Steady State Learning and the Code of Hammurabi

Drew Fudenberg and David K. Levine

5/10/04

Introduction

"If any one bring an accusation against a man, and the accused go to the river and leap into the river, if he sink in the river his accuser shall take possession of his house. But if the river prove that the accused is not guilty, and he escape unhurt, then he who had brought the accusation shall be put to death, while he who leaped into the river shall take possession of the house that had belonged to his accuser." [2nd law of Hammurabi] puzzling to modern sensibilities for two reasons

- ♦ based on a superstition that we do not believe to be true – we do not believe that the guilty are any more likely to drown than the innocent
- ♦ if people can be easily persuaded to hold a superstitious belief, why such an elaborate mechanism? Why not simply assert that those who are guilty will be struck dead by lightning?

we attack these puzzles from the perspective of the theory of learning in games

- ♦ which superstitions survive?
- ♦ Hammurabi had it exactly right: his law uses the greatest amount of superstition consistent with patient rational learning

Overview of the Model

- society consists of overlapping generations of finitely lived players
- indoctrinated into the social norm as children “if you commit a crime you will be struck by lightning”
- enter the world as young adults with prior beliefs that the social norm is true
- being young and relatively patient, having some residual doubt about the truth of what they were taught, and being rational Bayesians, young players optimally decide to commit a few crimes to see what will happen

The Hammurabi Games

Example 2.1: The Hammurabi Game

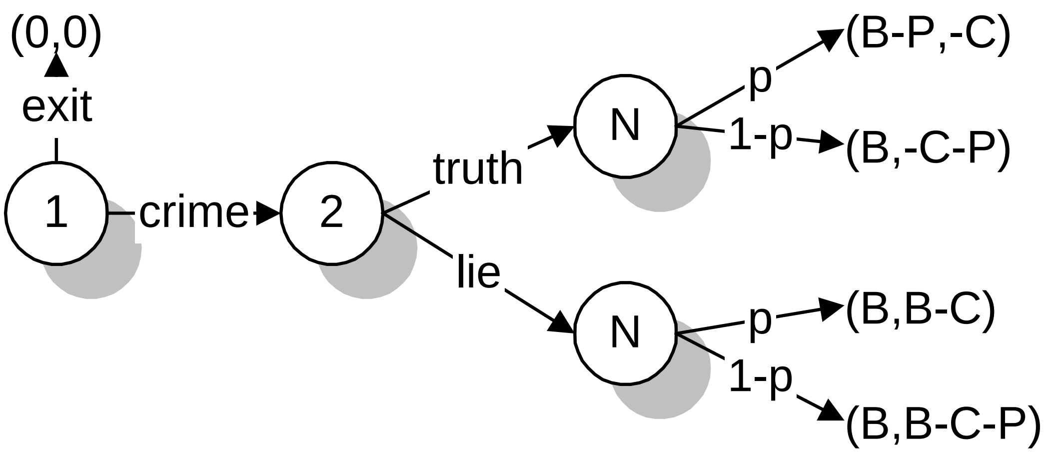

Player 1 is a suspect; player 2 an accuser

C social cost of the **crime**

benefit to accuser of a false accusation, or **lie**, B the same as the benefit of the **crime** to the suspect

the cost of punishment P same for both

assume that the probability of punishment sufficient to deter **crime**

Example 2.2: The Hammurabi Game Without a River

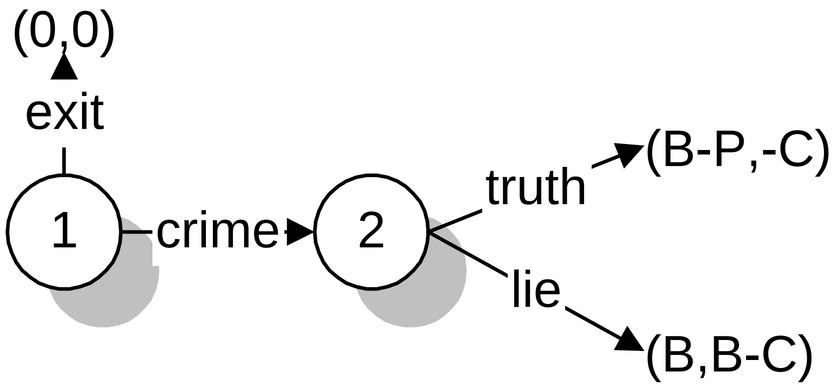

Example 2.3: The Lightning Game

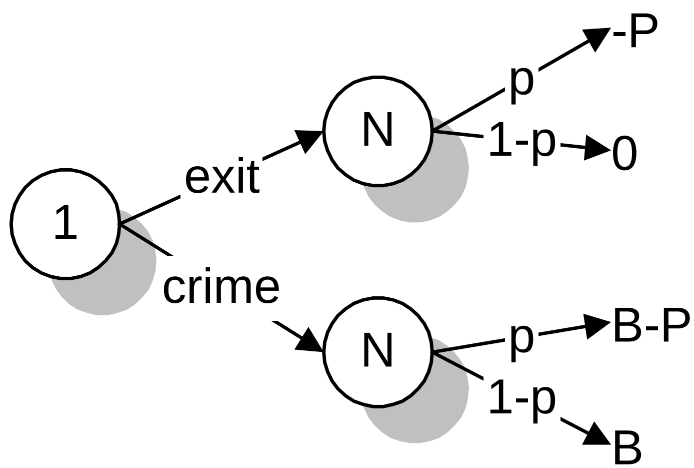

configurations in which there is no crime

Hammurabi game (Nash, but wrong beliefs about off-off path play)

- accuser tells the **truth** because he believes that if he **lie**s he will be punished with probability 1

Hammurabi game without a river (Nash, but not off-path rational)

- accuser tells the **truth,** and is indifferent (ex ante, not ex post)

lightning game (self-confirming, but not Nash)

- everyone believes that if they commit a **crime** they will be punished with probability 1, and that if they **exit** they will be punished with probability p

Simple Games

a *simple game*

- ♦ perfect information (each information set is a singleton node)
- ♦ each player has at most one information set on each path through the tree. (may have more than one information set, but once he has moved, he never gets to move again)

generic condition: no own ties

- ♦ weaker than no ties – allows the Hammurabi games

The Model

nodes in game tree $x \in X$, terminal nodes $z \in Z \subset X$

feasible actions at information sets $A(x)$

pure strategies $s_i \in S_i$, mixed σ_i, the state is θ a mixed profile interpreted as fraction of population playing different pure strategies

payoffs $u_i : Z \to \Re$

I players plus Nature ($I + 1$)

Nature plays a fixed and given mixed strategy σ^0_{I+1}

reachable nodes $Z(s_i), X(s_i)$, $X(\sigma_i)$

nodes reached $\bar{X}(\sigma)$ (the "equilibrium path")

behavior strategies π_i

beliefs about his opponents' play

μ_i a probability measure over Π_{-i}, the set of other players' behavior strategies

beliefs are *independent*: players do not believe that there is a correlation between how an opponent plays at different information sets, or how different opponents play

$p_i(x \mid \mu_i)$ marginal induced by beliefs

preferences:

$$u_i(s_i, \mu_i) \equiv u_i(s_i, p_i(\cdot \mid \mu_i)) \equiv \sum_{z \in Z(s_i)} p_i(z \mid \mu_i) u_i(z).$$

when μ_i is has a continuous density g_i we write $p_i(x \mid g_i), u_i(s_i, g_i)$.

Static Equilibrium Notions

Self-Confirming Equilibrium

Definition 4.1 : $\bar{\theta}$ *is a* self-confirming equilibrium *if for each player* i *and for each* s_i *with* $\bar{\theta}_i(s_i) > 0$ *there are beliefs* $\mu_i(s_i)$ *such that*

(a) s_i *is a best response* $\mu_i(s_i)$ *and*

(b) $\mu_i(s_i)$ *is correct at every* $x \in \bar{X}(s_i, \bar{\theta}_{-i})$,

Note also that *Nash equilibrium* strengthens (b) to hold at all information sets.

Subgame Confirmed Nash Equilibrium

Definition 4.2:

(a) In a simple game, node x is one step off the path *of π if it is an immediate successor of a node that is reached with positive probability under π.*

(b) Profile π is a subgame-confirmed Nash equilibrium *if it is a Nash equilibrium and if, in each subgame beginning one step off the path, the restriction of π to the subgame is self-confirming in that subgame.*

In a simple game with no more than two consecutive moves, self-confirming equilibrium for any player moving second implies optimal play by that player, so subgame-confirmed Nash equilibrium implies subgame perfection.

can fail when there are three consecutive moves.

Example 4.1: The Three Player Centipede Game

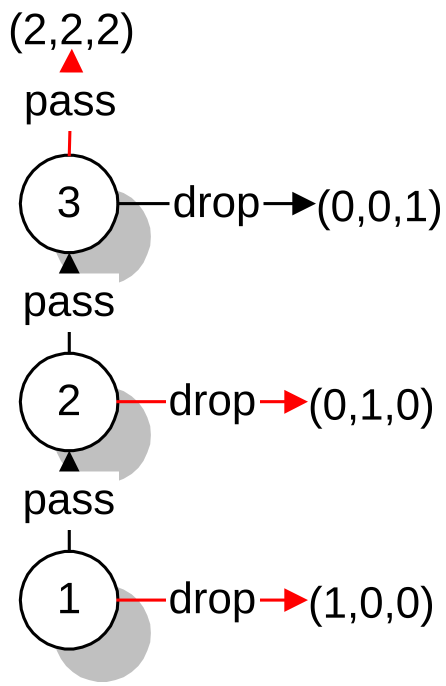

unique subgame-perfect equilibrium: all players to **pass**

(**drop**, **drop**, **pass)** is subgame-confirmed

Rational Steady-State Learning

The Agent's Decision Problem

"agent" in the role of player *i* expects to play game T times wishes to maximize

$$\frac{1-\delta}{1-\delta^T} E\sum_{t=1}^{T} \delta^{t-1} u_t$$

u_t realized stage game payoff

agent believes that he faces a fixed time invariant probability distribution of opponents' strategies, unsure what the true distribution is

Definition 5.1: *Beliefs* μ_i *are* non-doctrinaire *if* μ_i *is given by a continuous density function* g_i *strictly positive at interior points.*

Note that allow priors can go to zero on the boundary, as is the case for many Dirichlet priors assume non-doctrinaire prior g_i^0

$g_i(\cdot \mid z)$ posterior starting with prior g_i after z is observed

agents are assumed to play optimally

(dynamic programming problem defined in the paper)

histories are Y_i

optimal policy a map $r_i : Y_i \to S_i$ (may be several)

Steady States in an Overlapping generations model

- ♦ a continuum population
- ♦ doubly infinite sequence of periods
- ♦ generations overlap
- ♦ $1/T$ players in each generation
- ♦ $1/T$ enter to replace the $1/T$ player who leave
- ♦ each agent is randomly and independently matched with one agent from each of the other populations

each population assumed to use a common optimal rule r_i

given population fractions of each population playing pure strategies $\overline{\theta}_i(s_i)$

Using r we work out the fraction of the population with each experience $\theta_i(y)$

then recompute the fractions playing different strategies

$$f_i[\overline{\theta}](s_i) = \sum_{\{y_i | r_i(y_i) = s_{i\}}} \theta_i(y_i)$$

This is a polynomial map from the space of mixed strategy profiles to itself

a fixed point exists, and these fixed points are *steady states*.

Patient Stability

a sequence of steady states $\lim_{T\to\infty} \overline{\theta}^T \to \overline{\theta}$ we say that $\overline{\theta}$ is a g^0, δ-*stable state*

If $\overline{\theta}(\delta)$ are g^0, δ-*stable states* and $\lim_{\delta\to 1} \overline{\theta}(\delta) \to \overline{\theta}$, we say that $\overline{\theta}$ is a *patiently stable state.*

Theorem 5.1: *(Fudenberg and Levine [1993b])* g^0, δ-steady states are self-confirming; patiently stable states are Nash.

Patient Stability in Simple Games

two profiles $\overline{\theta}, \overline{\theta}'$ are *path equivalent* if they induce the same distribution over terminal nodes.

Theorem: *In a simple game, a patiently stable state* $\overline{\theta}$ *is path equivalent to a subgame-confirmed Nash equilibrium.*

corollary of a more general theorem; note that it is straightforward to show that a patiently stable state in a simple game must be Nash in weakly undominated strategies, which eliminates the "bad equilibrium" in Hammurabi without the river

key result is a converse for simple games a profile is *nearly pure* if Nature does not randomize on the equilibrium path, and no player except Nature randomizes off the equilibrium path

our proposed Hammurabi game profile is nearly pure – only Nature randomizes, and only off the equilibrium path

Theorem: *In simple games with no own ties, a subgame-confirmed Nash equilibrium that is nearly pure is path equivalent to a patiently stable state.*

This answers the Hammurabi puzzle: the Hammurabi equilibrium with the river is patiently stable; without the river it is not, nor is the lightning equilibrium stable

Games with Length at Most Three

a game has “length at most three” if no path through the tree hits more than three information sets

Theorem *In simple games with no own ties, no Nature’s move and length at most three, a subgame-confirmed Nash equilibrium is path equivalent to a patiently stable state.*

because in these games all equilibria are nearly pure

Lemma: *In simple games with no own ties, no Nature’s move and length at most three, a subgame-confirmed Nash equilibrium is path equivalent to a subgame-confirmed Nash equilibrium in which players play pure strategies.*

in turn follows from

Lemma: *In simple games with no own ties, no Nature’s move and length at most two, every self confirming equilibrium is path equivalent to a public randomization over Nash equilibria.*