Evolution & Learning in Games
Econ 243B

Jean-Paul Carvalho

Lecture 1.
The Rational Agent and its Limitations

The Knowledge Problem

What is the problem we wish to solve when we try to construct a rational economic order?

On certain familiar assumptions the answer is simple enough. If we possess all the relevant information, if we can start out from a given system of preferences and if we command complete knowledge of available means, the problem which remains is purely one of logic. That is, the answer to the question of what is the best use of the available means is implicit in our assumptions. The conditions which the solution of this optimum problem must satisfy have been fully worked out and can be stated best in mathematical form: put at their briefest, they are that the marginal rates of substitution between any two commodities or factors must be the same in all their different uses.

The Knowledge Problem

This, however, is emphatically not the economic problem which society faces. And the economic calculus which we have developed to solve this logical problem, though an important step toward the solution of the economic problem of society, does not yet provide an answer to it. The reason for this is that the "data" from which the economic calculus starts are never for the whole society "given" to a single mind which could work out the implications, and can never be so given.

The Knowledge Problem

The peculiar character of the problem of a rational economic order is determined precisely by the fact that the ***knowledge of the circumstances of which we must make use never exists in concentrated or integrated form, but solely as the dispersed bits of incomplete and frequently contradictory knowledge which all the separate individuals possess****. The economic problem of society is thus not merely a problem of how to allocate "given" resources if "given" is taken to mean given to a single mind which deliberately solves the problem set by these "data." It is rather a problem of how to secure the best use of resources known to any of the members of society, for ends whose relative importance only these individuals know. Or, to put it briefly, it is a problem of the utilization of knowledge not given to anyone in its totality.*

F.A. Hayek, The Use of Knowledge in Society, *American Economic Review*, 1945.

The Rational Agent

This lecture is an introduction to the rational agent in modern decision theory and economics.

Since there is much confusion as to what 'rationality' is, it is useful to review the prevailing model before moving on to various forms of *bounded rationality* in our treatment of evolution and learning in games.

I draw upon the following texts:

- Rubinstein (2006): *Lecture Notes in Microeconomic Theory – The Economic Agent*, Ch. 1, 2, 3, 8.
- Binmore (2009): *Rational Decisions*, Ch. 1, 3, 7.
- Gintis (2009): *Game Theory Evolving*, Ch. 2.

Preferences

A rational agent must have <u>well-defined</u> and <u>consistent</u> preferences (over consequences; note instrumental vs intrinsic preferences).

Definition. A *preference* on a set X is a binary relation $\succeq$ on X satisfying:

- *Completeness*: For any $x, y \in X$, either $x \succeq y$ or $y \succeq x$.
- *Transitivity:* For any $x, y, z \in X$, if $x \succeq y$ and $y \succeq z$, then $x \succeq z$.

Intransitivity due to Aggregation

Consider a multidimensional comparison of three occupations A, B, C:

1. *Salary*: $A \succ_1 B \succ_1 C$.
2. *Interest*: $C \succ_2 A \succ_2 B$.
3. *Autonomy*: $B \succ_3 C \succ_3 A$.

Suppose preferences are aggregated over dimensions, such that $x \succ y$ if x is preferred to y in at least two dimensions. Then we have intransitivity:

$$A \succ B \succ C \succ A.$$

Utility

To apply optimization techniques, we can define a utility function $U : X \to \mathbb{R}$.

The original interpretation of utility was an hedonic one (Bentham 1863).

In modern decision theory, it is simply a representation of preferences.

Definition. $U : X \to \mathbb{R}$ represents the preference $\succeq$ if for all $x, y \in X$, $x \succeq y$ if and only if $U(x) \geq U(y)$.

Equivalent Utility Representations

Proposition. If U represents $\succeq$, then for any strictly increasing function $f : \mathbb{R} \to \mathbb{R}$, $V(x) = f\big(U(x)\big)$ represents $\succeq$ as well.

Proof. $a \succeq b$
$\iff U(a) \geq U(b)$ (since U represents $\succeq$)
$\iff f\big(U(a)\big) \geq f\big(U(b)\big)$ (since f is strictly increasing)
$\iff V(a) \geq V(b)$.

Existence of Utility Representation

Proposition. If X is countable, then any preference relation on X has a utility representation with a range $(0, 1)$.

When there exists a utility representation, the most preferred option can be found by solving the following (utility) maximization problem:

$$max_{x \in X} U(x).$$

Note: some preferences cannot be represented by a utility function, e.g. lexicographic preferences.

Choice

Now we come to choice.

- Let D be a set of subsets of X (not necessarily $\mathcal{P}(X)$).
- A *choice function* C assigns to each set $A \in D$ a subset of A, with the choices denoted by $C(A)$.
- An *induced choice function* $C_{\succeq}$ is one that is consistent with the preference $\succeq$ (which is independent of the choice set A).

Revealed Preference

The choice function is unobservable, but some choice data may be available.

Revealed preference dispenses with trying to predict what is going on in agents' heads and restricts attention to choice data.

In particular, if choice is consistent , we can use data on an agent's choice from A to deduce what it will do when faced with some other $A' \in D$.

Definition. The *weak axiom of revealed preferences* (WARP): C satisfies WARP if whenever $x, y \in A \cap B$, $x \in C(A)$ and $y \in C(B)$, then $x \in C(B)$.

In words, if y is chosen when x is available, then x will never be chosen without y when both are available (an *independence* condition).

Money Pump or 'Dutch Book' Argument

Evolution (broadly conceived) and learning should eliminate intransitivity.

Suppose for person i:

$$apple \succ banana \succ cantaloupe \succ apple.$$

Person j could give i an apple, then charge i 1 cent to exchange the apple for a banana, then charge i 1 cent to exchange the banana for a cantaloupe, then charge i 1 cent to exchange the cantaloupe for an apple, ...

If this continues, j acquires i's entire wealth.

Risk

No distinction so far between actions and consequences, because deterministic association between action and consequence.

Now make stochastic:

- Let Z be a finite set of consequences (prizes).
- A lottery p is a probability measure on Z.
- $p(z)$ is the objective probability of receiving the prize z given lottery p (*risk*).
- The lottery in which prize x is received with prob. α and y with complementary probability is denoted by $\alpha x \oplus (1-\alpha)y$.
- $L(Z)$ is the space containing all lotteries with prizes in Z (simplex in Euclidean space).

Preferences over Lotteries

When there is risk, actions induce lotteries, so we have to define preferences over lotteries.

Now make stochastic:

- Let Z be a finite set of consequences (prizes).
- A lottery p is a probability measure on Z.
- $p(z)$ is the objective probability of receiving the prize z given lottery p (*risk*).
- The lottery in which prize x is received with prob. α and y with complementary probability is denoted by $\alpha x \oplus (1-\alpha)y$.
- $L(Z)$ is the space containing all lotteries with prizes in Z (simplex in Euclidean space).

Expected Utility

Von Neumann-Morgenstern Axioms

- Consider K lotteries, $p^1, p^2, \ldots p^K$.
- A compound lottery $\oplus_{k=1}^{K} \alpha_k p^k$ is generated as follows:

  Stage 1. One of the K lotteries is realized, with α_k being the probability that p^k is realized.

  Stage 2. The final prize is drawn from the lottery determined in stage 1.

Independence (I): For any $p, q, r \in L(Z)$ and any $\alpha \in (0,1)$,

$$p \succeq q \quad \text{iff} \quad \alpha p \oplus (1-\alpha) r \succeq \alpha q \oplus (1-\alpha) r.$$

Expected Utility

Continuity (C): If $p \succ q$, then there exist neighborhoods $B(p)$ of p and $B(q)$ of q such that

$$\text{for all } p' \in B(p) \text{ and } q' \in B(q),\ p' \succ q'.$$

Theorem (vNM). Let $\succeq$ be a preference over $L(Z)$ satisfying I and C. There are numbers $\big(v(z)\big)_{z \in Z}$ such that $\succeq$ has the following expected utility representation:

$$U(p) = \sum_{z \in Z} p(z) v(z).$$

- Description of choice, conveniently linear,
- $v(z)$ are called the Bernoulli numbers or vNM utilities,
- The probabilities are uniquely determined and the utilities are unique up to increasing affine transformations.

This means $p \succeq q$ iff $U(p) \geq U(q)$ as defined above.

Subjective Expected Utility

The probabilities $p(z)$ were assumed to be *objective* (e.g. roulette), which is a very special choice setting.

Theory of subjective probabilities developed by Ramsay (1931) and de Finetti (1937) reduces (Knightian) uncertainty to risk.

Savage (1951) adapted vNM expected utility theory to case of subjective probabilities. This is known as *Bayesian decision theory*.

- Requires additional consistency requirements on choice.
- Prior probabilities are updated in response to new information using Bayes' Rule.

Choice in Small Worlds

Dutch book arguments suggest that evolution and learning will lead individuals to be (subjective) expected utility maximizers.

But there are problems:

- Experimental anomalies include Allais (1953) and Ellsberg (1961) paradoxes.
- Bayes' theorem originally meant for frequencies not probabilities. What about rare events?
- Savage recognized that the consistency requirements and 'rational' formation of prior beliefs were only plausible in a "small world".

Considering Large Worlds

Technology adoption, macroeconomics, finance, etc. are large worlds with deep uncertainty.

Many other large-world choices:

- What occupation should I choose?
- What share of profits should I demand?
- Whom should I marry?

These choices are made all the time based on economic/social conventions/norms which are the product of evolution and learning.

Strategic Choice in Dark Worlds

The Centipede Game (Rosenthal 1981):

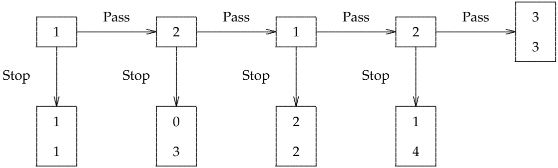

In classical game theory, solved via backward induction…

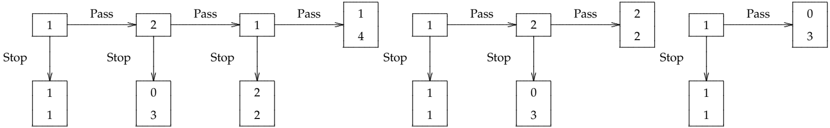

Only $\{(\text{Stop}, \text{Stop}), (\text{Stop}, \text{Stop})\}$ survives backward induction.

But how could players learn what will happen later in the game when (in equilibrium) the game stops immediately?