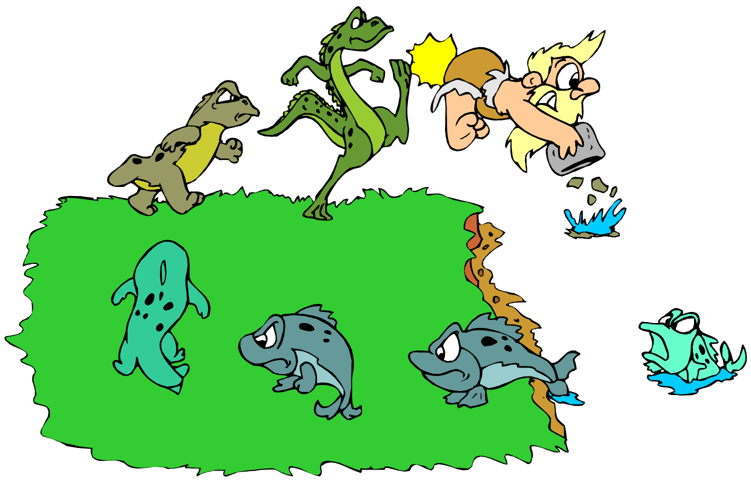

Computational Evolutionary Game Theory

and why I'm never using PowerPoint for another presentation involving maths ever again

Enoch Lau – 5 September 2007

Outline

- What is evolutionary game theory?
- Why evolutionary game theory?
- Evolutionary game theory concepts
- Computational complexity of evolutionary stable strategies
- Evolutionary game theory and selfish routing
- Evolutionary game theory over graphs
- Selection strategies
- Finite populations

What is evolutionary game theory?

Not creationism game theory

Evolutionary game theory (EGT)

- An infinite number of agents in 2-player symmetric games
- Payoffs calculate a fitness used for replication or imitation
- Similarities with conventional game theory
  - Both concerned with the decisions made by agents in a game
  - Equilibria are important concepts for both
- Differences from conventional game theory
  - Rationality of agents not assumed
  - Strategies selected by some force (evolution, cultural factors)
- Higher fitness means more (asexual) reproduction
- Other assumptions: complete mixing

Approaches to evolutionary game theory

- Two approaches
  1. Evolutionary stable strategy: derives from work of Maynard Smith and Price
  2. Properties of evolutionary dynamics by looking at frequencies of change in strategies

Evolutionary stable strategy (ESS)

- Incumbents and mutants in the population
- ESS is a strategy that cannot be invaded by a mutant population
- In an ESS, mutants have lower fitness (reproductive success) compared with the incumbent population
- ESS is more restrictive than a Nash equilibrium
- Not all 2-player, symmetric games have an ESS
- Assumptions very important:
  - If we have a finite number of players, instead of an infinite number, different ESS

Evolutionary stable strategy (ESS)

Finite population simulations on the Hawk-Dove game

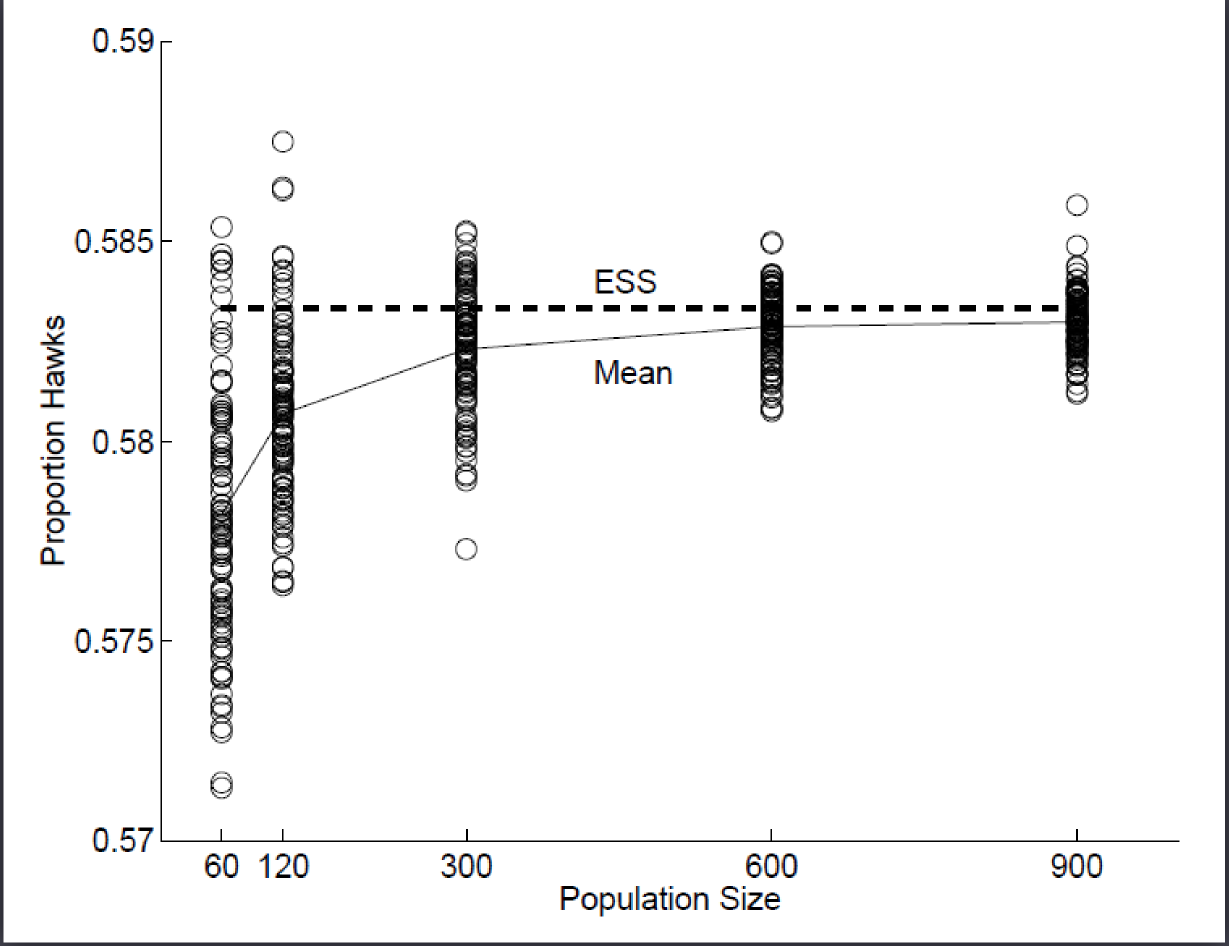

History

- First developed by R.A. Fisher in *The Genetic Theory of Natural Selection* (1930)
  - Attempted to explain the sex ratio in mammals
  - Why is there gender balance in animals where most males don't reproduce?
- R.C. Lewontin explicitly applied game theory in *Evolution and the Theory of Games* (1961)
- Widespread use since *The Logic of Animal Conflict* (1973) by Maynard Smith and Price
- Seminal text: *Evolution and the Theory of Games* (1984) by Maynard Smith

Example: hawks & doves

- Two organisms fighting over a resource, worth V
- Hawks: will fight for the resource, fighting costs C
- Doves: will retreat from aggressive hawks, share resource with other doves
- Example payoff matrix:

	H	**D**
H	-25	50
D	0	15

- Nash equilibrium and ESS given by mixed strategy of (7/12, 5/12)

Why evolutionary game theory?

Why not?

Equilibrium selection problem

- Problems with using Nash equilibria:
  - Not all games have pure Nash equilibria
  - Prisoner's Dilemma: sub-optimality of equilibria
  - Multiple Nash equilibria
- How to choose between different Nash equilibria?
  - Introduce refinements to the concept of Nash equilibria
  - Then how to choose between refinements?

Hyper-rational agents

- Humans sometimes prefer A to B, B to C, and C to A
- EGT can predict behaviour of animals, where strong rationality assumptions fail
- EGT better able to handle weaker rationality assumptions?

Lack of dynamical theory

- Traditional game theory, which is static, lacks the dynamics of rational deliberation
- Could use extensive form (tree form) instead of normal form
  - Quickly becomes unmanageable
  - Presupposes hyper-rational agents
  - Will not learn from observing opponent’s behaviour

Philosophical problems

- Objections to EGT, mainly from application to human subjects
- Measure of fitness in cultural evolutionary interpretations
- Explanatory irrelevance of evolutionary game theory
  - Does EGT simply reinforce existing values and biases?
  - EGT does not provide sufficient evidence for the origin of phenomena
  - Historical records more useful?

Evolutionary game theory concepts

This is where your head is meant to start hurting

Classical model

- Infinite population of organisms
- Each organism assumed equally likely to interact with each other organism
- Fixed, 2-player, symmetric game
- Fitness function F
- A is set of actions
- $\Delta(A)$ is set of probability distributions
- $F: \Delta(A) \times \Delta(A) \rightarrow R$
- $F(s|t)$ = fitness of s playing t
- ε proportion are mutants, $1 - \varepsilon$ are incumbents

Evolutionary stable strategy

- s is an incumbent, t is a mutant
- Expected fitness of an incumbent: $(1 - \varepsilon)\, F(s|s) + \varepsilon\, F(s|t)$
- Expected fitness of mutant: $(1 - \varepsilon)\, F(t|s) + \varepsilon\, F(t|t)$
- s is an ESS if there exists an ε_t such that for all $0 < \varepsilon < \varepsilon_t$, fitness of incumbent > fitness of mutant
- Implies:
  1. $F(s|s) > F(t|s)$, or
  2. $F(s|s) = F(t|s)$ and $F(s|t) > F(t|t)$
- A strategy s is an ESS for a 2-player, symmetric game given by a fitness function F, iff (s, s) is a Nash equilibrium of F, and for every best response t to s, $t \neq s$, $F(s|t) > F(t|t)$

Example: hawks & doves

- Generalised payoff matrix:

	H	D
H	(V – C) / 2	V
D	0	V / 2

- Note that (D, D) is not a Nash equilibrium
  - Cannot be an ESS either
- If V > C:
  - H is an ESS $F(t|H) = p\frac{V-C}{2} + (1-p)0 < \frac{V-C}{2}$
- If V ≤ C:
  - Mixed strategy: H with prob V/C, D with prob 1 – V/C is ESS

Example: hawks & doves

Map of proportions for Hawk-Dove game. Note that where the curve meets the straight line at a gradient of less than 1 (the middle point), that is a stable equilibrium. Where it meets it at a gradient greater than 1, it is an unstable equilibrium.

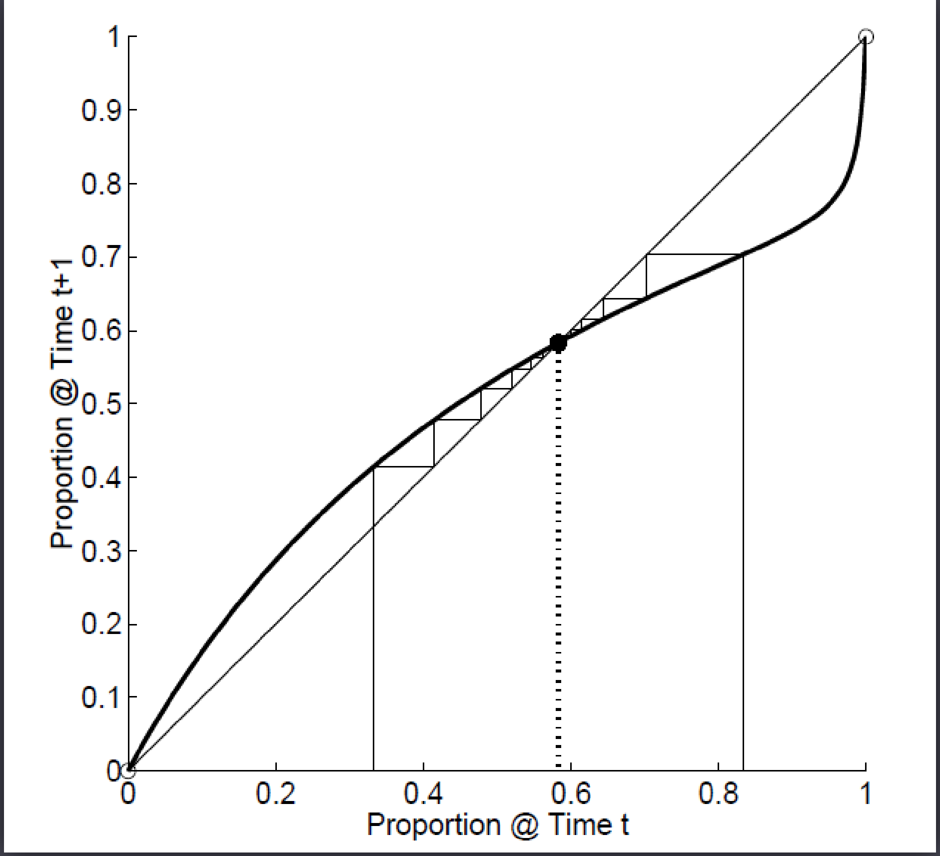

Replicator dynamics

- Continuous dynamics for EGT
- Find differential equations for the change in the proportion of each strategy over time
- In some cases, such as the Prisoner's Dilemma, stable states of replicator dynamics occur when everyone in the population follows the ESS
  - Roughly, true when only two pure strategies exist
  - Can fail to be true with more than two pure strategies

Example: Prisoner's Dilemma

- Generalised payoff matrix

	C	D
C	(R, R')	(S, T')
D	(T, S')	(P, P')

- with T > R > P > S and T' > R' > P' > S'
- Fitness functions

$$F_C = F_0 + p_c \Delta F(C, C) + p_d \Delta F(C, D)$$

$$F_D = F_0 + p_c \Delta F(D, C) + p_d \Delta F(D, D)$$

Example: Prisoner's Dilemma

- Proportion of C and D in next generation:

$$p'_c = \frac{p_c W_c}{W} \qquad p'_d = \frac{p_d W_d}{W}$$

- where W is the overall fitness of population (weighted by proportion)
- Leads to differential equations:

$$\frac{dp_c}{dt} = \frac{p_c(W_c - W)}{W} \qquad \frac{dp_d}{dt} = \frac{p_d(W_d - W)}{W}$$

- Use payoff matrix to show that $p'_d > 0$ and $p'_c < 0$

Computational complexity of evolutionary stable strategies

No good news here

Results and proof outline

- Finding an ESS is both NP-hard and coNP-hard
- Reduction from the problem of checking if a graph has a maximum clique of size exactly k
- Recognising whether a given strategy is an ESS is also coNP-hard
- Transform a graph G into a payoff matrix F, which will have an ESS iff the size of the largest clique in G is not equal to k
  - Transform adjacency matrix: replace all diagonal entries with the value ½, inserting 0^{th} row and 0^{th} column with entries $1 - 1/(2k)$

Proof idea

- For a mixed strategy s to be an ESS, incumbents should receive a relatively high payoff when playing other incumbents
  - When s plays itself, it must guarantee that the pure strategies chosen will correspond to two adjacent vertices
  - Mixed strategy with support over a clique will achieve this
- When max clique is greater than k, uniform mixed strategy corresponding to clique will be an ESS
- When max clique is less than k, get pure strategy ESS
- No ESS in the case where max clique is exactly k

Technical lemma

- If s is a strategy with $s_0 = 0$, then $F(s|s) \leq 1 - 1/(2k')$, where k' is the size of the maximum clique in G. This holds with equality iff s is the uniform distribution over a k'-clique.
- Proof idea
  - By induction over the number of non-edges between the vertices in G
  - Inductive step: Find two non-adjacent vertices u and v, and construct a new strategy s' by moving the probability in s from v to u

Lemmas

1. If C is a maximal clique in G of size k' > k, and s is the uniform distribution on C, then s is an ESS
2. If the maximum size clique in G is of size k' < k, then the pure strategy 0 is an ESS
3. If the maximum size clique of G is at least k, then the pure strategy 0 is not an ESS
4. If the maximum size clique of G is at most k, then any strategy for F that is not equal to the pure strategy 0, is not an ESS for F

Proof of Lemma 1

- By technical lemma, $F(s|s) = 1 - 1/(2k')$
- Any best response to s must have support over only C
  - $F(0|s) = 1 - 1/(2k) < F(s|s)$ by construction
  - Take a u not in C:
    - u is connected to at most $k' - 1$ vertices in C (since max clique size is k')
    - $F(u|s) \leq 1 - 1/k'$ (sum up the entries in the payoff matrix)
    - $F(u|s) < F(s|s)$
- Also by technical lemma, payoff of s is maximised when s is uniform distribution over C
- Hence, s is a best response to itself

Proof of Lemma 1

- Now, need to show that for all best responses t to s, $t \neq s$, $F(s|t) > F(t|t)$ (note: t has support over C)
- By technical lemma, $F(t|t) < 1 - 1/(2k')$ (note: no equality here since $t \neq s$)
- Using F, we can show that $F(s|t) = 1 - 1/(2k')$ (C is a clique, s and t are distributions with support over C)
  - You can get this by summing up the values in the payoff matrix
  - $(k' - ½)/k' = 1 - 1/(2k')$
- Hence, $F(s|t) > F(t|t)$

Proof of Lemma 2

- Mutant strategy t
- F(t|0) = 1 – 1/(2k) = F(0|0)
  - 0 is a best response to itself
- So need to show F(0|t) > F(t|t)
- Form t* by setting the probability of strategy 0 in t to zero and then renormalising
- Applying the technical lemma:
  - F(t*|t*) ≤ 1 – 1/(2k') < 1 – 1(2k) = F(0|t)

Proof of Lemma 2

- Expression for F(t|t):

$$F(t|t) = (2t_0 - t_0^2)\left(1 - \frac{1}{2k}\right) + (1 - 2t_0 + t_0^2)F(t^*|t^*)$$

- By expanding out expressions for F(t|t) and F(t*|t*):
  - F(0|t) > F(t|t) iff F(0|t) > F(t*|t*)

Evolutionary game theory and selfish routing

Ah, something related to my thesis topic

The model

- Each agent assumed to play an arbitrary pure strategy
- Imitative dynamics – switch to lower latency path with probability proportional to difference in latencies
- Recall: at a Nash flow, all s-t paths have the same latency
  - If we restrict the latency functions to be strictly increasing, then Nash flows are essentially ESS
- Paths with below average latency will have more agents switching to them than from them
- Paths with above average latency will have more agents switching from them than to them

Convergence to Nash flow

- As t → ∞, any initial flow with support over all paths in P will eventually converge to a Nash flow
- Use Lyapunov's direct method to show that imitative dynamics converge to a Nash flow
  - General framework for proving that a system of differential equations converges to a stable point
  - Define a potential function that is defined in the neighbourhood of the stable point and vanishes at the stable point itself
  - Then show that the potential function decreases with time
  - System will not get stuck in any local minima

Convergence to approximate equilibria

- ε-approximate equilibrium: Let P_ε be the paths that have latency at least $(1 + \varepsilon)l^*$, and let x_ε be the fraction of agents using these paths. A population is at ε-approximate equilibrium iff $x_\varepsilon < \varepsilon$
  - Only a small fraction of agents experience latency significantly worse than the average latency
- Potential function

$$\Phi(x) = l^* + \sum_e \int_0^{x_e} l_e(u)\,du$$

- Measures the total latency the agents experience
  - Integral: sums latency if agents were inserted one at a time

Convergence to approximate equilibria

- Theorem: the replicator dynamics converge to an ε-approximate equilibrium time $O(\varepsilon^{-3} \ln(l_{max}/l^*))$
  - Proof: see handout

Evolutionary game theory over graphs

Did you know? I am my neighbour's neighbour.

The model

- No longer assume that two organisms are chosen uniformly at random to interact
- Organisms only interact with those in their local neighbourhood, as defined by an undirected graph or network
- Use:
  - Depending on the topology, not every mutant is affected equally
  - Groups of mutants with lots of internal attraction may be able to survive
- Fitness given by the average of playing all neighbours

Mutant sets to contract

- We consider an infinite family $G = \{G_n\}$ (where G_n is a graph with n vertices)
  - Examine asymptotic (large n) properties
- When will mutant vertex sets contract?
  - Let M_n be the mutant subset of vertices
  - $|M_n| \geq \varepsilon n$ for some constant $\varepsilon > 0$
  - M_n contracts if, for sufficiently large n, for all but o(n) of the j in M_n, j has an incumbent neighbour i such that $F(j) < F(i)$
- ε-linear mutant population: smaller than invasion threshold ε'n but remain some constant fraction of the population (isn't a vanishing population)

Results

- A strategy s is ESS if given a mutant strategy t, the set of mutant strategies M_n all playing t, for n sufficiently large, M_n contracts
- Random graphs: pairs of vertices jointed by probability p
  - If s is classical ESS of game F, if $p = \Omega(1/n^c)$, $0 \leq c < 1$, s is an ESS with probability 1 with respect to F and G
- Adversarial mutations: At an ESS, at most o(n) mutants can be of abnormal fitness (i.e. outside of a additive factor τ)

Selection methods

The art of diplomacy

Role of selection

- Dynamics of EGT not solely determined by payoff matrix
- Let the column vector p represent strategy proportions
- F(p) is a fitness function
- S(f, p) is the selection function
  - Returns the state of the population for the next generation, given fitness values and current proportions
- $p_{t+1} = S(F(p_t), p_t)$
- Different selection strategies result in different dynamics
- Any S that maintains stable fixed points must obey $p^{fix} = S(c\ 1, p^{fix})$, and show convergence around p^{fix}

Selection methods

- Some selection methods commonly used in evolutionary algorithms:
  - Truncation
  - (μ, λ)-ES
  - Linear rank
  - Boltzmann selection

Example: Truncation selection

- Population size n, selection pressure k
- Sort population according to fitness
- Replace worst k percent of the population with variations of the best k percent

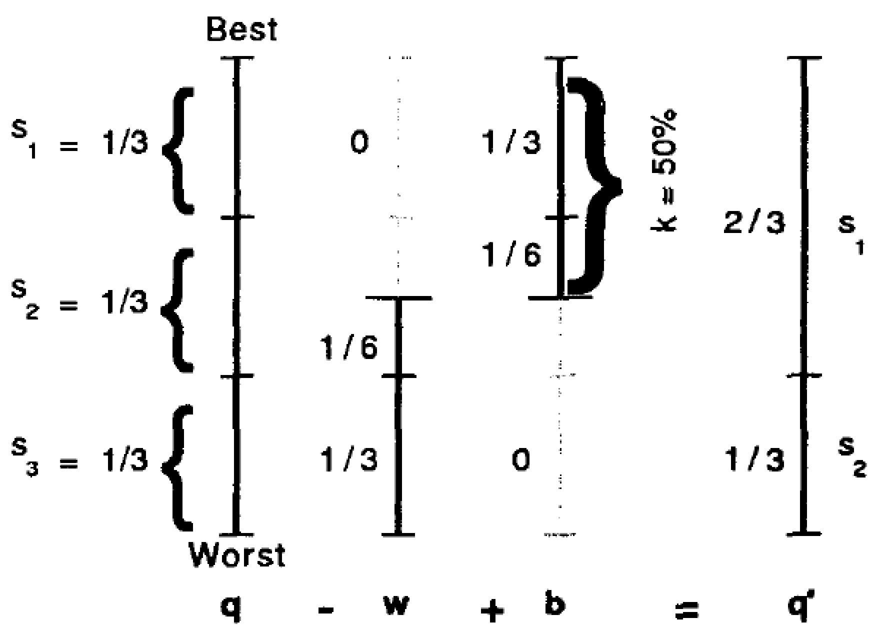

Example: Linear rank selection

- Often used in genetic algorithms
- Agents sorted according to fitness, assigned new fitness values according to rank
- Create roulette wheel based on new fitness values, create next generation
- Useful for ensuring that even small differences in fitness levels are captured

References

Just to prove I didn't make the whole talk up.

References (not in any proper format!)

- Suri S. *Computational Evolutionary Game Theory*, Chapter 29 of *Algorithmic Game Theory*, edited by Nisan N, Roughgarden T, Tardos E, and Vazirani V.
- Ficici S, and Pollack J. *Effects of Finite Populations on Evolutionary Stable Strategies*
- Ficici S, Melnik O, and Pollack J. *A Game-Theoretic Investigation of Selection Methods Used in Evolutionary Algorithms*