Bounded Rationality
Lecture 4

Individual Learning

Mikhail Anufriev

EDG, Faculty of Business, University of Technology Sydney (UTS)

European University at St.Petersburg

Faculty of Economics

January, 2014

Outline

Defense of unbounded rationality

1. The as if defense: people do not possess the information gathering and processing abilities to make rational decisions, but they act as if they do:
   1. People learn to behave rationally.
   2. Boundedly rational behavior is driven out by evolutionary pressure or ‘arbitrage opportunities’.
2. Assumption of rational choice disciplines economic modeling, whereas there is a “wilderness of bounded rationality” (with typically an *ad hoc* nature)

Outline

Models of Individual learning

- Given their perceived set of alternatives and perceived consequences agents select the most preferred action...
- ...agents observe the outcome...
- ...agents use this new information and may change their action next time.
  - they learn about set of alternatives and consequences
  - they learn about the behavior of an opponent

In other words, learning means that economic decision makers have certain beliefs about their economic environment and adapt these beliefs as new information comes along.

Learning in a static environment

- Consider a monopolist producing a certain commodity against zero costs.
- Let $P(q)$ be the inverse demand function the monopolist faces, $P(q) \geq 0$ and $P'(q) \leq 0$.
- Solution of individual rationality:

$$q^M = \arg\max_q \{P(q)\, q\}$$

[Write down the f.o.c. and the s.o.c.!]

Learning in a static environment

Assumption: Monopolist knows the entire inverse demand function (or profit function).

Question: How realistic is it to assume that the monopolist has global knowledge of the inverse demand function?

Learning in a static environment

Assumption: Monopolist knows the entire inverse demand function (or profit function).

Question: How realistic is it to assume that the monopolist has global knowledge of the inverse demand function?

Alternative assumption: The monopolist has only local information about the inverse demand function.

Gradient learning

- In period t monopolist produces q_t and observes his profit and associated marginal profit

$$\pi(q_t) \quad \text{and} \quad \frac{\mathrm{d}\pi}{\mathrm{d}q}(q_t)\,.$$

- Monopolist adjusts his output in the direction in which profits increase

$$\begin{aligned} q_{t+1} - q_t > 0 \quad &\text{if} \quad \frac{\mathrm{d}\pi}{\mathrm{d}q}(q_t) > 0 \\ q_{t+1} - q_t < 0 \quad &\text{if} \quad \frac{\mathrm{d}\pi}{\mathrm{d}q}(q_t) < 0 \end{aligned}$$

Gradient learning

- In period t monopolist produces q_t and observes his profit and associated marginal profit

$$\pi(q_t) \quad \text{and} \quad \frac{\mathrm{d}\pi}{\mathrm{d}q}(q_t)\,.$$

- Monopolist adjusts his output in the direction in which profits increase

$$q_{t+1} - q_t > 0 \quad \text{if} \quad \frac{\mathrm{d}\pi}{\mathrm{d}q}(q_t) > 0$$

$$q_{t+1} - q_t < 0 \quad \text{if} \quad \frac{\mathrm{d}\pi}{\mathrm{d}q}(q_t) < 0$$

- For example:

$$q_{t+1} = q_t + \lambda \frac{\mathrm{d}\pi}{\mathrm{d}q}(q_t)\,, \qquad \text{with } \lambda > 0\,.$$

Analyze the resulting dynamics!

Gradient learning

Main results for Gradient Learning:

1. The equilibrium steady-states q^* of the dynamics satisfy the first-order condition.
2. The stable steady-states satisfy the second-order condition.

Gradient learning

Main results for Gradient Learning:

1. The equilibrium steady-states q^* of the dynamics satisfy the first-order condition.
2. The stable steady-states satisfy the second-order condition.
3. If λ is not too high, the gradient learning will converge to a local maximum.
4. If $\lambda > |\frac{2}{\pi''(q^*)}|$, the process does not converge to the steady-state.

Gradient learning

Main results for Gradient Learning:

1. The equilibrium steady-states q^* of the dynamics satisfy the first-order condition.
2. The stable steady-states satisfy the second-order condition.
3. If λ is not too high, the gradient learning will converge to a local maximum.
4. If $\lambda > |\frac{2}{\pi''(q^*)}|$, the process does not converge to the steady-state.
5. The process might converge to the local maximum, which is not global (e.g. for $P(q) = 1 - 10(q-1)^3$)

Econometric learning

- Monopolist's perception of demand is linear:

$$\tilde{P}(q) = a - bq\,,$$

where a and b are positive parameters estimated from the information available.

Econometric learning

- Monopolist's perception of demand is linear:

$$\tilde{P}(q) = a - bq\,,$$

where a and b are positive parameters estimated from the information available.

- At the beginning of time $t+1$, given the perception $a_t - b_t q$, monopolist maximizes profit by producing

$$q_{t+1} = \frac{1}{2} \cdot \frac{a_t}{b_t}\,.$$

Econometric learning

- Monopolist's perception of demand is linear:

$$\tilde{P}(q) = a - bq\,,$$

  where a and b are positive parameters estimated from the information available.

- At the beginning of time $t+1$, given the perception $a_t - b_t q$, monopolist maximizes profit by producing

$$q_{t+1} = \frac{1}{2} \cdot \frac{a_t}{b_t}\,.$$

- The actual price, $P\left(\frac{a_t}{2b_t}\right)$, is typically different from the perceived price, $\frac{a_t}{2}$.

Econometric learning

- Monopolist's perception of demand is linear:

$$\tilde{P}(q) = a - bq\,,$$

  where a and b are positive parameters estimated from the information available.

- At the beginning of time $t+1$, given the perception $a_t - b_t q$, monopolist maximizes profit by producing

$$q_{t+1} = \frac{1}{2} \cdot \frac{a_t}{b_t}\,.$$

- The actual price, $P\left(\frac{a_t}{2b_t}\right)$, is typically different from the perceived price, $\frac{a_t}{2}$.
- The monopolist's perception can be revised.

Econometric learning. Case 1.

- In period t the monopolist produces q_t and observes

$$\big(q_t, P(q_t)\big) \quad \text{and} \quad P'(q_t) < 0\,.$$

Econometric learning. Case 1.

- In period t the monopolist produces q_t and observes

$$\big(q_t, P(q_t)\big) \quad \text{and} \quad P'(q_t) < 0\,.$$

- From this the monopolist estimates a linear demand curve as follows

$$p_t(q) = a_t - b_t\, q = P(q_t) + P'(q_t)(q - q_t)$$

Econometric learning. Case 1.

- In period t the monopolist produces q_t and observes

$$\big(q_t, P(q_t)\big) \quad \text{and} \quad P'(q_t) < 0\,.$$

- From this the monopolist estimates a linear demand curve as follows

$$p_t(q) = a_t - b_t\, q = P(q_t) + P'(q_t)(q - q_t)$$

- ...and the updating of beliefs leads to a new quantity

$$q_{t+1} = \frac{1}{2} \cdot \frac{a_t}{b_t} = -\frac{1}{2} \cdot \frac{P(q_t) - q_t P'(q_t)}{P'(q_t)} = \frac{1}{2} q_t - \frac{1}{2} \frac{P(q_t)}{P'(q_t)}\,.$$

Analyze the resulting dynamics!

Econometric learning. Case 1.

Main results for Econometric Learning I:

1. Steady states of learning model are local optima of profit function.

Econometric learning. Case 1.

Main results for Econometric Learning I:

1. Steady states of learning model are local optima of profit function.
2. Local minima are unstable.

Econometric learning. Case 1.

Main results for Econometric Learning I:

1. Steady states of learning model are local optima of profit function.
2. Local minima are unstable.
3. Local maxima are locally stable if curvature of the inverse demand curve is relatively mild:

$$P''(q^*) < 2\frac{P'(q^*)}{q^*}$$

Econometric learning. Case 2.

- From the first t periods the monopolist observed

$$\Big\{\big(q_0, P(q_0)\big), \big(q_1, P(q_1)\big), \ldots, \big(q_t, P(q_t)\big)\Big\}$$

Econometric learning. Case 2.

- From the first t periods the monopolist observed

$$\Big\{ \big(q_0, P(q_0)\big), \big(q_1, P(q_1)\big), \ldots, \big(q_t, P(q_t)\big) \Big\}$$

- The monopolist now estimates $P(q) = a_t - b_t\, q$ by OLS regression, that is,

$$\begin{pmatrix} a_t \\ -b_t \end{pmatrix} = (X'X)^{-1}X'y\,,$$

where

$$X = \begin{pmatrix} 1 & \cdots & 1 \\ q_0 & \cdots & q_t \end{pmatrix}' \quad \text{and} \quad y = \begin{pmatrix} P(q_0) & \cdots & P(q_t) \end{pmatrix}'\,.$$

Econometric learning. Case 2.

- From the first t periods the monopolist observed

$$\Big\{\big(q_0, P(q_0)\big), \big(q_1, P(q_1)\big), \ldots, \big(q_t, P(q_t)\big)\Big\}$$

- The monopolist now estimates $P(q) = a_t - b_t\, q$ by OLS regression, that is,

$$\begin{pmatrix} a_t \\ -b_t \end{pmatrix} = (X'X)^{-1}X'y\,,$$

where

$$X = \begin{pmatrix} 1 & \cdots & 1 \\ q_0 & \cdots & q_t \end{pmatrix}' \quad \text{and} \quad y = \begin{pmatrix} P(q_0) & \cdots & P(q_t) \end{pmatrix}'\,.$$

- Estimates (a_t, b_t) let monopolist to determine q_{t+1} and $P(q_{t+1})$ which gives new estimates (a_{t+1}, b_{t+1}), etc.

Econometric learning. Case 2.

Main questions for Econometric Learning II:

1. Will such process converge?
2. Does this process converge (only) to local/global maxima?
3. What is the role of initial conditions?

Remarks on Learning Model of Monopolist

- Many different learning models are possible.
  What is the most reasonable one?

Remarks on Learning Model of Monopolist

- Many different learning models are possible.
  What is the most reasonable one?
- Which are possible extensions?

Remarks on Learning Model of Monopolist

- Many different learning models are possible.
  What is the most reasonable one?

- Which are possible extensions?
  1. Monopolist may face costs when switching between production levels
  2. Inverse demand function might be time varying, or subject to random shocks.
  3. ...

Remarks on Learning Model of Monopolist

- Many different learning models are possible.
  What is the most reasonable one?
- Which are possible extensions?
  1. Monopolist may face costs when switching between production levels
  2. Inverse demand function might be time varying, or subject to random shocks.
  3. ...
- The proper objective function for the monopolist should be

$$\max_{\{q_t\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \delta^t \pi(q_t) ,$$

which shows a trade-off between experimenting with production levels to learn about the inverse demand function, and the forgone profits of this experimenting.

Models of Individual Learning: Summary

Important questions:

1. Does learning support the outcomes of the rational choice model?
2. Do economic agents learn the true set of alternatives and consequences?
3. Can models of learning give a better description of certain economic phenomena than the rational choice model does?
4. How should learning be modeled?