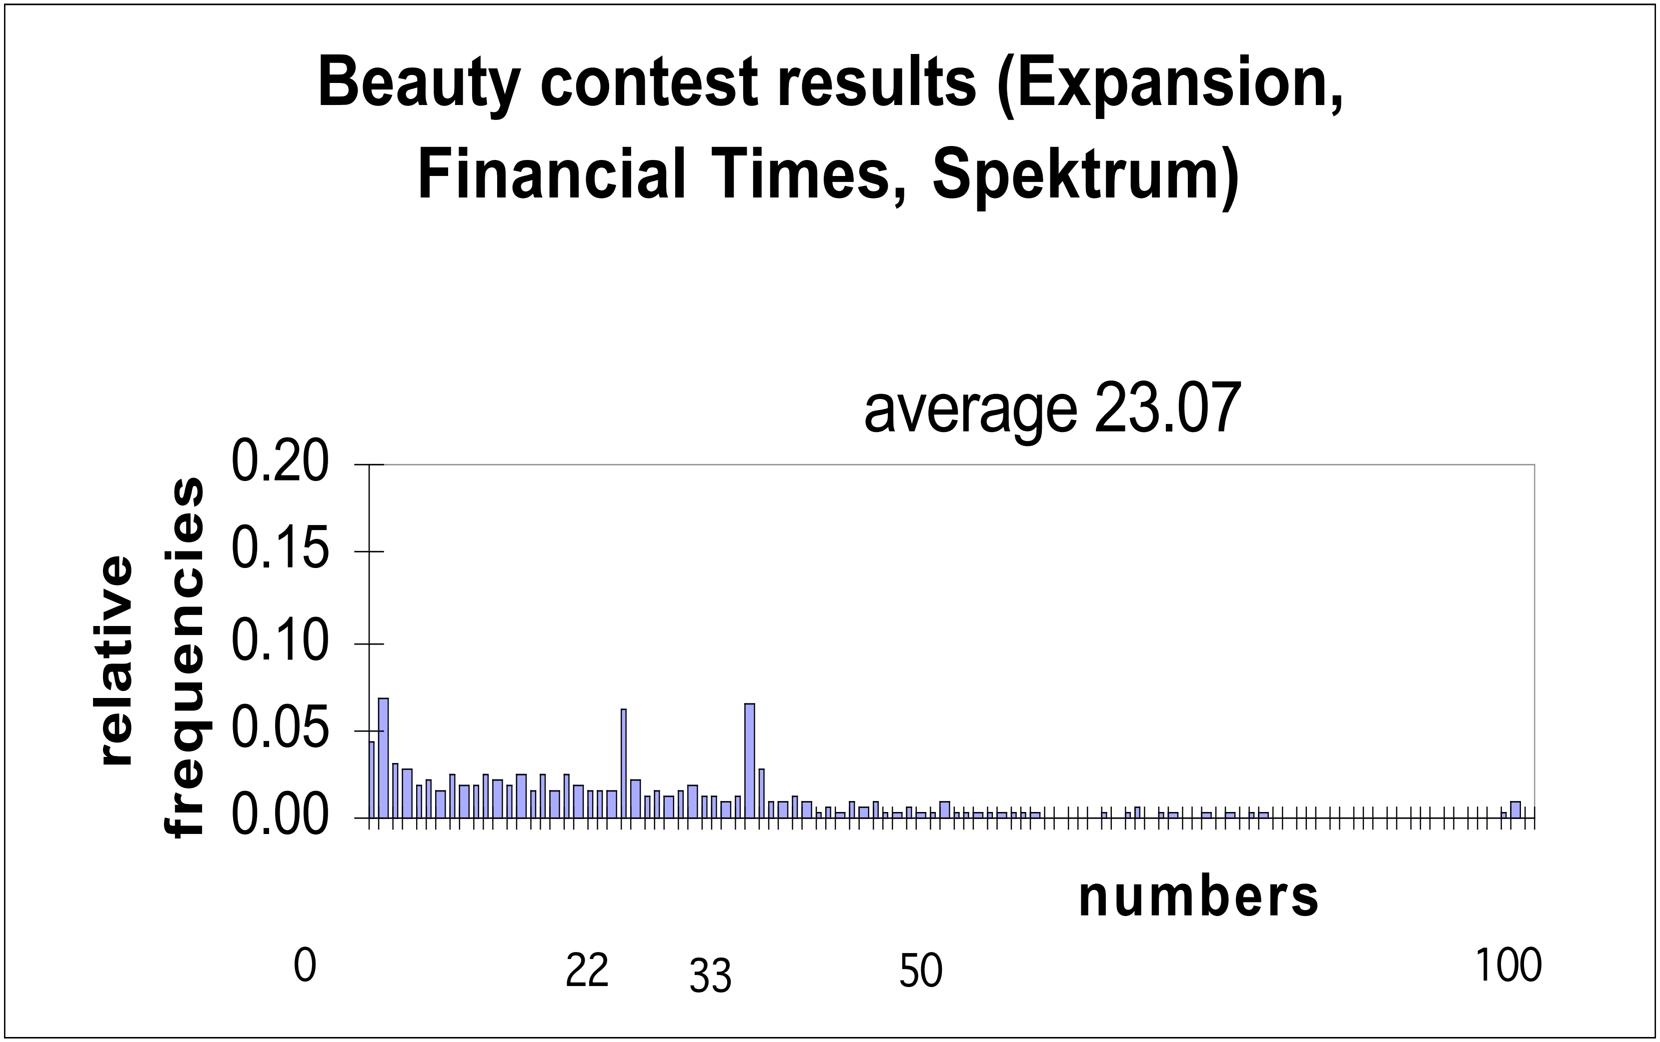

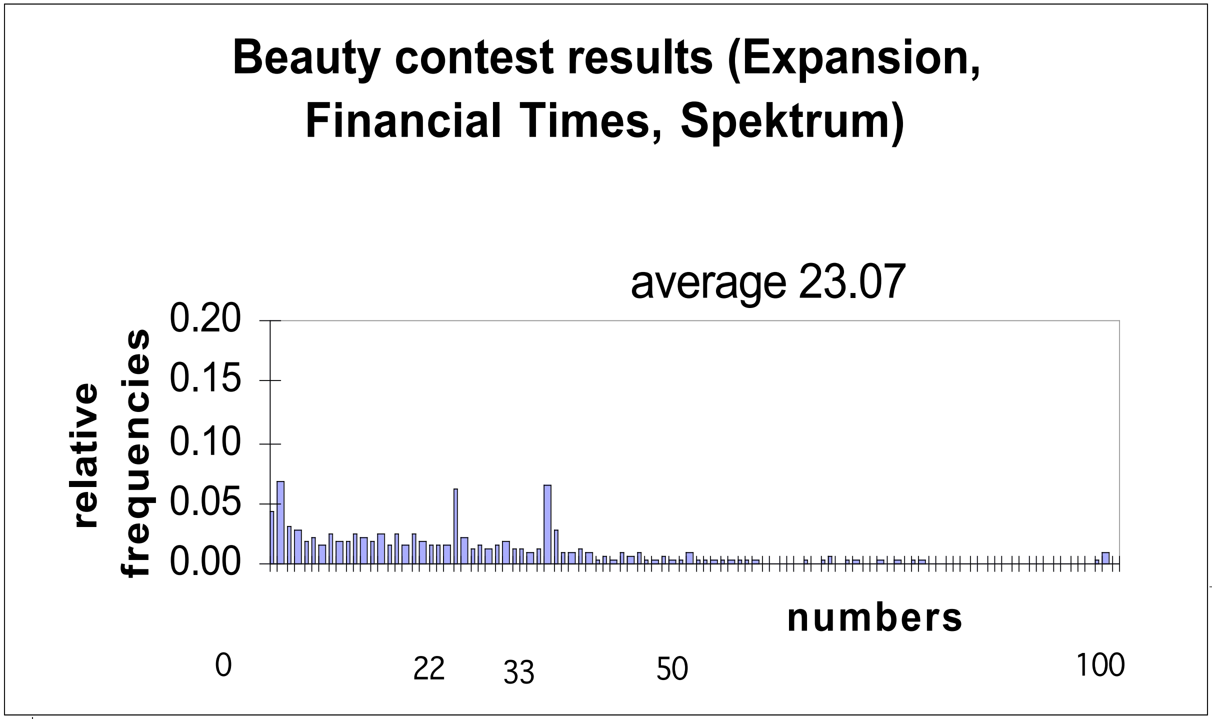
Beauty contest results (Expansion, Financial Times, Spektrum)
average 23.07
0.20
0.15
0.10
0.05
0.00
relative frequencies
numbers
0
22
33
50
100

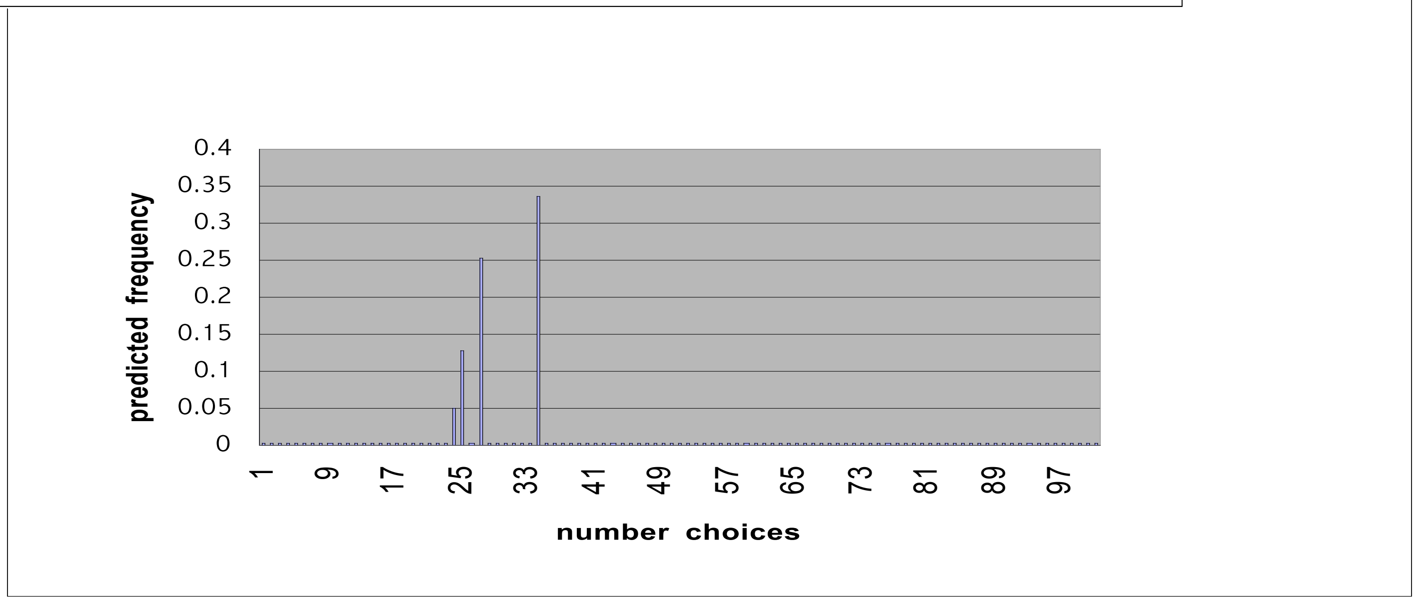
0.4
0.35
0.3
0.25
0.2
0.15
0.1
0.05
0
predicted frequency
1
9
17
25
33
41
49
57
65
73
81
89
97
number choices

Estimates of τ in pBC games

Table 1: Data and estimates of τ in pbc games (equilibrium = 0)			
	Mean		Steps of
subjects/game	Data	CH Model	Thinking
game theorists	19.1	19.1	**3.7**
Caltech	23.0	23.0	**3.0**
newspaper	23.0	23.0	**3.0**
portfolio mgrs	24.3	24.4	**2.8**
econ PhD class	27.4	27.5	**2.3**
Caltech g=3	21.5	21.5	**1.8**
high school	32.5	32.7	**1.6**
1/2 mean	26.7	26.5	**1.5**
70 yr olds	37.0	36.9	**1.1**
Germany	37.2	36.9	**1.1**
CEOs	37.9	37.7	**1.0**
game p=0.7	38.9	38.8	**1.0**
Caltech g=2	21.7	22.2	**0.8**
PCC g=3	47.5	47.5	**0.1**
game p=0.9	49.4	49.5	**0.1**
PCC g=2	54.2	49.5	**0.0**
		mean	**1.56**
		median	**1.30**

pBC estimation: Gory details

Table 3: Data and CH estimates of τ in various p-beauty contest games

subject pool or game	source[1]	group size	sample size	Nash equil'm	pred'n error	data: mean	data: std dev	data: mode	fit of CH model: τ	fit of CH model: mean	fit of CH model: error	fit of CH model: std dev	fit of CH model: mode	bootstrapped 90% c.i.
p=1.1	HCW (98)	7	69	200	47.9	152.1	23.7	150	0.10	151.6	-0.5	28.0	165	(0.0,0.5)
p=1.3	HCW (98)	7	71	200	50.0	150.0	25.9	150	0.00	150.4	0.5	29.4	195	(0.0,0.1)
high $	CHW	7	14	72	11.0	61.0	8.4	55	4.90	59.4	-1.6	3.8	61	(3.4,4.9)
male	CHW	7	17	72	14.4	57.6	9.7	54	3.70	57.6	0.1	5.5	58	(1.0,4.3)
female	CHW	7	46	72	16.3	55.7	12.1	56	2.40	55.7	0.0	9.3	58	(1.6,4.9)
low $	CHW	7	49	72	17.2	54.8	11.9	54	2.00	54.7	-0.1	11.1	56	(0.7,3.8)
.7(Mean+18)	Nagel (98)	7	34	42	-7.5	49.5	7.7	48	0.20	49.4	-0.1	26.4	48	(0.0,1.0)
PCC	CHC (new)	2	24	0	-54.2	54.2	29.2	50	0.00	49.5	-4.7	29.5	0	(0.0,0.1)
p=0.9	HCW (98)	7	67	0	-49.4	49.4	24.3	50	0.10	49.5	0.0	27.7	45	(0.1,1.5)
PCC	CHC (new)	3	24	0	-47.5	47.5	29.0	50	0.10	47.5	0.0	28.6	26	(0.1,0.8)
Caltech board	Camerer	73	73	0	-42.6	42.6	23.4	33	0.50	43.1	0.4	24.3	34	(0.1,0.9)
p=0.7	HCW (98)	7	69	0	-38.9	38.9	24.7	35	1.00	38.8	-0.2	19.8	35	(0.5,1.6)
CEOs	Camerer	20	20	0	-37.9	37.9	18.8	33	1.00	37.7	-0.1	20.2	34	(0.3,1.8)
German students	Nagel (95)	14-16	66	0	-37.2	37.2	20.0	25	1.10	36.9	-0.2	19.4	34	(0.7,1.5)
70 yr olds	Kovalchik	33	33	0	-37.0	37.0	17.5	27	1.10	36.9	-0.1	19.4	34	(0.6,1.7)
US high school	Camerer	20-32	52	0	-32.5	32.5	18.6	33	1.60	32.7	0.2	16.3	34	(1.1,2.2)
econ PhDs	Camerer	16	16	0	-27.4	27.4	18.7	N/A	2.30	27.5	0.0	13.1	21	(1.4,3.5)
1/2 mean	Nagel (98)	15-17	48	0	-26.7	26.7	19.9	25	1.50	26.5	-0.2	19.1	25	(1.1,1.9)
portfolio mgrs	Camerer	26	26	0	-24.3	24.3	16.1	22	2.80	24.4	0.1	11.4	26	(2.0,3.7)
Caltech students	Camerer	17-25	42	0	-23.0	23.0	11.1	35	3.00	23.0	0.1	10.6	24	(2.7,3.8)
newspaper	Nagel (98)	3696, 1460, 2728	7884	0	-23.0	23.0	20.2	1	3.00	23.0	0.0	10.6	24	(3.0,3.1)
Caltech	CHC (new)	2	24	0	-21.7	21.7	29.9	0	0.80	22.2	0.6	31.6	0	(4.0,1.4)
Caltech	CHC (new)	3	24	0	-21.5	21.5	25.7	0	1.80	21.5	0.1	18.6	26	(1.1,3.1)
game theorists	Nagel (98)	27-54	136	0	-19.1	19.1	21.8	0	3.70	19.1	0.0	9.2	16	(2.8,4.7)
								mean	1.30					
								median	1.61					

2. Approximate equilibration in entry games

- Entry games:

  N entrants, capacity c

  Entrants earn \$1 if n(entrants)$\leq$c;

  earn 0 if n(entrants)>c

  Earn \$.50 by staying out

  All choose simultaneously
- Close to equilibrium in the 1st period:

  Close to equilibrium prediction n(entrants) $\approx$ c

  "To a psychologist, it looks like magic"-- D. Kahneman '88
- How? Pseudo-sequentiality of CH → "later" entrants smooth the entry function

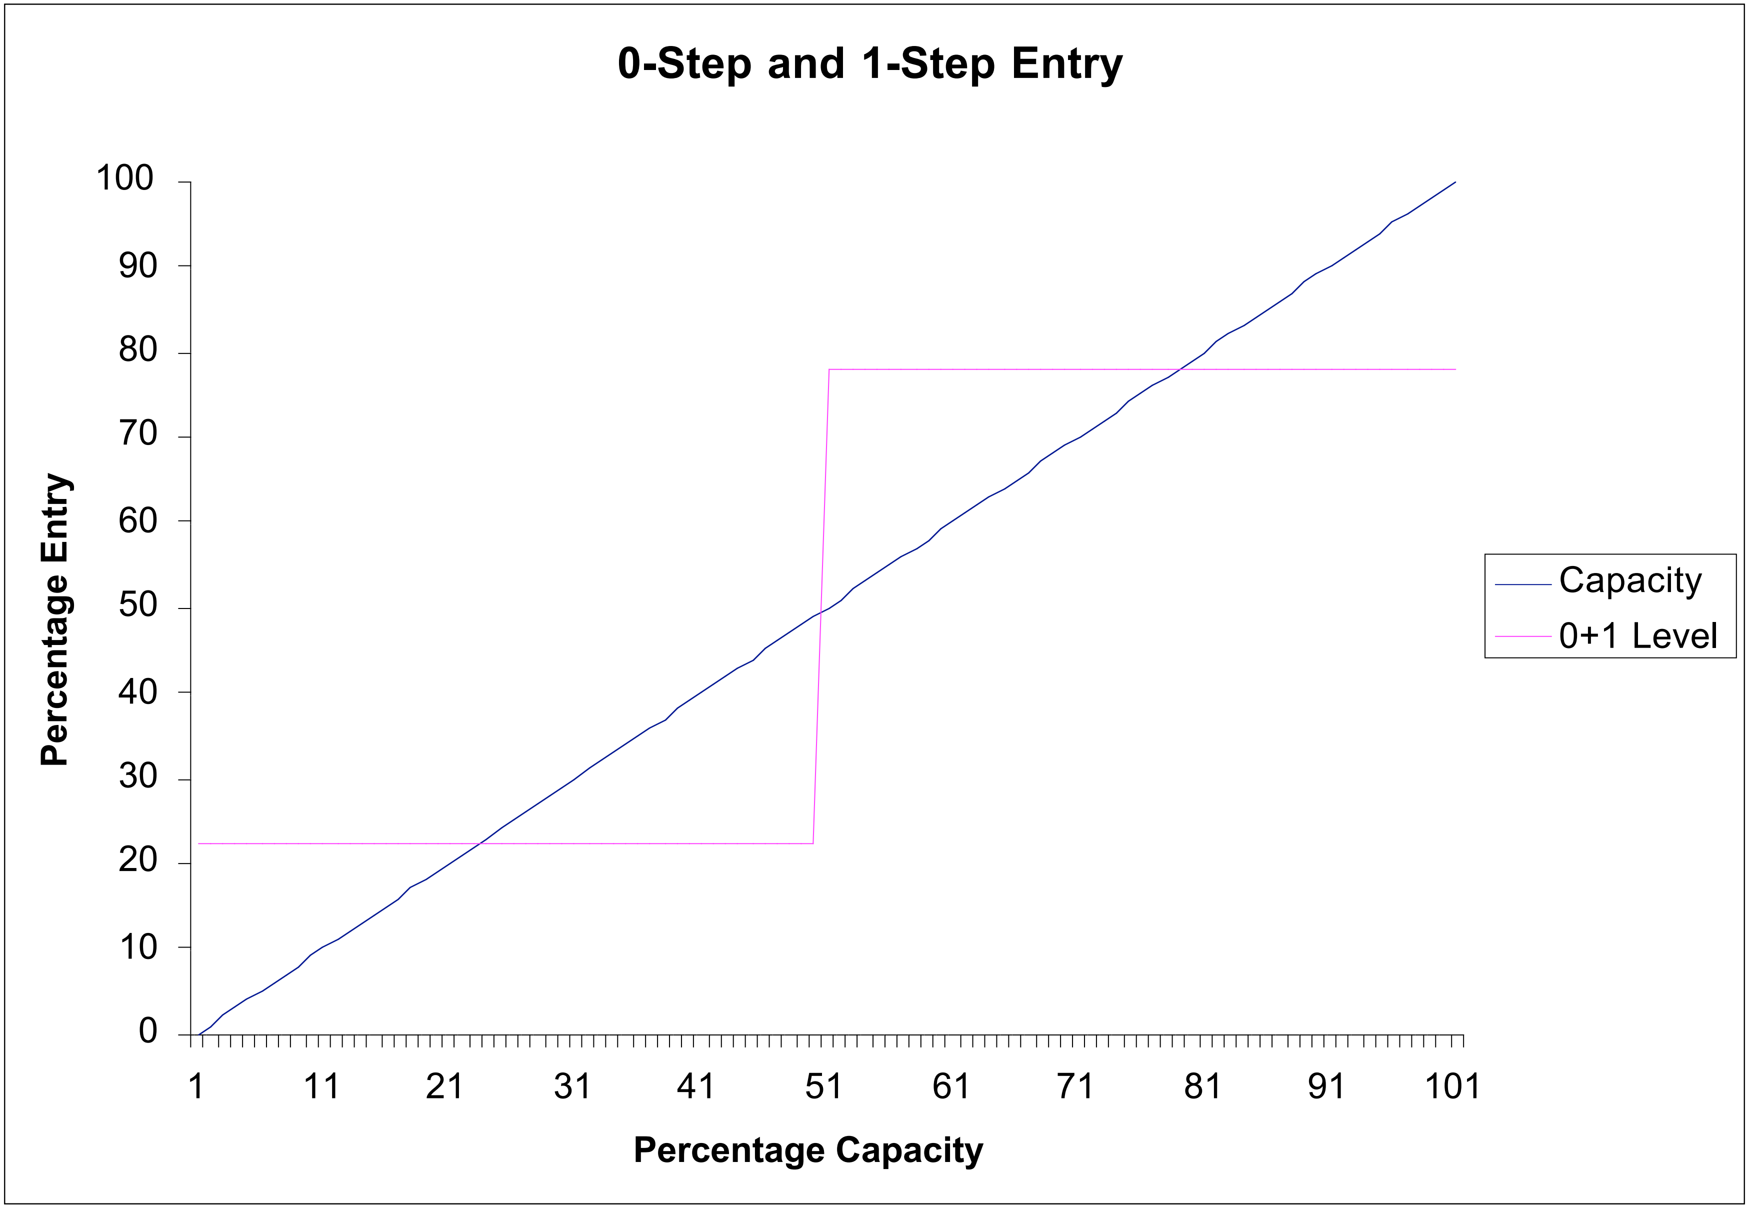
0-Step and 1-Step Entry
100
90
80
70
60
50
40
30
20
10
0
Percentage Entry
1
11
21
31
41
51
61
71
81
91
101
Percentage Capacity
Capacity
0+1 Level

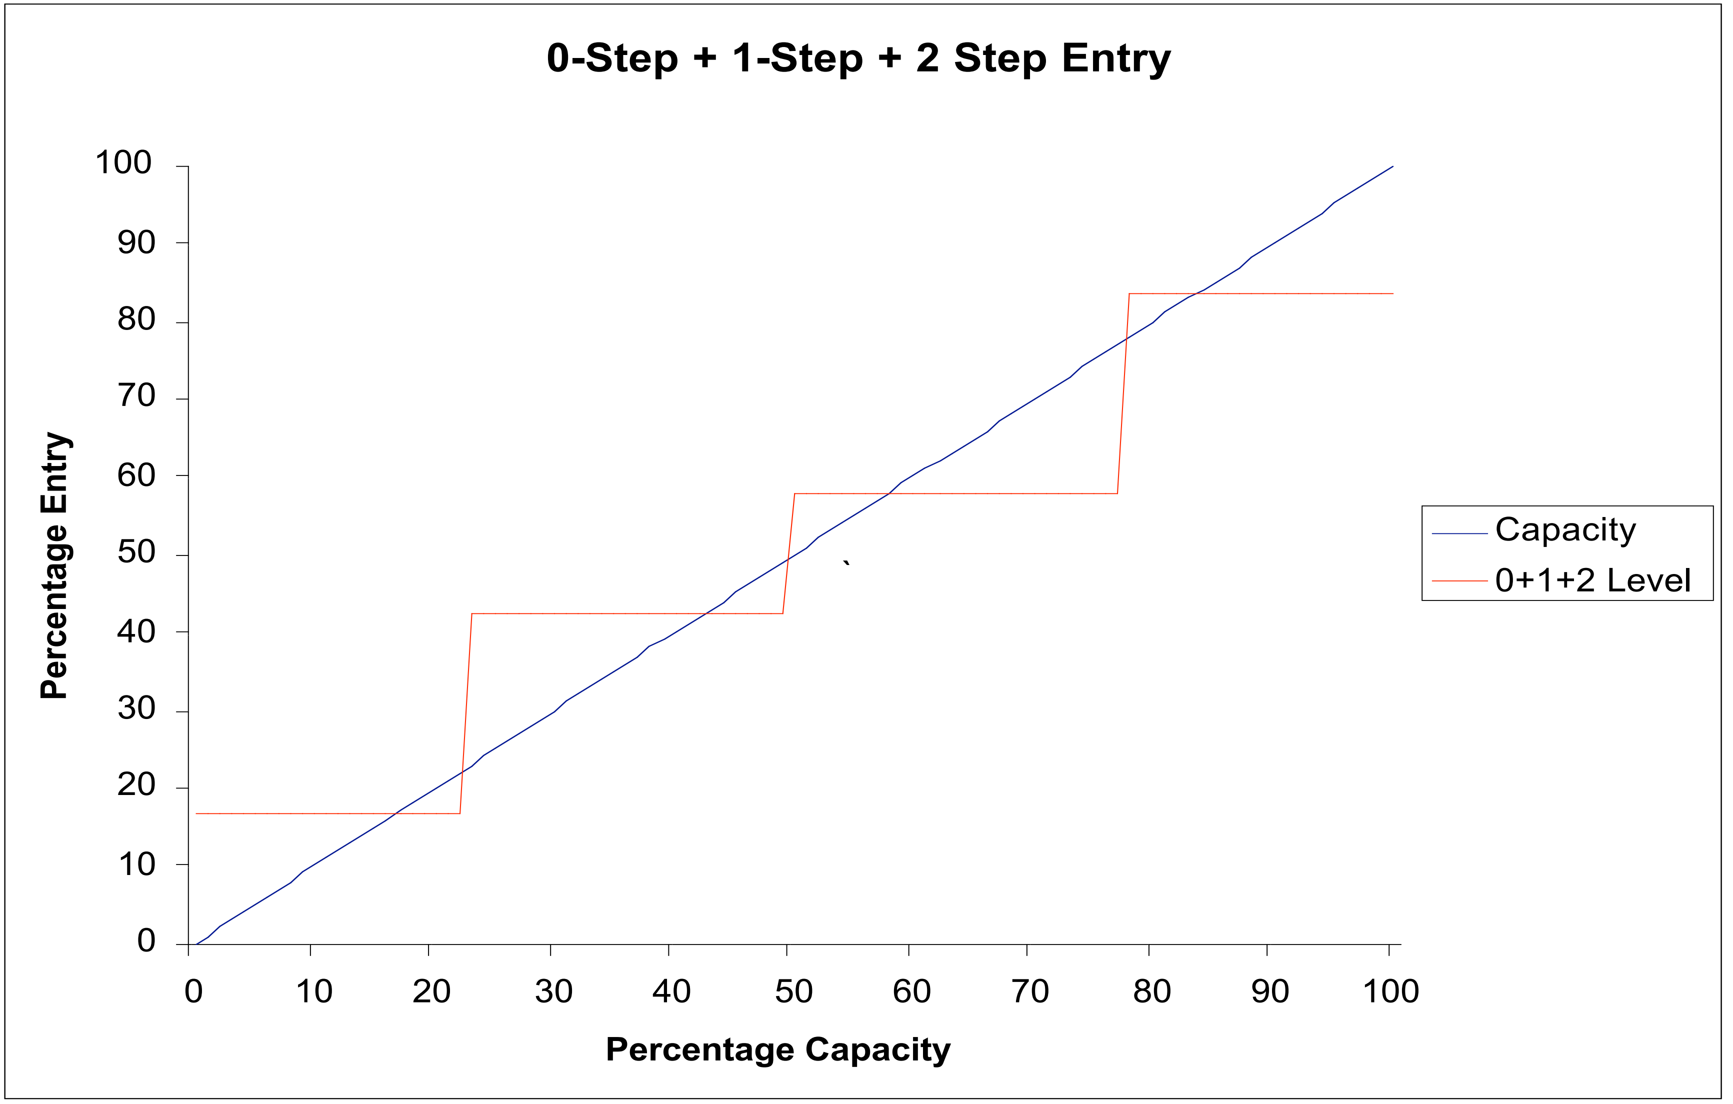

0-Step + 1-Step + 2 Step Entry
Percentage Entry
Percentage Capacity
0
10
20
30
40
50
60
70
80
90
100
Capacity
0+1+2 Level

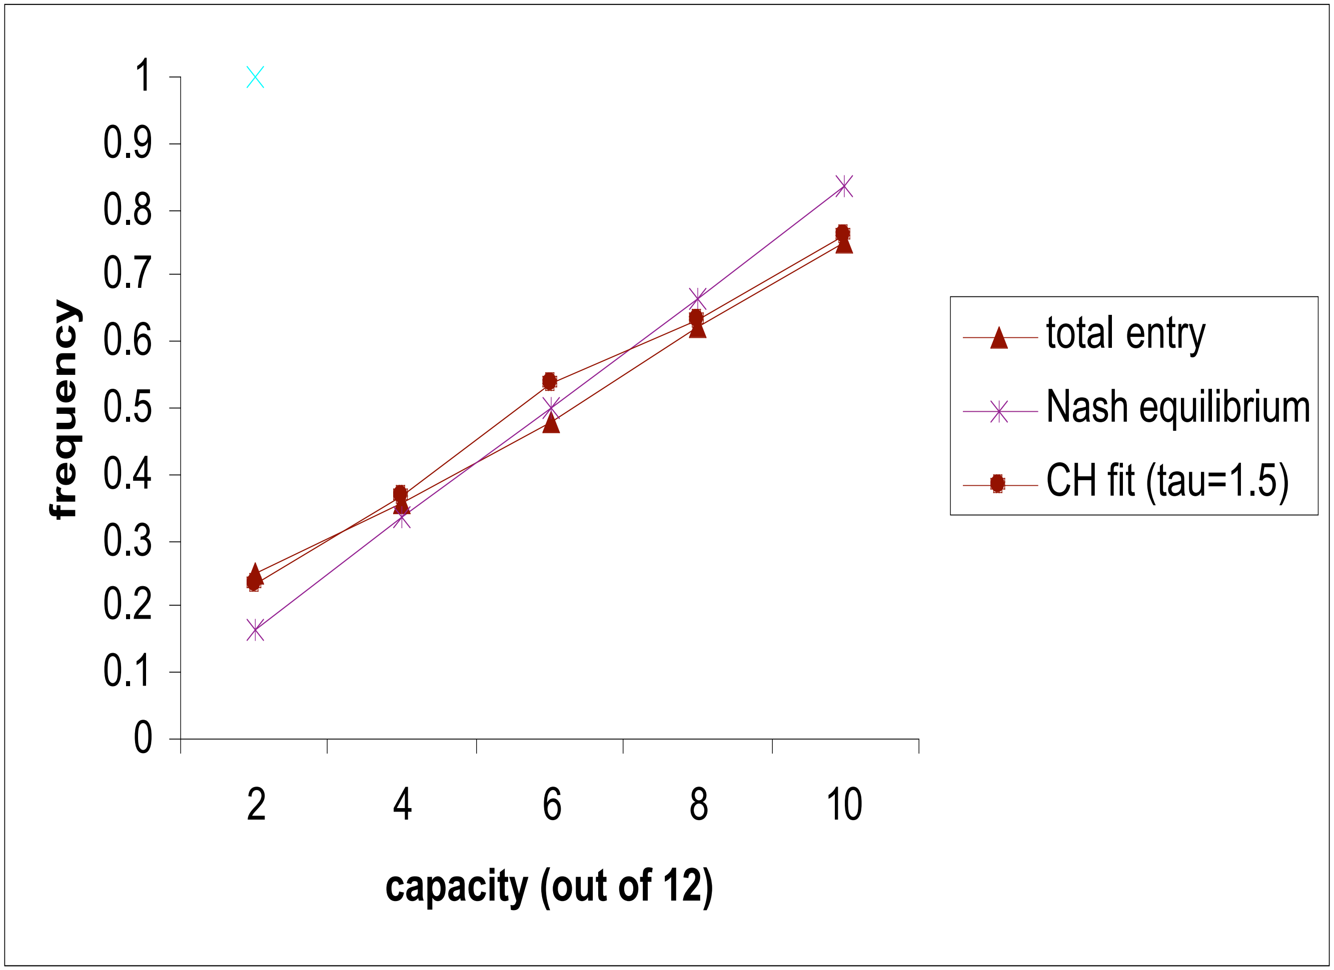
frequency
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
2
4
6
8
10
capacity (out of 12)
total entry
Nash equilibrium
CH fit (tau=1.5)

The Mating Game

- Mate-for-life Game

Male \ Female	Commit	Pass (search)
Commit	(M_m, M_f)	$(S_m - R_m, S_f)$
Pass (search)	$(S_m, S_f - R_f)$	(S_m, S_f)

Males and Females evaluate each other via a costly (possibly multi-stage) search process. When both commit, they get payoffs (M_m, M_f). If neither commits, then both get the cost of continued search. If only one commits, then there is also a rejection cost.

Mate Payoff related to mate quality

$$M_f = f_q(q_m)$$

$$q_m = g(attract, Intell, status, altruism, \ldots)$$

Search cost related to aspiration level

$$S_j = f_s(aspiration_j)$$

Rejection cost, a function of
(Own perceived quality, other's quality)

$$R_j = f_r(\hat{q}_j, q_{other})$$

Player j's type (what defines strategy)

$$T_j = Strategy(\{\hat{q}_j, M_f(q_f), S_j, R_j\})$$

Perceiving Mate Quality

$$p(q_{other} \mid \{cue_1, \cdots, cue_n\}) = \prod_{k=1}^{n} p(cue_k \mid q_{other})\, p(q_{other})$$

$$\begin{aligned} \log p(q_{other} \mid \{cue_1, \cdots, cue_n\}) &= \sum_{k=1}^{n} \log p(cue_k \mid q_{other}) + \log p(q_{other}) \\ &\approx \sum_{k=1}^{n} w_k \hat{q}^k_{other} + \log p(q_{other}) \end{aligned}$$

where $\hat{q}^k_{other} = \arg\max_{q_{other}} \left(\log p(cue_k \mid q_{other}) \right)$

$$p(q_{other})$$

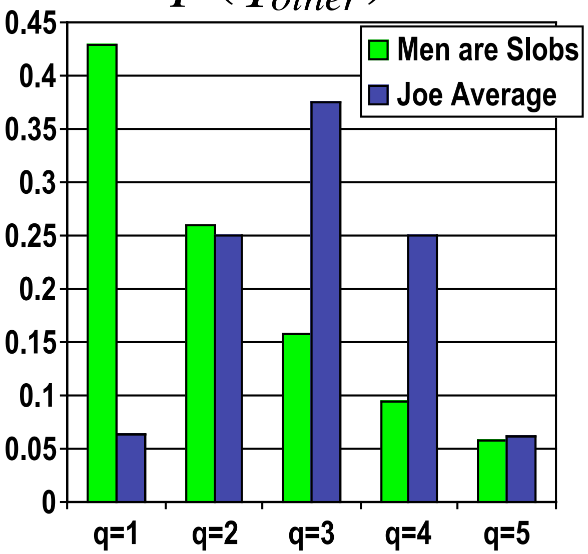

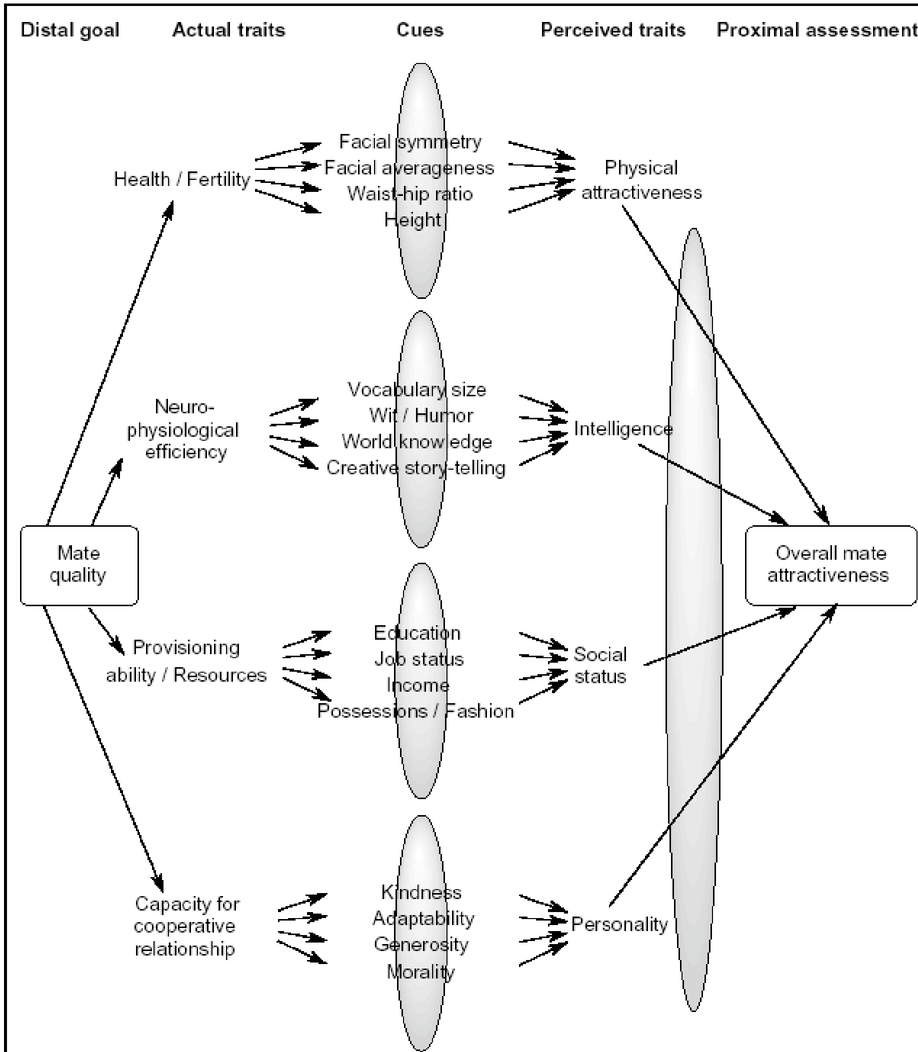

Fig. 1 A two-stage lens model (expanded from ideas in Refs 21,47) for assessing distal mate quality by integrating perceived sexual cues. The 16 cues and 4 traits depicted here are just possible examples. A sexual prospect's actual mate value ('Distal goal') is segmented into objective trait values ('Actual traits'), each of which might be advertised by several sexual cues (middle). Each cue has some correlation with the underlying trait, and hence with actual mate value. According to this kind of model, when judging a sexual prospect, an individual integrates these cue values into trait attributions ('Perceived traits'), and then integrates these into an assessment of overall attractiveness ('Proximal assessment'). While the traits are likely to be evolutionarily stable, the cues may be culture-specific. Each stage of judgment could use various linear or nonlinear cue-integration mechanisms.

Setting Aspiration

Sequential Mating Game

Commit to:

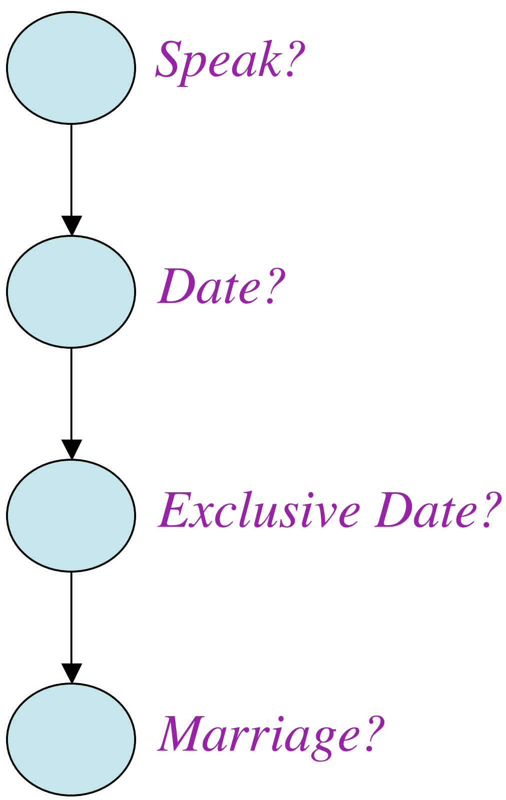

Aspiration can be set differently for each subgame

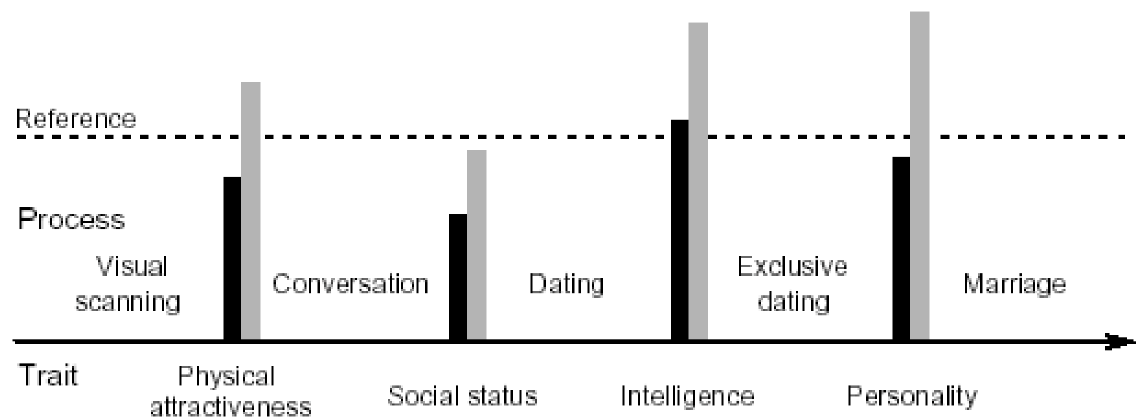

Fig. 2 A sequential aspiration-level model for mate choice. Time proceeds from left to right, over days or weeks. An individual assesses a sexual prospect's trait values in sequence, from those most easily perceived through those that take longest to judge. Individuals assign an aspiration level to each of four traits (the same examples used in Fig. 1), and pursue further courtship only with prospects who exceed each successive aspiration level. These aspiration-level 'hurdles' can be set at different heights depending on the individual's sexual strategy and mating goals, and can be set relative to one's own trait values, or relative to a population percentile value (as in Ref. 22). Two aspiration levels are shown for each trait, based on Kenrick's data (Ref. 22). Typical male aspiration levels for a short-term dating partner (shorter, black bars) and a long-term marriage partner (taller, grey bars) are shown relative to a reference value (broken line). Note that, unlike the lens model of Fig. 1, no integration of trait values into an overall mate value assessment is necessary. The assessment order here is idealized; in real life, some aspects of personality and intelligence can be judged before some aspects of bodily attractiveness (because of clothing) or social status (because of deceptive advertising).

Strategy search: How to set aspiration?

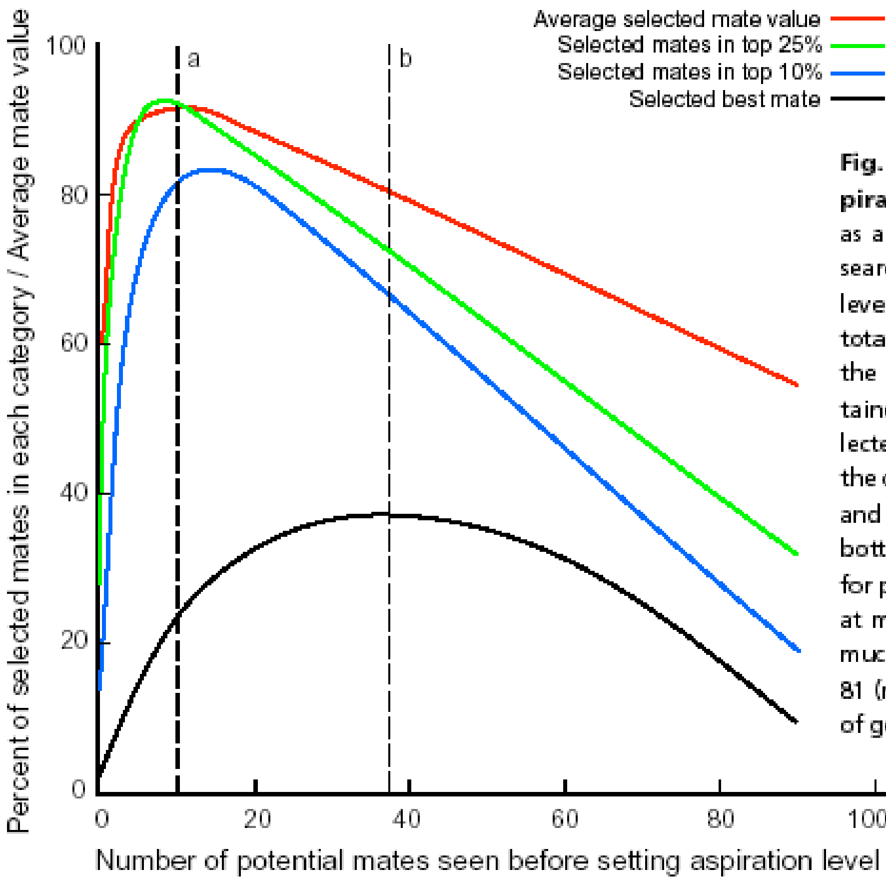

Fig. 3 Search performance based on number of individuals sampled to form an aspiration level for further mate search. Performance on four different criteria is plotted as a function of number of individuals sampled initially (assuming one-sided, non-mutual search), with the highest-value individual from that initial sample used as the aspiration level for future search. These smoothed results are based on simulations (Refs 64,65) with a total mating population of 100, each individual having a mate value between 1 and 100. If the initial sample comprises 10 individuals (x=10, marked a), the average mate value obtained given the resulting aspiration level will be 92 (red curve, top); the chance that the selected mate will be in the top 25% of the population is 92% (green curve, second from top); the chance that the selected mate will be in the top 10% is 82% (blue curve, third from top); and the chance that the very highest-value mate will be selected is 23% (black curve, bottom). Sampling about 10 individuals is also a fast way to form a useful aspiration level for population sizes larger than 100. In contrast, the '37% rule' (x=37, marked b) does better at maximizing the chance of obtaining the very highest-valued mate (to 37%), but has a much higher search cost, and does worse on the other three criteria [average mate value of 81 (red); only a 72% chance of getting a mate in the top 25% (green); only a 67% chance of getting a mate in the top 10% of the population (blue)].

How long to sample before you know the mate quality distribution?

Other Games

- Trust game:

Players: Investor (I) & Trustee (T)

I & T both receive \$S.

I can invest \$y =\$0 to \$S with T. The experimenter then triples the amount sent, so gets 3y

T then returns \$z=\$0 to 3y

Payoffs:

$P_I = S - y + z$

$P_T = S + 3y - z$

Nash Eq?

$P_I(y)$; y>0 only if z>y

$P_T(z)$; if y>0 then best is z=0

Hence N.E. is 0,0

Which y amount is empirically best for investor

Average: y = \$5.16

z = \$4.66

Figure 3: Amounts sent and paid back in trust experiment, (1995) Berg, Dickhaut & McCabe

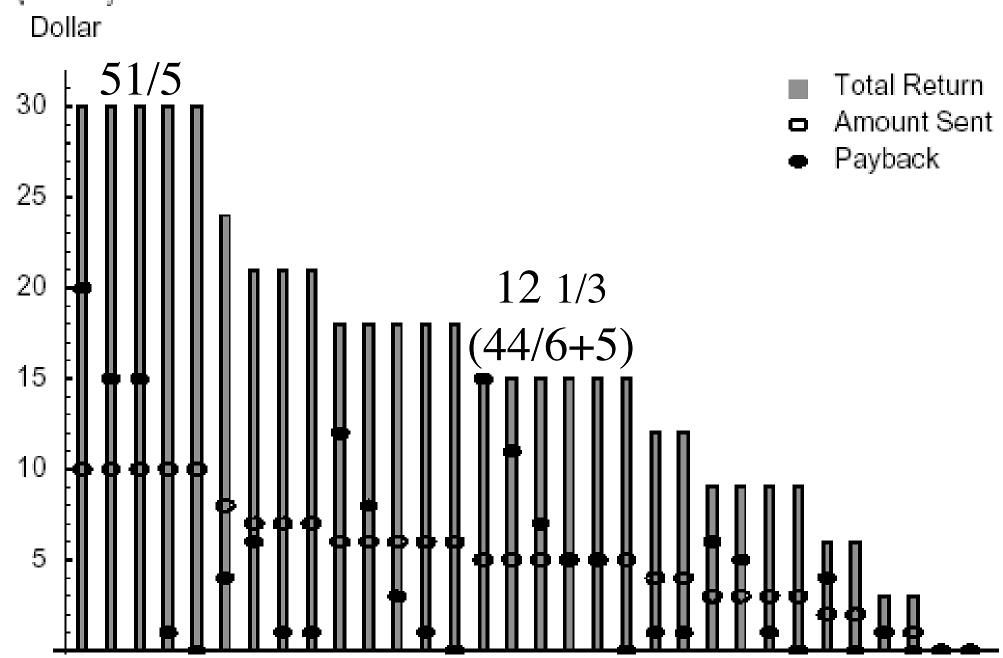

Trust games

- Typically 10-15% trustees give nothing back
- 5-15% invest nothing
- Typically trust is underpaid by about 10% (Bolle,1995)
- However,
  - Koford(1998) Bulgaria study (country with low trust in authority and high fraud rates, with most students cheating on exams, and Professors accepting bribes for grades).
  - Investors average 69% investment
  - Trustees return 50% over investment
  - Bulgarians trust each other?

Trust across countries

countries	direct	group (foursome)	society	mean
	fraction invested			
American-Chinese	.76	.49	.49	.54
Japanese-Korean	.51	.48	.28	.41
mean	.64	.48	.39	.47
	fraction returned			
American-Japanese	.28	.13	.11	.15
Chinese-Korean	.41	.25	.18	.25
mean	.35	.19	.15	.20

Americans - give trust but not reciprocated
Chinese - do the best over all
Japanese/Korean - do worse than expected from sociological speculations about their family structure.

What's going on?

Binary choice variant: Snider and Kerens, 1996

Investor \ Trustee	Repay trust	Don't repay
Don't trust	*(P,P)*	*(P,P)*
Trust	*(R,R)*	*(S,T)*

P = initial payment
R = reciprocity payment
T = Selfish Trustee payoff
S = Sucker!
where S<P<R<T

Varied payoffs, looked at two variables derived from social Preference model with "guilt" and "regret" factors. Key theoretical variables should be:

Trustee Temptation: (T-S)/(T-R) *(high when pays to keep)*
Investor Risk: (P-S)/(R-S) (*high when loss is large relative to gain)*

These two variables appear to account for all changes in subject behavior with changes in payoff.

Taking trust games to the workplace

Work or Shirk!

6 firms
8 workers

Firms post offers anonymously in random order and workers accept or reject them.

If accept, then worker chooses an effort level that is fixed, and both get their payoffs. Proceed for 10-20 rounds.

Incentive condition:

Fines for e < criterion, detected with p=1/3

Worker j chooses effort level e_j=[0.1-1]

Firm i offers wage w_i

Payoffs: $P_{firm}=(q-w)e$

$P_{firm}=(w)c(e)$, $c(e)$ convex

Figure 4: Actual effort-rent relation in the absence and presence of explicit performance incentives

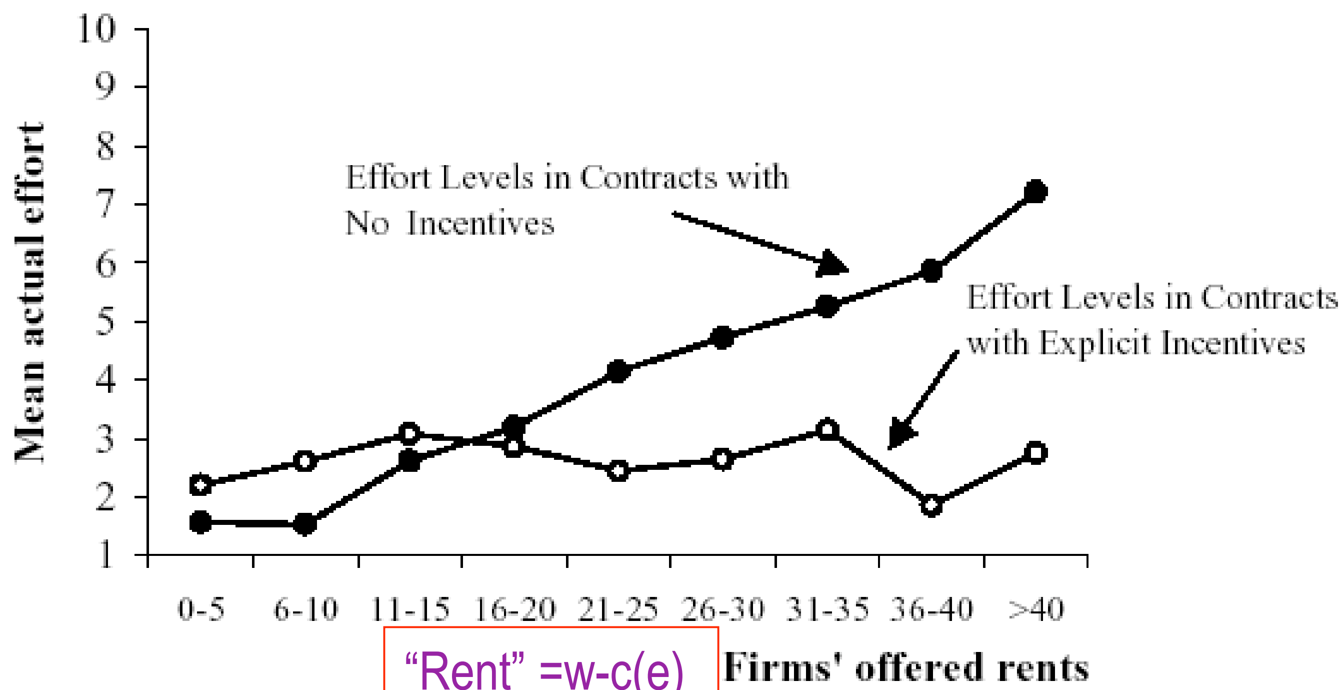

Source: Fehr and Gächter (2000).