Foraging for Trust: Exploring Rationality and the Stag Hunt Game*

Steven O. Kimbrough
565 JMHH
kimbrough@wharton.upenn.edu
215-898-5133

INFORMS, 16 November 2005

These foils available online at `http://opim-sun.wharton.upenn.edu/~sok/sokpapers/2006/informs-2005-trust-foils.pdf`

*File: informs-2005-trust-foils.tex/pdf.

Too Many Foils. Basic Points:

1. Trust and cooperation are prevalent and fundamentally important social phenomena.

2. Their prevalence is inadequately explained with classical game-theoretic (and economic) models, suggesting consideration of different modeling/explanatory perspectives.

3. Agent-based models, built upon a philosophy of exploring rationality are plausible and promising in this regard.

4. Three very different models—a gridscape model, a Markov model, and a linear operator model (and there are others)—each easily and amply explain the emergence of trust and cooperation among agents.

Available Online

"Foraging for Trust: Foraging Rationality and the Stag Hunt Game,"

PDF at http://opim-sun.wharton.upenn.edu/~sok/sokpapers/2005/, then iTrust2005asPublished.pdf, Steven O. Kimbrough, in P. Herrmann et al. (Eds.): iTrust 2005, LNCS 3477, pp. 1–16, 2005.

Foils accompanying that paper: http://opim-sun.wharton.upenn.edu/~sok/sokpapers/2005/itrust/trust-foils.pdf

Cooperation: "The Cement of Society"

- Cooperation: key to social felicity (whether a good thing or not)
- Trust: a means to achieve cooperation. Assumption of risk, depending on behavior of counter-parties. One trusts that the other guy will do the right thing.
- Is it like altruism, perhaps: more apparent than real?
- Must ask: Cooperation or trust for what?
  - Usually: for mutual advantage. Pareto optimality.
  - Other or additional goals? Fairness? Justice? Team purposes? Agreed goals? Game theory has always abstracted these away.

Honor among Thieves

". . . drug trafficking still employs ethnic networks for efficiency and trust, but not exclusively. In contrast to a Mafia model, in which all transactions take place among members of a crime 'family,' the drug trade gives rise to specialties that take advantage of location, language, local knowledge, or ability to melt into the crowd. Some drug transactions rely on the trust and mutual recognition that a common ethnic background implies; others are enforced by the threat of violence. But in a world of increasing sources of supply and product destinations, a great many drug transactions are simply that—transactions." [2, page 73]

Models for Trust

- Actions that do-or-do-not involve trust are essentially strategic, and so addressable by the theory of games. (Decisions: parametric versus strategic.)
- On the strategy of modeling (in this context)
- What are the games?
- KISS: Prisoner's Dilemma and Stag Hunt

Puzzle: why is there so much of it?

Prisoner's Dilemma

	Cooperate (C)	Defect (D)
Cooperate (C)	3, 3	0, 4
Defect (D)	4, 0	1, 1

Required: $T > R > P > S$ and $2R > T + S$.

	Cooperate (C)	Defect (D)
Cooperate (C)	R, R′	S, T′
Defect (D)	T, S′	P, P′

Stag Hunt

	Hunt stag (S)	Chase hare (H)
Hunt stag (S)	4, 4	0, 3
Chase hare (H)	3, 0	1, 1

Generally: $R > T \geq P > S$. Three Nash equilibria in the one-shot game.

	Hunt stag (S)	Chase hare (H)
Hunt stag (S)	R, R′	S, T′
Chase hare (H)	T, S′	P, P′

Widely Shared Worries about Neoclassical Models

Among the main ones:

- Heroic, implausible assumptions (e.g., common knowledge, unlimited computational capacity)

- Too many equilibria (Folk Theorem), and neglect of repetition/iteration

- Poor track record in predicting behavior

  E.g., behavioral game theory

Background: Rationality Box

Hooked up to the decision maker and guarantees rationality (à la Rational Choice Theory). Given a decision problem, it

- Obtains a consistent set preferences from the decision maker
- Collects all relevant data
- Calculates what are the rational choices

In strategic contexts (games), it assumes that the counter-players all have and use Rationality Boxes.

Suppose Such Rationality Boxes Existed

Is there anything left to investigate in decision making, parametric or strategic?

1. Plants, bacteria, and lower animals

   Even if *homo economicus* is built with an innate Rationality Box, "Birds do it, bees do it, even monkeys up in trees do it" without Rationality Boxes. "It," of course, being engage in strategic interaction. We want to understand them and the properties that emerge from their micro-behavior.

   How would we study this, absent Rational Choice Theory? Investigate plausible and otherwise interesting rules for governing strategic behavior.

Coping without Rationality Boxes (con't.)

As always in science, look for testable explanations and for general principles.

In short, the agent-based modeling paradigm is broadly applicable here. Agent-based models typically:

(a) Are procedural (and constructive) rather than equational
(b) Are bottom up rather than top down
(c) Are distributed rather than centralized
(d) Are stochastic and probing rather than deterministic
(e) Are adaptive and learning (meliorizing) rather than maximizing (optimizing)

NB. Simulation is typical, by necessity, but analytic results are always welcome!

Coping without Rationality Boxes (con't.)

2. Given rationality violations in humans . . .

   By now amply documented experimentally and observationally, perhaps we best think of humans as lower animals vis à vis the race of Rationality Box owners.

   Empirical question: which affords better explanations of human behavior: Rationality Boxes or agent-based models akin to those that work for plants, bacteria, and lower animals?

Coping without Rationality Boxes (con't.)

3. Interactions between Rationality Box owners and those without it

   Rationality Boxes, recall, don't work unless everyone has one.

   How would one investigate? Again, agent-based models would be a prime candidate.

Coping without Rationality Boxes (con't.)

4. Failures of the Rationality Box to provide optimal results

   In a "rationally hurtful" situation (Schick, *Making Choices*) Rationality Box players will get hurt. They will be in Sen's phrase, "rational fools."

   There will be circumstances in which clever (think: Odysseus) players will jointly prefer to do without their Rationality Boxes. Some cases:

   (a) Repeated games in which the Rationality Box can make no useful recommendation because (Nash) equilibria proliferate.
   (b) Definitely repeated Prisoners' Dilemma, for which Rationality Boxes counsel universal defection.
   (c) Generally, situations in which the achievement of trust would be mutually beneficial

Coping without Rationality Boxes (con't.)

5. When Rationality Boxes are not fully functional

   Think of decision theory, and game theory, as attempts to design the Rationality Box. When the theory is incomplete (as in practice it will always be), investigation by other means of areas not covered is apt.

Evaluation criteria for Agent-Based Models of Individual Learning and Behavior

1. Does well against itself
2. Does well against other regimes that do well against themselves
3. Is not catastrophically exploitable
4. Robust under perturbations of its parameters
5. Can be parameterized in such a way that an agent can learn (easily) profitable, well-performing settings
6. Is computationally tractable (ideally, simple)
7. Relies on plausibly available information

Now to some example models.

Aside. . .

Why is it even attractive to use equilibrium as an—let alone *the*—evaluation criterion (solution concept) in games?

Given at least a degree of rationality by the players, why not use maximization (or at least meliorization) as a (main) solution concept? Predict that the agents will learn to maximize their returns, individually.

That's really what's up here. The above evaluation criteria articulate and adorn this idea.

Suggestion: If the context is maximum taking, then classical game theory applies and equilibrium analysis pertains. If the context is maximum seeking, then strategic maximization or meliorization pertains, and equilibrium may or may not result.

2D Models: The lattice or gridscape

6$\times$6 gridscape:

	1	2	3	4	5	6
1						
2		NW	N	NE		
3		W	X	E		
4		SW	S	SE		
5						
6						

X plays eight neighbors (Moore neighbohood). Update rule: *imitate-the-best.* As Skyrms reports [3, chapter 3], S typically takes over.

Materials and results from the conference paper

	1	2	3	4	5	6	7	8
1								
2		S	S	S				
3		S	S	S				
4		S	S	S				
5								
6								
7								

(a)
Generation
x

	1	2	3	4	5	6	7	8
1			S					
2		S	S	S				
3	S	S	S	S	S			
4		S	S	S				
5			S					
6								
7								

(b)
Generation
x+1

Continuing

	1	2	3	4	5	6	7	8
1		S	S	S				
2	S	S	S	S	S			
3	S	S	S	S	S			
4	S	S	S	S	S			
5		S	S	S				
6								
7								

(c) Generation x+2

Analytically

- Given imitate-the-best, then
- if we get a 3×3 or larger block of Ss, and
- if $5(R - P) > 3(T - S)$ (true when R=4, T=3, P=1, S=0), then
- the chance of takeover by S is quite high. (See paper for details.)

In sum on spatial (network) models for Stag Hunt

- There are known variations and much to be investigated, but . . .
- At least some models show that under reasonable conditions trust in Stag Hunt can spontaneously arise and be maintained.
- Key factors are:

  (a) payoff/reward structure — which is neglected in the classical theory

  (b) social structure — also classically neglected
- Results for Prisoner's Dilemma, too.

Stag Hunt: Initialization (black=hunt stag)

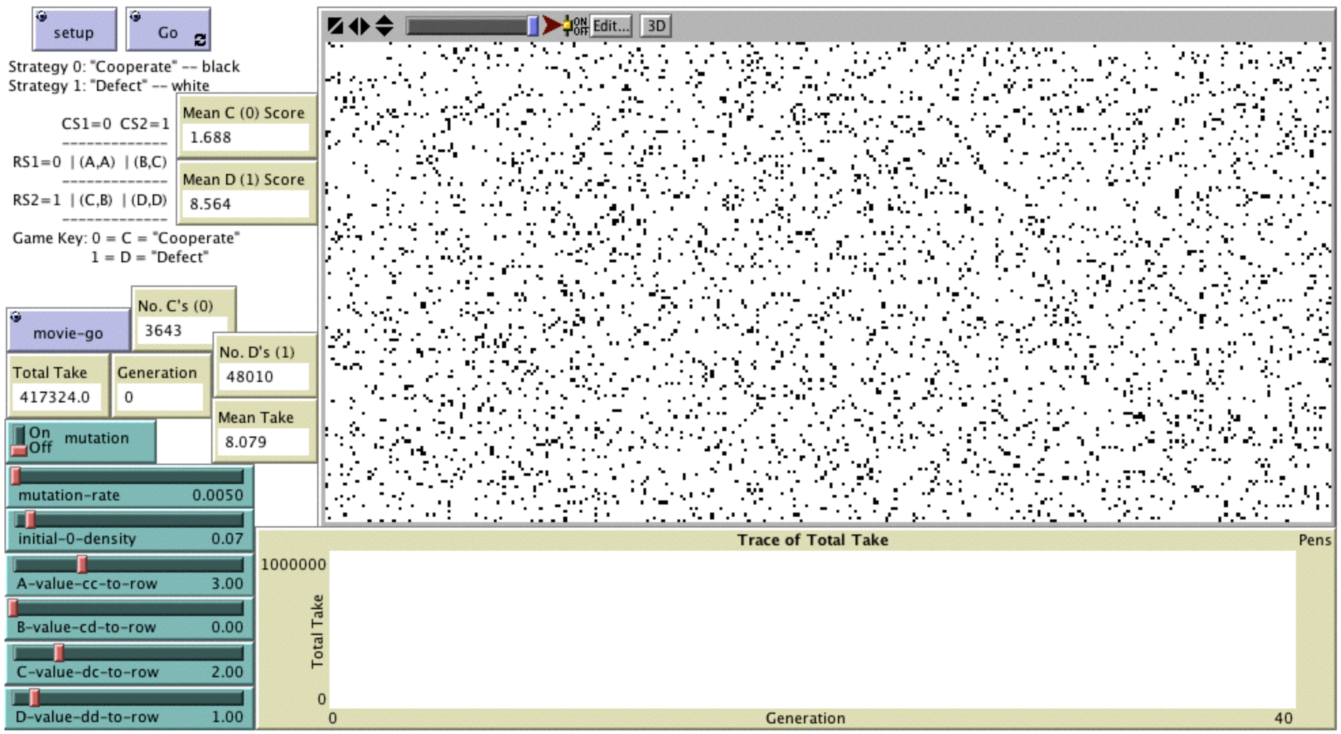

Stag Hunt: Take-Over (black=hunt stag)

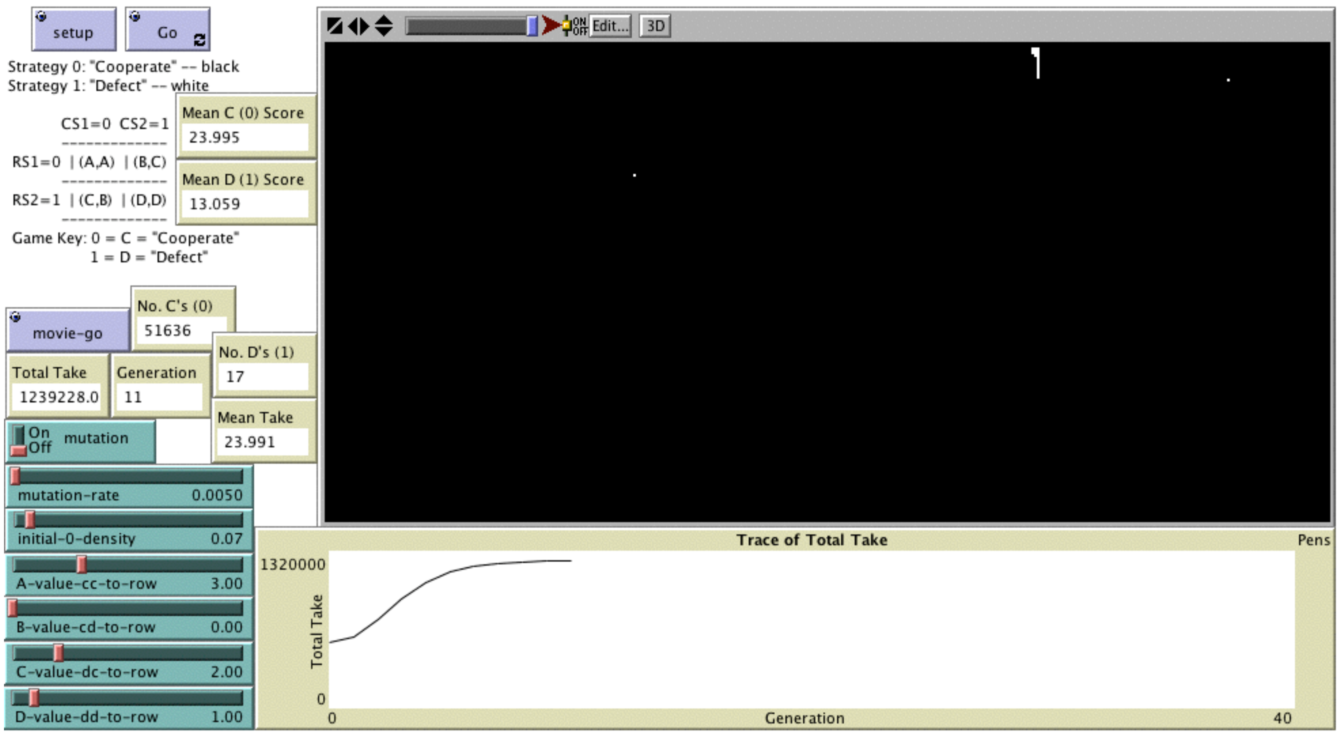

Prisoners' Dilemma: Take-Over (black=cooperate)

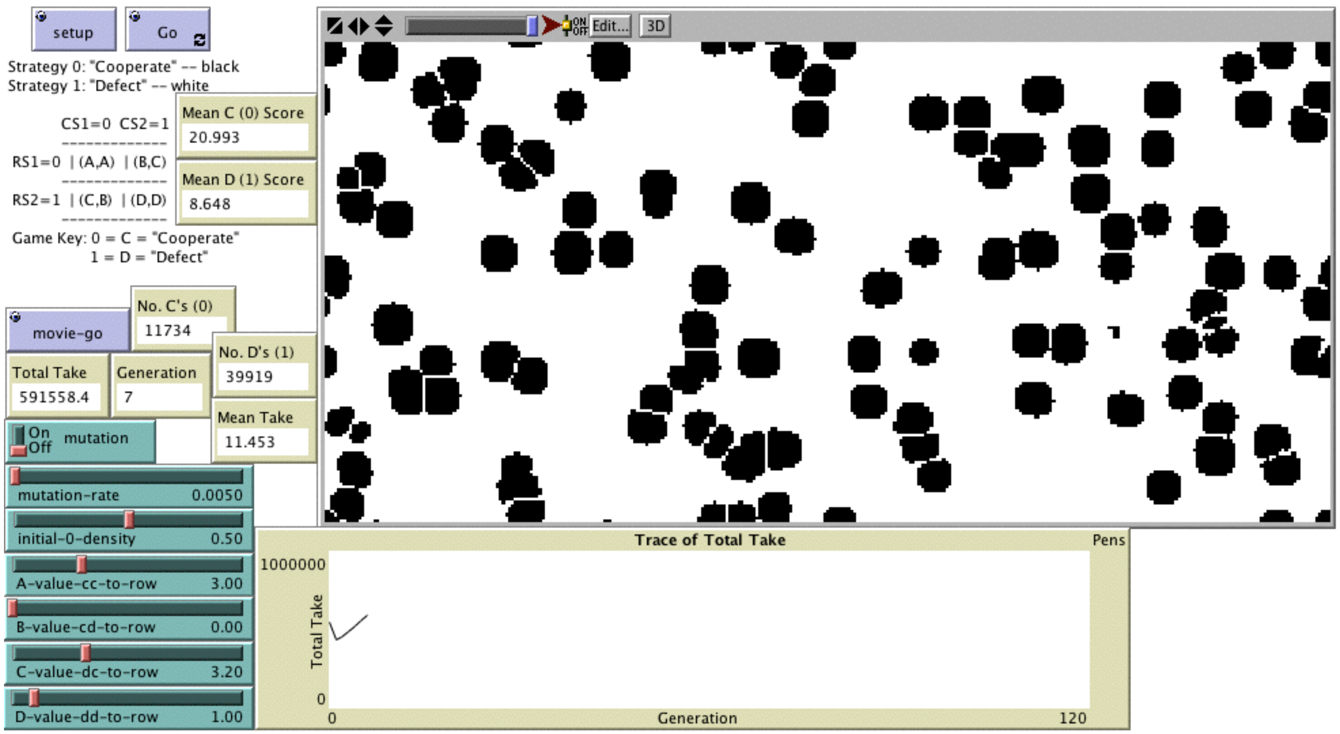

Prisoners' Dilemma: Take-Over (black=cooperate)

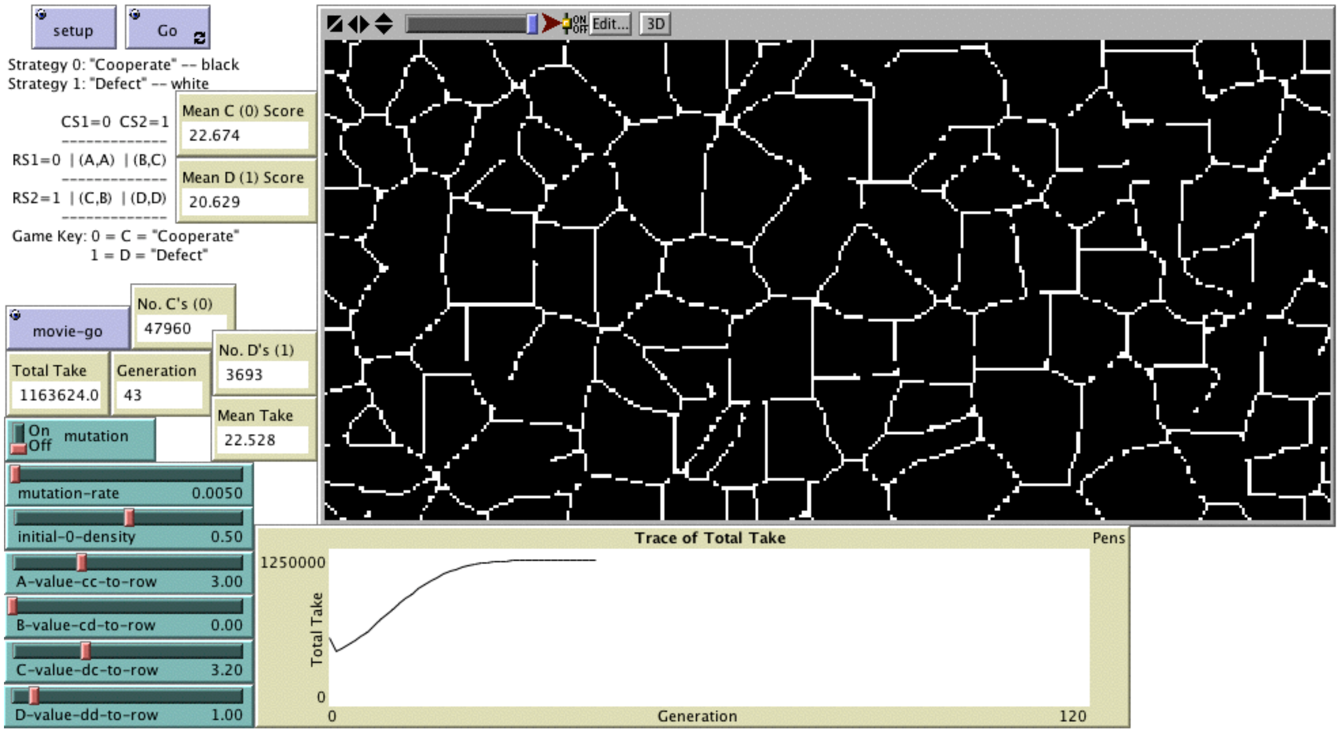

Now, models of individual learning

- Replicator dynamics (and most other mainstream models): adaptation, or evolution, but not learning.

- Network models: imitation learning; limited but not to be neglected. Richer models have begun to be explored.

  Upcoming, LPS models: generalize and abstract earlier models:

- Reinforcement learning

- MLPS, a Markov model

- Models of matching and other adaptive learning processes

Utility theory; rational choice theory

Begins by assuming a set Ω of outcomes and a preference relation $\succeq$ on Ω.

Posits for all $a, b, c \in \Omega$

1. Totality: $a \succeq b$ or $b \succeq a$.

2. Transitivity: if $a \succeq b$ and $b \succeq c$, then $a \succeq c$.

Sounds good. Why not?

Foraging context

n patches each of which may be tried for a reward. At the next time step the agent may stay or move to another patch. NB: multi-armed bandits.

Problem: where does $\succeq$ come from?

It is determined by context-specific conditions. The agent's challenge is less to reason about $\succeq$ than to discover it, to ascertain it with sufficient precision to support effective action.

$\Rightarrow$ Utility theory, rational choice theory, is a tool of minimal use in a foraging context.

Rational choice theory is appropriate for choosing from Ω given a well-articulated and supported $\succeq$.

Exploring rationality

- A *theory of exploring rationality* is needed and appropriate for choice when $\succeq$ is poorly known.

- What would that look like?

  Compare maxims of rationality:

  1. "It is rational to have transitive preferences."
  2. "In the presence of incomplete knowledge, is it rational to probe one's environment, sampling for useful information."

  Each maxim commands consideration, but in distinct contexts.

- Illustration: a multi-armed bandit problem in which a casino customer must sample the slot machines for the best chance of net positive payoff.

Contrasting summary

- EU or Rational Choice theory is a theory of *maximum taking* (MT) among presented and well-evaluated options.

- A theory of exploring rationality is a theory of *maximum seeking* (MS) among incompletely known options.

  Inherently a trial-and-error, probing, exploratory process.

- Rational Choice theory and exploring rationality theories address different problems.

Q: Why do you hate Hollywood? A: [Robert Altman] I don't hate Hollywood. They make shoes. I make gloves. We're in different businesses.

LPS: Learning in Policy Space, pseudo code

repeat forever:

1. Select a policy $\pi_i \in \Pi$, where Π is the consideration set of policies.
2. Pick a length of play, l, for policy π_i.
3. Play the next l rounds of the game using π_i.

   Note: At each round, π_i will observe the current state, s_t, take an action a and obtain a reward r_t.
4. Update V^{π_i} based on the individual-round rewards, r_ts, obtained during the l rounds of play of policy π_i. Loop.

MLPS: A model for exploring rationality

- MLPS: Markov learning in policy space. See the conference paper for details.
- Basic setup:
  - The supergame: individual rounds of play of Stag Hunt repeated indefinitely, between two players.
  - The supergame is divided into games: rounds of play of Stag Hunt, each of fixed length.
  - The games are divided into epochs: rounds of play of length le. Each game consists of n_e epochs.

Individual play in the MLPS model

- Each player has a *consideration set of policies for play,* $\mathcal{S}$.

  E.g., Always hunt stag, Always hunt hare, Tit for Tat, Suspicious Tit for Tat

- At the beginning of each game, each player independently picks a *focal strategy* from its $\mathcal{S}$. Choice is made by fitness-proportional selection on returns from the last game for policies in $\mathcal{S}$.

- At the beginning of each epoch in a game, each player independently picks a *policy-in-use* for the duration of the epoch. With probability $1-\varepsilon$, a player picks its current focal policy. With probability ε the player randomly chooses among the non-focal policies.

Example (from the paper)

- Each player has two policies in its $\mathcal{S}$: Always Defect and Tit for Tat. $l_e = 10$, $\varepsilon = 0.1$.

- Assume: n_e large enough, or mechanism-based coordination strong enough that expected values of returns to policies are exactly realized.

- We have then a Markov process, with state transitions occurring a the end of games. There are 4 states: (1) Row focuses on Tit for Tat & Colum focuses on Tit for Tat, (2) Row focuses on Tit for Tat & Colum focuses on Always Defect, . . .

  Note: Neither player knows what state the system is in.

The transition matrix

	s(1)=(1,1)	s(2)=(1,2)	s(3)=(2,1)	s(4)=(2,2)
s(1)	$0.7577 \cdot 0.7577$ $= 0.5741$	$0.7577 \cdot 0.2423$ $= 0.1836$	$0.2423 \cdot 0.7577$ $= 0.1836$	$0.2423 \cdot 0.2423$ $= 0.0587$
s(2)	$0.5426 \cdot 0.7577$ $= 0.4111$	$0.5426 \cdot 0.2423$ $= 0.1315$	$0.4574 \cdot 0.7577$ $= 0.3466$	$0.4574 \cdot 0.2423$ $= 0.1108$
s(3)	$0.7577 \cdot 0.5426$ $= 0.4111$	$0.7577 \cdot 0.4574$ $= 0.3466$	$0.2423 \cdot 0.5426$ $= 0.1315$	$0.2423 \cdot 0.4574$ $= 0.1108$
s(4)	$0.5426 \cdot 0.5426$ $= 0.2944$	$0.5426 \cdot 0.4574$ $= 0.2482$	$0.4574 \cdot 0.5426$ $= 0.2482$	$0.4574 \cdot 0.4574$ $= 0.2092$

Table 5: Stag Hunt transition matrix data assuming fitness proportional policy selection by both players. Numeric example for $\varepsilon = 0.1 = \varepsilon_1 = \varepsilon_2$.

At convergence of the Markov process

Pr(s(1))	Pr(s(2))	Pr(s(3))	Pr(s(4))
0.4779	0.2134	0.2134	0.0953

So 90%+ of the time at least one agent is playing $\text{T}\textsc{f}\text{T}$. They learn to trust (a lot). Note the expected take for Row per epoch by state:

1. $(1-\varepsilon)(40-31\varepsilon)+\varepsilon(12-2\varepsilon)=34.39$

2. $(1-\varepsilon)(9+31\varepsilon)+\varepsilon(10+2\varepsilon)=11.91$

3. $\varepsilon(40-31\varepsilon)+(1-\varepsilon)(12-2\varepsilon)=14.31$

4. $\varepsilon(9+31\varepsilon)+(1-\varepsilon)(10+2\varepsilon)=10.39$

The players do better playing this way

- Row (and Column) gets 2.302 per round of play (on average).
- At the (H,H) Nash equilibrium, each player gets 1 per round of play.
- If both players play AllD with ε-greedy exploration to Tit for Tat, each gets 1.039 per round of play.
- If both players play the mixed Nash equilibrium, each gets on average 2 per round of play.

An intuitive, very simple, cognitively undemanding model yields substantial trust and cooperation as an emergent property. Compare with 'rational fools'.

How robust and realistic is the MLPS model?

- Very robust, across a variety of settings and games. See paper and working paper cited.
- In detail, not realistic at all, but
  - It captures some important features (policy space, trial and error exploring)
  - Speedy convergence
  - Simulation with relaxed assumptions leads to qualitatively similar results.
- It demonstrates that simple learning procedures of a certain type can reliably produce a degree of trust and cooperation.

SLO (simple linear operator) Learning

This learning setup is fundamentally simple and natural. I will begin by describing the system for parametric models, as originally articulated by Bush and Mosteller [1], and will subsequently generalize and abstract it in various ways.

The general mathematical system, and the particular models that conform to it, may be described as simple linear (stochastic) operator (SLO) models. In a SLO model, time is discrete. At each choice point (or time), the agent has a collection of distinct actions, A_i, that it may take. The agent also has a probability, p_i, for taking each action. One of the actions is taken, according the the p_is, and an event, E_j, results. Depending on which event occurs, the agent's action probabilities are updated, and the process cycles.

The setup may be formalized as follows using stochastic linear operators for updating the action probabilities. Let:

A_i A series of r distinct actions, one of which is taken by a agent at each choice point. $i = 1, 2, \ldots, r$.

$\mathbf{p}_t$ A column vector of r elements, $p_1, p_2, \ldots, p_r$. p_i is the probability that the agent will take action A_i at the next choice point. $\mathbf{p}_0$ is the initial distribution of the action probabilities.

E_j A series of n event types, whose occurrences typically depend upon which action was taken by the agent. $j = 1, 2, \ldots, n$.

$e_{j,t}$ The event of event type E_j occurring at time t. Again: time is discrete.

$\mathbf{T}_j$ An $r \times r$ stochastic matrix (columns are probabilities adding to 1). $j = 1, 2, \ldots, n$.

Given $\mathbf{p}_t$ followed by action A_i and (crucially) event E_j, $\mathbf{p}_{t+1}$ is generated as follows:

$$\mathbf{p}_{t+1} = \mathbf{T}_j \mathbf{p}_t \tag{1}$$

$\mathbf{T}_j$ *transforms* $\mathbf{p}_t$ to yield $\mathbf{p}_{t+1}$, and does so in a linear fashion. $\mathbf{T}_j$ is said to be a *linear operator.* The full model, with n events and corresponding $\mathbf{T}_j$'s is then:

$$\mathbf{p}_{t+1} = \begin{cases} \mathbf{T}_1 \mathbf{p}_t & \text{if } E_1 \\ \mathbf{T}_2 \mathbf{p}_t & \text{if } E_2 \\ \vdots & \vdots \\ \mathbf{T}_n \mathbf{p}_t & \text{if } E_n \end{cases} \tag{2}$$

For example, if upon starting E_2 occurs, then E_3, then E_2, then E_1, we have

$$\mathbf{p}_4 = \mathbf{T}_1 \mathbf{T}_2 \mathbf{T}_3 \mathbf{T}_2 \mathbf{p}_0 \tag{3}$$

The case in which $r = 2$—when there are two possible event types that may occur, given an action A_i—is particularly simple. We let

$$\mathbf{p}_t = \begin{pmatrix} p \\ q \end{pmatrix} \tag{4}$$

Note: p and q are probabilities and $p = 1 - q$. $\mathbf{T}_j$ has the form

$$\mathbf{T}_j = \begin{pmatrix} 1 - b_j & a_j \\ b_j & 1 - a_j \end{pmatrix} \tag{5}$$

Note: both a and b are in $[0, 1]$. Thus

$$\mathbf{T}_j \mathbf{p}_t = \begin{pmatrix} (1 - b_j)p + a_j q \\ b_j p + (1 - a_j)q \end{pmatrix} \quad (j = 1, 2, \ldots n) \tag{6}$$

Define $Q_j p$ as the first element in $\mathbf{T}_j \mathbf{p}_t$:

$$Q_j p = (1 - b_j)p + a_j q \tag{7}$$

Define $\tilde{Q}_j q$ as the second element in $\mathbf{T}_j \mathbf{p}_t$:

$$\tilde{Q}_j q = 1 - Q_j p = b_j p + (1 - a_j)q \tag{8}$$

Focusing now on $Q_j p$, expression (7), we define α_j as

$$\alpha_j = 1 - a_j - b_j \tag{9}$$

Since $b_j = 1 - a_j - \alpha_j$ and $q = (1 - p)$,

$$Q_j p = (1 - b_j)p + a_j q = (1 - (1 - a_j - \alpha_j))p + a_j(1 - p) \tag{10}$$

Simplifying and rearranging the rightmost expression, we get

$$Q_j p = a_j + \alpha_j p \tag{11}$$

Expression (11) is the *slope-intercept form* for $Q_j p$. Notice that since a_j and α_j are constants, this is an expression for a straight line in variable p, which of course is a probability and lies in $[0,1]$.

A second form, called the *fixed-point form,* will be especially useful to us. We define λ_j with

$$a_j = (1 - \alpha_j)\lambda_j \tag{12}$$

Equivalently,

$$\lambda_j = \frac{a_j}{(a_j + b_j)} \tag{13}$$

We get the *fixed-point form* by using expression (12), substituting $(1-\alpha_j)\lambda_j$

for a_j in expression (11), yielding

$$Q_j p = \alpha_j p + (1 - \alpha_j)\lambda_j \tag{14}$$

which is the *fixed-point form* for $Q_j p$. Note that when $p = \lambda_j$

$$Q_j p = \alpha_j p + (1 - \alpha_j) p = p \tag{15}$$

assuming (as we shall) that $\alpha_j \in [0, 1]$.

This notation, and indeed the entire system, was developed to model a single agent engaged in parametric decisions. Our purpose in this chapter is to explore this system in the context of 2×2 repeated games between two players. A simple adjustment will suffice notationally: for each distinct symbol we add a subscript to indicate the player or agent. Thus, p for the row player becomes p_R, q for the column player becomes q_C. Similarly, we have λ_{jR} for the row player and α_{jC} for the column player.

Generalized to the Strategic Case

The new setup—for strategic learning by two agents—remains fundamentally simple and natural. Time proceeds discretely. At each choice point (or time), agents R and C each have a collection of distinct actions, A_{iR} and A_{iC}, that they may take. Each agent also has a probability, p_{iR} and p_{iC}, for taking each action. Each agent independently chooses and takes of its actions,[1] according to its p_{iR}s or p_{iC}s, and events $E_{j,R}$ and $E_{j,C}$ result. Depending on which event occurs, each agent's action probabilities are updated, and the process cycles.

[1]Hence Pr$(A_{iR} \cap A_{iC}) = p_{iR}p_{iC}$.

SLO (Simple Linear Operator) Learners

- NetLogo program `bushandmosteller1.nlogo`

Stag Hunt

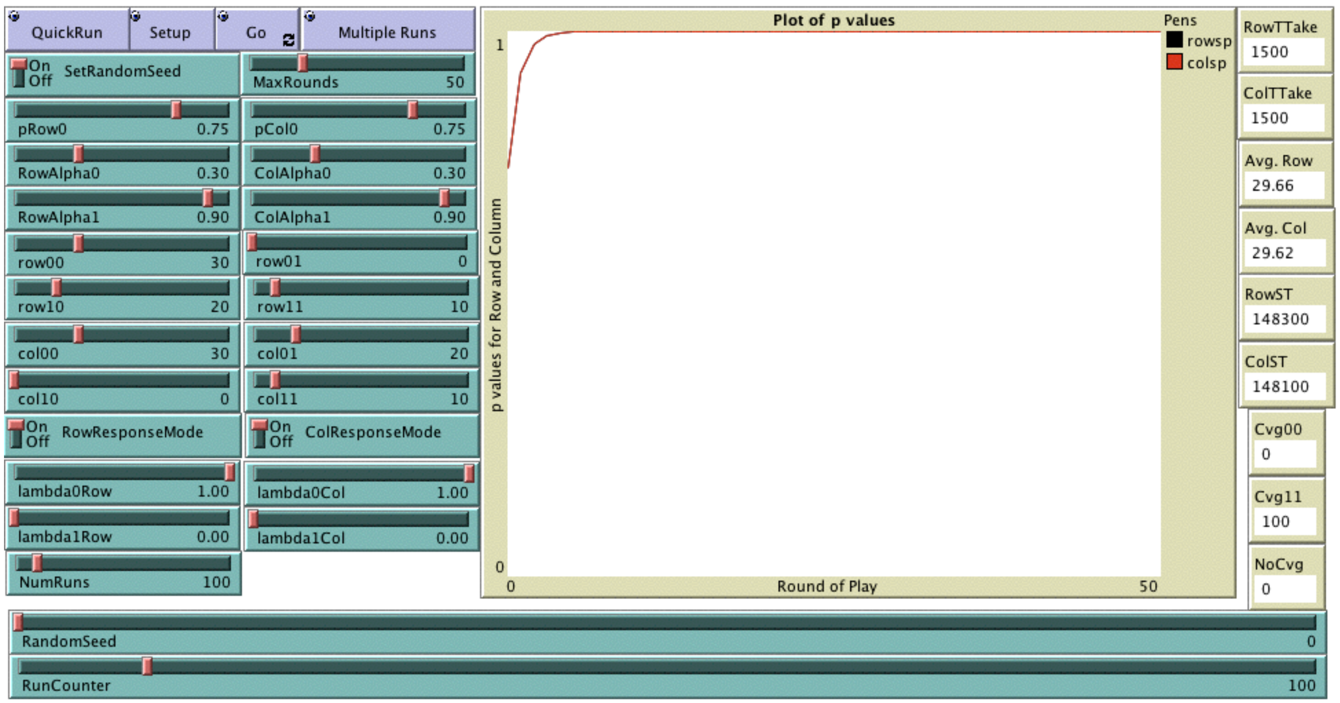
QuickRun
Setup
Go
Multiple Runs
On Off SetRandomSeed
MaxRounds 50
pRow0 0.75
pCol0 0.75
RowAlpha0 0.30
ColAlpha0 0.30
RowAlpha1 0.90
ColAlpha1 0.90
row00 30
row01 0
row10 20
row11 10
col00 30
col01 20
col10 0
col11 10
On Off RowResponseMode
On Off ColResponseMode
lambda0Row 1.00
lambda0Col 1.00
lambda1Row 0.00
lambda1Col 0.00
NumRuns 100
Plot of p values
Pens
rowsp
colsp
p values for Row and Column
Round of Play
1
0
0
50
RowTTake 1500
ColTTake 1500
Avg. Row 29.66
Avg. Col 29.62
RowST 148300
ColST 148100
Cvg00 0
Cvg11 100
NoCvg 0
RandomSeed 0
RunCounter 100

Prisoners' Dilemma

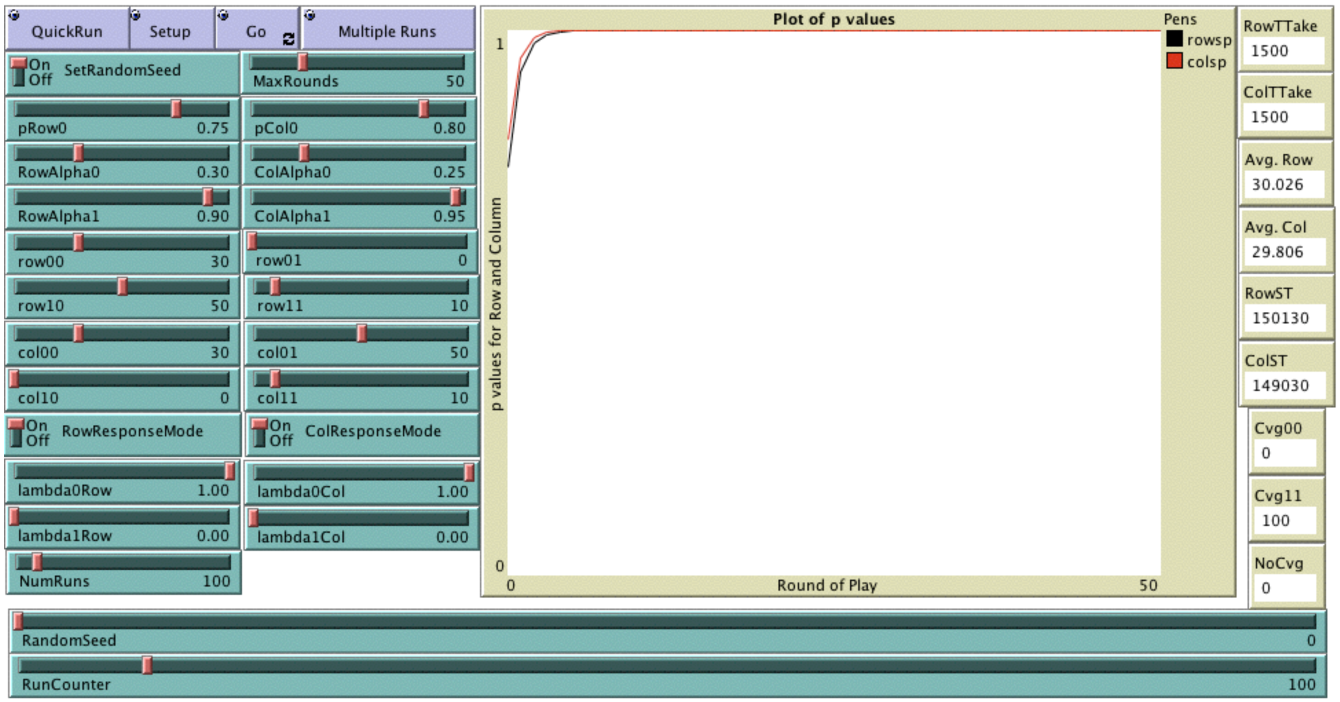

Too Many Foils. Basic Points:

1. Trust and cooperation are prevalent and fundamentally important social phenomena.

2. Their prevalence is inadequately explained with classical game-theoretic (and economic) models, suggesting consideration of different modeling/explanatory perspectives.

3. Agent-based models, built upon a philosophy of exploring rationality are plausible and promising in this regard.

4. Three very different models—a gridscape model, a Markov model, and a linear operator model (and there are others)—each easily and amply explain the emergence of trust and cooperation among agents.

Conclusion: Melioration not equlibrium!

These are the concluding 2 paragraphs from John Dewey. "The Influence of Darwin on Philosophy", Chapter 1 in *The Influence of Darwin on Philosophy and Other Essays.* New York: Henry Holt and Company (1910) : 1-19.

Old ideas give way slowly; for they are more than abstract logical forms and categories. They are habits, predispositions, deeply engrained attitudes of aversion and preference. Moreover, the conviction persists—though history shows it to be a hallucination—that all the questions that the human mind has asked are questions that can be answered in terms of the alternatives that the questions themselves present. But in fact intellectual progress usually occurs through sheer abandonment of questions together with both of the alternatives they assume, an abandonment that results from their decreasing vitality and a change of urgent interest. We do not solve them: we get over them.

Old questions are solved by disappearing, evaporating, while new questions corresponding to the changed attitude of endeavor and preference take their place. Doubtless the greatest dissolvent in contemporary thought of old questions, the greatest precipitant of new methods, new intentions, new problems, is the one effected by the scientific revolution that found its climax in the "Origin of Species."

References

[1] R. R. Bush and F. Mosteller. *Stochastic Models for Learning*. Wiley, New York, NY, 1955.

[2] Moisés Naím. *Illicit: How Smugglers, Traffickers, and Copycats Are Hijacking the Global Economy*. Doubleday, New York, NY, 2005.

[3] Brian Skyrms. *The Stag Hunt and the Evolution of Social Structure*. Cambridge University Press, Cambridge, UK, 2004.

`$Id: informs-2005-trust-foils.tex,v 1.5 2005/11/16 00:32:58 sok Exp $`