On the emergence of social conventions

Michael Repplinger

January 24, 2010

What we are interested in

- We want to model systems of agents trying to achieve different goals
- In such systems, rules helps promote cooperative behavior, increasing efficiency
- Rules could be agreed upon beforehand, or emerge during the process

How to model it

- Classical (static) framework of game theory; social laws in this framework
- Rationality as individual advantage
- Stochastic setting in which certain conventions emerge through repeated interaction
- Analytic and empirical considerations regarding emergence of conventions

The static game

- k-person games, as Normal-form game
- Only symmetric games:
  - Agents have same strategies
  - Identical play, identical payoff
  - Invariant under permutation, i.e. payoff does not depend on role
- Further restriction: 2-person 2-choice games

2-person 2-choice game, general form

$$\begin{pmatrix} x, x & u, v \\ v, u & y, y \end{pmatrix}$$

Coordination game (e.g driving conventions)

$$\begin{pmatrix} 1, 1 & -1, -1 \\ -1, -1 & 1, 1 \end{pmatrix}$$

Cooperation game (Prisoner's dilemma)

$$\begin{pmatrix} 1, 1 & -3, 3 \\ 3, -3 & -2, -2 \end{pmatrix}$$

Strategies and results

- Rational agents are maximizers of individual utility. (*)
- Maximin strategy: highest minimal payoff (find worst payoff, decide on best response). Makes sense if other players might be imbeciles.
- Nash equilibrium is sensible if other players know the game as well and are rational as defined above (as they will not chose a bad payoff for themselves).
- Pareto optimality: if no joint action exists that improves one player while the others stay at least as good.

Applied to our examples

- Coordination: maximin value is -1, both strategies are maximin strategies; 1,1 and 2,2 are Nash equilibria and Pareto optimal.
- Cooperation: maximin value is -2, 2,2 is maximin strategy and Nash equilibrium. Everything *but* 2,2 is Pareto optimal.

Central authority

- Without a central authority, it seems plausible that players defect (even though they suffer).
- Constraints that benefit society as a whole, or each individual could be introduced by an authority.
- The interesting question: can such (globally beneficial) constraints emerge from individual interaction? If so, only if all (or most) players accept a constraint as individually rational.

Definitions

- A social law is a restriction on the set of actions of the agents.
- Game variables: $V(g)$ is the value (or set of values) that can be achieved by a certain strategy wrt. a game g.
- Rational social law wrt. g and V: if $V(g) < V(g_{sl})$. NB: Ordering simple for maximin.
- Social convention: a social law that restricts players to one particular strategy.

Applied to our examples

- Coordination: two rational conventions wrt. maximin value: 1,1 and 2,2.
- Cooperation: unique rational convention: cooperate.

Outline of the model

- Repeated games of individual interaction, with limited information. (*)
- Despite restrictions, potential convergence to global, rational state as if players had complete information.

n-k-g stochastic social game

- n agents, k-person game g, unbounded sequence of tuples of k selected players (out of the n agents).

Selection function

- Which information can agents rely on? Two principles:
  - Obliviousness: identity of agents or names of agents are inaccessible.
  - Locality: selection function based on personal history.
- Local selection function: based on the history of actions and payoffs of the agent alone.
- Semi-local selection function: based on the history of actions and payoffs of the agent and the agents he encountered.

Intuitions behind the formal properties

- Coordination problem could be trivially solved, ahead of time, by reference to "left" or "right", if we would not demand *obliviousness*.
- Better example: specifying assembly robot behavior in general terms, without reference what they should use, but only how they should react given their (payoff) observations.

Highest cumulative reward rule

- An agent switches to a new action iff the total payoff from that action in the last m iterations was greater than the total payoff of the current action in the same time period.
- NB: cannot be pure *locality*, otherwise agents would never switch (since there is no random mutation or error). According to their own definition, it should be considered *semi-local*, I believe, or, equivalently, with some knowledge of the game (i.e. payoff for actions not taken).

Social agreement games

(The class of games to which their results apply)

$$\begin{pmatrix} x, x & u, v \\ v, u & y, y \end{pmatrix}$$

with $x, y, u, v \neq 0$; $x > 0$ or $y > 0$; $u < 0$ or $v < 0$; if both $x > 0$ and $y > 0$, then $x = y$.

Convergence result

- Probability of social convention being reached approaches 1.
- This social convention is stable.
- Every agent receives at least initial maximin value (i.e. maximin in the unrestricted game)
- If a *rational* social convention exists wrt. the maximin value, it will be reached.

What it means

- Local rule brings about emergence of globally reasonable stable convention. (*)
- Emergence of *cooperation* and *coordination* by local update rule.

Efficiency, analytic

- Section is a complete mess, but probably: lower bound (= minimal number of iterations) in the order of $n * log(n)$ (where n is the total number of agents).
- Upper bound is not stated explicitly at all (as far as I can see), but probably in the order of x^n.

Efficiency, empirical

- Despite identical framework, coordination is much faster than cooperation. In fact, cooperation seems to be untenable.
- Coordination game convergence approaches lower bound.
- Assumptions for the following results (unless noted otherwise): unlimited memory, 100 agents

Update delay

95% goal. 1600 iterations. Out of 4000 trials.

Result: high update frequency is good.

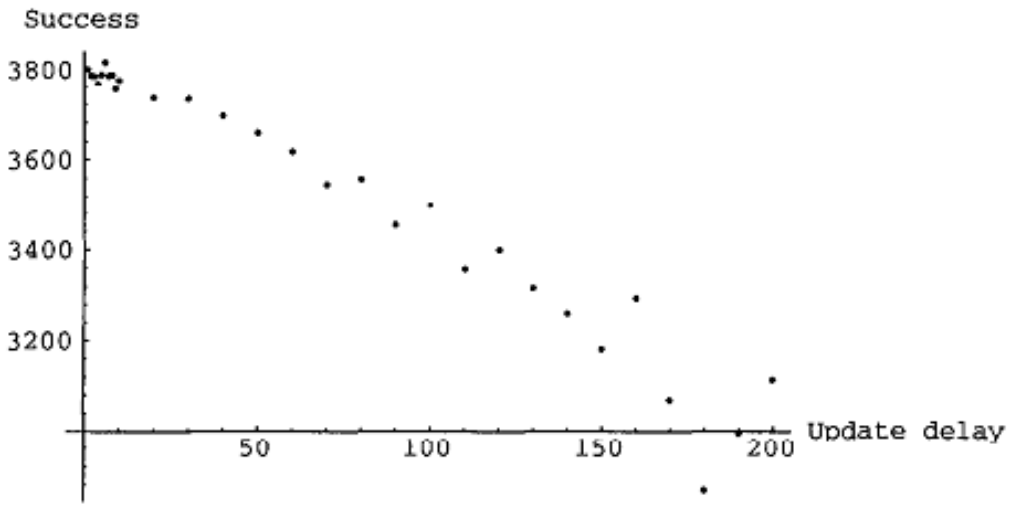

Fig. 1. The effects of update frequency.

Memory restart

85% goal. 800 iterations. Out of 4000 trials.

Result: At least for now, full memory is good.

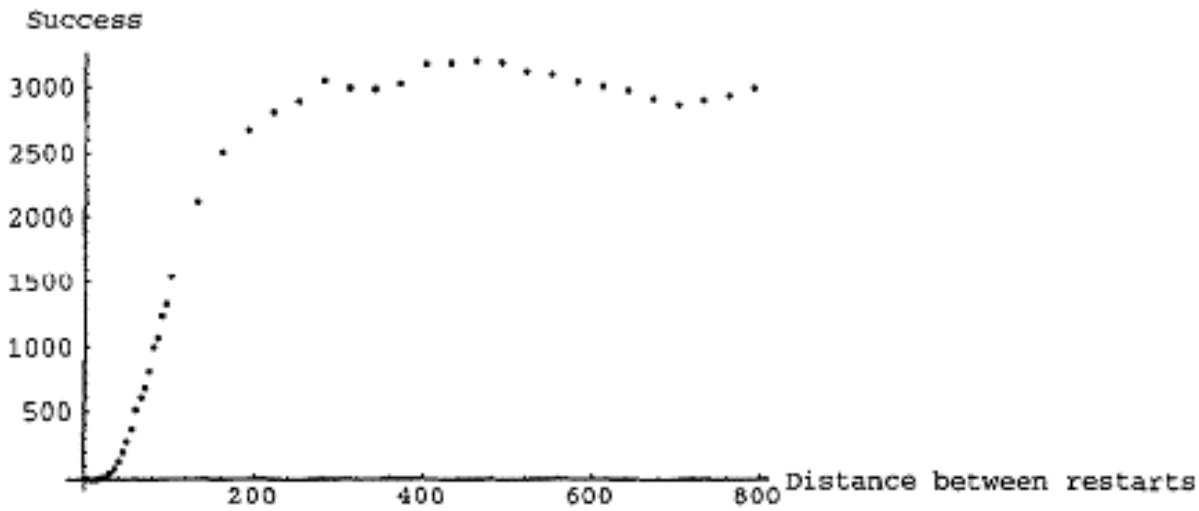

Fig. 2. The effects of memory restarts.

Combined delay and restart

95% goal. 1600 iterations. Out of 4000 trials.

Result: linked restart and update are better than full memory and update delays.

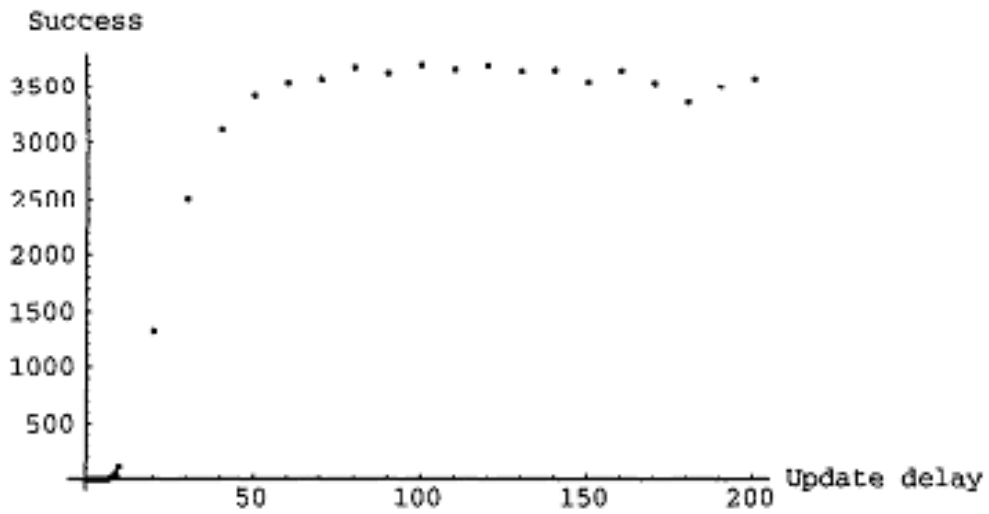

Fig. 3. The case in which update frequency = memory restart frequency.

Limited memory window, observations

85% goal. 800 iterations. Out of 4000 trials.

Result: (Somewhat) limited (continuous) memory is beneficial.

NB: Difference to Fig. 1 probably due to number of iterations.

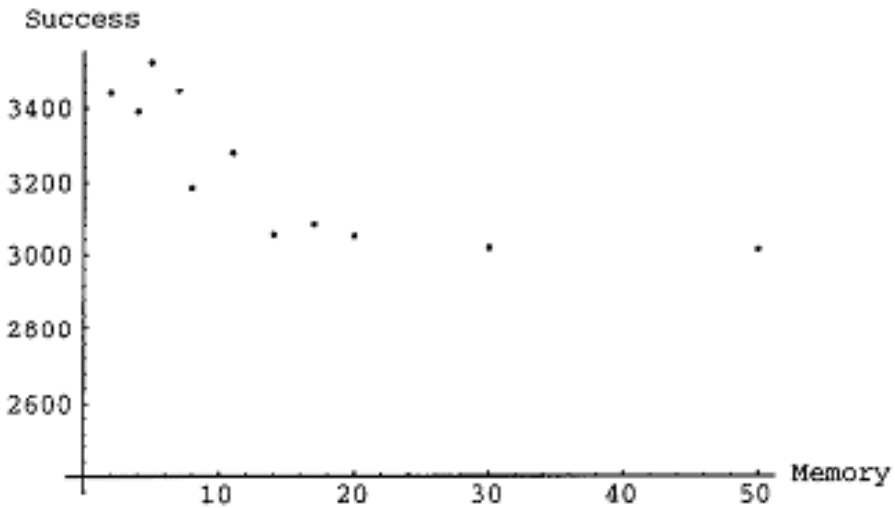

Fig. 4. Limited memory (latest observations).

Limited memory window, iterations

85% goal. 800 iterations. Out of 4000 trials.

Result: Same as for Fig. 4.

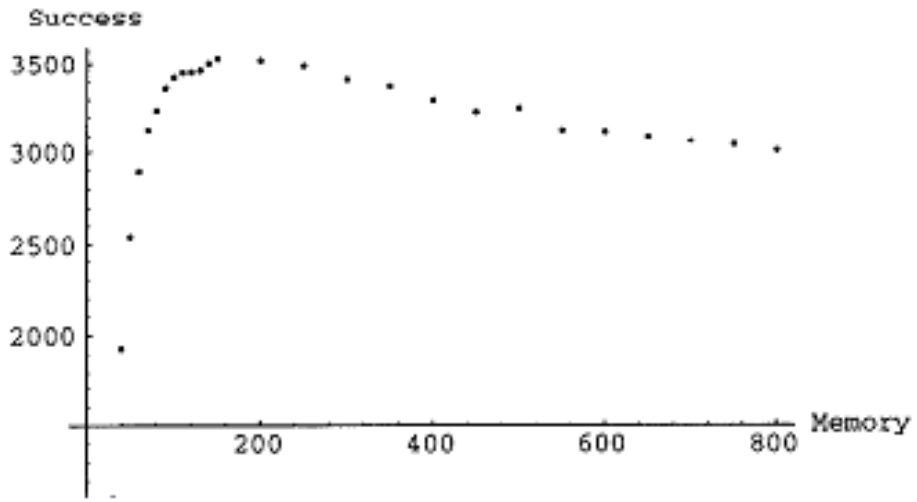

Fig. 5. Limited memory (latest iterations).

Convergence behavior

Full memory.

Result: Slow in the beginning and the end.

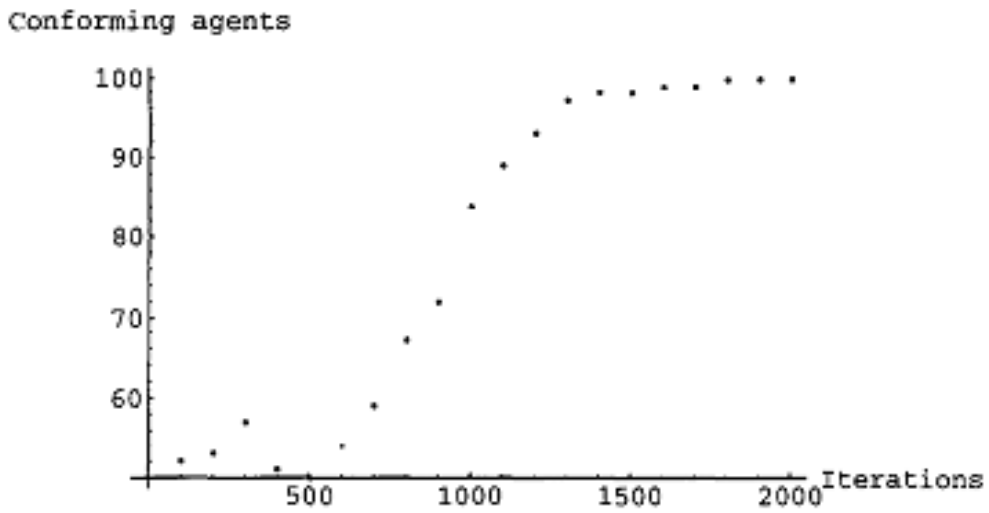

Fig. 6. The shape of convention evolution.

Extended coordination game

Definition: payoff for both agents is $x > 0$ iff the coordinate, $-x$ otherwise.

Result: Number of potential conventions (i.e. more possible actions for players) decreases efficiency, but relatively slowly.

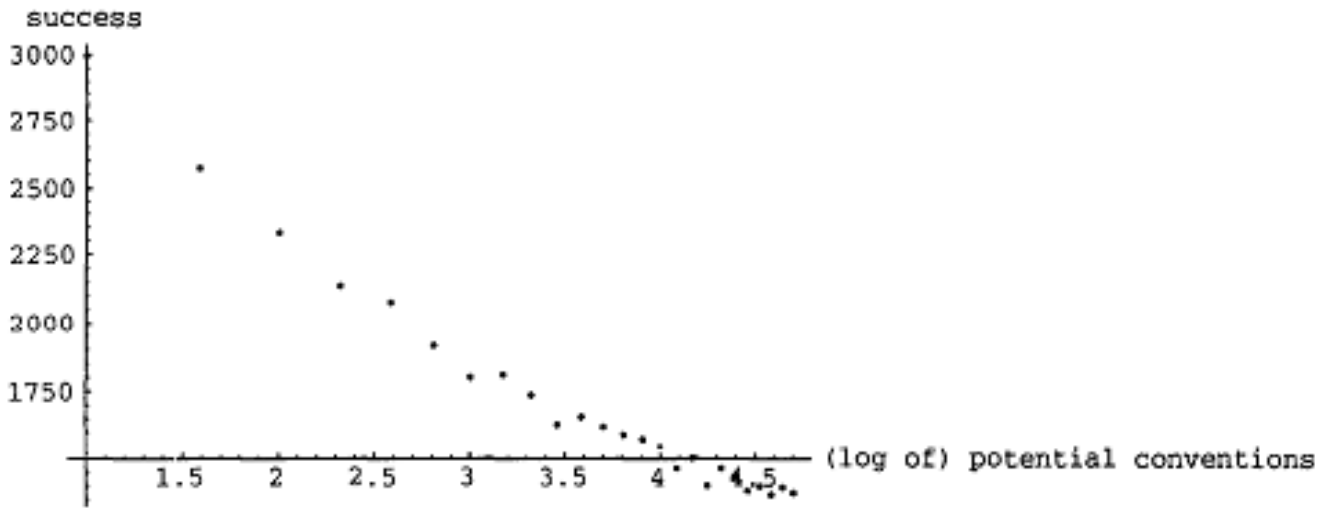

Fig. 7. The effects of the number of potential conventions.

Results, according to the authors

- Interesting because local update rules plus global interaction yields global advantages through stable conventions.
- "Some of our results refer to the emergence of social conventions that are not Nash equilibria". That's great, why didn't you show us those (practical) results?

Comparison to previous frameworks

- Young, "Evolution of Conventions"
  - Similarities: stochastic process (repeated games); classic game theoretical setting
  - Differences: Nash equilibria; best response (global knowledge, but error and limited memory)
- Axelrod "Evolutionary Approach to Norms":
  - Similarities: stochastic process; evolutionary setting
  - Differences: no analytic results (stable states are neither named nor proven), perhaps because they are more complex

Suggestions for discussion

- "Rational agents are maximizers of individual utility" (Always true?)
- "Repeated games of individual interaction, with limited information" (How little do they know; the game, other players' decisions, etc?)
- "Local rule brings about emergence of globally reasonable stable convention." (But only wrt. maximin value.)
- What is the idea behind the class of *social agreement games*?
- No mutation or error assumed in this approach.