AN EXPERIMENTAL STUDY OF COGNITION AND BEHAVIOR IN NORMAL-FORM GAMES

Miguel Costa-Gomes (UC San Diego and Harvard)
Vincent Crawford (UC San Diego)
Bruno Broseta (University of Arizona and OPVI)

February 2000

● Two (too many!) noncooperative game theories: equilibrium analysis (deductive) and adaptive learning models (inductive)

● These theories differ mainly in assumptions about *strategic sophistication*, the extent to which players' decisions reflect attempts to predict others' decisions, taking incentives into account: equilibrium analysis assumes it is unlimited, while adaptive learning models assume it is nonexistent or severely limited

● The persistence of such different theories suggests the importance of combining theory with empirical evidence; the differences might not affect the set of possible limiting outcomes (Nash equilibria), but they affect equilibrium selection, speed of convergence, and responses to changes in the environment

● Empirical studies of strategic behavior need to control the environment, for which experiments have important advantages

● In studying sophistication experiments have a further advantage, making it possible to study sophistication more directly by observing subjects' searches for hidden payoff information

● Camerer et al. studied backward induction in extensive-form alternating-offers bargaining games, using computer interface called MouseLab that allows subjects to search for hidden pie sizes

● We study sophistication in normal-form games, using MouseLab to allow subjects to look up own and partner's payoffs for each possible decision combination as often as desired, one at a time

● There are close connections between equilibrium analyses of *strategies* in normal- and extensive-form games, but their cognitive foundations are very different; the different presentation of payoff information in normal-form games yields a much larger space of possible look-up patterns, which allows us to study different aspects of sophistication in series of games with various patterns of dominance, iterated dominance, and unique pure-strategy equilibria

● Main goals:

To use subjects' information searches, in the light of the cognitive implications of alternative theories of their behavior, to identify subjects' decision rules more precisely

To learn the extent to which individual subjects' deviations from the information search implications of equilibrium analysis help to predict their deviations from equilibrium decisions

● Main results:

Subjects' decisions are highly heterogeneous, too varied to describe by any single decision rule, even allowing for errors

Subjects' information searches are even more heterogeneous than their decisions, and generally confirm the interpretation of behavior suggested by their decisions, with some differences

Subjects reveal a surprising degree of sophistication, but not enough to justify unqualified use of equilibrium analysis; most sophistication is best described by boundedly rational rules

Incorporating the cognitive implications of decision rules into the analysis allows a unified account of decisions and information search, and reveals systematic relationships between them that allow better estimates of rules and better predictions of decisions

Experimental design

● Structure follows largely from the goal of studying sophistication:

As always, the design must control the strategic environment so that the results can be interpreted for clearly identified games

In studying sophistication it is important to observe subjects' *initial* responses to each game they play, because learning can mimic sophistication and we would lose control over information search if subjects could recall current payoffs from earlier plays

Varying the games also helps to prevent preconceptions about their strategic structures, and to identify subjects' decision rules

Main (Baseline or "B") treatment (2 runs, 21 and 24 subjects)

● Game-theoretically naïve subjects randomly grouped into Rows and Columns from "large" population, repeatedly, anonymously paired to play a common, randomly ordered series of 18 two-person games, with different games and partners each period

● Each subject faced a mix of games with varying structures: 5 with simple dominance, 9 with kinds of iterated dominance, and 4 with unique pure-strategy equilibria but no (pure-strategy) dominance

● To control preferences, each subject was paid in proportion to his payoff in one randomly selected game; payments averaged about $15 per hour; no face-to-face interaction; and identities confidential

● To control information, the structure of the environment was made public at the start except for the game payoffs, to which subjects were given unlimited access via MouseLab screen display

● To suppress learning, subjects were given no feedback about partners' decisions or own or partners' payoffs during experiment

● All subjects framed as Row players and called "You"; abstract decision labels, random orders of decisions, and lack of observable differences in player roles limit framing effects (still left-right bias)

● Payoffs are numbers of "points" from 1-99, normally hidden in "closed boxes"; move cursor into box and left-click to look up payoff; must close box by right-clicking before opening next box; decisions made and confirmed in boxes below payoffs; cursor automatically moves to top-center at start of each game

	S/He: &	**S/He: @**	**S/He: &**	**S/He: @**
You: #				
You: *				
	Your Points		**Her/His Points**	
	You: #		**You: ***	

● To ensure comprehension, subjects given instructions on screen and via handout, allowed questions, and required to participate in four unpaid practice rounds and pass an Understanding Test

● MouseLab automatically records subjects' look-up sequences (boxes opened), *gaze times* (look-up durations), and decisions

Maintained assumptions about effects of the design

● Public knowledge justifies focus on complete-information theories

● Lack of feedback and variations in games suppress learning, and random pairing in large population suppresses repeated-game effects, justifying separate, static analyses of game-subject pairs

● Framing allows focus on structural theories of behavior, like traditional game theory

Subsidiary treatments

Open Boxes or "OB" (1 run, 27 subjects)

- OB identical to Baseline except that the games were presented via Mouselab with all payoffs continuously visible (in "open boxes")

- Comparing OB and Baseline results reveals whether presentation via MouseLab affects decisions *per se*; we usually find somewhat higher equilibrium compliance in OB, but not statistically significant

Trained Subjects or "TS" (1 run, 15 subjects)

- TS identical to Baseline except: subjects were taught dominance, iterated dominance, and pure-strategy equilibrium-checking (but *not* information search) on the screen; did not interact; and were rewarded only for correctly identifying their equilibrium decisions

- Comparing TS and Baseline results reveals if Baseline deviations from equilibrium are due to cognitive limitations and provides a benchmark for information searches; in TS we find high equilibrium compliance even in complex games and information search different than Baseline subjects who make equilibrium decisions

Decision rules ("types")

- We structure our analysis by assuming each subject's behavior is determined in all games by one of nine possible decision rules or "types"; these are general principles of decision-making chosen for appropriateness as possible descriptions of subjects' behavior, theoretical interest, and separation of behavioral implications

- Each type is naturally associated with algorithms that describe how to process payoff information into decisions; using these as models of cognition allows us to describe subjects' behavior in a comprehensible way without overfitting or over-constraining the analysis, and links their decisions and information searches so we can identify relationships between them

● We allow four nonstrategic types and five strategic types:

Altruistic, which maximizes the sum of its own and other's payoffs over all possible decision combinations

Pessimistic ("Maximin"), which (without randomizing) maximizes its minimum payoff over other's possible decisions

Naïve (Stahl and Wilson's *L1*), which best responds to beliefs that assign equal probabilities to other's decisions

Optimistic ("Maximax"), which maximizes its maximum payoff over other's decisions

L2 (a relative of S&W's *L2*), which best responds to *Naïve*

D1 (*D2*), which does one round (two rounds) of deleting decisions that are dominated by pure decisions and then best responds to a uniform prior over other's remaining decisions

Equilibrium (a relative of S&W's *Naïve Nash*), which makes its equilibrium decision, unique in all of our games

Sophisticated (a relative of S&W's *Rational Expectations*), which best responds to the probability distributions of others' decisions, as estimated from the observed population frequencies in our experiment (depends on the data)

● All of our strategic types exhibit some sophistication, in that their decisions reflect attempts to predict others' decisions; *Sophisticated* represents the ideal of a player who can predict others' decisions, included to learn if any subjects have a prior understanding that transcends mechanical rules like *L2*, *D1*, *D2*, and *Equilibrium*

● Our 18 games were chosen to separate nonstrategic from strategic types as much as possible, give strong incentives to follow types' decisions, and avoid artificial clarity of payoffs and structures

Figure II. Games

2A (1,2)	A,P,N	D12,L2,E,S
A	72,93	31,46
D	84,52	55,79

2B (1,2)	A,P,N	D12,L2,E,S
D	94,23	38,57
A	45,89	14,18

3A (2,1)	D	A
D12,L2,E,S	75,51	42,27
A,P,N	48,80	89,68

3B (2,1)	D	A
A,P,N	21,92	87,43
D12,L2,E,S	55,36	16,12

4A (2,1)		D	A
A,P,N	59,58	46,83	85,61
D12,L2,E,S	38,29	70,52	37,23

4B (2,1)	A	D
D12,L2,E,S	31,32	68,46
P	72,43	47,61
A,N	91,65	43,84

4C (1,2)	D12,L2,E,S	A,P,N
	28,37	57,58
A	22,36	60,84
D	51,69	82,45

4D (1,2)	D12,L2,E,S	P	A,N
D	42,64	57,43	80,39
A	28,27	39,68	61,87

5A (3,2)	A,P,N	D12,L2,E,S
A	53,86	24,19
P,N,D1,L2,S	79,57	42,73
D2,E	28,23	71,50

5B (3,2)	A,P,N	D12,L2,E,S
A	76,93	25,12
D2,E	43,40	74,62
P,N,D1,L2,S	94,16	59,37

6A (2,3)	A	D2,E,S	P,N,D1,L2
D12,L2,E,S	21,26	52,73	75,44
A,P,N	88,55	25,30	59,81

6B (2,3)	D2,E	A	P,N,D1,L2,S
A,P,N	42,45	95,78	18,96
D12,L2,E,S	64,76	14,27	39,61

7A (∞,∞)	N,D12,L2,S	A,P	E
L2,E,S	87,32	18,37	63,76
A,P,N,D12	65,89	96,63	24,30

7B (∞,∞)	N,D12,L2,S	A,P	E
A,P,N,D12	67,91	95,64	31,35
L2,E,S	89,49	23,53	56,78

8A (∞,∞)	L2,E,S	A,P,N,D12
E	72,59	26,20
A,P	33,14	59,92
N,D12,L2,S	28,83	85,61

8B (∞,∞)	L2,E,S	A,P,N,D12
A,P	46,16	57,88
E	71,49	28,24
N,D12,L2,S	42,82	84,60

9A (1,2)	D12,L2,E,S	A,P,N
	22,14	57,55
	30,42	28,37
A	15,60	61,88
D	45,66	82,31

9B (2,1)			A	D
A,P,N	56,58	38,29	89,62	32,86
D12,L2,E,S	15,23	43,31	61,16	67,46

Aggregate analysis of decisions

● No significant differences between aggregate decisions in Baseline and OB, or for Rows and Columns in isomorphic games; highly significant differences between Baseline and TS

● In Baseline and OB, high equilibrium compliance in simplest games, falling below random in most complex games, where it depends on 2-3 rounds of iterated dominance or equilibrium logic

● In TS, high equilibrium compliance even in most complex games, so low equilibrium compliance in Baseline is not due to cognitive limitations or the difficulty of looking up payoffs via MouseLab

Table II. Percentages of Decisions that Comply with Equilibrium by Type of Game

Type of Game (rounds of dominance)	Baseline	OB	B+OB	TS
2x2 with dominant decision (1) (2A, 2B for Rows; 3A, 3B for Cols.)	85.6% (77/90)	92.6% (50/54)	88.2% (127/144)	100.0% (24/24)
2x3 with dominant decision (1) (4D for Rows; 4B for Cols.)	82.2% (37/45)	100.0% (27/27)	88.9% (64/72)	100.0% (12/12)
3x2 with dominant decision (1) (4C for Rows; 4A for Cols.)	86.7% (39/45)	92.6% (25/27)	88.9% (64/72)	100.0% (12/12)
4x2 with dominant decision (1) (9A for Rows; 9B for Cols.)	88.9% (40/45)	96.3% (26/27)	91.7% (66/72)	100.0% (12/12)
2x2, partner has dominant decision (2) (3A, 3B for Rows; 2A, 2B for Cols.)	61.1% (55/90)	79.6% (43/54)	68.1% (98/144)	95.8% (23/24)
2x3, partner has dominant decision (2) (4A for Rows; 4C for Cols.)	62.2% (28/45)	63.0% (17/27)	62.5% (45/72)	100.0% (12/12)
3x2, partner has dominant decision (2) (4B for Rows; 4D for Cols.)	60.0% (27/45)	55.6% (15/27)	58.3% (42/72)	83.3% (10/12)
2x4, partner has dominant decision (2) (9B for Rows; 9A for Cols.)	73.3% (33/45)	70.4% (19/27)	72.2% (52/72)	100.0% (12/12)
2x3 with 2 rounds of dominance (2) (6A, 6B for Rows; 5A, 5B for Cols.)	62.2% (56/90)	68.5% (37/54)	64.6% (93/144)	100.0% (24/24)
3x2 with 3 rounds of dominance (3) (5A, 5B for Rows; 6A, 6B for Cols.)	11.1% (10/90)	22.2% (12/54)	15.3% (22/144)	87.5% (21/24)
2x3, unique equilibrium, no dominance (7A, 7B for Rows; 8A, 8B for Cols.)	50.0% (45/90)	51.9% (28/54)	50.7% (73/144)	91.7% (22/24)
3x2, unique equilibrium, no dominance (8A, 8B for Rows; 7A, 7B for Cols.)	17.8% (16/90)	27.8% (15/54)	21.5% (31/144)	91.7% (22/24)

Maximum likelihood error-rate analysis of decisions

● We wish to infer individual subjects' types by comparing their decisions over all 18 games with our types' decisions

● *Naïve* and *Optimistic* always make the same decisions, so we lump them together for now; *L2* and *Sophisticated* decisions are separated, weakly, in only one game for Columns, but we include both because we pool the data for Rows and Columns; any two other types make different decisions in at least 2/18 games for each player role, with strategic and nonstrategic types strongly separated

● Maximum-likelihood error rate analysis (El-Gamal and Grether, Harless and Camerer), using a mixture model in which subjects' types are drawn from a common prior distribution; in each game a subject's type determines his information search, with error, and type and information search then determine his decision, with error

● A type-*k* subject normally makes type *k*'s decision, but in each game there is a probability $\varepsilon_k \in [0, 1]$, type *k*'s *error rate*, that he makes an error, and makes each of *c* decisions with probability $1/c$; probability of type-*k* decision is $1 - (c - 1)\varepsilon_k/c$ and of any non-type *k* decision is ε_k/c; given *k*, errors are i.i.d. across games and subjects

● $k = 1,\ldots,K$ and $i = 1,\ldots,N$ index types and subjects; c=2,3,4 is the number of subject *i* 's decisions; T^c is the number of games with *c* decisions (the same for all subjects in our design); $p \equiv (p_1,\ldots,p_K)$ are the prior type probabilities; $x_k^i \equiv (x_{2k}^i, x_{3k}^i, x_{4k}^i)$, $x^i \equiv (x_1^i,\ldots,x_K^i)$, and $x \equiv (x^1,\ldots,x^N)$, where x_k^{ic} is the number of subject *i*'s decisions that equal type *k*'s in games where he has *c* decisions; and $\varepsilon \equiv (\varepsilon_1,\ldots,\varepsilon_K)$ are the types' error rates

● The probability of observing a particular sample with $x_k^i \equiv (x_{2k}^i, x_{3k}^i, x_{4k}^i)$ type-*k* decisions when subject *i* is type *k* is:

$$L_k^i(\varepsilon_k \mid x_k^i) = \prod_{c=2,3,4} [1-(c-1)\varepsilon_k / c]^{x_k^{ic}} [\varepsilon_k / c]^{T^c - x_k^{ic}}$$

● Weighting by p_k, summing over *k*, taking logarithms, and summing over *i* yields the log-likelihood for the entire sample:

$$\ln L(p,\varepsilon \mid x) = \sum_{i=1}^{N} \ln \left[\sum_{k=1}^{K} p_k \prod_{c=2,3,4} [1-(c-1)\varepsilon_k / c]^{x_k^{ic}} [\varepsilon_k / c]^{T^c - x_k^{ic}} \right]$$

● The model has 15 independent parameters: 7 type probabilities and 8 type-dependent error rates; maximum likelihood estimation yields consistent parameter estimates

● A type-*k* decision is evidence for type *k*, but only to the extent that estimated ε_k suggests it was more likely than a non-type *k* decision

Econometric results for decisions

● Subjects' decisions are highly heterogeneous, with a great deal of sophistication: 72-80% of Baseline-OB subjects estimated to be strategic types; many are also estimated to be *Naïve/Optimistic*

● Few are estimated to be *Equilibrium* and none *Sophisticated* (*L2* is weakly separated from *Sophisticated*, but its better predictions in one game are amplified into a large lead by lower error rates)

● Most sophistication is best described by boundedly rational types *L2* and *D1*, which respect two rounds of pure-strategy dominance (*L2* also respects two rounds of mixed-strategy dominance), because their decisions tend to switch from *Equilibrium* to *Naïve* in our most complex games, like most of our subjects' decisions

Table III. Parameter Estimates for OB and Baseline Subjects (— vacuous)

	Decisions Alone (*Naïve* and *Optimistic* Parameters Constrained Equal)		Decisions Alone with Compliance-conditional Error Rates	Decisions & Information Search
Treatment (log-likelihood)	**OB (-246.44)**	**B (-446.39)**	**B (-433.23)**	**B (-852.02)**
Type				
Altruistic p_k	**0.000**	**0.044**	**0.089**	**0.022**
ζ_{kj}, *j=H,M,L,0*	—	—	(0.04,0.02,0.36,0.57)	0.89,0.00,0.00,0.11
ε_k or ε_{kj}, *j=H,M,L,0*	—	0.253	0.26,0.63,0.79,0.82	0.00,—,—,0.66
Pessimistic p_k	**0.000**	**0.000**	**0.000**	**0.045**
ζ_{kj}, *j=H,M,L,0*	—	—	—	0.47,0.00,0.00,0.53
ε_k or ε_{kj}, *j=H,M,L,0*	—	—	—	0.60,—,—,1.00
Naïve p_k	**0.199**	**0.240**	**0.227**	**0.448**
ζ_{kj}, *j=H,M,L,0*	—	—	(0.97,0.02,0.01,0.01)	0.95,0.01,0.01,0.02
ε_k or ε_{kj}, *j=H,M,L,0*	0.285	0.286	0.24,0.43,0.58,0.81	0.50,0.39,0.47,0.85
Optimistic p_k	**0.199**	**0.240**	**0.000**	**0.022**
ζ_{kj}, *j=A,0*	—	—	—	1.00,0.00
ε_k or ε_{kj}, *j=A,0*	0.285	0.286	—	0.29,0.50
L2 p_k	**0.344**	**0.496**	**0.442**	**0.441**
ζ_{kj}, *j=H,M,L,0*	—	—	(0.88,0.07,0.02,0.03)	0.87,0.04,0.03,0.06
ε_k or ε_{kj}, *j=H,M,L,0*	0.233	0.203	0.18,0.35,0.21,0.21	0.25,0.61,0.16,0.22
D1 p_k	**0.298**	**0.175**	**0.195**	**0.000**
ζ_{kj}, *j=H,M,L,0*	—	—	(0.44,0.12,0.06,0.38)	—
ε_k or ε_{kj}, *j=H,M,L,0*	0.276	0.704	0.43,0.63,0.15,1.00	—
D2 p_k	**0.000**	**0.000**	**0.000**	**0.000**
ζ_{kj}, *j=H,M,L,0*	—	—	—	—
ε_k or ε_{kj}, *j=H,M,L,0*	—	—	—	—
Equilibrium p_k	**0.160**	**0.044**	**0.052**	**0.000**
ζ_{kj}, *j=H,M,L,0*	—	—	(0.00,0.08,0.75,0.17)	—
ε_k or ε_{kj}, *j=H,M,L,0*	0.165	0.163	—,0.41,0.00,0.97	—
Sophisticated p_k	**0.000**	**0.000**	**0.000**	**0.022**
ζ_{kj}, *j=H,M,L,0*	—	—	—	0.00,0.00,0.71,0.29
ε_k or ε_{kj}, *j=H,M,L,0*	—	—	—	—,—,0.54,1.00

● Except for *D1* in the Baseline, whose error rate is 70%, error rates for the five types with positive estimated probabilities range from 16-29%, low for initial responses to abstractly framed games; probability of types' decisions (except *D1*'s) ranges from 0.79-0.92

● Posterior probabilities for an uninformative prior over parameter vector place at least 0.90 on one type for 74-84% of subjects

● Widespread sophistication does not justify unqualified use of equilibrium analysis: equilibrium compliance in complex games is 11-52%, far less than estimated 72-80% sophisticated subjects

● High equilibrium compliance in simple games suggests this isn't irrationality, and TS and OB results suggest it isn't due to cognitive limitations or presentation via Mouselab; our analysis of decisions suggests that it is due to the prevalence of boundedly rational types *L2* and *D1*; a possible alternative explanation is a high frequency of *Sophisticated* subjects that is not common knowledge: a puzzle for analysis of information search and further experiments to resolve

Cognition and information search

● Each type is naturally associated with algorithms that describe how to process payoff information into decisions; using these as models of subjects' cognitive processes links their decisions and information searches so we can identify relationships between them

● The algorithms require mainly pairwise payoff comparisons; we call single payoffs *look-ups* and operations on pairs *comparisons*

● Different types require different look-ups and comparisons, so observing information search may allow inferences about types; but the inferences depend on how cognition affects information search

● Problems:

Very little theory, many possible look-up sequences, and our subjects' sequences are noisy and highly heterogeneous

We cannot directly observe comparisons, and if a subject scanned and memorized all payoffs before thinking, the order of look-ups would be unrelated to cognition; thus inferences depend on subjects using repeated look-ups rather than memory

● We make two simple assumptions about how cognition affects information search, suggested by Camerer et al.'s control and our TS treatment; they impose minimal restrictions on a type's look-ups and comparisons, to avoid arbitrarily imputing inconsistency to subjects whose cognitive processes we cannot observe

Occurrence: For a given type in a given game, each look-up in a minimal set needed to identify the type's decision appears at least once in the subject's look-up sequence

Adjacency: For a given type in a given game, Occurrence is satisfied and each comparison in a minimal set needed to identify the type's decision is represented by an adjacent look-up pair (or group) at least once in the subject's look-up sequence

● Both Occurrence and Adjacency are conditional on type

● Occurrence is uncontroversial, of limited use because subjects make so many look-ups it is too likely to be satisfied by chance

● Adjacency is satisfied if subjects perform comparisons one at a time via adjacent look-ups, relying on repeated look-ups rather than memory; more controversial, but has more discriminatory power

● Illustrate types' implications under Occurrence and Adjacency in game 3A (Column has dominant decision, dominance-solvable in two rounds, nonstrategic Rows pick B and strategic Rows pick T)

	S/He: L	S/He: R	S/He: L	S/He: R
You: T	**75**	**42**	**51**	**27**
You: B	**48**	**89**	**80**	**68**
	Your Points		**Her/His Points**	
	You: T		**You: B**	

● *Altruistic* compares totals of own and other's payoffs for each decision combination; Occurrence requires all own and other's look-ups and Adjacency requires comparisons (75,51), (42,27), (48,80), and (89,68)

● *Naïve* compares expected payoffs of own decisions given uniform prior over other's, via a set of either up-down or left-right own payoff comparisons; Occurrence requires look-ups 75, 48, 42, and 89 and Adjacency requires either the set of comparisons {(75,42), (48,89)} *or* the set of comparisons {(75,48), (42,89)}

● *Optimistic* compares maximal payoffs of own decisions, scanning own payoffs in any order; can eliminate some if find payoff higher than the maximum for previously checked decisions; Occurrence requires look-ups 75, 42, and 89 and Adjacency is vacuous

● *Pessimistic* compares minimal payoffs of own decisions by left-right comparisons of own payoffs; can eliminate some if find payoff lower than previously identified minimum for another decision; Occurrence requires look-ups 48, 89, and 42 and Adjacency requires comparison (48,89)

	S/He: L	S/He: R	S/He: L	S/He: R
You: T	75	42	51	27
You: B	48	89	80	68
	Your Points		Her/His Points	
	You: T		You: B	

● If *Equilibrium* has a dominant decision it needs only to identify it; if not it can use iterated dominance or equilibrium-checking, either decision combination by combination or "best-response dynamics"; Occurrence requires look-ups 51, 27, 80, 68, 75, and 48 and Adjacency requires comparisons (51,27), (80,68), and (75,48)

● If *Sophisticated* has a dominant decision it needs only to identify it; if not it needs to form beliefs (identifying the strategic structure) and compare the expected payoffs of undominated decisions; Occurrence requires all own and other's look-ups and Adjacency requires the comparisons *Equilibrium* requires plus (42,89) to check for own dominance and compare own decisions' expected payoffs

● *L2* needs to identify other's *Naïve* decision and *L2*'s best response to it; Occurrence requires all other's look-ups plus 75 and 48, the own look-ups for other's *Naïve* decision, and Adjacency requires either the set of comparisons {(51,27), (80,68)} or the set of comparisons {(51,80), (27,68)} to identify other's Naïve decision, plus the comparison (75,48) to identify *L2*'s best response

● In games solvable by two rounds of iterated dominance, as here, Occurrence and Adjacency are the same for *D1* and *D2* as for *Equilibrium*, because iterated dominance is the best way to identify *Equilibrium* decisions; this yields almost the same look-ups and comparisons as for *L2*; by contrast, in our games without pure-strategy dominance *D1* and *D2* Occurrence and Adjacency require the look-ups and comparisons needed to check for other's dominance plus those for *Naïve*, which are very different from *L2*

Aggregate analysis of information search

● We incorporate types' Occurrence and Adjacency implications by evaluating compliance for each game-subject pair and sorting it into four categories, indexed by $j = H$, M, L, or 0: B_H (100% Occurrence, 67-100% Adjacency), B_M (100% Occurrence, 34-66% Adjacency), B_L (100% Occurrence, 0-33% Adjacency), and ~A (<100% Occurrence); we now distinguish *Naïve* and *Optimistic*, and to avoid bias, we treat *Optimistic* A as the union of vacuous B_H, B_M, and B_L

● Table VI summarizes types' implications for 13 search measures, under Occurrence and Adjacency, and our experimental results for those measures, aggregated for TS and Baseline and for Baseline subjects sorted by most likely type estimated from decisions alone

● There is strong separation of implications across three groups of types: *Altruistic*; *Pessimistic*, *Naïve*, or *Optimistic*; and *L2*, *D1*, *D2*, *Equilibrium*, or *Sophisticated*; there is also some separation within groups, e.g. *L2* from *D1* for Own Look-Ups versus Other Look-Ups

● TS subjects have more look-ups, longer string lengths, and shorter gaze times than Baseline, suggesting more systematic analyses; TS subjects also have more own up-down and other's left-right transitions; TS and Baseline both have more other's left-right than own up-down look-ups, suggesting some left-right bias

● The observed search measures for Baseline subjects are higher than the theoretical lower bounds in the top of Table VI, but they vary across estimated types in rough proportion to the bounds

● *Altruistic* has more own and other's look-ups, shorter string lengths, fewer own-to-own and other's-to-other's transitions, and more Altruistic own-to-other's transitions

● All types but *Altruistic*, particularly *Naïve/Optimistic*, have more own than other's look-ups; and all have longer own than other's gaze times; own payoff first exceeds 58% for every type, and 70% for all but *L2*; and own payoff last exceeds 54% for all but *Altruistic*

● Table VII summarizes aggregate compliance with Occurrence and Adjacency in TS and Baseline, and in Baseline with subjects sorted by most likely type estimated from decisions alone

● TS and Baseline subjects differ very sharply in compliance with *Equilibrium*, *D1*, and *D2* Occurrence and Adjacency; *Sophisticated* and *Pessimistic* Adjacency; and, less sharply, *L2* Occurrence

● Occurrence alone doesn't discriminate well among types, because most subjects usually comply with it for most types; but Adjacency discriminates well even in the aggregate: note approximate diagonal dominance in bottom part of Table VII (minor exceptions for *Naïve/Optimistic* and larger exceptions for *D1*)

● The contrast with the sharp separation of TS from Baseline subjects in the top of Table VII suggests important differences between TS subjects and "naturally occurring" *Equilibrium* Baseline subjects (Baseline subjects whose most likely type is *Equilibrium*)

Table VI. Implications of Types and Aggregate Look-up Measures for TS and Baseline Subjects, and Baseline Subjects by Most Likely Type Estimated from Decisions Alone (— vacuous)

Type	Own Look-Ups	Other Look-Ups	Own String Length	Other String Length	Own Gaze Time	Other Gaze Time	Own Payoff First	Own Payoff Last	Own-Own Trans.	Other-Other Trans.	Own Up-Dn. Trans.	Other L.-Rt. Trans.	Altr. Own-Oth. Trans.
Implications of Types													
Altruistic	≥5.8	≥5.8	≤1.82	≤1.82	Long	Long	—	—	≤45%	≤45%	—	—	≥10%
Pessimistic	≥3.9	—	≥1.82	—	Long	Short	≥50%	≥50%	≥45%	—	≤31%	—	≤10%
Naïve	≥5.8	—	≥1.82	—	Long	Short	≥50%	≥50%	≥45%	—	≈31%	—	≤10%
Optimistic	≥5.8	—	≥1.82	—	Long	Short	≥50%	≥50%	≥45%	—	—	—	≤10%
L2	≥2.4	≥5.8	≥1.82	≥1.82	Long	Long	≤50%	≥50%	≥45%	≥45%	≥31%	≈31%	≤10%
D1	≥4.6	≥2.0	≥1.82	≥1.82	Long	Long	—	≥50%	≥45%	≥45%	≥31%	≥31%	≤10%
D2	≥4.5	≥2.4	≥1.82	≥1.82	Long	Long	—	≥50%	≥45%	≥45%	≥31%	≥31%	≤10%
Equilibrium	≥4.1	≥3.6	≥1.82	≥1.82	Long	Long	—	≥50%	≥45%	≥45%	≥31%	≥31%	≤10%
Sophisticated	≥5.8	≥4.2	≥1.82	≥1.82	Long	Long	—	≥50%	≥45%	≥45%	≥31%	≥31%	≤10%
Random	—	—	1.82	1.82	—	—	50%	50%	45%	45%	31%	31%	10%
Aggregate Look-up Measures													
TS	19.0	15.7	6.88	7.33	0.60	0.45	68.3%	83.9%	84.2%	81.6%	63.3%	69.3%	5.1%
All Baseline	16.8	14.6	5.46	5.95	0.67	0.60	72.8%	78.5%	79.7%	77.5%	31.6%	42.9%	8.5%
Altruistic	24.4	26.5	2.20	2.26	0.48	0.44	88.9%	33.3%	33.5%	38.0%	21.0%	60.0%	36.8%
Pessimistic	—	—	—	—	—	—	—	—	—	—	—	—	—
Naïve/Optim.	13.7	8.4	6.76	6.03	0.82	0.69	96.0%	77.3%	84.9%	80.5%	21.1%	48.3%	4.9%
L2	18.0	17.2	5.80	7.13	0.59	0.52	58.5%	87.9%	84.7%	83.0%	39.4%	30.3%	6.2%
D1	14.6	12.8	3.74	3.73	0.81	0.76	70.6%	54.8%	70.4%	68.3%	28.3%	61.7%	14.5%
D2	—	—	—	—	—	—	—	—	—	—	—	—	—
Equilibrium	18.4	13.4	4.05	3.67	0.55	0.51	100.0%	72.2%	72.0%	69.5%	21.5%	79.0%	5.3%
Sophisticated	—	—	—	—	—	—	—	—	—	—	—	—	—

Table VII
Aggregate Rates of Compliance with Types' Occurrence and Adjacency for TS and Baseline Subjects, and for Baseline Subjects by Most Likely Type Estimated from Decisions Alone, in percentages (— vacuous)

Treatment (# subjects)	Altruistic *j*=H,M,L,0	Pessimistic *j*=H,M,L,0	Naïve *j*=H,M,L,0	Optimistic *j*=A,0	L2 *j*=H,M,L,0	D1 *j*=H,M,L,0	D2 *j*=H,M,L,0	Equilibrium *j*=H,M,L,0	Sophisticated
TS (12)	3,10,50,27	44,7,36,13	83,2,0,15	86,14	76,2,0,22	92,3,1,5	92,3,1,5	96,1,1,3	75,1,1,24
Baseline (45)	14,11,51,24	74,2,11,14	78,4,4,14	85,15	67,14,5,14	52,19,15,14	50,19,15,14	42,23,19,16	39,21,20,21
Altruistic (2)	78,6,11,6	56,8,33,3	53,3,42,3	97,3	47,8,39,6	36,6,56,3	33,8,56,3	31,11,56,3	28,14,56,3
Pessimistic (0)	—,—,—,—	—,—,—,—	—,—,—,—	—,—	—,—,—,—	—,—,—,—	—,—,—,—	—,—,—,—	—,—,—,—
Naïve/Opt. (11)	9,5,53,33	85,1,9,5	89,5,3,4	96,4	42,24,3,31	45,22,20,13	43,18,23,16	26,24,28,23	23,23,27,27
L2 (23)	8,12,58,22	72,2,9,17	78,3,0,18	80,20	85,6,3,6	57,20,9,15	54,21,10,15	49,24,12,15	46,22,12,20
D1 (7)	23,21,26,29	59,3,16,23	63,7,6,23	77,23	53,21,6,21	48,17,14,20	45,19,15,21	42,20,17,21	38,14,21,27
D2 (0)	—,—,—,—	—,—,—,—	—,—,—,—	—,—	—,—,—,—	—,—,—,—	—,—,—,—	—,—,—,—	—,—,—,—
Equilibrium (2)	6,8,86,0	100,0,0,0	97,3,0,0	100,0	64,36,0,0	69,17,14,0	67,19,14,0	56,25,19,0	53,19,28,0
Sophisticated (0)	—,—,—,—	—,—,—,—	—,—,—,—	—,—	—,—,—,—	—,—,—,—	—,—,—,—	—,—,—,—	—,—,—,—

Maximum likelihood error-rate analysis of decisions and information search

● Because type determines information search with error, and type and search determine decision with error, we allow a subject's deviations from his type's decisions and searches in a given game to be correlated, but we assume that, given *c* and *k*, the deviations are i.i.d. across games and subjects

● Our model of decisions and information search allows a general joint probability distribution, except that it constrains how subjects' deviation probabilities vary with *c* and it assumes that *j* is the only aspect of search compliance that matters; our model of decisions with compliance-contingent error rates ignores information search, except that it makes error rates depend on compliance

● Let ζ_{kj} be the *un*conditional probability that a type-*k* subject has type-*k* information search compliance *j* in a game

● A subject of type *k* normally makes type *k*'s decision, but in each game there is a conditional probability $\varepsilon_{kj} \in [0, 1]$, type *k*'s *error rate with compliance j*, that he makes an error, and makes each of his *c* decisions with probability 1/*c*; the probability of type *k*'s decision is $1-(c-1)\varepsilon_{kj}/c$ and of any non-type *k* decision is ε_{kj}/c

● $\zeta_k \equiv (\zeta_{kH}, \zeta_{kM}, \zeta_{kL}, \zeta_{k0})$, $\zeta \equiv [\zeta_{kj}]$, $\varepsilon_k \equiv (\varepsilon_{kH}, \varepsilon_{kM}, \varepsilon_{kL}, \varepsilon_{k0})$, and $\varepsilon \equiv [\varepsilon_{kj}]$. T_{kj}^{ic} is the number of games in which subject *i* has *c* decisions and type-*k* compliance *j*, x_{kj}^{ic} is the number of such games in which subject *i* also makes type *k*'s decision, $\bar{T}_k^i \equiv [T_{kj}^{ic}]$, $\bar{T}^i \equiv (\bar{T}_1^i, \ldots, \bar{T}_K^i)$, $\bar{T} \equiv (\bar{T}^1, \ldots, \bar{T}^N)$, $\bar{x}_k^i \equiv [x_{kj}^{ic}]$, $\bar{x}^i \equiv (\bar{x}_1^i, \ldots, \bar{x}_K^i)$, and $\bar{x} \equiv (\bar{x}^1, \ldots, \bar{x}^N)$.

● The probability of observing a particular sample with T_{kj}^{ic} and x_{kj}^{jc} when subject *i* is type *k* is:

$$L_k^i(\varepsilon_k,\zeta_k \mid \bar{x}_k^i,\bar{T}_k^i)=\prod_j \prod_{c=2,3,4} \zeta_{kj}^{T_{kj}^{ic}}[1-(c-1)\varepsilon_{kj}/c]^{x_{kj}^{ic}}[\varepsilon_{kj}/c]^{T_{kj}^{ic}-x_{kj}^{ic}}$$

● The log-likelihood function for the entire sample is:

$$\ln L(p,\varepsilon,\zeta \mid \bar{\mathrm{x}},\bar{T})=\sum_{i=1}^{N}\ln\left[\sum_{k=1}^{K} p_k \prod_j \prod_{c=2,3,4} \zeta_{kj}^{T_{kj}^{ic}}[1-(c-1)\varepsilon_{kj}/c]^{x_{kj}^{ic}}[\varepsilon_{kj}/c]^{T_{kj}^{ic}-x_{kj}^{ic}}\right]$$

● In this log-likelihood, information search and decisions work together to distinguish types; the model of decisions with compliance-contingent error rates simply omits the ζ_{kj} terms

● The model of decisions and information search has 67 independent parameters: 8 type probabilities, 34 (4 compliance categories for each of 9 types, less 2 for *Optimistic*) compliance-conditional error rates ε_{kj}; and 25 unconditional compliance probabilities ζ_{kj} (34 categories less 1 $\sum_j \zeta_{kj}=1$ restriction for each type); the model with compliance-contingent error rates has 42; both yield consistent parameter estimates

● For both models, a type-*k* decision is evidence for type *k* only to the extent that ε_{kj} suggests it was more likely than a non-type *k* decision; the ε_{kj} terms in the log-likelihood favor types *k* for which the x_{kj}^{ic} are concentrated on particular *j* values, since concentration lowers estimated error rates; theory suggests that concentration should be on *high j* values, the estimates confirm that, so they favor types whose decisions occur with the "right" searches

● In the model of decisions and information search, the ζ_{kj} terms in the log-likelihood favor types *k* for which the T_{kj}^{ic}, and hence the estimated ζ_{kj}, are concentrated on particular levels of type-*k* compliance *j*, whether high or low; the theory suggests *high*, but we don't assume that: again, the estimates confirm it

● To avoid arbitrarily favoring *Optimistic* because it has fewer categories, we equalize its likelihood with that of a hypothetical type whose compliance is random across the missing categories, so that all types get equal credit for actual success predicting subjects' information search compliance

● Even so, the model of decisions and information search assigns the same meaning to a given pattern of compliance for each type; because this is a strong, untested distributional assumption, when the estimates differ we are more confident in those from our model of decisions with compliance-contingent error rates

● *Naïve* and *Optimistic* have the same decisions and almost the same Occurrence, but *Naïve* Adjacency is more restrictive; this makes the above effects favor *Naïve* to the extent that subjects' searches satisfy *Naïve* Adjacency more than randomly, and that subjects' *Naïve/Optimistic* decisions are more concentrated on high *Naïve* compliance (because *Naïve* Adjacency is then more useful than vacuous *Optimistic* Adjacency in predicting decisions)

● Adjacency and Occurrence also help to separate *L2* and *D1*, even though both respect two rounds of iterated pure-strategy dominance and make very similar decisions in our games

Econometric results for decisions and information search

● Subjects' information searches are even more heterogeneous than their decisions, and generally confirm the interpretation of behavior suggested by their decisions, with some differences

● Table III reports maximum likelihood parameter estimates for the Baseline treatment, taking search into account; for the model with compliance-contingent error rates it also reports estimates of ζ_{kj} conditional on p_k estimated from decisions, to indicate the compliance frequencies on which error rate estimates are based

● The estimates for decisions with compliance-contingent error rates are close to those for decisions alone: *L2* has the largest frequency, 44%, followed by *Naïve* at 23% and *D1* at 20%; the main difference from estimates based on decisions alone is the strong separation of *Naïve* and *Optimistic*, accomplished via differences in their Adjacency restrictions as explained above

● For this model, types' estimated error rates tend to decrease as compliance increases, but a likelihood ratio test cannot reject the hypothesis that $\varepsilon_{kj} \equiv \varepsilon_k$ for all j and k (p-value 0.20), or that ε_{kj} is weakly increasing in j for all k; this (weakly) supports the theory's implication that subjects with higher compliance make their types' decisions more often, suggesting a systematic relationship between subjects' deviations from search patterns associated with equilibrium analysis and from equilibrium decisions

● *Naive* and *L2* have high compliance, error rates that generally decrease with higher compliance, and low error rates when compliance is high; *D1* has fairly high compliance and high error rates that usually decrease with compliance; and *Altruistic* and *Equilibrium* have low compliance and error rates that usually decrease with compliance

● *Naive* and *L2* compliance is usually high enough that the implied noncompliance frequencies are lower than the corresponding error rates, supporting the interpretation that subjects made *Naive* and *L2* decisions intentionally, except for errors, and suggesting that those types' estimated frequencies are reliable; but lower compliance gives less reason for confidence in the estimated frequency of *D1* and still less in those of *Altruistic* or *Equilibrium*

● Most sophistication appears to be best described by boundedly rational types, *L2* and *D1*, which respect two rounds of dominance

● The estimates for the model of decisions and information search generally confirm the view of behavior from the other models, with some changes: *Naïve* and *L2* now have the largest frequencies, each around 45%, and *D1* has disappeared; the shift toward *Naïve*, mainly at the expense of *Optimistic* and *D1*, reflects the fact that *Naïve* compliance explains more of the variation in subjects' searches and decisions than *Optimistic*'s, which is too unrestrictive to be useful in our sample, or *D1*'s, which is more restrictive than *Naïve*'s, but less correlated with subjects' behavior

● Types' error rates again tend to decrease as compliance increases, but a likelihood ratio test cannot reject $\varepsilon_{kj} \equiv \varepsilon_k$ for all *j* and *k* (*p*-value 0.99) or that ε_{kj} is weakly increasing in *j* for all *k*

● *Altruistic*, *Naïve*, and *Optimistic* have high compliance, error rates that decrease with compliance, and low to moderate error rates when compliance is high; *L2* has high compliance but widely varying error rates; and *Pessimistic* and *Sophisticated* have high error rates that decrease with compliance

● *Altruistic*, *Naïve*, and *Optimistic* compliance is generally high enough to support the interpretation that subjects made their types' decisions intentionally, except for errors, suggesting that their estimated frequencies are reliable; but compliance gives less reason for confidence in the estimated frequency of *L2* and little reason for confidence in those of *Pessimistic* and *Sophisticated*

● With these qualifications, our analysis suggests that there are very large frequencies of *Naïve* and *L2* subjects, together making up at least five-sixths of the population; depending on confidence in our model of decisions and information search, there might also be some *D1* subjects, but there are at most traces of other 6 types

● Most sophistication still appears to be best described by boundedly rational types, *L2* and perhaps *D1*

● Posterior probabilities for an uninformative prior over the parameter vector based on model of decisions and information search place at least 0.90 on one type for 43 of 45 Baseline subjects: 19 *L2*, 19 *Naïve*, 2 *Pessimistic*, 1 *Altruistic*, 1 *Optimistic*, and 1 *Sophisticated*; observing searches allows us to assign most subjects to a single type with more confidence than decisions

● Focusing on the modal type probabilities, observing search identifies 11 *Naïve*/*Optimistic* subjects as *Naïve* (10) or *Optimistic* (1); changes 6 subjects from *L2*, 2 from *D1*, 1 from *Equilibrium*, and 1 from *Altruistic* to *Naïve*; changes 2 from *D1* and 1 from *Equilibrium* to *L2*; changes 2 from *D1* to *Pessimistic*; and changes 1 from *D1* to *Sophisticated*; sharpens the identification of 11 *L2* subjects; clouds the identification of 3 *L2* subjects; and leaves 4 posteriors (3 *L2* and 1 *Altruistic*) unchanged at 1.000

● Incorporating the cognitive implications of decision rules into an error rate analysis allows a coherent, unified account of subjects' decisions and information search, which allows better predictions