Equilibrium Dynamics

I. Introduction

Non-cooperative game theory has made great strides over past 25 years. Game theory has swept through economics leaving virtually no field in economics untouched. Game theory provides a powerful set of tools and approaches. Suite of equilibrium concepts to apply to games of various sorts:

- Nash equilibrium (static game of complete information)
- Subgame perfect equilibrium (dynamic games of complete information)
- Bayesian Nash equilibrium (static games of incomplete information)
- Perfect Bayesian equilibrium – intuitive criterion – and various refinements (dynamic games of incomplete information)

Often the application of equilibrium concepts generates a clear prediction of how players should play the game: equilibrium prediction. Equilibrium is by definition internally consistent.

But what if you start with players who don't necessarily know the equilibrium strategies? How do they get to equilibrium?

Question of equilibrium dynamics:

- Would players learn to play equilibrium?
- How quickly would they do so?
- What is the path to equilibrium?
- Is an equilibrium stable in the sense that a small perturbation would not destroy equilibrium?
- In cases with multiple equilibria, can we predict anything about which equilibrium is likely to be selected? Notion of equilibrium selection.

There is a set of closely related questions regarding the sophistication of players.

- Do people have the requisite cognitive abilities to get to equilibrium? What is required of people in the solution concepts, especially in dynamic games of incomplete information, is often quite difficult.
- Perhaps people use simple rules of thumb. For example, they might tend to increase the probability of playing a strategy that has given a good payoff and avoid strategies yielding low payoffs (reinforcement learning).
- If people play with rules of thumb or with various learning strategies, how does play evolve? Does it approach an outcome that matches equilibrium analysis in non-cooperative game theory?

These are important questions that non-cooperative game theory is not well adept at answering and until recently have not been addressed in a satisfactory way. This has given rise to interest in evolutionary game theory, learning and equilibrium selection: something we will call "equilibrium dynamics."

Need a theory about path to equilibrium – dynamics prior to equilibrium being achieved. Promising strategy is to look explicitly at dynamic analysis where players start with an initial strategy and modify strategy based experience and beliefs about how play will evolve.

Note: such theory is necessarily one of limited rationality. Players do not immediately find equilibrium but grope their way toward equilibrium (perhaps). This approach is useful and realistic for novel situations or inexperienced players.

II. Replicator Dynamics

A. The Basics

There is a large range of possible dynamics. Here we will start with simple dynamics story based on a simple idea inherent in natural selection that strategies that yield higher than average payoffs will get played more often in future rounds and those with below average payoffs will get played less often in future rounds.

Application of dynamics will initially be in cases where there is a large population of players and where payoffs are symmetric.

Symmetric Games: payoffs are the same to all players who play the same strategy when facing the same strategies of other players. Two player case: $\Pi(s_i, s_j)$ is the payoff to either player when they play strategy s_i and other plays s_j. Alternative way of saying this: $\Pi_i(a,b) = \Pi_j(b,a)$ Ex: Prisoner's dilemma. Only need one set of payoffs.

N_{it} = number of players playing strategy i
N_t = total number of players
m different strategies
$x_{it} = N_{it}/N_t$
x_t is vector $(x_{1t}, x_{2t}, \ldots, x_{mt})$ – a mixed strategy for the population

Evolution of x: discrete and continuous time version

Continuous time version: $\dot{x}_{it} = x_{it}(\Pi_i(x_t) - \overline{\Pi}(x_t))$, where $\overline{\Pi}(x_t)$ is the average payoff for the population. Note that we do not need to normalize this to get the dynamics to work out correctly. We must have:

$$\sum_{i=1}^{I} \dot{x}_{it} = 0$$

$$\sum_{i=1}^{I} x_{it}(\Pi_i(x_t) - \overline{\Pi}(x_t))$$

$$= \sum_{i=1}^{I} x_{it}(\Pi_i(x_t)) - \overline{\Pi}(x_t)\sum_{i=1}^{I} x_{it}$$

$$= \overline{\Pi}(x_t) - \overline{\Pi}(x_t) = 0$$

Discrete time version : $x_i(t+1) = x_i(t)\dfrac{\Pi_i(x_t)}{\overline{\Pi}(x_t)}$

Start at some initial set of beliefs, x_0, and then watch the system evolve. Will it reach a resting point (dynamic equilibrium)? Or a limit cycle? Or not converge?

Note: players here have a simple adjustment process. They move towards the strategy that gives higher payoffs. In some sense this is very naïve behavior. It ignores the effect that this reasoning is likely to have on where the system will evolve and what this will imply about likely future payoffs. This approach is like adaptive expectations rather than fully rational expectations. Not a fully rational (or hyper-rational) equilibrium approach.

B. Definitions of Dynamic Equilibrium

x^* in X is a *stationary point* of the differential equation $dx/dt = f(x)$ if $f(x^*) = 0$.

x^* is *stable* if it is a stationary point with the property that for every neighborhood V of x^*, there exists a neighborhood U contained in V such that for x_0 in U then $x(x_0,t)$ contained in V for all $t > 0$.

x^* is *asymptotically stable* if it is stable and there exists a neighborhood W of x^* such that x_0 in W implies $\lim_{t\to\infty} x(x_0,t) = x^*$.

C. Examples of Dynamic Evolution of Simple Games

1. One-dimensional games: 2x2 symmetric games.

Because there are only two strategies, we only need a single number to summarize the vector x. Let p be the proportion of the population playing strategy 1 and $(1-p)$ be the proportion of the population playing *strategy 2*. Knowing p then tells us how the population splits among the two strategies.

a. Hawk-dove game
In the following table we report the payoffs to the rows player as a function of

	H	*D*
	(p)	*(1-p)*
H	-1	2
D	0	1

Payoff to playing *H* as a function of p: $\Pi_H(p) = -1p + 2(1-p) = 2-3p$
Payoff to playing *D* as a function of p: $\Pi_D(p) = 0p + (1-p) = 1 - p$

The average payoff: $$\overline{\Pi}(p) = p\Pi_H(p) + (1-p)\Pi_D(p) = p(2-3p) + (1-p)(1-p)$$

$$\dot{p} = dp/dt = p(\Pi_H(p) - \overline{\Pi}(\text{p}))$$
$$= p\{2-3p - [p(2-3p) + (1-p)(1-p)]\} = p[(1-p)(2-3p) - (1-p)(1-p)]$$
$$= p(1-p)(1-2p)$$

There are three equilibrium points: $p = 0$, $p = 0.5$, $p = 1$. Note: you can't come back from extinction ($p = 0$, or $p = 1$). For interior cases: then $dp/dt > 0$ for $p < ½$, and $dp/dt < 0$ for $p > ½$. $p = ½$ is asymptotically stable. The equilibrium points at the ends $p = 0$, $p = 1$, are unstable.

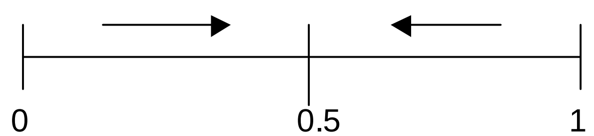

b. Stag-hunt game

The story for this game originates from the writings of Jean-Jacques Rousseau. Rousseau described a situation where two hunters are deciding what to hunt. Hunting stag (a big animal) requires the cooperation of both hunters to be successful, while hunting hare (a small animal) can be done successfully without cooperation – but yields the hunter a lower payoff. The payoffs for this game are summarized below.

	Stag	Hare
	(*p*)	(1-*p*)
Stag	5	0
Hare	3	3

Payoff to playing *S* as a function of p: $\Pi_S(p) = 5p + 0(1-p) = 5p$
Payoff to playing *H* as a function of p: $\Pi_H(p) = 3p + 3(1-p) = 3$
$dp/dt = p\{5p - [p5p + (1-p)3]\} = p(1-p)(5p-3)$

For $dp/dt > 0$ for $p > 3/5$. $dp/dt < 0$ for $p < 3/5$.

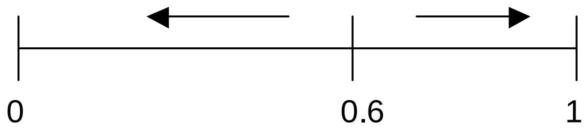

There are three stationary states (Nash equilibria) : $p = 0$, $p = 3/5$, $p = 1$; but only two stable stationary states: 0 and 1. So, what is the prediction for this game? We can get to either equilibrium at $p = 0$ or $p = 1$ depending upon where you start. The basin of attraction is larger for $p = 0$. This equilibrium ($p = 0$), which is to choose *Hare*, is the risk-dominant equilibrium. You never do worse than a payoff of 3 with this choice. This equilibrium though is Pareto dominated. If players are reasonably confident of cooperation in hunting Stag then they will get a higher payoff choosing Stag ($p = 1$). What one should predict in this game depends on how risk averse or cooperative the population is.

c. General 2x2 symmetric game: switch to simple dynamics so that direction of change in strategy is toward strategy giving higher payoff.

Payoff Matrix for game: $A = \begin{bmatrix} a_{11} \, a_{12} \\ a_{21} \, a_{22} \end{bmatrix}$

Current state $s = (p, 1\text{-}p)$, i.e. proportion of population playing first strategy is p, and second strategy is $1\text{-}p$.

Direction of change:

$$\begin{aligned} D(p) &= U(s_1) - U(s_2) \\ &= pa_{11} + (1-p)a_{12} - pa_{21} - (1-p_{22})a \\ &= (1-p)(a_{12} - a_{22}) - p(a_{21} - a_{11}) \\ &= (1-p)a - pb \\ &\text{where } a = (a_{12} - a_{22}) \text{ and } b = (a_{21} - a_{11}) \end{aligned}$$

Three cases:

- $a, b > 0$, unique root of $D(p)$ is $D(p^*) = 0$; $p^* = a/(a+b)$: Mixed strategy Nash eq, stable evolutionary equilibrium

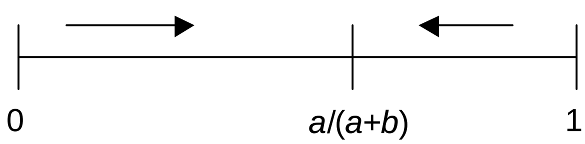

- $a, b < 0$, $D(p^*) = 0$; $p^* = a/(a+b)$: but this is an unstable equilibrium. Two pure strategy Nash equilibrium – movement is toward pure strategies

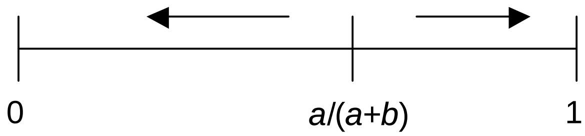

- a, b of opposite signs – then no intermediate point for $D(p) = 0$. Will go to either one or other pure strategy.

2. Two-dimensional games (3x3 symmetric game or 2x2 asymmetric game)

   a. Battle of the Sexes example

   $$A = \begin{bmatrix} 1 & 0 \\ 0 & 2 \end{bmatrix} \; B = \begin{bmatrix} 2 & 0 \\ 0 & 1 \end{bmatrix}$$

   Let p = proportion of A players who play strategy 1; (1-p) is the proportion who play strategy 2
   Let q = proportion of B players who play strategy 1; (1-q) is the proportion who play strategy 2

   For player A:
   the expected payoff from playing strategy 1 is: q
   the expected payoff from playing strategy 2 is: 2(1-q)

   $$\begin{aligned} \dot{p} = dp/dt &= p(\Pi(p) - \overline{\Pi}(\mathrm{p})) \\ &= p(\Pi(p) - [p(\Pi(p) + (1-p)\Pi(1-p)]) \\ &= p(q - pq - 2(1-p)(1-q)) \\ &= p(1-p)(3q-2) \end{aligned}$$

   $dp/dt > 0$ for $q > 2/3$; $dp/dt < 0$ for $q < 2/3$.

   For player B:
   the expected payoff from playing strategy 1 is: $2p$
   the expected payoff from playing strategy 2 is: 1(1-p)

   $$\begin{aligned} \dot{q} = dq/dt &= q(\Pi(q) - \overline{\Pi}(\mathrm{q})) \\ &= q(\Pi(q) - [q(\Pi(q) + (1-q)\Pi(1-q)]) \\ &= q(2p - 2pq - (1-p)(1-q)) \\ &= q(1-q)(3p-1) \end{aligned}$$

   $D(p) = 0$ for $p = 1/3$; $dq/dt > 0$ for $p > 1/3$; $dq/dt < 0$ for p $< 1/3$.

Phase diagram:

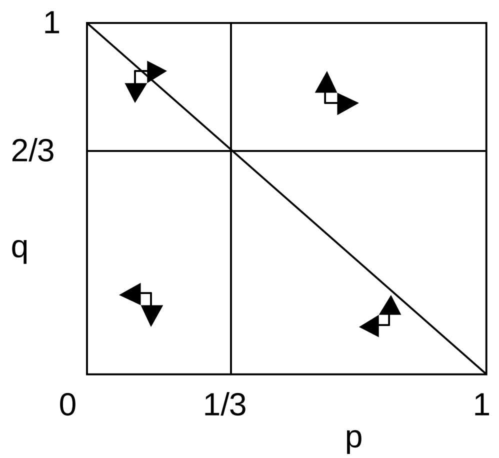

Two stable equilibrium points (0,0) and (1,1). Third Nash equilibrium, mixed strategy Nash eq. (1/3, 2/3), is not stable.

D. Summary

The study of equilibrium dynamics can show which equilibria are stable and which are likely to emerge as equilibrium selection.

Replicator dynamics: simple rules on evolution of strategy choices. Good theory of biological evolution – fitness rules. Replicator dynamics is not necessarily a good theory for describing how strategy of thinking rational people will evolve. For this, we really want a theory of learning – based on results but also on beliefs on how play will evolve. So: on to theories of learning and experimental results on learning.