האוניברסיטה העברית בירושלים
THE HEBREW UNIVERSITY OF JERUSALEM

EVOLUTIONARY STABILITY IN
FINITE STOPPING GAMES UNDER A
FAST BEST-REPLY DYNAMICS

By

ZIBO XU

Discussion Paper # 632   **Jan 2013**

מרכז לחקר הרציונליות

CENTER FOR THE STUDY
OF RATIONALITY

Evolutionary stability in finite stopping games under a fast best-reply dynamic

Zibo Xu[1]

Center for the Study of Rationality

The Hebrew University of Jerusalem

17th December 2012

Abstract

We consider a fast evolutionary dynamic process on finite stopping games, where each player at each node has at most one move to continue the game. A state is evolutionarily stable if its long-run relative frequency of occurrence is bounded away from zero as the mutation rate decreases to zero. The fast dynamic process allows each individual in each population to change its strategy at every stage. We define a robustness index of backward induction and show examples where the backward induction equilibrium component is not evolutionarily stable for large populations. We show some sufficient conditions for evolutionary stability, which are different from the ones for the conventional evolutionary model. Even for this fast dynamic process,

the transition between any two Nash equilibrium components may take very long time.

1   Introduction

In [6] and [7], we have discussed the equilibrium refinement of evolutionary stability in two-player extensive-form games of perfect information. In an evolutionary process of selection and mutation, we say that a state is *evolutionary stable* if its long-term relative frequency of occurrence is bounded away from zero as the mutation rate decreases to zero. The model is adapted from Hart [2] and Gorodeisky [1]. Given an extensive-form game, we give an associated population game where for each player there is a population playing the game in that player's role. Our model in [6] and [7] requires that, at every stage, one random individual in each population is chosen. With very small probability, a mutation happens and the individual picks a strategy randomly. When a selection happens with high probability, the individual picks a best-reply strategy against the current distribution of all other populations.

To check the evolutionary stability of each Nash equilibrium component, we study the transition under the best-reply dynamics from each component triggered by a mutation. If, from an evolutionarily stable component, with probability bounded from below when populations approach infinity, it transits to another component, then the latter component is also evolutionarily stable. We call such best-reply process triggered by a single mutation an *one-mutation transition*.

We however have not given a characterisation theorem of the evolutionary stability in extensive-form games. Indeed, as we comment in [6] and [7], the one-mutation transition can be subtle and the dynamics of distribution of populations needs various combinatorial results. One may wonder whether a simplified model can shed light on the general rules in evolutionary process on extensive-form games of perfect information.

We turn to the classic paper of Kandori, Mailath and Rob [4]. They consider a *Darwinian* property, where the proportion of a population using a currently best-reply strategy increases at each stage. Their conclusion is robust with respect to any possible increase of individuals currently in best reply. In our model, if we consider the extreme case that at each stage every individual can adopt their currently best-reply strategy when selection takes effect, can we reach a simple and clear characterisation result? Under what conditions, does the stability result coincide with the one in [6] and [7]?

This is the main topic we study in this paper. Note that every evolutionary dynamics process corresponds to a Markov chain. We define the probability of the backward induction equilibrium component in the invariant distribution to be a *robustness index of backward induction*. We concentrate on finite stopping games in this paper (see Definition 3.3), as they capture the basic tree property, while being free of dynamics of strategies involving moves at multiple branches. The extreme evolutionary process above can be viewed as a fast learning process disturbed by small errors. It essentially eliminates the inertia propositions in our previous model in [6] and [7]. So one should expect some different evolutionary results (see Section 4.1). Therefore, the fast evolutionary dynamics cannot be simply regarded as analogous to those addressed by the result of the previous model but involving less transition time. This fast learning model may explain the interaction in a society where everyone has the same information and rationality at the same time. Moreover, they are impatient and can move quickly. We show that, in the case of large populations, after the initial stage, with asymptotic probability 1, every population is either with all individuals playing a pure strategy or in an approximate uniform distribution of all (reduced) strategies terminated in a subgame.

The result still holds that only Nash equilibria can be evolutionarily stable for any population size. We can further show that, from any state that is not a Nash equilibrium, it takes only finitely many steps, independent of population size, to enter a state of Nash equilibrium. So we only need to study the transitions between some specific equilibria (populations in pure strategies or uniform distribution) included in the corresponding components. When each transition between components is an one-mutation transition, the dynamic process can be simplified by a stationary Markov chain with states of those specific equilibria (see irreducibility condition and the definition of recurrent class in Section 3.3), and the robustness index of backward induction follows easily. We shall give examples where the backward induction equilibrium component is not evolutionarily stable for large populations.

While we can find some sufficient conditions for evolutionary stability in this simplified model, we still need to face the same problem discussed in [7]: the stability result can be very sensitive to payoff vectors. There exists an extensive-form game such that when populations increase to infinity the probability of a best-reply transition triggered by any mutant between any two components is approaching zero. Thus there is no one-mutation transition between any two components.

We show such an example in Section 3.5. The source of transitions between components in this game is the event that the distribution of a population at a disconnected node is uneven enough to trigger some further transition. We can show each such strategy is selected with equal probability. So, when populations increase to infinity, the probability of such event decreases to zero. We can view this small-probability event as an inconsistency of the strong law of large numbers. The evolutionarily stable component is the one with smaller probability that the distribution of the population at

disconnected nodes (triggered by a mutant) passes the threshold and leads to a deviation transition. Because such probability is so small, the transition between components takes very long time. (A simple estimate shows it take much more than $e^m \mu^{-1}$ stages where m is the population size and μ is the mutation rate.) This example shows that, even for fast evolutionary dynamics, the evolutionarily stability result can be intrinsically dependent on payoff vectors, as the payoffs determine the threshold in the ratio of population proportions that must be reached in each direction of transition.

We discuss the variation of the model in Section 4.

2   Model

We adopt the model used in [6] and [7] but with a slight change on the number of individuals who can change their strategies at each stage.

We consider a generic finite N-player game Γ in extensive form with perfect information. In the generic assumption, no player obtains the same utility at any two terminal vertices. Thus, the backward induction equilibrium in Γ is unique. Without loss of generality, we assume that Γ does not contain any chance move. Therefore, a vertex in the game is either a (decision) node or a terminal vertex attached with a payoff vector. We denote the strategy set of player i to be A^i, $i = 1, ..., N$. Put $A := \prod_{1 \leq i \leq N} A^i$. For every a in A players receive payoffs determined by a in Γ. Note that in this paper, we only consider *reduced-form strategies*. That is, we view two strategies of a player the same strategy if it is impossible to reach different payoff to this player from these two strategies. For instance, given the game in Figure 1, conventionally, for a strategy of player I where she plays the action a_1 at the top node, that strategy must also specify the action she would play at the

bottom node. We, however, do not specify it, as it is impossible to reach the bottom node in that case. Hence, in this one-player game, there are only three strategies a_1, a_2 and a_3 in our definition.
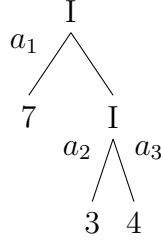

Figure 1: Example to the definition of strategies

For each player i, the set of mixed strategies of player i is a unit simplex on A^i, and we denote it by

$$X^i := \triangle(A^i) = \{x^i \in \mathbb{R}_+^{A^i} : \sum_{a^i \in A^i} x_{a^i}^i = 1\}.$$

We then denote the set of mixed strategy profiles by $X := \prod_{1 \le i \le N} X^i$. Let $u^i : A \to \mathbb{R}$ be the payoff function of player i.

We consider the following *population game* associated to Γ. Assume that for each i with $1 \le i \le N$ there is a non-empty population $M(i)$ of individuals playing the game in the role of player i. Without loss of generality, we put $|M(i)| = m$ for all i. We define a state w to be the collection of pure strategies of each individual in each population, i.e., $w = (w^i)_{1 \le i \le N}$ where $w^i := (w^i(q))_{q \in M(i)}$ and $w^i(q) \in A^i$ for each i. Let the state space W be $\prod_{1 \le i \le N} (A^i)^{M(i)}$. At a state w, for each i, let $m_{a^i}(w)$ be the number of individuals playing strategy $a^i \in A^i$, and let $x_{a^i}^i(w)$ be the proportion of population $M(i)$ that plays the strategy a^i, i.e., $x_{a^i}^i(w) = m_{a^i}(w)/m$. We put $x^i(w) := (x_{a^i}^i(w))_{a^i \in A^i}$, and thus we may view each state w as an N-tuple mixed strategy $x(w) = (x^i(w))_i$ in X.

We define a discrete-time Markov chain $(Z_t)_{t \geq 0}$ such that each Z_t is a random variable valued in W. Note that for a rigorous analysis, one needs to define first the sample space $\Omega := W^{\mathbb{N}}$, which is the necessity for the definition of this Markov chain and its transition probabilities. We shall not refer to it explicitly, when the context allows. Given the population size m and mutation rate μ, the transition probabilities $P_{\mu,m}$ of this Markov chain specify the probability that Z_{t+1} equals a state $\tilde{w}$ given a history $Z_1 = w_1, ..., Z_t = w_t$, i.e., $P_{\mu,m}(Z_{t+1} = \tilde{w} | Z_1 = w_1, \ldots Z_t = w_t)$, for each $t \geq 0$. We further assume $(Z_t)_{t \geq 0}$ to be a stationary Markov chain dictated by an one-step transition probability matrix $Q_{\mu,m}$, i.e.,

$$P_{\mu,m}(Z_{t+1} = \tilde{w} | Z_1 = w_1, \ldots Z_t = w_t) = Q_{\mu,m}(\tilde{w} | w_t),$$

for every $w_1, ..., w_t, \tilde{w}$ in W and $t = 1, 2, \ldots$. If context allows, we drop the subscript μ and m of P and Q; for a strategy a of some player i, at stage t, we also abbreviate $m_a(Z_t)$ by $m_a(t)$ or simply $A(t)$, which is the random number of individuals in population $M(i)$ playing the strategy a in the Markov chain at stage t. We view Q as a transition rule concerning 'selection' and 'mutation' described in the model below consistent with Darwinian property.

Let $\mu > 0$ and $\sigma > 0$ be given such that $\mu + \sigma \leq 1$. Let k be a natural number no more than m. We suppose that the current state is w. At each stage, for each i with $1 \leq i \leq N$, a set S^i of k individuals in $M(i)$ is randomly chosen with probability $\binom{m}{k}^{-1}$. For each i, all individuals in $M(i)$ except those in S^i do not change their strategies. The following process is performed independently for each chosen individual $q^i \in S^i$ and for each population $M(i)$.

- The chosen q^i undergoes mutation, selection or no change, with probability μ, σ and $1 - \mu - \sigma$, respectively.

- Conditional on mutation, the individual q^i chooses a random strategy, i.e., $\tilde{w}^i(q^i) = a^i$ with probability $1/|A^i|$ for each a^i in A^i.

- Conditional on selection, we define a set of 'best strategies', namely

$$BS^i(w) := \arg \max_{a^i \in A^i} u^i(a^i, w^{-i}). \tag{2.1}$$

  where w^{-i} indicates the collection of pure strategies of each individuals in all populations other than $M(i)$. If no ambiguity, we abbreviate the notation $BS^i(w)$ by BS^i. If $w^i(q^i) \notin BS^i$, then the new strategy $\tilde{w}^i(q^i)$ of q^i is a randomly chosen best strategy, i.e., $\tilde{w}^i(q^i) = a^i$ with probability $1/|BS^i|$ for each $a^i \in BS^i$. If $w^i(q^i) \in BS^i$, then there is no change in the strategy of q^i.

- For the case of no change, the strategy of q^i does not change: $\tilde{w}^i(q^i) = w^i(q^i)$.

Thus, the one-step transition probabilities $Q(\tilde{w}|w)$ is well defined.

The combined results from Hart [2] and Gorodeisky [1] are restricted within the case that each player can only play at one node in a game, and at each stage only one individual can change its strategy, i.e., $k = 1$. We have studied the case of an extensive-form game where a player can move at multiple nodes in [6] and [7]. In this paper, we turn to the problem of $k > 1$, in particular the extreme case of $k = m$. This corresponds to a kind of fast adaptive learning, and the evolutionary result can be different from what we have seen in [6] and [7].

3 Results

3.1 Definitions and previous results

To give the formal definition of evolutionary stability, we study the behaviour of the steady-state analysis and limiting distribution of the Markov chain $(Z_t)_{t \geq 0}$ with the one-step transition probability matrix $Q_{\mu,m}$ generated from the transition process in Section 2. It is straightforward to check that, for every finite m and every $\mu > 0$, the induced Markov chain is irreducible and aperiodic. Hence, for every pair of population size m and mutation rate μ, there exists a unique invariant distribution $\pi_{\mu,m} \in \triangle(W)$ on W such that $\pi_{\mu,m} = \pi_{\mu,m} Q_{\mu,m}$, or

$$\pi_{\mu,m}[\tilde{w}] = \sum_{w \in W} \pi_{\mu,m}[w] Q_{\mu,m}(\tilde{w}|w)$$

for every $\tilde{w} \in W$.

Lemma 3.1 *Given an one-step transition probability matrix $Q_{\mu,m}$ defined in Section 2, for a fixed population size m and any k with $0 < k \leq m$, the limit distribution $\lim_{\mu \to 0} \pi_{\mu,m}$ exists.*

Proof. It immediately follows from Theorem 3.1 in [8]. □

For a fixed population size m, a state w in W is called *m-evolutionarily stable* if its invariant probability $\pi_{\mu,m}[w]$ does not go to zero as μ decreases to zero, i.e., $\lim_{\mu \to 0} \pi_{\mu,m}[w] > 0$. We call a state w *evolutionarily stable for large populations* if $\liminf_{m \to \infty} \lim_{\mu \to 0} \pi_{\mu,m}[w] > 0$. Given a limiting process that m increases to infinity, μ decreases to zero, and that μm is always in a set $R \subseteq \mathbb{R}^+$, we call a state w *evolutionarily stable in this limiting process* if $\pi_{\mu,m}[w]$ is bounded away from zero throughout this process. All definitions above can be naturally extended to subsets included in the state

space. For every set of mixed strategy profiles $Y \subseteq X$, if no ambiguity, we abbreviate $\pi_{\mu,m}[w \in W : x(w) \in Y]$ by $\pi_{\mu,m}[Y]$. Recall that the game we consider is generic. Denote the unique backward induction equilibrium by $BI = (s^i)_{1 \leq i \leq N}$, and denote the ϵ-neighbourhood of BI by

$$BI_\epsilon := \{x \in X : x^i_{s^i} \geq 1 - \epsilon \; \forall 1 \leq i \leq N\}.$$

It is straightforward to see that Theorem 4.4 in [6] can be generalised for any $k \leq m$, and we have the following corollary.

Corollary 3.2 *Under the best-reply selection mechanism, given a finite N-player extensive-form game of perfect information, for any $k \leq m$, it has*

$$\lim_{\mu \to 0} \pi_{\mu,m}[NE] = 1$$

for all m in $\mathbb{N}$, and hence

$$\lim_{m \to \infty} \lim_{\mu \to 0} \pi_{\mu,m}[NE] = 1.$$

One can prove the unique evolutionary stability of BI for large populations under the condition that each player can only move at one node by the approach presented in [5] and then applied in [4]. That is,

$$\liminf_{m \to \infty} \lim_{\mu \to 0} \pi_{\mu,m}[BI] = 1.$$

We omit the proof here.

We only consider *finite stopping games* in this paper, as it is the first step to understand general extensive-form games. See Section 4.2 for the discussion on general extensive-form games.

Definition 3.3 *A finite stopping game is a finite extensive-form game of perfect information where each node has at most one immediate succeeding node and all his other immediate children are terminal vertices.*

The main contribution of this paper is to show the conditions when a Nash equilibrium component is m-evolutionarily stable for large m.

We introduce the central definition of the robustness index of backward induction from the evolutionary stability of the backward induction equilibrium component when $k = m$.

Definition 3.4 *Given an evolutionary dynamic process of large populations on an extensive-form game, for $k = m$ and $1 - \mu = \delta$, we define the robustness index I_b of backward induction to be the probability of the backward induction equilibrium component in the invariant distribution for large populations, i.e.,*

$$I_b = \liminf_{m \to \infty} \lim_{\mu \to 0} \pi_{\mu,m}[BC].$$

3.2 Preliminary results

In the model introduced in Section 2, when $k = m$ & $1 - \mu - \sigma = 0$, then at each stage in each population all individuals move to the current best strategy if no mutation happens. We only consider two-player finite stopping games in such fast dynamic process before Section 4.

Given a two-player finite stopping game Γ, we denote the deepest node in Γ as $\bar{N}$, and we say that a player stops at some node $N \neq \bar{N}$, if he move to a terminal vertex at N. We also say the corresponding strategy terminates at N. If the player move towards $\bar{N}$ at every node before $\bar{N}$ where she should move, we then say that player continues until the end of the game Γ.

For each node N in Γ, we denote $d(N)$ and $p(N)$ to be the set of all successor nodes and all predecessor nodes of n, respectively. Note that both $d(N)$ and $p(N)$ can be empty sets. We denote the set of immediate successor nodes and the set of immediate predecessor nodes of N by $id(N)$ and $ip(N)$, respectively.

Lemma 3.5 *Given a finite two-player generic stopping game Γ, each Nash equilibrium in Γ has a unique Nash equilibrium path.*

Proof. Suppose that there are more than one path $(p_1, p_2, ..., p_n)$ for a Nash equilibrium e in Γ. We denote the deepest node on all these paths to be N. Without loss of generality, we suppose that player I moves at node N. It is straightforward to see that N is not the bottom node, and at least one supporting pure strategy of player II terminates at a node deeper than N. Player II's strategy in equilibrium e cannot be a pure one, otherwise player I's best response in e can only sustain one equilibrium path. We suppose it is a mixed strategy of n supporting pure strategies of player II, and that they terminate at nodes $N_1, N_2, ..., N_n$, respectively. Suppose the deepest terminate node of all supporting strategies of player I is $\bar{N}$. We denote N_1^m and N_2^m as the nodes in $(N_1, N_2, ..., N_n)$ that are immediate before $\bar{N}$ and after $\bar{N}$, respectively. We may further infer that there is no more strategy of player I sandwiched between N_1^m and N_2^m.

Now consider two strategies of player II which terminate at N_1^m and N_2^m, respectively. They cannot be both the best reply against the strategy of player I in e. To see this, check that they lead to the same payoff for all supporting strategies of player I terminating before N_1^m, but differently for the strategy of player I terminating at $\bar{N}$. Contradiction.

$\square$

Given a node N in an extensive-from game Γ, denote Γ_N to be the subgame rooted at N.

Definition 3.6 *Given a node N in Γ, if a mixed strategy s of a player requires equal probability on each her possible pure strategy stopping after N, then we say s is a uniform strategy of that player at node N.*

For instance, in the game in Figure 2, the uniform strategy of player II at the second node is a mixed strategy with equal probability on strategies t_1, t_2 and t_3.

Under the assumption of large populations and best-reply dynamics with no mutation, when a player's best reply is a uniform strategy, we apply the strong law of large numbers: at the next stage, it is an approximate uniform distribution of all moving individuals on the set of all supporting pure strategies of that uniform strategy.

Given a game Γ and a strategy s of a player i, denote the set of the best reply strategies of the other player by β_s.

Definition 3.7 *Suppose a two-player extensive-form game Γ is given. When m is large, there exists a $\lambda > 0$ such that, for a uniform strategy s of player i with support size $|s|$ in Γ, if the number of individuals playing at each supporting strategy of s is at least $m((1/|s|) - \lambda)$, then the set of best-reply strategies of the other player is still β_s. We fix such a λ for the game Γ, and call any distribution of population $M(i)$ above λ-consistent with s.*

If the best reply is a uniform strategy s, and no individual plays any supporting strategy of s before, then the distribution of the according population will be λ-consistent with s at the next stage, with probability approaching 1 as the population size increases to infinity. If no ambiguity, we skip λ when we mean λ-consistency.

On the one hand, given a sufficiently large population size m in the dynamic process applied to Γ, we can see that if all individuals in one population are playing some pure strategy or the distribution of that population is consistent with a uniform strategy, the best reply of the other population is also some pure strategy or a mixed strategy with supporting strategies in a uniform strategy. The latter case means the distribution of all moving

12

individuals in that population is consistent with the uniform strategy with asymptotic probability 1. On the other hand, we can find some particular m such that the best reply of one population against any distribution of the other population is alway either some pure strategy or a uniform strategy. In fact, for any game Γ, there is an infinite increasing sequence $(m_1, m_2, ...)$ with this proposition. For simplicity, we only consider an m in this sequence. So, we only need to consider an initial state where both populations are playing a pure strategy or consistent with a uniform strategy. If i goes to infinity in the sequence $(m_i)_{i \geq 1}$, we may abuse the notation and simply write $m \to \infty$.

Comment: A close look at the analysis later will see that the exact number of m is not crucial, as most of the time the distribution of any population is in a pure strategy or consistent with a uniform strategy. This is why we can abuse the notation. We assume $m \in (m_i)_{i \geq 1}$ simply for avoiding discussion on the issue how the distribution approaches to one in pure strategy or consistent with a uniform strategy, which is not important in the long run.

From Corollary 3.2, we know that the set of all non-Nash-equilibrium states is not evolutionarily stable for any m. We can also show below that, under the best-reply dynamics with no mutation, any non-Nash-equilibrium state will finally transit to a a state in NE in boundedly many stages. The *cycle* below means returning to any state the process has stayed before.

Lemma 3.8 *Under the best-reply dynamics with no mutations, there is no cycle of states in $W \setminus NE$. To be specific, from any state not in NE, with positive probability, the transition leads to a state in NE in boundedly many stages for any population size m_i in $(m_1, m_2, ...)$. When i goes to infinity, that probability goes to 1.*

Proof. Since m is in $(m_1, m_2, ...)$, we only need to consider a transition within the states where the whole population is playing a pure strategy or it is consistent with a uniform strategy. For simplicity, we say that a population is playing a strategy if its population is consistent with that strategy.

Suppose that at the initial stage n, $M(1)$ and $M(2)$ are stopping at node N and N', respectively. (They may play a pure strategy or a uniform strategy there.) We further suppose that $N \in p(N')$. All length of the process mentioned below is only dependent on the game tree, not on m.

Claim 1: If both populations are not playing the best reply against the other population at stage n, then after finitely many stages, either the process moves to a state in NE, or one of the population is playing a best reply strategy against the other population, at the start of some stage $n' > n$.

Suppose both populations are always not playing the best reply at the start of each stage greater than n. We firstly note that $M(2)$'s best reply is a pure strategy terminated at a node $N_1 \in p(N)$ at stage $n+1$. Then, $M(1)$'s best reply is a pure strategy terminated at a earlier node $N_2 \in p(N_1)$ at stage $n+2$. In this way, the pure strategy terminated at the earlier node of two populations moves towards the root of Γ. When N_i is exactly the root, the population who does not stop at the root trivially plays a best reply. The population stopping at the root has no incentive to deviate as well. So the process moves to a state in NE. Contradiction.

Claim 2: If at the start of stage n, $M(1)$ stops earlier than $M(2)$ and $M(1)$ is playing a best reply against $M(2)$. Then after finitely many stages, the process moves to a stage in NE.

Note that $M(1)$ is playing a pure strategy at a node N at stage n. Suppose the above claim is not true. At stage n, $M(2)$ is not playing a best reply against $M(1)$. Thus it will play a pure strategy terminates at a node $N_1 \in$

$p(N)$ at stage $n + 1$. With the similar argument as in the proof of Claim 1, we can reach a contradiction.

Claim 3: If at the start of stage n, $M(2)$ is playing a uniform strategy at a node N' later than where $M(1)$ stops and $M(2)$ is playing a best reply against $M(1)$, then after finitely many stages, the process moves to a stage in NE.

Note that $M(1)$ will move to a best reply at stage $n + 1$, while $M(2)$ keeps unchanged. If $M(1)$ plays a pure strategy terminating at some node $N_1 \in p(N')$, we can apply the conclusion of Claim 2.

If $M(1)$ stops at a node N_2 later than N' (either with a pure or uniform strategy at N_2), then we check the movement of $M(2)$ at stage $n+2$. If $M(2)$ unchanged, then it is in an NE state. If $M(2)$ moves to a pure strategy terminating at a node $N_3 \in p(N_2)$, then we are back to Claim 2. If $M(2)$ stops at a node $N_4 \in d(N_2)$, we then check the movement of $M(1)$. Since Γ is finite, if it's always not in a state in NE and not back to the cse in Claim 2, then at some stage, some population will continue until the end of Γ, and that surely leads to a state in NE. $\square$

We may view a state in NE as a stable state. From Corollary 3.2, we know that the process stays in NE for most of the time. When a rare mutation happens, it may have no effect on the other population and the mutant will move back to its previous position at the next stage. Only if the mutant triggers the best reply of the other population, it may lead to a new stable state. We know that, when the mutation rate is very low, the transition time between two stable states is short compared with the time of staying at a stable state. So we can roughly view the dynamic process as a sequence of blocks where each block corresponds to the time spent at a stable state. The movement to the next block is triggered by a mutation which happens at the

end of that block. If the relative proportion of the blocks corresponding to a component C is bounded from below when time goes to infinity, then C is evolutionarily stable. (For details and rigorous analysis, see Section 8.2 in [6].)

We can define a directed graph G to describe the transitions between states in NE. If the distribution of each population is in one pure strategy or consistent with a uniform strategy, then the current state can be described by a vertex in G. Thus under our assumptions, each equilibrium component corresponds to a finite set of vertices in the dynamic process of large populations. For our purpose, we view a component as the set of these corresponding vertices. Suppose that a state is described by a vertex V and, after at most one mutation in some population, the state moves to a new vertex V' under the best-reply mechanism (without further mutation in the transition) with asymptotic probability 1. We then say that V is *one mutation away from* V'. In G, we draw a directed edge from V to V'. We allow multiple edges on a pair of vertices (for the case of multiple transition paths). If there is a sequence of vertices $(V_1, V_2, ..., V_k)$ such that V_i is one mutation away from V_{i+1} for all $1 \leq i < k$, we then say that there is an *one-mutation chain* from V_1 to V_k.

3.3 One-mutation transition process

Without loss of generality, we assume that there are multiple equilibrium components in Γ. We first consider a condition of no non-trivial partition of the set of all equilibrium components in graph G: there exists a component C (maybe not unique) such that, for any other component C' and any vertex $V' \in C'$, there is a directed path (one-mutation chain) from V' to C. We call a graph G *irreducible* in this paper if it satisfies this condition. We shall show

16

an example in Section 3.5 where such irreducibility condition is not satisfied.

A *recurrent class* R in G is a set of vertices with no outbound edges from any vertex in R to any vertex not in R. One can check that, if G is irreducible, then there is only one recurrent class in G.

Theorem 3.9 *Suppose a dynamic process of large populations on a game Γ and the associated directed graph G are given. If G is irreducible, then the evolutionarily stable equilibrium components for large populations are those with vertices included in the recurrent class in G.*

Proof. For a vertex not included in the recurrent class R in G, it will transit to R in boundedly many stages with asymptotic probability 1. From any vertex in R, every outbound edge (to a vertex in R) represents an one-mutation transition. Once the process is at a vertex in R, with asymptotic probability 1, at the next stage it will not escape from R. Furthermore, it will transit to any other vertex in R in $c\mu^{-1}$ stages in expectation, where c is independent to m. By the proposition of the asymptotic behaviour in stationary Markov chain, the state (or equilibrium component) represented in any vertex in the recurrent class is evolutionarily stable. $\square$

In a two-player finite stopping game, for a strategy profile, we say that player I terminates this game at node N if player I moves to a terminal vertex at N with probability 1 and both players continue the game before node N with probability 1. Note that, in a state in NE, only the mutation which happens in the population terminating the game can lead to a transition to other equilibrium component. (For the population which plays a pure strategy or is consistent with a uniform strategy at a later node, the possible mutation cannot provide enough incentive to the other population.)

We briefly show how to calculate I_b by a new (simpler) Markov chain defined on the recurrent class R when G is irreducible. Suppose at a state in

17

NE, all individuals in population $M(1)$ are playing a pure strategy s, and s terminates the game. Suppose this state is denoted as V_1 in G, and we denote the number of pure strategies of player I in Γ by ς_1. We further suppose that V_1 is one mutation away from a vertex V_2 in R and the transition can be triggered by $n_{V_1 \to V_2}$ different mutations. Then, we simply assign the transition probability from V_1 to V_2 as $n_{V_1 \to V_2}/\varsigma_1$ in the new Markov chain to be defined on R. In this way, we can assign the transition probability of all pair of vertices in R which is connected by an one-mutation transition in the original dynamic process. We then calculate the invariant distribution in this new Markov chain. The probability of a component C in the invariant distribution in the original system is simply the sum of the invariant probabilities of all vertices included in both C and R in the new simplified Markov chain.

The simplest case is that only two Nash equilibrium components BC and NC exist in Γ, and each component contains only one vertex in R. When both components are evolutionarily stable, i.e., $I_b \neq 1$ or 0, we check the number n_b of mutations from BC that lead to the other equilibrium component, and the number n_c of mutations for transition from NC to BC. Suppose that there are ς_b and ς_c reduced pure strategies for the player who terminates the path of BC and NC in Γ, respectively. (If one player terminates both paths, then $\varsigma_b = \varsigma_c$.) Then the relative probability in the Markov chain from BC to NC is n_b/ς_b, and the probability from NC to BC is n_c/ς_c. From the definition of invariant distribution, we find

$$I_b = \frac{\frac{n_c}{\varsigma_c}}{\frac{n_c}{\varsigma_c} + \frac{n_b}{\varsigma_b}}.$$

Here is an example of a game with $0 < I_b < 1$. For convenience, all backward induction moves are arrowed and the terminating moves in the alternative pure Nash equilibrium are double arrowed in this paper. In game
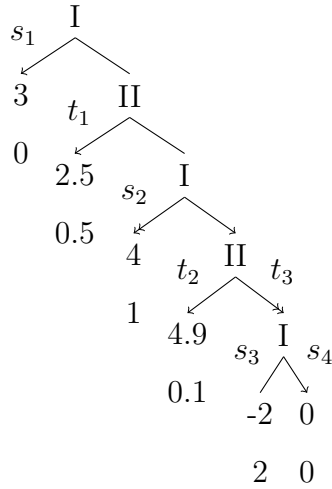
18

Figure 2: Γ_1

Γ_1 in Figure 2, the backward induction equilibrium is (s_1, t_1). We shall see from the game that two (stable) vertices in BC in the corresponding directed graph G have the same one-mutation transition to NC, so we can assume virtually only one vertex associated to BC here. (This is analogous to the trick used in a stationary Markov chain.) An alternative pure Nash equilibrium is (s_2, t_3). On one hand, from NC, a mutation in $M(1)$ from s_2 to s_4 will lead to BC. (The best reply of $M(2)$ to s_4 is t_1, which pushes $M(1)$ to s_1.) The probability of such mutation conditional on happening in $M(1)$ is $1/4$. On the other hand, from BC, a mutation in $M(1)$ from s_1 to either s_2 or s_3 leads to NC, and its total conditional probability is $1/2$. Note that in the long run, in most of the time that the state is in BC, all individuals in $M(2)$ play t_1. This can be seen from the one-mutation transition from NC to BC. (The mutation at s_2 will lead to the distribution of $M(2)$ consistent with a uniform strategy at the second bottom node, which sustains s_2 as the best reply of $M(1)$; the mutation at s_3 means the best reply of $M(2)$ is t_3, which then pushes $M(1)$ to s_2.) So $I_b = 1/3$.

19

The game Γ_2 in Figure 3 is an example that $I_b = 0$. The backward induction equilibrium in this game is $e_b = (s_2, t_2)$ and an alternative pure Nash equilibrium is $e_n = (s_1, t_1)$. Suppose that the initial state is in e_b. After a mutation from s_2 to s_3, $M(2)$ moves to an approximate uniform distribution of t_3 and t_4 at the next stage. Then $M(1)$ is attracted to s_4, which further moves $M(2)$ to s_2 and hence a transit to NC. Suppose the initial state is in e_n, any mutation will either makes $M(2)$ unaffected or lead to $M(1)$ playing s_4. With the similar argument above, we can show that the process will be back to NC. Furthermore, it is back to exactly e_n after the movement of $M(2)$ triggered by a mutation.
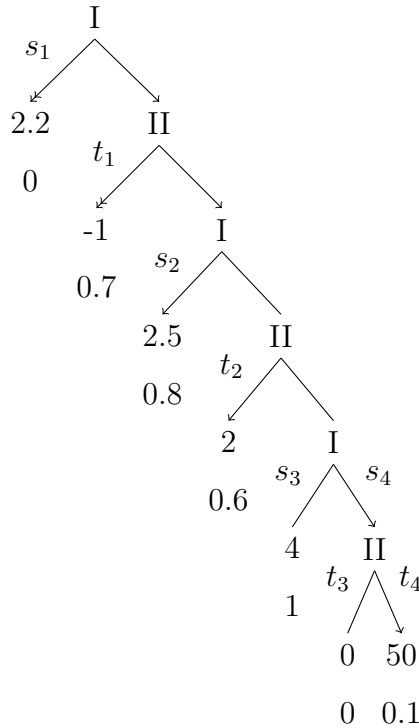

Figure 3: Γ_2

Comment: We have shown that there is no one-mutation chain between BC and NC (G not irreducible), and BC is not evolutionarily stable for large

populations in the example in Section 5.1, [7]. However, we do not know whether a game exists with an one-mutation transition from BC to NC and BC not evolutionarily stable for large populations, under the condition that only one individual in each population can change strategy at each stage.

3.4 A sufficient conditions for evolutionary stability for games with only two equilibrium components

We consider the simplest case of two Nash equilibrium components BC and NC in Γ, and show a sufficient conditions for evolutionary stability for large populations.

Lemma 3.10 *Suppose there are two equilibrium components C_1 and C_2 in Γ, and the paths of these two components terminate at node N_1 and N_2, respectively. We suppose that N_2 is a descendant node of N_1. Without loss of generality, we assume that in C_2 player II moves at node N_2 and the pure strategy of player II in C_2 is s_2. If s_2 is the best reply of player II against the uniform strategy of player I at node N_2, then C_2 is evolutionarily stable for large populations.*

Proof.

   Case 1: the paths of two equilibrium components are terminated by different player. Then, the path of C_1 is terminated by player I. In C_2, we fix one pure strategy s of player I such that (s, s_2) is in C_2. Note that s terminates after N_2. Suppose Z_0 is in C_1. Consider a mutation in population $M(1)$ from its original position to s at stage 1. The best reply for $M(2)$ is s_2, and the whole $M(2)$ moves to s_2 at stage 2. At stage 3, $M(1)$ is under the selection pressure and moves to a distribution consistent with the uniform

21

strategy at N_2. As we have assumed that s_2 is the best reply to that uniform strategy, the state is in C_2 at stage 3.

Case 2: the paths of two equilibrium components are terminated by the same player, player II. Denote the set of all nodes earlier than N_2 where player II moves by $\bar{S}$. Denote the payoff to player II in C_2 as u_2. From the definition of equilibrium component, we know that, any strategy of player II stopping at any node in $\bar{S}$ cannot give a payoff to player II as high as u_2, if the strategy of player I terminates later than N_2.

Suppose that the initial state is in C_1 at stage 0. Check whether there exists a strategy x_1 of player II terminating in Γ_{N_2} whose best reply of player I is to terminate earlier than N_2.

1. If so, consider a mutation in $M(2)$ from its original position to x_1, then the whole $M(1)$ moves outside Γ_{N_2} as the best reply. Furthermore, it is consistent with a pure strategy of player I, say strategy x_2 terminating at N_3.

   (a) If $M(2)$ has no incentive to deviate from its original position, then it is in NE. consider the next mutation in $M(2)$ to s_2, and we find that the state will transit to C_1.

   (b) If $M(2)$ is attracted to a new strategy by x_2, then it must be a uniform strategy at N_3. This is not in a equilibrium state, so $M(1)$ is under the selection pressure, and his best reply is a pure strategy x_3 in Γ_{N_3}. (The definition of x_2 requires that $M(1)$ cannot terminates at an earlier node than the current node.) Once again, this is not in an equilibrium state, and $M(2)$ is moving at the next stage. This process continues until some time $M(1)$ terminates later than N_2. (Note that $M(1)$ can only terminates

22

later and later in the process before it adopts a strategy terminating in Γ_{N_2}. Otherwise, it will convert back to x_2, which gives a equilibrium path terminating at N_3. But there are only two equilibrium components. Contradiction.) From then, to land on an equilibrium state, $M(2)$ has to stop at node N_2, which gives a higher payoff than stopping anywhere earlier. So it transits from C_1 to C_2.

2. If no such x_1 exists, then given any mutation in $M(2)$ to a strategy terminating in Γ_{N_2}, the best reply of $M(1)$ also terminates in Γ_{N_2}. With similar arguments in 1(b), we can confirm a transition from C_1 to C_2.

☐

Comment: If both C_1 and C_2 are evolutionarily stable for large populations, from Case 2 1.(a), one may argue that, although it is possible to transit from C_1 to C_2, C_1 is not always one transition away from C_2. In fact we can prove that, in the dynamic process, the transition from C_1 to C_2 only needs one mutation with asymptotic probability 1: the strategy of the whole $M(1)$ in C_1 is almost always the same pure strategy terminating before N_2.

3.5   One-mutation transition in G does not always exist

Theorem 3.9 concerns the situation that G is irreducible. This is not always true for all finite stopping games.

In the following game Γ_3 in Figure 4, the backward induction equilibrium is $e_b = (s_2, t_2)$; an alternative pure Nash equilibrium is $e_n = (s_1, t_1)$. Suppose $Z_0 = e_b$. The mutation in $M(1)$ from s_2 to either s_3 or s_5 will push the whole $M(2)$ to t_2 and hence no transition to NC. For a mutation from s_2 to s_4 at

stage 1, $M(2)$ is under the selection pressure to strategy t_4 and t_5. So it will be an approximate uniform distribution of $M(2)$ between t_4 and t_5. Then, At stage 2, $M(1)$ is unaffected by this distribution of $M(2)$. Now suppose $Z_0 = e_n$. Only the mutation in $M(1)$ from s_1 to s_2 will trigger the movement of $M(2)$ to a uniform distribution at t_2, t_3, t_4 and t_5, at stage 2. Then, the best reply of $M(1)$ is s_5 at stage 3, which pushes the whole $M(2)$ to t_1 at stage 4. Hence, the process is back to e_n at stage 5.

Figure 4: Γ_3

One may argue that, if the initial state is in BC, then it does not have

to be in e_b. For example, if $Z_0 \in BC$ but a significant proportion of $M(2)$ is at t_5, then a mutation from s_1 to s_4 may give more than $3/5$ of $M(2)$ at t_5, which triggers $M(1)$ to move to s_5 and thus the state transits to NC. Similar argument could also applies to a state in NC. A close look could reveal the hidden problem of this argument. Once it transits to the other equilibrium component, say C, due to the contribution of the initial distribution, with asymptotic probability 1 the distribution of $M(2)$ at the current state strongly favours the stability of the current component, and then it will convert to the pure Nash equilibrium in C after further mutations in $M(1)$.

So what is the invariant distribution of the Markov chain corresponding to Γ_3? One possibility for a component transition is that there are multiple mutations before it enters a Nash equilibrium. However, the probability of this event is approaching 0 when the mutation rate decreases to 0. The other possibility is the distribution of a population corresponding to a uniform strategy may be significantly deviated from the uniform distribution. By strong law of large numbers, the probability of this event is also low. However, compared with the even smaller mutation rate, it is the source of component transition in Γ_3. (Recall that the definition of evolutionary stability requires the mutation rate approaching to 0 for a fixed population size.) For the possible transitions between NC and BC, we are essentially comparing the probability inconsistent with the strong law of large numbers in two corresponding cases.

In Γ_3 in Figure 4, for the transition from BC to NC, without loss of generality, we can assume $Z_0 = e_b$. After the mutation from s_2 to s_4, $M(2)$ moves to strategy t_4 and t_5. If the proportion of t_5 can be greater than $3/5$ at stage 1, then $M(1)$ is under the selection pressure to s_5, and finally

the process transits to NC. On the other direction, if the initial state is at $Z_0 = e_n$, then a mutation from s_1 to s_2 will make $M(2)$ move to strategies t_2, t_3, t_4 and t_5. If the proportion of t_3 is enough small (some number smaller than $1/30$) and the total proportion of t_2 and t_4 is big enough, then the best reply of $M(1)$ is strategy s_2, which completes the transition from NC to BC. It is straightforward to see that in this game it is 'easier' to transit from BC to NC than the opposite direction. Given a population size m, we denote p_m^1 and p_m^2 to be the probability that the distribution of $M(2)$ is inconsistent with the Strong Law of Large Numbers and the transition takes place from BC to NC and from NC to BC, respectively. It is straightforward to see that

$$\lim_{m \to \infty} \frac{p_m^2}{p_m^1} = 0.$$

We can then conclude that NC is the only evolutionarily stable result for large populations in Γ. So $I_b = 0$ in Γ.

One should note that, since the event of a transition between BC and NC is inconsistent with the Strong Law of Large Numbers, the probability of a successful transition triggered by one mutation decreases to zero as the population size increases to infinity. A simple estimate shows that in expectation it will take $O(e^m \mu^{-1})$ stages for a component transition. The best-reply dynamic applied in this population game is fast, but the transition time can be exponentially long with respect to the population size.

From this example, one may find that, for the case of no one-mutation chain in G, the evolutionary stability can (heavily) depends on payoff vectors. There seems no simple way to calculate the exact I_b for each such game, and it is usually either $I_b = 0$ or 1 for a generic game.

Comment: The stability result of the game Γ_3 in Figure 4 may be different for different limiting process discussed in Section 6 in [7]. For example,

26

when $c_1 < m\mu < c_2$ as $m \to \infty$ and $\mu \to 0$, the possibility of multiple mutations before the state arrives at the next equilibrium cannot be ignored. (We may need to prove first that only Nash equilibria can be evolutionarily stable in this limiting process for any population size.)

4 Variation of the model

So far, we have focused on population games of large populations. Due to the strong law of large numbers, with asymptotic probability 1, the distribution of each population at the next stage is consistent with a pure or a uniform strategy. We study some variation of the model below.

4.1 Different results for dynamic parameters

From the results in Section 3.3 and 3.4, we find it is straightforward to calculate I_b for a game whose graph G is irreducible. One might hope that under the irreducibility condition, the stability result for the process with $k = m$ is a good approximation of the process with $k = 1$, and we could then only refer to the fast dynamic process. It turns out that their results are not completely identical.

In the game Γ_4 below in Figure 5, we showed in Section 7 in [7] that BC is the only evolutionarily stable result for large populations when $k = 1$. However, by Lemma 3.10, when $k = m$ and $\mu = 1 - \sigma$, NC is also m-evolutionarily stable for large populations.

Here is another game Γ_5 in Figure 6. In this game, the backward induction equilibrium $e_b = (a_1, b_3)$ and an alternative pure Nash equilibrium is $e_n = (a_2, b_1)$. We showed in Section 4 in [7] that both BC and NC are m-evolutionarily stable for any m when $k = 1$. However, when $k = m$ and
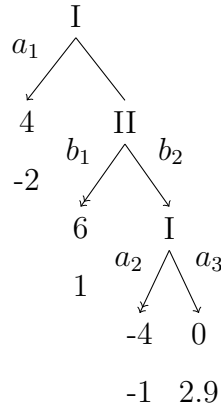
27

Figure 5: Γ_4

$\mu = 1 - \sigma$, only BC is evolutionarily stable for large populations. To see this, we suppose that the initial state is in BC. If a mutation happens in $M(1)$ and it adopts a strategy a_3 at stage 1, then $M(2)$ will be attracted to b_2, which confirms the best reply of a_1 for $M(1)$. If the mutants changes from a_1 to a_2 at stage 1, then $M(2)$ would move to b_1 at stage 2. At stage 3, $M(1)$ prefers the uniform strategy at the second node from the top. However, that approximate uniform distribution (when m is large) puts $M(2)$ under the selection pressure to strategy b_2, which gives him payoff $4/3$ rather than 1. With the similar argument as in the earlier case, the process converts back to BC. So there is no one-mutation transition from BC to NC. The other direction from NC to BC can be triggered by a mutation from b_1 to b_3.

Let's turn to the parameter of selection rate σ. If we set $k = m$ but $\sigma > 0$, then we can show the results for the above two examples are the same as the case when $k = 1$. Moreover, if $k < m$ then the same results hold for whatever σ. It is an open question to give a characterisation when the evolutionary stability result is always the same for the two cases of (1) $k < m$ and (2) $k = m$ but $\sigma = c > 0$ for some constant c, for generic extensive-form games
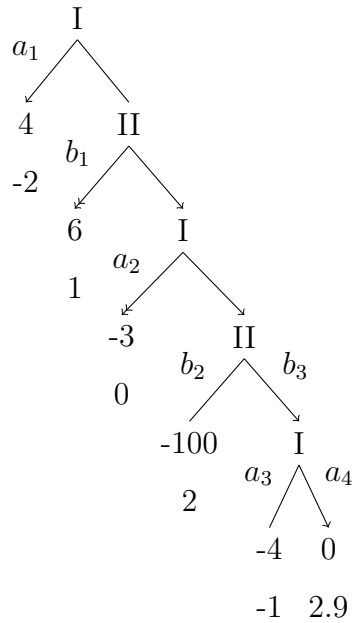
28

Figure 6: Γ_5

of perfect information with large populations.

4.2 Beyond finite stopping games

For all results above, we have constrained ourselves within finite stopping games. An immediate question is what we can say for general extensive-form games of perfect information. In finite stopping games, there is only one branch (or maximum path) in the whole game tree, however in general extensive-form games, it contains multiple branches. The first difficulty is the definition of mutation. Should we assign the same probability for the mutation that an individual changes his move in only one branch and the mutation on multiple branches at the same time? One should also consider the definition of the best-reply strategy: when an individual moves to his so called best-reply strategy, shall we require that individual's moves unchanged

at nodes which are disconnected in both the states before and after the change? We have made some assumptions on these issues in Section 4.2 in [6], and one should expect that the stability results depend on such definitions and conditions.

For the dynamic process of the population game with $k < m$, one should anticipate the complexity of multiple equilibrium components on multiple branches. When a population is moving to a new equilibrium on the other branch, the distribution of the remaining (number decreasing) individuals on the starting branch is not deterministic in the transition. Furthermore, the distribution on this branch is toward more even (by the strong law of large numbers), which may reverse the transition direction of the whole population at some stage.

It would be difficult to give a succinct characterisation result on this more general case.

4.3 Games played by more than two players

It would be interesting to see the stability result for extensive-form games with more than two players, which is one of my research directions in the future.

Acknowledgements

References

[1] Z. Gorodeisky, Evolutionary stability for large populations and backward induction, Mathematics of Operations Research 31 (2006) 369–380.

[2] S. Hart, Evolutionary dynamics and backward induction, Games Econ. Behav. 41 (2002) 227–264.

[3] G. Noldeke and L. Samuelson, An evolutionary analysis of backward and forward induction, Games Econ. Behav. 5 (1993) 425–454.

[4] M. Kandori, G. Mailath and R. Rob, Learning, mutation, and long-run equilibrium in games, Econometrica, 61 (1993) 29–56.

[5] M. I. Freidlin and A. D. Wentzell, Random Perturbations of Dynamical Systems, New York: Springer Verlag, 1984.

[6] Z. Xu, Evolutionary stability in general extensive-form games of perfect information, Center for the Study of Rationality, Hebrew University of Jerusalem, Mimeo.

[7] Z. Xu, The instability of backward induction in evolutionary dynamics, Center for the Study of Rationality, Hebrew University of Jerusalem, Mimeo.

[8] H. P. Young, Individual Strategy and Social Structure, Princeton University Press, 1998.