On Concepts of Rationality in Games*

Steven O. Kimbrough
Robert L. Axtell
University of Pennsylvania   The Brookings Institution
kimbrough (à) wharton.upenn.edu
Web page:
`http://opim-sky.wharton.upenn.edu/~sok/`
This paper (Kimbrough & Axell), these foils:
`http://opim-sky.wharton.upenn.edu/~sok/comprats/`

SEP, 19 May 2006, San Diego, CA

*File: rationality-concepts-foils.tex/pdf.

Acknowledgements

The ideas in part 1 occurred during a period of intense conversation and joint work with Rob Axtell, fall 2005. I thank him for this and for much else.

The ideas in part 2 emerged from discussions and joint work with Fred Murphy, beginning several years ago, and continuing through the fall of 2005. I thank him for this and for much else.

The Talk Is in Two Parts

1. The ‘philosophical’ part.

   No new facts. Comment on facts already available, indeed well known, to us. New distinctions and observations, however. See paper at `http://opim-sky.wharton.upenn.edu/~sok/sokpapers/2006/two-concepts-rationality.pdf`.
   Here and in paper, emphasize paradoxes.

2. The ‘scientific’ part. (Begins on page 36.)

   New results. Presentation of models and report on their properties and behaviors. Necessarily brief, however. Supporting programs at `http://opim-sky.wharton.upenn.edu/~sok/agebook/applications/nlogo/`.

Context

- Rationality is a contested concept (in game theory and economics)
- In some quarters, a deep unease with using equilibrium concepts to explain economic outcomes and to ‘solve’ games.
- Here, will present a new problem for rationality as standardly conceived (in game theory), and sketch an alternative, with equilibrium worries in mind.

Rationality

- Ordinary language senses (see Wikipedia, e.g.).

- Special meaning in decision theory, games, economics. Roughly: self-consistent with exogenous preferences.

  But is this correct or sensible?

- I will discuss more nuanced concepts and name them: fundamental rationality, IER (individual economic rationality), and effective rationality. All applicable to strategic (game-theoretic) contexts.

Some senses of ‘rationality’ (Wikipedia)

- “A logical argument is sometimes described as rational if it is logically valid.”

- “In economics, sociology, and political science, a decision or situation is often called rational if it is in some sense optimal, and individuals or organizations are often called rational if they tend to act somehow optimally in pursuit of their goals;”

- “Rationality is a central principle in artificial intelligence, where a rational agent is specifically defined as an agent which always chooses the action which maximises its expected performance, given all of the knowledge it currently possesses.”

Fundamental rationality

An agent is said to be *fundamentally rational* (with respect to a set of payoff vectors Ω) if the agent has a preference ordering, $\succeq$, on Ω such that for every $a, b \in \Omega$:

1. The *totality* condition obtains:

   $a \succeq b$ or $b \succeq a$ (or both, in which case we say that the agent is indifferent between a and b and we write $a \sim b$), and

2. The *transitivity* condition obtains:

   If $a \succeq b$ and $b \succeq c$ then $a \succeq c$.

Rational Choice: Fundamental rationality

Definition 1. [Fundamentally rational choice] *Given* $\Omega' \subseteq \Omega$, *a fundamentally rational preference ordering on* Ω, *then choosing* $\omega \in \Omega'$ *is a* fundamentally rational choice *iff there is no* $\omega' \in \Omega'$ *such that* $\omega' \succ \omega$.

- Comment: Standard stuff. I've merely given it a name (for the purposes of making distinctions later). Fundamental rationality can be, and has been, challenged. I'll give it a pass for now.

- Now to games, contexts of strategic interaction (CSIs). First, the standard setup, which I call a *type A* setup.

Type A Setup and Analysis

For type A analysis of a game (or a type A game), we need to specify the following items as constituting the setup of the game:

1. The players.

   For the examples to hand there are two players, R and C. In general there may be any finite number of players.

Type A Setup and Analysis

2. The pure strategy sets, Σ^i for each player, i.

   A strategy (for player i) is a complete set of instructions for play of the game (by player i). When the game is presented in strategic form, the pure strategies for the row (column) player are the rows (columns) in the table. In addition to its pure strategies, each player also has *mixed strategies*. These are the probability-weighted combinations of the pure strategies. We denote mixed strategies with a tilde. For example, Σ^C denotes player C's set of pure strategies and $\tilde{\Sigma}^C$ denotes C's mixed strategies. Since the pure strategies are a special case of probabilistic combination of pure strategies (one has a weight of 1, the others have 0), $\tilde{\Sigma}^C$ denotes all of C's strategies.

Type A Setup and Analysis

3. For each outcome, a payoff vector giving payoffs in that outcome for each player.

   An *outcome* of a (type A) game is a strategy vector, giving the played strategy of each player.[1] Thus for Standard Prisoners' Dilemma, there are four possible outcomes: (r_1, c_1), (r_1, c_2), (r_2, c_1), and (r_2, c_2). The payoff vectors, $\omega(\cdot)$, for these outcomes are: $\omega(r_1, c_1) = (3, 3)$, $\omega(r_1, c_2) = (0, 5)$, $\omega(r_2, c_1) = (5, 0)$, and $\omega(r_2, c_2) = (1, 1)$. Our convention is that in payoff vector (x, y), R gets x and C gets y.

[1] My terminology here is nonstandard. Usually, *outcome* is used for what I am calling the payoff vector. Also, see below for the distinction between the *play* of a game and the *outcome* of play.

Type A Setup and Analysis

4. Rules of play for the game.

   Standardly, "The rules of a game must tell us *who* can do *what* and *when* they can do it. They must also indicate who gets *how much* when the game is over." [Bin92, page 25]. I am handling the *how much* aspect of a game separately, as the payoff vectors ω (previous item).

Type A Setup and Analysis: Summary

1. Two or more players. Each player, i, has a set of pure strategies, Σ^i.

2. The total strategies available to each player is $\tilde{\Sigma}^i$, the set of mixed strategies on its Σ^i.

3. The play of the game: (mixed) strategy choices made by each player. E.g., for A, rock, paper, scissors, each with probability $\frac{1}{3}$.

4. The outcome of the game: pure strategies obtained by resolving all the mixed strategies. E.g. A plays rock, B plays scissors.

5. The payoff of the game: the reward to each player given the outcome of the game. E.g., A beats B.

Rational Choice in Games

- In games (of type A), we assume that the possible payoffs for each player are in its Ω^i and that each player is fundamentally rational with regard to its Ω^i. Every player i has a fundamentally rational preference ordering on Ω^i. (I neglect uncertainty but without loss of generality.)

- Note that a player doesn't pick $\omega \in \Omega^i$. Rather, a player picks a strategy $s^i \in \tilde{\Sigma}^i$.

  Hence the need to distinguish the *play* of a game, from the *outcome* of (play of) a game, from the *payoff* resulting from the outcome of a game.

Rational Choice: IER

We are now in position to define rationality for type A games. Let $\tilde{s}^i \in \tilde{\Sigma}^i$, that is, $\tilde{s}^i$ denotes a (possibly mixed) strategy available to agent i. Numbering the players from 1 to n, we can denote the play of a game by *Play* = $(\tilde{s}^1, \tilde{s}^2, \ldots, \tilde{s}^i, \ldots, \tilde{s}^n)$.

Definition 2. [IER choice] *Agent* i *is* individually economically rational *(IER) in the play of a game, if there is no strategy other than the one played by* i *that would yield in expectation a superior payoff for* i*, assuming the play is otherwise unchanged.*

Compare with Definition 1, page 7, for fundamentally rational choice. Similar structure.

Core idea: Optimal or maximizing because there is nothing better.

NB. *Not* "trying to maximize". Rather, actually doing so.

Rational Choice: IER

Formally, given *Play* = $(\tilde{s}^1, \tilde{s}^2, \ldots, \tilde{s}^i, \ldots, \tilde{s}^n)$, agent i is IER (makes an IER choice) iff there is no $\tilde{t}^i \in \tilde{\Sigma}^i$ such that E*Play*$(i, (\tilde{s}^1, \tilde{s}^2, \ldots, \tilde{t}^i, \ldots, \tilde{s}^n) \succ$ E*Play*$(i, (\tilde{s}^1, \tilde{s}^2, \ldots, \tilde{s}^i, \ldots, \tilde{s}^n)$.

Put otherwise, given i played $\tilde{s}^i$, then i is IER if i has no better strategy (in expectation), given how the other players played. Again, i is IER if the strategy i played is on average a *best response* to what the other players played.

Note: An agent that is individually economically rational (IER) is also said to be *consistent* with its fundamental preference ordering on Ω. This is how remarks like the following should be interpreted.

> When game theorists describe players as “rational”, they mean no more than that they make choices *consistently.* [Bin92, page 309]

Nash equilibrium

- The *Play* of a game is a (Nash) equilibrium iff every player is IER.
- Solution theory for (classical, type A) games: Every player will be IER, and play will be at a Nash equilibrium.

Accessibility

- **Definition 3. [Accessibility]** *Suppose that algorithm (or procedure or rule) α accepts inputs β and produces γ. Let us then say that γ* is accessible from β via α*.*

- **Definition 4. [Practical accessibility]** *Suppose that algorithm (or procedure or rule) α accepts inputs β and produces γ. Let us then say that γ* is practically accessible from β via α *if this is done at reasonable cost and in reasonable time.*

  Formal treatments can be added in specific cases.

Problems with IER

- Practical accessibility

  Zermelo and checkers. What since 1913?

- Accessibility

  Coordination: when there are multiple equilibria, how can the players find them?

  The problem arises in one-shot games (see below). It gets massively worse in indefinitely repeated games (Folk Theorem).

Consider: Stag Hunt

	c_1	c_2
r_1	(3, 3)	(0, 2)
r_2	(2, 0)	(1, 1)

Figure 1: Standard Stag Hunt (SSH). Player R has $\Sigma^R = \{r_1, r_2\}$. Player C has $\Sigma^C = \{c_1, c_2\}$. A good label for r_1 and c_1 is 'Hunt Stag' and for r_2 and c_2, 'Hunt Hare'.

This game has two Nash equilibria in pure strategies, play of (r_1, c_1) and play of (r_2, c_2), yielding a payoff to each player of 3 and 1, respectively. There is a third, mixed strategy.

NB: The paper discusses 'games of type B.'

Stag Hunt (con't.)

	Hunt stag (S)	Chase hare (H)
Hunt stag (S)	R [NP] R	T S
Chase hare (H)	S T	P [N] P

Figure 2: Generic Stag Hunt (aka: Assurance game): $R > T > P > S$

Calculating the mixed equilibrium we have

$$E(S) = Rx + (1-x)S \quad (1)$$

$$E(H) = Tx + (1 - x)P \quad (2)$$

Equating and solving for x gives us

$$x = \frac{P + S}{R + P - T - S} \quad (3)$$

For the specific game in Figure 1

$$x = \frac{1 + 0}{3 + 1 - 2 - 0} = \frac{1}{2} \quad (4)$$

The expected return received by a player of the game at this equilibrium is

$$G_{x=0.5} = \frac{1}{4}(3 + 0 + 2 + 1) = 1.5$$

So, there are three strategies that could be IER

	c_1	c_2	$(c_1, \frac{1}{2}; c_2, \frac{1}{2})$
r_1	$(3, 3)$	$(0, 2)$	$(1.5, 2.5)$
r_2	$(2, 0)$	$(1, 1)$	$(1.5, 0.5)$
$(r_1, \frac{1}{2}; r_2, \frac{1}{2})$	$(2.5, 1.5)$	$(0.5, 1.5)$	$(1.5, 1.5)$

Table 1: Expected payoffs for row and column, each playing one of its strategies supported by a Nash equilibrium (could be IER). Game of Figure 1, page 19.

Coordination is along the diagonal. By what a priori principle of rationality will these agents be able to arrive at a position on the diagonal?

Changing the game, there are still three strategies that could be IER, and are Nash equilibria

	c_1	c_2	$(c_1, \frac{1}{50}; c_2, \frac{49}{50})$
r_1	$(51, 51)$	$(0, 2)$	$(1.02, 2.98)$
r_2	$(2, 0)$	$(1, 1)$	$(1.02, 0.98)$
$(r_1, \frac{1}{50}; r_2, \frac{49}{50})$	$(2.98, 1.02)$	$(0.98, 1.02)$	$(1.02, 1.02)$

Table 2: Expected payoffs for row and column, each playing one of its strategies supported by a Nash equilibrium (could be IER). Game of Figure 1, page 19, but with $R = 3$ changed to $R = 51$, making mutual hunting of stag very attractive.

Coordination is still along the diagonal. If (r_1, c_1) looks prominent now, why not before? Notice decline of payoff for the mixed equilibrium. Principle: if both players are better off with a play, they should be more likely to reach it. Violated here.

The Gift of the Magi

- Short story by O. Henry, about 1900
- Available on Project Gutenberg

  `http://www.gutenberg.org/dirs/etext05/magi10h.htm`
- Captures the point through fictional dramatization.
- Della and Jim are a young married couple, very much in love, very poor. Christmas is coming and they have a problem. . . .

The story begins. . .

One dollar and eighty-seven cents. That was all. And sixty cents of it was in pennies. Pennies saved one and two at a time by bulldozing the grocer and the vegetable man and the butcher until one's cheeks burned with the silent imputation of parsimony that such close dealing implied. Three times Della counted it. One dollar and eighty- seven cents. And the next day would be Christmas.

. . . and ends. . .

Jim tumbled down on the couch and put his hands under the back of his head and smiled.

“Dell,” said he, “let’s put our Christmas presents away and keep ‘em a while. They’re too nice to use just at present. I sold the watch to get the money to buy your combs. And now suppose you put the chops on.”

. . . here.

The magi, as you know, were wise men—wonderfully wise men—who brought gifts to the Babe in the manger. They invented the art of giving Christmas presents. Being wise, their gifts were no doubt wise ones, possibly bearing the privilege of exchange in case of duplication. And here I have lamely related to you the uneventful chronicle of two foolish children in a flat who most unwisely sacrificed for each other the greatest treasures of their house. But in a last word to the wise of these days let it be said that of all who give gifts these two were the wisest. Of all who give and receive gifts, such as they are wisest. Everywhere they are wisest. They are the magi.

In Summary: Widely Shared Worries about Neoclassical Models

Among the main ones:

- Heroic, implausible assumptions

  Common knowledge, unlimited computational capacity, selfishness and unreality of *homo economicus*

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

  How can we predict which equilibrium will occur? (This is *not* the coordination problem just discussed.)

In Summary: Widely Shared Worries about Neoclassical Models

- Poor track record in predicting behavior

> Ironically, game theory is often hoisted on its own pétard: many of its most fundamental predictions—predictions that would have been too vague to test with any confidence in the pre-game-theoretic era—are *decisively and repeatedly disconfirmed,* in laboratory settings, with substantial agreement among experimenters, regardless of their theoretical priors. [Gin00, page xxiv] (emphasis in original)

Documented extensively in the behavioral game theory literature.

And, new here. . .

- The problems of accessibility and coordination, discussed above.

So, it makes implausible assumptions, is generally indefinite in its predictions, when it makes predictions they are at variance with experience, and the solution concept is oblivious to implementation; otherwise, it's a pretty good theory.

What to do?

One alternative

It is better to drop the term “rational” altogether, ...

In the same vein, we do not follow classical game theory in asking how agents “learn” to play optimal strategies, because the cognitive processes involved in “learning” are probably, under most conditions, much less important than the forms of imitation underlying the replicator dynamic. . . and cultural transmission. . . . In short, evolutionary game theory replaces the idea that games have “solutions” that agents “learn,” with the idea that games are embedded in natural and social processes that produce agents who play effectively.

Dispensing with the rationality postulate does not imply that people are *irrational* (whatever that means). The point is that the concept of “rationality” does not help us understand the world. [Gin00, pages xxv-xxvi]

Another Alternative: Develop a Concept of Effective Rationality

Approximately, roughly:

Definition 5. [Effective Rationality] *Suppose that* γ *is practically accessible via procedure* α *applied to inputs* β*. If* α *has a preponderance of attractive properties, then* α *is effectively rational.*

More of a definition schema. Think: β as a history of play in a game, γ as a move in a game, and α as a procedure that picks a move, given the history. What are those attractive properties? (see next foil)

See the paper and the discussion of type B games. Idea: α does *learning in policy space*.

Some attractive properties pertaining to α

1. Performance against self.

   Does the apparatus do well against itself? If all agents use the apparatus, do the agents as a group prosper relatively well in the ambient environment?

2. Performance against others.

   Does the apparatus do well playing against other regimes that do well against themselves?

Some attractive properties pertaining to α

3. Exploitability.

   Is the apparatus catastrophically exploitable? Does it have weaknesses that may be discovered by another apparatus?

4. Robustness.

   Is the apparatus robust under perturbations of its parameters? Does the apparatus perform well against a field of others?

5. Learnability.

   Can the apparatus be parameterized in such a way that an agent can (easily) learn profitable, well-performing settings?

Some attractive properties pertaining to α

6. Computational cost.

   Is the apparatus computationally tractable? Is it simple or does it require excessive computational resources from the agent?

7. Informational requirements.

   Does the apparatus rely on plausibly available information? Or does it require information not likely to be available in the actual system being modeled?

Begin Part 2

- Will develop and explore a model, ProbeAndAdjust1.
- An example of *learning in policy space* (LPS).
- ProbeAndAdjust1 will be used for acquiring values for a continuous parameter.

ProbeAndAdjust1 Model/Learning Procedure

1. Initialize `currentValue`, δ, ϵ, and `history`
2. `upValues` $\leftarrow$ [ ], `downValues` $\leftarrow$ [ ], `episodeOver` $\leftarrow$ false
3. `while (not episodeOver):`
   (a) Select `probe` randomly from $[\texttt{currentValue} - \delta, \texttt{currentValue} + \delta]$
   (b) If $\texttt{probe} \geq \texttt{currentValue}$ append $\texttt{Return}(\texttt{probe})$ to `upValues`; else append to `downValues`.
   (c) If $\texttt{Length}(\texttt{upValues}) + \texttt{Length}(\texttt{downValues}) \geq \texttt{history}$, `episodeOver` $\leftarrow$ true
4. If $\texttt{Mean}(\texttt{upValues}) \geq \texttt{Mean}(\texttt{downValues})$, $\texttt{currentValue} \leftarrow \texttt{currentValue} + \epsilon$, else $\texttt{currentValue} \leftarrow \texttt{currentValue} - \epsilon$.
5. Go to (2).

Why Does This Matter? And where?

One example: microeconomics. Bear with me. To begin: monopoly.

Introductory textbooks in microeconomics will tell the story roughly as follows. Suppose, for the sake of simplicity that the product in question can be produced at 0 cost. (Extensions of the model can be made for non-zero costs. The results will not be terribly sensitive to this.) The market's demand, Q (think: quantity demanded), for our product is a linear function of its price, P:

$$Q = c - b \cdot P \tag{5}$$

Assume: $b, c > 0$. Here, c is a constant, representing the quantity demanded when price is 0 and, since we are linear, the price point at which demand disappears when the price is too high (when $b \cdot P = c$).

Monopoly (con't.)

Rearranging (5) we get:

$$P = a - slope \cdot Q \tag{6}$$

where a = c/b and slope = 1/b. (We assume $slope > 0$.)

The profit, π, made by the monopolist is $P \cdot Q$ (since the cost of production is 0, we need only account for revenue, which is defined here as $P \cdot Q$).

$$\pi = P \cdot Q = (a - slope \cdot Q) \cdot Q = a \cdot Q - slope \cdot Q^2 \tag{7}$$

Monopoly (con't.)

The monopolist will seek to maximize π, which can be done by a simple exercise with the calculus.

$$\frac{d\pi}{dQ} = a - 2 \cdot slope \cdot Q \tag{8}$$

Setting $a - 2 \cdot slope \cdot Q$ to zero and solving for Q yields

$$Q = \frac{a}{(2 \cdot slope)} \tag{9}$$

So Q in (9) is Q^*, the optimal quantity for the monopolist to put on the market.

Monopoly (con't.)

Checking that $\frac{d^2\pi}{dQ^2} = -2 \cdot slope < 0$ verifies that we indeed have found a maximum. So,

$$Q^* = \frac{a}{(2 \cdot slope)} \tag{10}$$

Important lesson: monopolists make extra profits. The value of Q under perfect competition is 0: marginal costs = marginal revenues.

Even though the model is exceedingly simple, clearly it teaches a valuable lesson.

Nagging questions

- How did the monopolist discover the demand function in the first place?
- What if demand varies randomly in some way?

The model ignores accessibility, something the argument above warns us against.

Note: Of course, in the monopoly case, we do not have a game. This is parametric decision making.

Can we find a model that will grant accessibility to the monopolist?

ProbeAndAdjust1 Model Applied to the Monopoly Problem

`http://opim-sky.wharton.upenn.edu/~sok/agebook/applications/nlogo/monopolyProbeAndAdjust.nlogo`

- Not surprisingly, ProbeAndAdjust1 takes our agent to the vicinity of the monopoly quantity (and hence profit). Note: this is probing with quantity. Works as well for probing with price.

- Perhaps surprisingly, ProbeAndAdjust1 can do very well at tracking a demand curve that is undertaking a random walk. Clearly, an agent using ProbeAndAdjust1 might *learn* good values for δ and ϵ. This takes us well beyond the standard account.

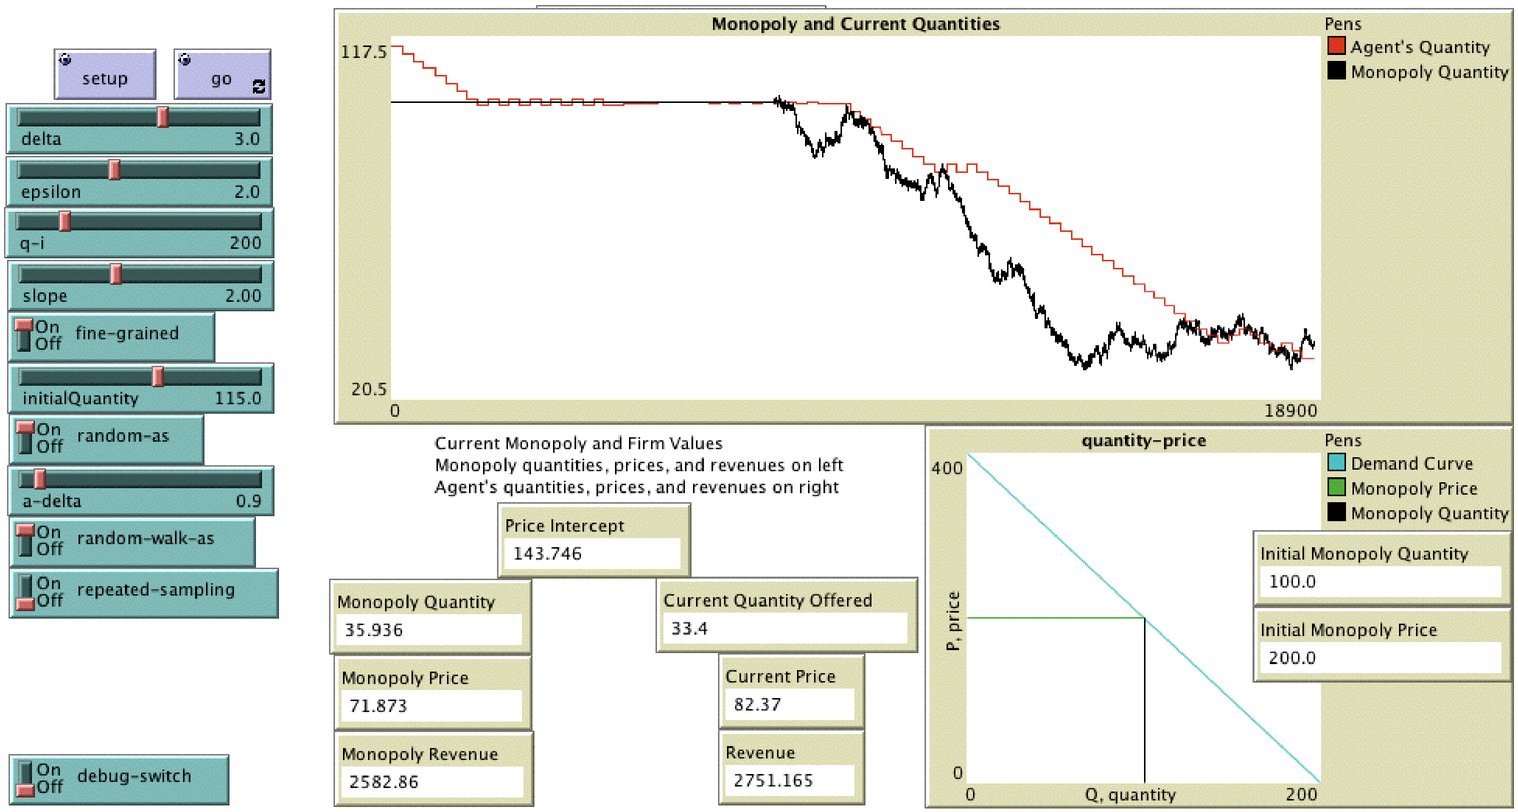

http://opim-sky.wharton.upenn.edu/~sok/agebook/applications/nlogo/monopolyProbeAndAdjust.nlogo

Monopoly Multiexamples

- At start $\epsilon = 2.0$, `random-as` is Off (so no random walk). Agent promptly moves to vicinity of optimal quantity.
- Upon getting to neighborhood of Q^*, I changed ϵ to 0.2, causing the agent to hover closer to Q^*.
- Then I turned `random-as` to On, initiating a random walk on the demand curve's Q intercept. As luck would have it, a steep drop in Q^* occurs. The agent quickly falls away.
- Then I reset ϵ to 2.0. The agent tracks pretty well.

Monopoly + 1: Duopoly

- The monopoly problem is not a game, is not strategic. Only one agent/player.

- Much attention in economics has been directed at duopoly, which now is a game. It is thought that any findings for duopoly will apply, approximately to oligopoly, a small number of firms.

Duopoly Models

- The picture is much less settled for duopoly (oligopoly). Three main models or types of competition:

  1. Bertrand. The two agents compete on price and drive the price down to the competitive equilibrium.
  2. Cournot. The two agents in effect find a partially collusive outcome, between the competitive solution and the monopoly solution. Agents choose simultaneously, on quantity, Q.
  3. Stackelberg. Leader-follower model, choices are on quantity, Q. Results in equilibrium between Cournot and Bertrand.

We'll focus on Bertrand and Cournot.

Bertrand

Price competition.

> If firm 1 really believes that firm 2 will charge a price $\hat{p}$ that is greater than the marginal cost, it will always pay firm 1 to cut its price to $\hat{p} - \epsilon$. But firm 2 can reason the same way! Thus any price higher than marginal cost cannot be an equilibrium; the only equilibrium is the competitive equilibrium. [Var03, page 488]

In terms of our assumption of 0 cost (viz., monopoly), this means that the equilibrium price is 0.

Cournot: Quantity competition

Generalizing Expression (5) from the monopoly case, page 38:

$$Q = Q_1 + Q_2 = c - b \cdot P \tag{11}$$

Q_i = quantity offered to the market by firm i. The profit of firm 1 is then

$$\pi_1 = P \cdot Q_1 = (a - slope \cdot (Q_1 + Q_2)) \cdot Q_1 \tag{12}$$

For firm 2 we have

$$\pi_2 = P \cdot Q_2 = (a - slope \cdot (Q_1 + Q_2)) \cdot Q_2 \tag{13}$$

Cournot: Quantity competition

Solving to maximize π_is:

$$\frac{d\pi_1}{dQ_1} = a - 2 \cdot slope \cdot Q_1 - slope \cdot Q_1 \cdot \frac{dQ_2}{dQ_1} - slope \cdot Q_2 \qquad (14)$$

$$\frac{d\pi_2}{dQ_2} = a - 2 \cdot slope \cdot Q_2 - slope \cdot Q_2 \cdot \frac{dQ_1}{dQ_2} - slope \cdot Q_1 \qquad (15)$$

Problem: How do we get values for $\frac{dQ_2}{dQ_1}$ and $\frac{dQ_1}{dQ_2}$?

Answer: Stipulate that they equal 0. This is tantamount to assuming that the other player will not change its quantity offered from last period to next period. You play your best response to what the other played last time, and hope he does so again.

Cournot: Quantity competition

Now we can get a solution, analytically.

$$\frac{d\pi_1}{dQ_1} = a - 2 \cdot slope \cdot Q_1 - slope \cdot Q_2 = 0 \tag{16}$$

$$\text{So, } \; Q_1 = \frac{a - slope \cdot Q_2}{2 \cdot slope} \tag{17}$$

$$\frac{d\pi_2}{dQ_2} = a - 2 \cdot slope \cdot Q_2 - slope \cdot Q_1 = 0 \tag{18}$$

$$\text{So, } \; Q_2 = \frac{a - slope \cdot Q_1}{2 \cdot slope} \tag{19}$$

Cournot: Quantity competition

Now, continuing to solve

$$Q_1 = \frac{a - slope \cdot Q_2}{2 \cdot slope} = \frac{a - slope \cdot \frac{a - slope \cdot Q_1}{2 \cdot slope}}{2 \cdot slope} \tag{20}$$

implies

$$Q_1 = \frac{a}{3 \cdot slope} \tag{21}$$

Similarly

$$Q_2 = \frac{a}{3 \cdot slope} \tag{22}$$

Compare Cournot profits with monopoly profits

Let our demand function be

$$P = 400 - 2Q \tag{23}$$

Then $Q^* = 100$ and $P^* = 200$. The monopoly profit is $P^* \cdot Q^* = 20000$.

Under the Cournot duopoly model,

$$Q = Q_1^* + Q_2^* = \frac{400}{3 \cdot 2} + \frac{400}{3 \cdot 2} = \frac{400}{3} \approx 133 \tag{24}$$

$$\pi = \pi_1 + \pi_2 = (400 - 2 \cdot 133) \cdot 133 \approx 17822 < 20000 \tag{25}$$

What algorithm should our 2 agents use?

Same as before: ProbeAndAdjust1. But: lower bid gets the deal, all of it.

Here, in `duopoly.nlogo`, the agents are bidding prices (instead of quantities).

The Bertrand model has them competing away all profits.

We just saw what the Cournot model predicts.

What will ProbeAndAdjust1 do?

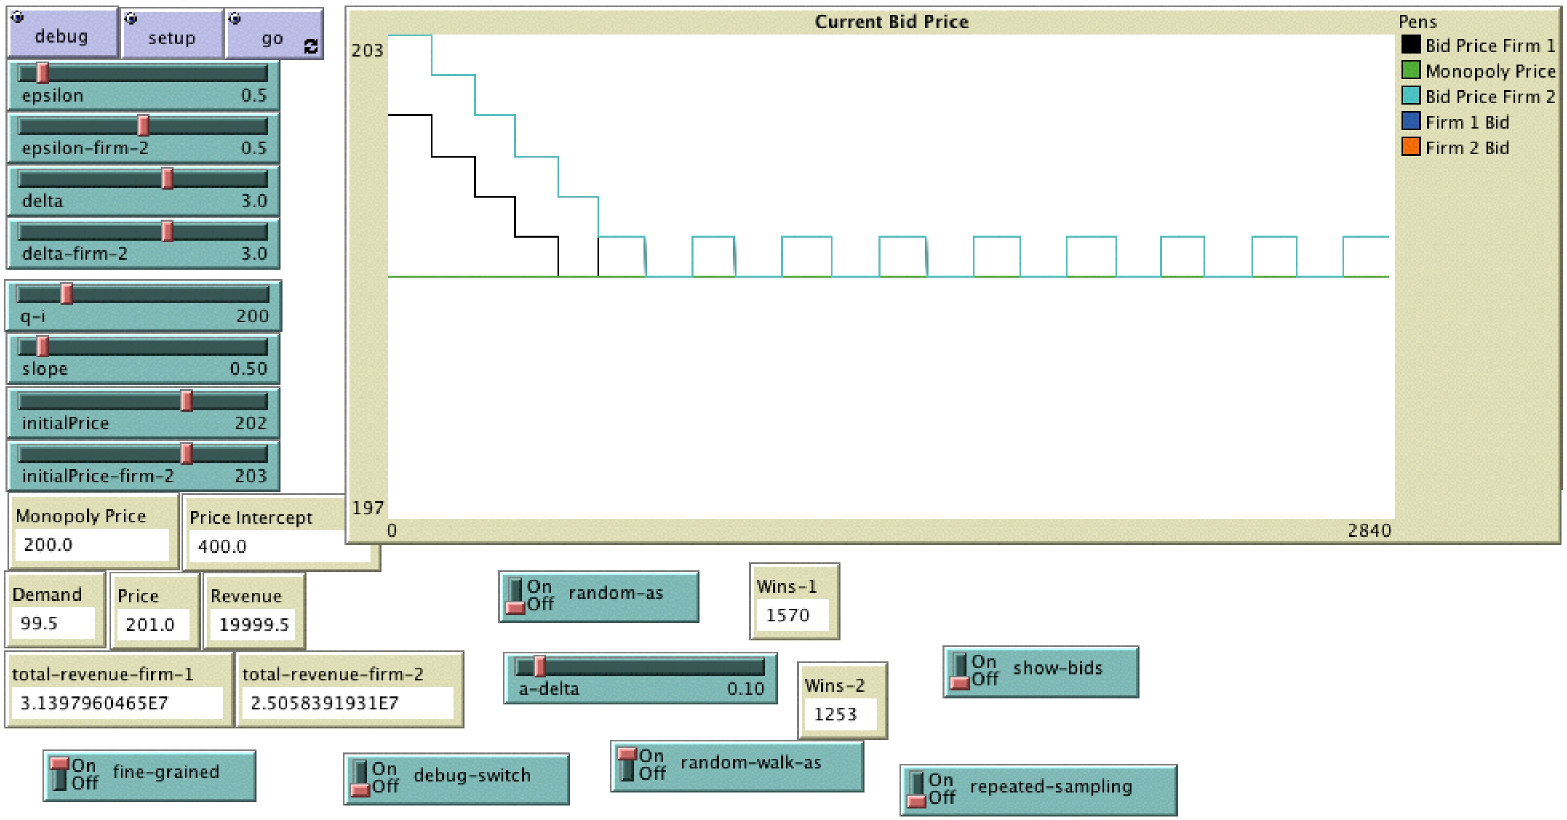

http://opim-sky.wharton.upenn.edu/~sok/agebook/applications/nlogo/duopoly.nlogo

Comments

- The agents start out with different initial prices and no knowledge of the demand function.

- They independently probe and adjust, according only to successful bids. They stochastically split the market.

- Convergence is quick (in this example), and rests in the neighborhood of the monopoly price!

  In the figure, the current price is 201, compared to the monopoly price of 200.

- And if the demand curve is a random walk?

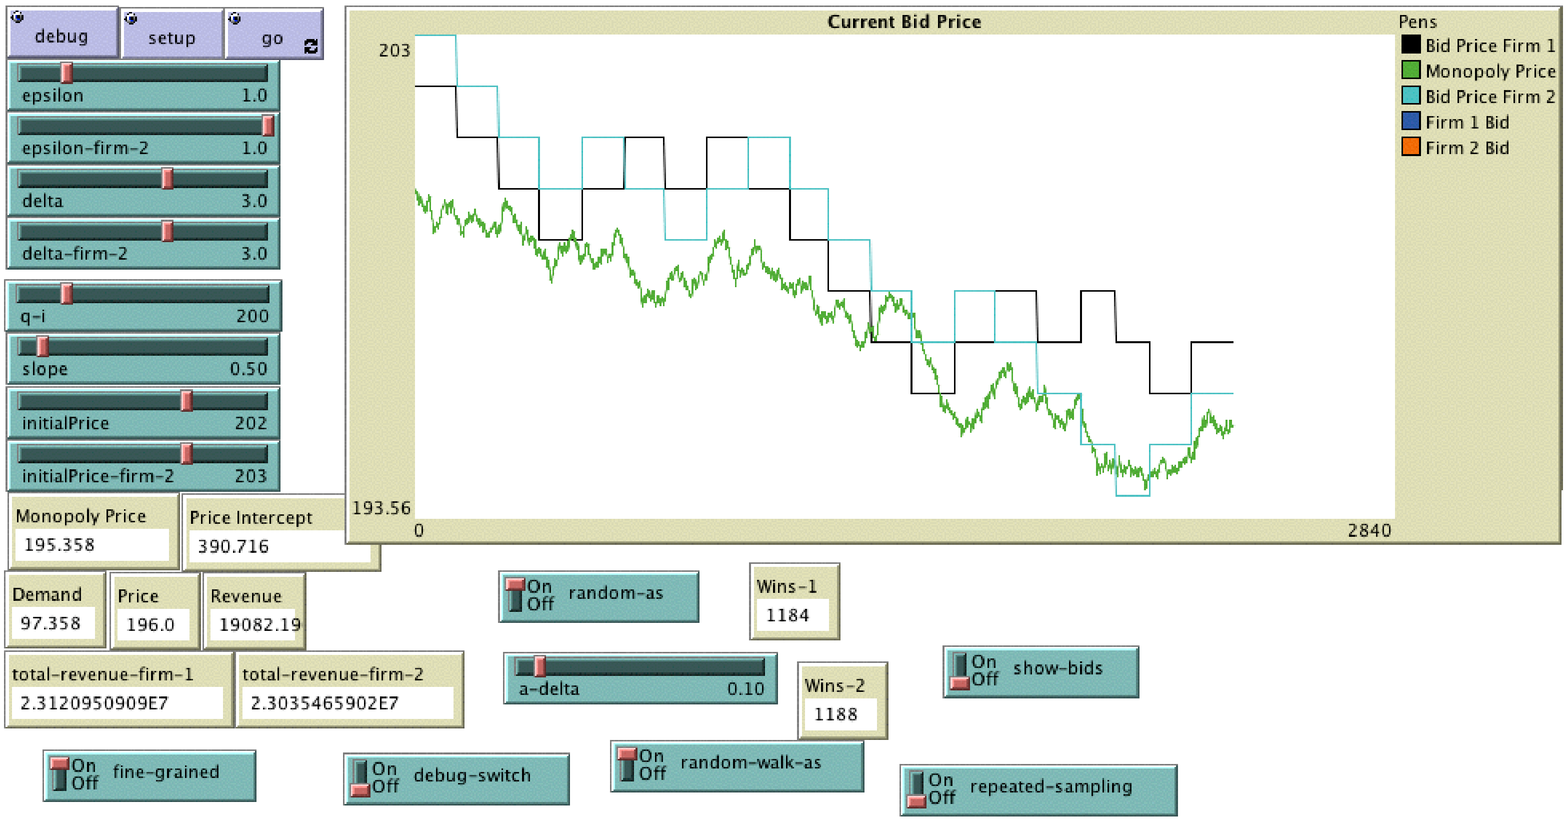

http://opim-sky.wharton.upenn.edu/~sok/agebook/applications/nlogo/duopoly.nlogo

And

What if the cost structures of the two firms are different?

Say, firm 1 has cost 0 and firm 2 cost 50.

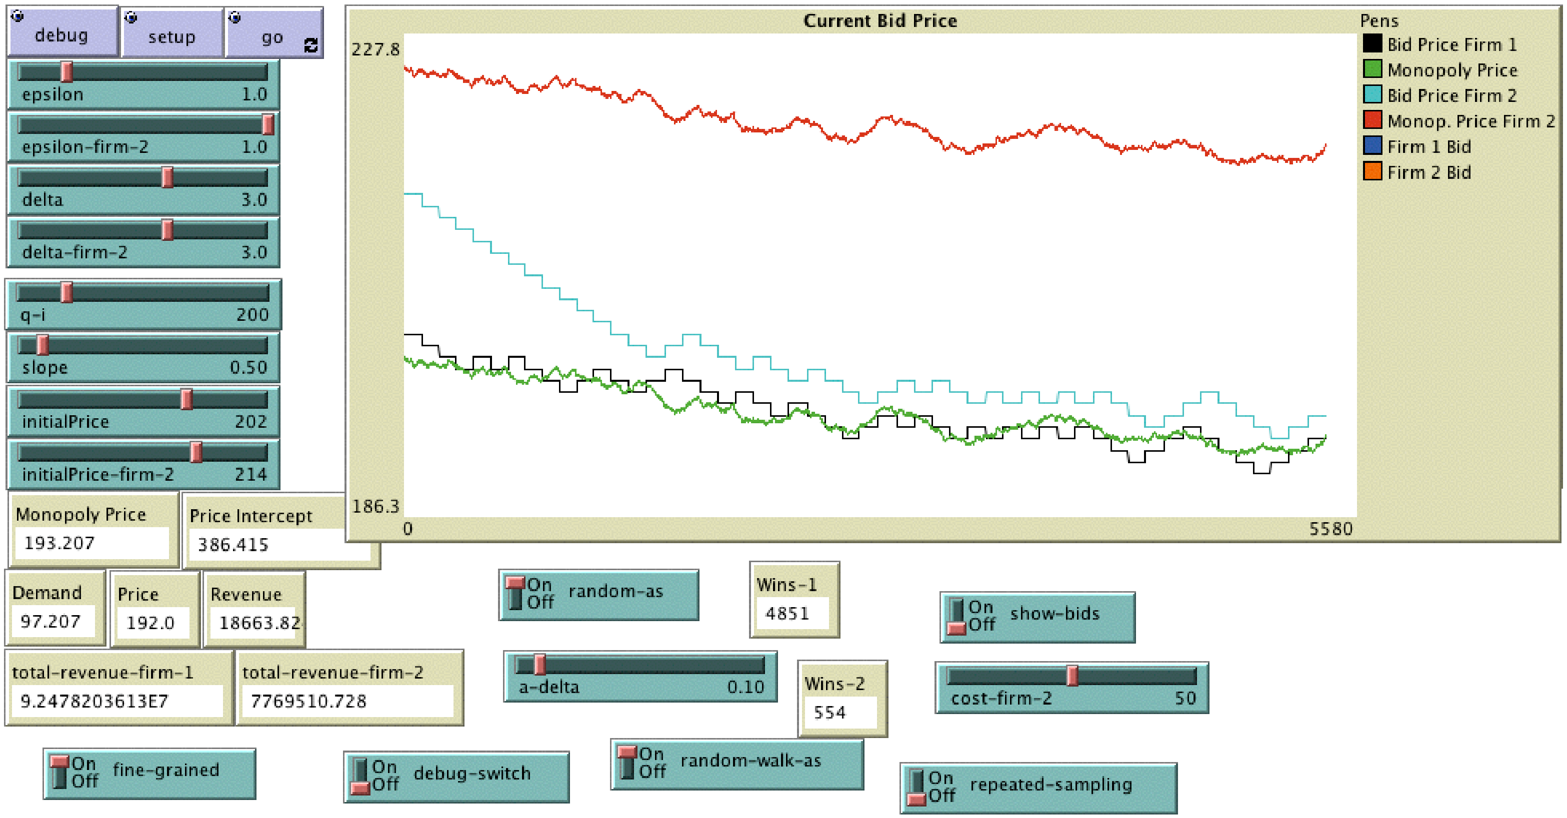

http://opim-sky.wharton.upenn.edu/~sok/agebook/applications/nlogo/duopoly1.nlogo

Comment

They converge to the monopoly price associated with the low-cost firm.

Discussion

- Classical duopoly theory ignores the fact that the game is indefinitely repeated (iterated). Every morning the firms open for business and are free to set prices and quantities available.

- Ignoring repetition permits finessing the strategy coordination problem. Or does it? Price competition, quantity competition or ?

- There is no effectively rational procedure of play in the classical literature.

- ProbeAndAdjust1 is an effectively rational procedure in both the monopoly and the duopoly cases.

Revisiting the α evaluation criteria

1. Performance against self.

   Does ProbeAndAdjust1 do well against itself? If all agents use the apparatus, do the agents as a group prosper relatively well in the ambient environment?

   Answer: Yes, it appears so in the present case.

2. Performance against others.

   Does the apparatus do well playing against other regimes that do well against themselves?

   Answer: No other candidates at this time. ProbeAndAdjust1 could obviously be improved. Also, against a fixed defector, ProbeAndAdjust1 will learn to defect.

Revisiting the α evaluation criteria

3. Exploitability.

   Is the apparatus catastrophically exploitable? Does it have weaknesses that may be discovered by another apparatus?

   Answer: See previous answer.

4. Robustness.

   Is the apparatus robust under perturbations of its parameters? Does the apparatus perform well against a field of others?

   Answer: Does well under changing demand function, random walk, etc.

5. Learnability.

   Can the apparatus be parameterized in such a way that an agent can (easily) learn profitable, well-performing settings?

   Answer: Yes. And this is important. I learned (just a little) to set the δ and ϵ parameters.

   During play, the agents will receive information from play that is useful for adjusting these parameters.

   For example, if at the end of all the recent episodes we decrease the `current-value`, this suggests that ϵ and perhaps δ should be increased. Conversely, if we are bouncing around up and down, this suggests that ϵ and perhaps δ should be decreased.

   Bacteria could learn this. So could MBAs.

Revisiting the α evaluation criteria

6. Computational cost.

   Is the apparatus computationally tractable? Is it simple or does it require excessive computational resources from the agent?

   Answer: Low cost.

7. Informational requirements.

   Does the apparatus rely on plausibly available information? Or does it require information not likely to be available in the actual system being modeled?

   Answer: Uses history of play, which is available.

Final comments

The flip side of the critique of IER is a critique of (Nash) equilibrium as a solution concept for games. This critique has be developed elsewhere, and extensively. I would only add this.

Why are we interested in equilibria (in games)? I imagine that a basic motivation has been that this is where things have to settle and if you look for something, your best bet is where it will settle. After all, it's usually there. Where to look for the water? In the valleys, in the basement. Down hill. That's where it is in equilibrium.

The argument I have tried to make is that while the equilibrium principle might be good for water, it is because we have a theory—gravity—of how and why the water will be driven to, will be caused to get to, equilibrium. We have no such theory or account in the case of games with multiple equilibria.

Really final (and speculative) comments

- We might distinguish theories of rationality by their objects. Roughly: individual choices versus policies that lead to individual choices.

  Choice versus habit theories of rationality.

- Other names: habits, principles, policies.

  These have been much addressed by philosophers, but are largely beyond the ken of Rational Choice Theory.

How to characterize the differences between them?

Two approaches to decision

Given a number of options...

Choice Measure the value of each option and pick one with the highest measure.

(e.g., Rational Choice Theory)

Habit Use some form of randomized search to sample the choices and learn which are the good ones.

Policies, habits, principles may be thought of as (considered) inductive biases, necessary for randomized search. Rationality in this sense is about conducting search, not about particular choices.

References

[Bin92] Ken Binmore, *Fun and games: A text on game theory*, D.H. Heath and Company, Lexington, MA, 1992.

[Gin00] Herbert Gintis, *Game theory evolving: A problem-centered introduction to modeling strategic interaction*, Princeton University Press, Princeton, NJ, 2000.

[Var03] Hal R. Varian, *Intermediate microeconomics: A modern approach*, sixth ed., W. W. Norton & Company, New York, NY, 2003.