

האוניברסיטה העברית בירושלים

THE HEBREW UNIVERSITY OF JERUSALEM

EVOLUTIONARY STABILITY FOR LARGE POPULATIONS AND BACKWARD INDUCTION

by

ZIV GORODEISKY

Discussion Paper # 312

April 2003

מרכז לחקר הרציונליות
**CENTER FOR THE STUDY
OF RATIONALITY**

Feldman Building, Givat-Ram, 91904 Jerusalem, Israel
PHONE: [972]-2-6584135 FAX: [972]-2-6513681
E-MAIL: ratio@math.huji.ac.il
URL: <http://www.ratio.huji.ac.il/>

Evolutionary Stability for Large Populations and Backward Induction

Ziv Gorodeisky*

July 20, 2005

Abstract

It has been shown (Hart [2]) that the backward induction (or subgame-perfect) equilibrium of a perfect information game is the unique stable outcome for dynamic models consisting of selection and mutation, when the mutation rate is low and the populations are large, under the assumption that the expected number of mutations per generation is bounded away from zero.

Here it is shown that one can dispense with this last condition. In particular, it follows that the backward induction equilibrium is evolutionarily stable for large populations.

Key words: evolutionary dynamics, evolutionary stability, Markov chains, transition times, backward induction equilibrium, large populations

1 Introduction

In this paper we follow the work of Hart [2]. We study the long-run behavior of evolutionary dynamics, we introduce a few notions of stability, and finally, we show the stability properties of the backward induction equilibrium (BIE).

As in Hart [2], the games we consider are generic finite games in extensive form with perfect information. In these games, there exists a unique *subgame-*

*Center for the Study of Rationality, Department of Mathematics, The Hebrew University of Jerusalem, 91904 Jerusalem, Israel. E-mail: zivg@math.huji.ac.il

perfect equilibrium, or *backward induction equilibrium*. For each such game, there is an associated *population game*: at each node there is a distinct population of individuals who play the game in the role of the corresponding player.

The evolutionary dynamic process is a Markov chain on the space of the mixed strategies of the game with unique invariant distribution. We are looking for stable strategies in this model. When the populations are fixed, a strategy profile is *evolutionarily stable* if its occurrence is positive independently of the mutation rate, i.e., if its invariant probability is bounded away from zero as the mutation rate goes to zero.¹ When the populations increase, the number of possible outcomes of the dynamics increases, and the invariant probabilities of the different strategy profiles change. Therefore, we define a strategy profile to be *evolutionarily stable for large populations* (ESLP) if its invariant probability is bounded away from zero as both the mutation rate goes to zero and the populations increase to infinity.

In Hart [2] it is shown that when the populations are fixed, the BIE is evolutionarily stable. The main result there is that in the limit the BIE becomes the only stable outcome as the mutation rate decreases to zero and the populations increase to infinity, provided that the expected number of mutations per generation is bounded away from zero. In this paper we show that in the models we consider, this additional proviso (on the expected number of mutations per generation) is not needed. Thus, the BIE is also ESLP.

Section 2 presents the model. Section 3 defines stability and presents

¹This is the “stochastic stability” of Kandori et al. [3] and Young [6].

the Main Theorem. Section 4 explains two major assumptions in our model, that of mutation rate and that of population size. Section 5 proves the Main Theorem. The Appendix proves several general propositions on Markov chains, needed for the proof of the Main Theorem.

2 The Model

The model is as in Hart [2], except for a somewhat less general class of dynamics.² We present a summary of the model below.

2.1 The Game

Let Γ be a finite extensive-form game with perfect information. Each non-terminal vertex corresponds to a *move*, and each move of one of the players is called a *node*. Let N be the set of nodes, and let \hat{N} be the set of nodes and terminal vertices (i.e., the set of all vertices that are not chance moves). For each node $i \in N$, let $N(i)$ be the set of nodes that are successors of i in the tree. Assume that the nodes are numbered $\{1, \dots, n\}$, where $n = |N|$, such that $j \in N(i)$ implies $j > i$.

The game is in *agent-normal form*: at each node $i \in N$ there is a different agent with a set of choices A^i , which is the set of outgoing branches at i . Let $A := \prod_{i \in N} A^i$, and let $u^i : A \rightarrow \mathbb{R}$ be the payoff function of agent i , which is extended multilinearly to *mixed* strategies; thus $u^i : X = \prod_{i \in N} X^i \rightarrow \mathbb{R}$, where $X^i := \Delta(A^i)$ is the set of probability distributions over A^i .

The classic result of Kuhn [4] states that there always exists a *pure* BIE.

²See Remarks 3 and 4 in Section 2.2.

We assume here that the game Γ has a *unique* BIE (this is true generically), which must therefore be pure; we denote it by $b = (b^i)_{i \in N} \in A$, and refer to b^i as the “backward induction strategy of i .”

We associate a population game³ to Γ : at each node $i \in N$ there is a nonempty finite population $M(i)$ of individuals playing the game in the role of agent i . We assume that the populations at different nodes are disjoint. For each individual $q \in M(i)$, let $\omega_q^i \in A^i$ be the pure strategy of q , and let $\omega^i = (\omega_q^i)_{q \in M(i)}$ and $\omega = (\omega^i)_{i \in N}$. The proportions of the different pure strategies in population i induce a mixed strategy of agent i . We therefore have a map x of each ω to a vector of mixed strategies, where $x_{a^i}^i(\omega)$ — the probability of a^i — is the proportion of individuals in populations i that play a^i .

2.2 The Dynamics

We now come to the dynamic model. We formulate the dynamics as *discrete-time stationary Markov chains*.

Since they are evolutionary, the dynamics are based on selection, i.e., changes towards better strategies, and on mutations, i.e., random changes. Each period of the dynamics is assumed to be small enough that the probability that more than one individual in each population will change his strategy is negligible. Thus, we assume that at each period, at most one individual may change his strategy, either due to selection or due to mutation.

For every *mutation rate* parameter $\mu > 0$, and *population size* m such

³This is called the *gene-normal form* of Γ in Hart [2].

that⁴ $|M(i)| = m$ for all nodes $i \in N$, the process is a stationary Markov chain on the state space $\Omega_m := \prod_{i \in N} (A^i)^{M(i)}$, where a *state* ω of the system specifies the pure strategy of each individual in each population, as defined in Section 2.1. The one-step transition probability of the process is given by a transition matrix $Q = (q_{\omega, \tilde{\omega}})_{\omega, \tilde{\omega} \in \Omega_m}$ that satisfies⁵:

- Conditional independence over $i \in N$, i.e.,⁶

$$Q[\tilde{\omega} \mid \omega] = \prod_{i \in N} Q[\tilde{\omega}^i \mid \omega]. \quad (2.1)$$

- For each $i \in N$, one individual $q(i) \in M(i)$ is chosen, such that there exist constants $\gamma_1, \gamma_2 > 0$ with

$$\frac{\gamma_1}{m} \leq Q[q(i) = q \mid \omega] \leq \frac{\gamma_2}{m} \text{ for each } q \in M(i), \text{ and} \quad (2.2)$$

$$Q[\tilde{\omega}_q^i = \omega_q^i \mid \omega] = 1 \text{ for each } q \in M(i) \setminus \{q(i)\}. \quad (2.3)$$

- There exists a constant $\sigma > 0$ such that, for each $i \in N$,

$$Q[\tilde{\omega}_{q(i)}^i = a^i \mid \omega] \geq \sigma \text{ for each } a^i \in B^i \cap S^i, \quad (2.4)$$

where $B^i \equiv B^i(q(i), \omega) := \left\{ a^i \in A^i : u^i(a^i, \omega^{-i}) > u^i(\omega_{q(i)}^i, \omega^{-i}) \right\}$ is

the set of “better strategies” — those strategies at node i that are

⁴See Remark 3.

⁵For an example of a dynamic that satisfies these conditions, see the *basic model* in Hart [2].

⁶For each $\omega \in \Omega$, take $Q[\cdot \mid \omega]$ to be a probability distribution over Ω , such that $Q[\Omega' \mid \omega] = \sum_{\omega' \in \Omega} q_{\omega, \omega'}$ for all $\Omega' \subseteq \Omega$; derived probabilities — like its marginals, etc. — will also be denoted by $Q[\cdot \mid \omega]$.

strictly better in Γ , against the populations at the other nodes, than the strategy $\omega_{q(i)}^i$ of the chosen individual $q(i)$, and $S^i \equiv S^i(\omega) := \{a^i \in A^i : \omega_q^i = a^i \text{ for some } q \in M(i)\}$ is the set of strategies at node i that are present in state ω .

- There exist constants $\alpha_1, \alpha_2 > 0$ such that, for each $i \in N$,

$$Q[\tilde{\omega}_{q(i)}^i = a^i \mid \omega] \geq \alpha_1 \mu \text{ for each } a^i \in A^i, \text{ and} \quad (2.5)$$

$$Q[\tilde{\omega}_{q(i)}^i = a^i \mid \omega] \leq \alpha_2 \mu \text{ for each } a^i \notin B^i, a^i \neq \omega_{q(i)}^i. \quad (2.6)$$

Remarks:

1. The game Γ and the constants γ_1 , γ_2 , σ , α_1 , and α_2 do not change while m and μ vary.
2. As in Hart [2, Section 3.2, Remark 2], we assume ((2.5) and (2.6)) that the probabilities of mutation in different populations are comparable. See Section 4.1.
3. We assume for simplicity that the sizes of the populations are the same, but it suffices that the sizes be comparable (i.e., the ratios $|M(i)| / |M(j)|$ be bounded away from zero and ∞). This assumption is not needed in Hart [2], but as will be seen in Section 4, it is a natural assumption.
4. We assume that the probability of switching to a better strategy by selection depends only on whether that strategy is currently present in the population and not on the proportion of individuals playing that

strategy (see (2.4)). Therefore, our class of dynamics is less general than that of Hart [2] (see Hart [2, (2.7)]).

5. A general model with a one-step transition matrix Q satisfying (2.1) – (2.6) yields a Markov chain that is irreducible and aperiodic. Hence there exists a unique invariant distribution on Ω_m that describes the long-run behavior of the process (see, e.g., Feller [1, Chapter 15, Theorem 2 and Corollary I]).

3 Stability

In this section we define stability and present the main result of this work.

3.1 Definitions

We are interested in the behavior of the process when the mutation rate is low, i.e., the limit of the invariant distribution π as $\mu \rightarrow 0$. We will look first at the case where the population size m is fixed, and then at the case where $m \rightarrow \infty$.

Let the game Γ and the constants α_1 , α_2 , γ_1 , γ_2 , and σ be fixed. For every mutation rate $\mu > 0$ and population sizes m , let $\Xi(m, \mu)$ be the set of all one-step transition matrices $Q = (q_{\omega, \tilde{\omega}})_{\omega, \tilde{\omega} \in \Omega_m}$ satisfying (2.1) – (2.6) with μ and m . For every $Q \in \Xi(m, \mu)$, let π_Q be the unique invariant distribution of Q , and let $\pi_{m, \mu}[\omega] = \inf_{Q \in \Xi(m, \mu)} \pi_Q[\omega]$ for all $\omega \in \Omega_m$.

Definition 3.1. A state $\omega \in \Omega_m$ is *m-evolutionarily stable* if

$$\liminf_{\mu \rightarrow 0} \pi_{m,\mu}[\omega] > 0.$$

Recall that each state $\omega \in \Omega_m$ may be viewed as an N -tuple of mixed strategies $x(\omega) = (x^i(\omega))_{i \in N} \in X$. We can therefore view the dynamic as a dynamic on the space of mixed strategies. Since the population is finite, and each agent plays a pure strategy, the number of mixed strategies that can arise in this dynamic is finite. The invariant distribution π_Q on Ω_m therefore induces a probability distribution $\hat{\pi}_Q := \pi_Q \circ (x)^{-1}$ over X ; i.e., $\hat{\pi}_Q[Y] := \pi_Q[\{\omega \in \Omega_m : x(\omega) \in Y\}]$ for every (measurable) $Y \subseteq X$. Let $\hat{\pi}_{m,\mu}[Y] = \inf_{Q \in \Xi(m,\mu)} \hat{\pi}_Q[Y]$.

Definition 3.2. An N -tuple of mixed strategies $x \in X$ is *m-evolutionarily stable* if

$$\liminf_{\mu \rightarrow 0} \hat{\pi}_{m,\mu}[x] > 0.$$

We now consider the case where the populations increase, i.e., $m \rightarrow \infty$, (while the game and constants remain fixed). When the population size increases, the size of the state space increases as well, and the invariant distribution of each fixed mixed strategy may converge to zero.⁷ To properly define evolutionary stability in this context, we consider probabilities of neighborhoods rather than the probability of a single point.

For every $\varepsilon > 0$ and N -tuple of mixed strategies $x \in X$, let x_ε be the ε -neighborhood of x , i.e., $x_\varepsilon := \{y \in X : \|x - y\| < \varepsilon\}$ (we will also use BI_ε instead of b_ε for consistency with the notations in Hart [2]).

⁷For an example of where the $\liminf_{\substack{\mu \rightarrow 0 \\ m \rightarrow \infty}} \hat{\pi}_{m,\mu}[b]$ is zero, see Hart [2, Footnote 49].

Definition 3.3. An N -tuple of mixed strategies $x \in X$ is *evolutionarily stable for large populations* (ESLP) if for every $\varepsilon > 0$

$$\liminf_{\substack{\mu \rightarrow 0 \\ m \rightarrow \infty}} \hat{\pi}_{m,\mu}[x_\varepsilon] > 0.$$

We will also use ESLP to denote the set of all N -tuples that are ESLP.

Thus, x is ESLP if, for any neighborhood of x , the relative frequency of visits at that neighborhood is bounded away from zero, for all large population sizes and all small mutation rates.

Remarks:

1. Conditions (2.1) – (2.6) are closed. Therefore, $\Xi(m, \mu)$ is a closed subset, and hence compact. Therefore, for every $\omega \in \Omega_m$, there exists a matrix $Q_\omega \in \Xi(m, \mu)$ such that $\pi_{m,\mu}[\omega] = \pi_{Q_\omega}[\omega]$. Moreover, as $\pi_Q[\omega] > 0$ for all $Q \in \Xi(m, \mu)$ and $\omega \in \Omega_m$, we have $\pi_{m,\mu}[\omega] > 0$; however, $\pi_{m,\mu}$ need not be a probability distribution over Ω_m . The same arguments hold for $\hat{\pi}_{m,\mu}$, and it need not be a probability distribution over X .
2. For each m , let ES_m be the set of all m -evolutionarily stable N -tuples of mixed strategies. By Hart [2, Theorem 3.1], the BIE is m -evolutionarily stable for all m ; therefore, $\{b\} \subset ES_m$. As seen in Hart [2, Section 3.1, Remark 2], equilibria other than the BIE may be m -evolutionarily stable, and hence we have $\{b\} \subsetneq ES_m$. Together with our Main Theorem below, we have $ESLP = \{b\} \subsetneq ES_m$.

3.2 The Main Result

In Hart [2] it is shown that the BIE is m -evolutionarily stable for all m . It is also shown that when the populations increase the probability of any neighborhood of the BIE goes to 1, as long as the populations increase fast enough.

Theorem 3.4 (Hart [2]). *For every $\varepsilon > 0$ and $\delta > 0$,*

$$\lim_{\substack{\mu \rightarrow 0 \\ m \rightarrow \infty \\ m\mu \geq \delta}} \hat{\pi}_{m,\mu}[BI_\varepsilon] = 1.$$

This theorem implies that only the BIE may be ESLP. In order to show that it is ESLP, we show that the double limit (without the restriction $m\mu \geq \delta > 0$) is positive — in fact, it equals one.

Theorem 3.5 (Main Theorem). *For every $\varepsilon > 0$,*

$$\lim_{\substack{\mu \rightarrow 0 \\ m \rightarrow \infty}} \hat{\pi}_{m,\mu}[BI_\varepsilon] = 1, \tag{3.1}$$

and, therefore, $ESLP = \{b\}$. Moreover, there exists a constant C , depending on the game, on the constants defining the dynamics $\gamma_1, \gamma_2, \sigma, \alpha_1$, and α_2 , and on ε , such that

$$\hat{\pi}_{m,\mu}[BI_\varepsilon] \geq 1 - C\left(\mu + \frac{1}{m}\right) \tag{3.2}$$

for all $\mu > 0$ and m .

Our proof of the Main Theorem is based on estimating transition times. Thus, besides the results on the long-run behavior of the populations (i.e., that most individuals play the backward induction strategy), we also learn the number of periods it takes to get there. Dealing with the case that $m\mu \leq D < \infty$ (which yields the Main Theorem by combining it with Theorem 3.4

for $\delta = D$), it takes on average of the order of at most $m + 1/\mu$ periods until most individuals play the backward induction strategy.

Proposition 3.6. *For every $\varepsilon > 0$, there exists a constant C' such that the expected number of periods to reach BI_ε , from any state, is at most⁸ $C'(m + 1/\mu)\exp(C'm\mu)$.*

Notice that the BIE is m -evolutionarily stable for all m , and therefore any neighborhood of it is also m -evolutionarily stable. Moreover, from the double limit (3.1) it follows that the iterative limit $\lim_m \lim_\mu$ is 1, and thus positive, which implies that BI_ε is m -evolutionarily stable *uniformly* in⁹ m .

4 Comparability between Populations

In this section we discuss the assumptions on the mutation rates $(\mu_i)_{i \in N}$ and on the sizes of the different populations $(m_i)_{i \in N}$, where $m_i := |M(i)|$.

Mutation is a “mechanical” process and, therefore, it is natural to assume that all genes have the same probability of mutation. Therefore, it is reasonable to assume that *the probability of mutation of each individual is the same* (or comparable if, for example, the number of genes in different populations is different). It follows that μ_i/m_i — the probability of mutation of each individual in population i — is the same for all i , or at least comparable.

⁸We use $\exp(x)$ for the exponent function.

⁹The other iterative limit, $\lim_\mu \lim_m$, is also 1, as already seen in Hart [2, Remark 1 in Section 3.2].




Figure 1: The Game Γ

4.1 Mutation Rates

In the models we consider, the probabilities of mutation in the different populations are assumed to be of the same order of magnitude — see (2.5) and (2.6). In the next example we will show that this assumption is indeed needed, and without it there is no convergence to the BIE, even if the probability of mutation of each individual is the same.

Example. Let us consider the game Γ in Figure 1. The BIE in this game is $b = (b^1, b^2)$. Consider the following dynamics on Γ : each period there is a probability of $1/m_i$ for each individual $1 \leq q \leq m_i$ to be chosen, and probabilities of $2\mu_i$ and $1 - 2\mu_i$ for the chosen individual to change his action by mutation (with equal probabilities to mutate to each of the two strategies) or selection respectively.

Let Y^i , for $i = 1, 2$, be the proportion of individuals in population i that play c^i , and let R^2 be an indicator random variable, defined as 1 if node 2 is reached (see also Section 5.1).

Notice first that there is always selection in node 1 towards b^1 regardless of population 2, and the transition probabilities of population 1 depend only

on Y^1 . Notice also that node 2 is reached iff Y^1 is positive. If node 2 is reached, then, on average, after Cm_1 periods (for some constant C) we have $Y^1 = 0$ (because of selection in node 1), and node 2 is no longer reached. Therefore, we have Cm_1 changes toward b^2 because of selection. On the other hand, once $Y^1 = 0$, node 2 will stay “not reached” until there is a mutation in node 1, which takes, on average, $1/\mu_1$ periods. In the $1/\mu_1$ periods node 2 is not reached, there are $\mu_2(1/\mu_1)$ mutations, in which population 2 gets closer to $1/2$ (moves from b^2 to c^2 if we are close to the BIE) by $\sqrt{\mu_2/\mu_1}$. By choosing μ_i and m_i such that $\sqrt{\mu_2/\mu_1} \geq Cm_1$ (while the probability of mutation of each individual is the same, i.e., $\mu_1/m_1 = \mu_2/m_2$), we will show that Y^2 does not converge to 0.

Let $m \in \mathbb{N}$ and let $m_1 = m$, $m_2 = m^3$, $\mu_1 = 1/m^3$, and $\mu_2 = 1/m$; then we have $\mu_1/m_1 = \mu_2/m_2 = 1/m^4$. Let π be the invariant distribution of the dynamic described above.

We start by finding the probability of node 2 to be reached, i.e.,¹⁰ $\pi[R^2 = 1] = \pi[Y^1 > 0]$. Let $H_k = \{\omega : Y^1(\omega) = k/m_1\}$ for $k = 0, \dots, m_1$. As there is selection in node 1 towards b^1 , we have $P[\omega_{t+1} \in H_{k+1} \mid \omega_t] = \mu_1(m_1 - k)/m_1$ for all $\omega_t \in H_k$, and $P[\omega_{t+1} \in H_k \mid \omega_t] = (1 - \mu_1)k/m_1$ for all $\omega_t \in H_{k+1}$. Therefore (see Hart [2, Footnote 70]), we have

$$\pi[H_k] \frac{\mu_1(m_1 - k)}{m_1} = \pi[H_{k+1}] \frac{(1 - \mu_1)k}{m_1}, \text{ or}$$

$$\pi[H_{k+1}] = \pi[H_k] \frac{\mu_1}{1 - \mu_1} \frac{m_1 - k}{k},$$

¹⁰As the chosen individual in population 1 has probabilities of μ_1 and $1 - \mu_1$ to play c^1 and b^1 respectively, we are able to show that Y^1 is distributed according to the binomial distribution with m_1 and μ_1 .

and for all k we get

$$\begin{aligned}\pi[H_k] &= \pi[H_0] \binom{m_1}{k} \left(\frac{\mu_1}{1-\mu_1}\right)^k, \text{ and} \\ \pi[R^2 = 0] &= \pi[H_0] = \left(1 + \frac{\mu_1}{1-\mu_1}\right)^{-m_1} = \left(1 - \frac{1}{m^3}\right)^m.\end{aligned}\quad (4.1)$$

To show that Y^2 does not converge to 0, let $\omega \in \Omega$ be distributed according to π , and let $\tilde{\omega} \in \Omega$ be the next state, given by the transition probabilities; then $\tilde{\omega}$ is also distributed according to π . Let $dY := Y^2(\tilde{\omega}) - Y^2(\omega)$; then we have $E[dY] = 0$.

When $R^2 = 1$ there is selection in node 2 towards b^2 , and when $R^2 = 0$ there are only mutations. Therefore, we have

$$\begin{aligned}E[dY \mid R^2 = 1] &= \frac{\mu_2}{m_2} - \frac{1}{m_2} E[Y^2(\omega) \mid R^2 = 1], \text{ and} \\ E[dY \mid R^2 = 0] &= \frac{\mu_2}{m_2} - \frac{2\mu_2}{m_2} E[Y^2(\omega) \mid R^2 = 0].\end{aligned}$$

Therefore, we get

$$\begin{aligned}0 = E[dY] &= E[dY \mid R^2 = 1] \pi[R^2 = 1] + E[dY \mid R^2 = 0] \pi[R^2 = 0] \\ &= \frac{\mu_2}{m_2} - \frac{2\mu_2}{m_2} E[Y^2(\omega)] - \frac{1-2\mu_2}{m_2} E[Y^2(\omega) \mid R^2 = 1] \pi[R^2 = 1].\end{aligned}$$

After rearranging and using (4.1), we get

$$\begin{aligned}E[Y^2(\omega)] &= \frac{1}{2} - \frac{1-2\mu_2}{2\mu_2} E[Y^2(\omega) \mid R^2 = 1] \pi[R^2 = 1] \\ &\geq \frac{1}{2} - \frac{1-2\mu_2}{2\mu_2} \pi[R^2 = 1] = \frac{1}{2} - \frac{m-2}{2} \left(1 - \left(1 - \frac{1}{m^3}\right)^m\right) \\ &\geq \frac{1}{2} - \frac{m-2}{2} \frac{1}{m^2} \xrightarrow{m \rightarrow \infty} \frac{1}{2}.\end{aligned}$$

4.2 Population Sizes

For simplicity we have assumed that all the populations are of equal size m . In general, we have a *population profile* $\mathbf{m} = (M(i))_{i \in N}$, with corresponding dynamic processes on the state space $\Omega_{\mathbf{m}}$.

Letting now $m := \min_{i \in N} m_i$, and $\bar{m} := \max_{i \in N} m_i$, and replacing (2.2) with

$$\frac{\gamma_1}{m_i} \leq Q[q(i) = q \mid \omega] \leq \frac{\gamma_2}{m_i} \text{ for each } q \in M(i), \quad (4.2)$$

our Main Theorem becomes:

Theorem 4.1. *For any $\varepsilon > 0$, there exists a constant C such that for any mutation rate μ , population profile \mathbf{m} , and dynamic system Q satisfying conditions (2.1), (4.2), (2.3) - (2.6), with invariant distribution π_Q , we have*

$$\hat{\pi}_Q[B I_\varepsilon] \geq 1 - C \left(\left(\frac{1}{m} + \mu \frac{\bar{m}}{m} \right) \exp(C\mu\bar{m}) \exp\left(C\frac{\bar{m}}{m}\right) \right). \quad (4.3)$$

Equation (4.3), together with Hart [2, (3.6)], imply (3.2) (and hence the Main Theorem) for models with comparable population sizes (i.e., $\bar{m}/m \leq D$, for some constant D).

Remarks:

1. As seen in Section 4.1, we must assume that the mutation rates are comparable. Together with the assumption that the probabilities of the mutations of different individuals are comparable, the assumption on the sizes of the population is not a further restriction on the model.
2. We did not manage to obtain a better bound than (4.3) on the prob-

ability of BI_ε , and hence we need the populations to be comparable.

We do not know whether this assumption is indeed needed.

5 The Proof

In this section we use a few general propositions on Markov chains to describe the long-run behavior of our dynamics, and to prove the Main Theorem.

5.1 Notations

We present some notations about transition times, and about the game and the dynamic system.

Given a finite state space Ω with a transition matrix $Q = (q_{\omega, \omega'})_{\omega, \omega' \in \Omega}$, we use the following notations.

- For each $\Omega' \subset \Omega$, define $T(\Omega') = \min \{t \geq 0 : X_t \in \Omega'\}$, the *first entrance time* to Ω' .
- For every $\omega \in \Omega$, define $U(\omega, \Omega') = E[T(\Omega') \mid X_0 = \omega]$; i.e., $U(\omega, \Omega')$ is the expected number of periods it takes to get to Ω' starting from ω .
For each $\Omega_1, \Omega_2 \subset \Omega$ let $U(\Omega_1, \Omega_2) = \max_{\omega \in \Omega_1} U(\omega, \Omega_2)$.
- For all $\omega \in \Omega$ and $\Omega' \subset \Omega$, let $Q(\omega, \Omega') = \sum_{\omega' \in \Omega'} q_{\omega, \omega'}$ — the probability of going from ω to Ω' in one step.
- For an event \mathbf{A} and $\omega \in \Omega$, we use $P_\omega[\mathbf{A}]$ for $P[\mathbf{A} \mid X_0 = \omega]$.
- For every $\omega \in \Omega$ let $\tilde{\omega}$ be the next state, i.e., for all $\omega' \in \Omega$, we have $P_\omega[\tilde{\omega} = \omega'] = P[X_{t+1} = \omega' \mid X_t = \omega] = q_{\omega, \omega'}$.

- For any partition $\{H_k\}_{k=0}^K$ of Ω , and for every $0 \leq k \leq K$, let $A_k = \bigcup_{l=0}^{k-1} H_l$ and $B_k = \bigcup_{l=k+1}^K H_l$.

Given a game Γ , and population profile \mathbf{m} , we use the following notations about Γ and $\Omega_{\mathbf{m}}$.

- For all $i \in N$, let $Y^i(\omega) = 1 - x_{b^i}^i(\omega)$. This is the proportion of population i that does not play the backward induction strategy.
- Given two vertices $i, j \in \widehat{N}$ such that i is a descendant of j , let $R^{j,i}(\omega)$ be an indicator random variable, defined as 1 if i is reached from j in state ω , and 0 otherwise. For every node $k \in N$ on the path from j to i , let $a^{k,i} \in A^k$ be the strategy that leads toward i . When j is the root we will use $R^i(\omega)$ instead of $R^{j,i}(\omega)$.¹¹
- Let $\lambda = \lambda(\Gamma) > 0$ be such that for each node $i \in N$, if for every $j \in N(i)$ we have $Y^j(\omega) < \lambda$, then b^i is the unique local best reply of i .
- For all $i \in N$, let $L^i(\omega)$ be an indicator random variable, defined as 1 if $Y^j(\omega) < \lambda$ for all $j \in N(i)$, and 0 otherwise.
- For all $i \in N$, let $s(i) \in \widehat{N}$ be the immediate successor of i such that $a^{i,s(i)} = b^i$.
- Let $\varepsilon > 0$. For all $j \in N$ define $G_j(\varepsilon) = \{\omega : Y^l(\omega) \leq \varepsilon, \forall l \geq j\}$, and define $G_{n+1}(\varepsilon) = \Omega_{\mathbf{m}}$. Notice that $G_1(\varepsilon) = x^{-1}(BI_\varepsilon)$. For all $j = 1, \dots, n+1$ and $i \in N$, define $G_{j,i}(\varepsilon) = \{\omega \in G_j(\varepsilon) : R^{s(i)}(\omega) = 1\}$.

¹¹The root may be a chance move, and hence not in \widehat{N} , but this does not change the definition of R^i .

I.e., in $G_{j,i}(\varepsilon)$, node i is reached, $s(i)$ is reached from i , and for all the nodes $l \geq j$, we have $Y^l \leq \varepsilon$.

We will now restate the conditions satisfied by our class of dynamics. These conditions hold for every μ , \mathbf{m} , $Q \in \Xi(\mathbf{m}, \mu)$, and $i \in N$.

All strategies have a positive probability of being chosen by mutation:

$$P_\omega[\tilde{\omega}_{q(i)}^i = a^i] \geq \alpha_1 \mu \text{ for every } a^i \in A^i. \quad (5.1)$$

If node i is not reached, i.e., $R^i(\omega) = 0$, then all strategies of i yield the same payoff and only mutation affects ω^i . Therefore:

$$\text{If } R^i = 0 \text{ then } P_\omega[\tilde{\omega}_{q(i)}^i \neq \omega_{q(i)}^i] \leq \alpha_2 \mu. \quad (5.2)$$

If $R^i(\omega) = 1$ and $L^i(\omega) = 1$ then b^i is the global best reply of i and thus certainly a “better strategy” for a “non- b^i individual” (i.e., $b^i \in B^i(q(i), \omega)$ when $\omega_{q(i)}^i \neq b^i$), and there isn’t any “better strategy” for a “ b^i individual” (i.e., $B(q(i), \omega) = \phi$ when $\omega_{q(i)}^i = b^i$). Therefore:

$$\text{If } L^i R^i = 1, \omega_{q(i)}^i \neq b^i \text{ and } b^i \in S^i \text{ then } P_\omega[\tilde{\omega}_{q(i)}^i = b^i] \geq \sigma. \quad (5.3)$$

$$\text{If } L^i R^i = 1 \text{ and } \omega_{q(i)}^i = b^i \text{ then } P_\omega[\tilde{\omega}_{q(i)}^i \neq b^i] \leq \alpha_2 \mu. \quad (5.4)$$

$$\text{If } L^i R^i = 1 \text{ then } P_\omega[Y^i(\tilde{\omega}) > Y^i(\omega)] \leq \alpha_2 \mu. \quad (5.5)$$

Using (5.2) and (5.5) we get:

$$\text{If } L^i = 1 \text{ then } P_\omega[Y^i(\tilde{\omega}) > Y^i(\omega)] \leq \alpha_2 \mu. \quad (5.6)$$

Only one individual may change his strategy (by (2.3)); therefore:

$$P_\omega[Y^i(\tilde{\omega}) \leq Y^i(\omega) + \frac{1}{m_i}] = 1. \quad (5.7)$$

If $R^{j,i}(\omega) = 1$, then for all the nodes $l \in N$ on the path from j to i , there must exist $q \in M(l)$ such that $w_q^l = a^{l,i}$. To get $R^{j,i}(\tilde{\omega}) = 0$, there must be some node l on the path where no one plays $a^{l,i}$, which means that the only individual who played $a^{l,i}$ must have been chosen. Therefore (by (4.2)):

$$\text{If } R^{j,i} = 1 \text{ then } P_\omega[R^{j,i}(\tilde{\omega}) = 0] \leq n\gamma_2 \frac{1}{m}. \quad (5.8)$$

If $\varepsilon < \lambda$, then for all $1 \leq j \leq n+1$, $\omega \in G_j(\varepsilon)$ and every $j \leq l \leq n$, we have $Y^l(\omega) \leq \varepsilon < \lambda$. By the assumption on the numbering of the nodes, for all $i \geq j-1$, we have $N(i) \subset \{j, \dots, n\}$, and therefore:

$$\text{If } \omega \in G_j(\varepsilon) \text{ then } L^i(\omega) = 1 \text{ for all } i \geq j-1. \quad (5.9)$$

If $j \leq n$ and $\tilde{\omega} \notin G_j(\varepsilon)$ then there is some node $i \geq j$ such that $Y^i(\tilde{\omega}) \geq \varepsilon > Y^i(\omega)$. As $G_{n+1}(\varepsilon) = \Omega_{\mathbf{m}}$, we have for all $1 \leq j \leq n+1$ (using (5.9) and (5.6) for $j \leq n$):

$$\text{If } \omega \in G_j(\varepsilon) \text{ then } P_\omega[\tilde{\omega} \notin G_j(\varepsilon)] \leq n\alpha_2\mu. \quad (5.10)$$

If $\varepsilon < \lambda$ and $\omega \in G_{j,i}(\varepsilon)$, then to get $\tilde{\omega} \notin G_{j,i}(\varepsilon)$, we need either $\tilde{\omega} \notin G_j(\varepsilon)$ or $R^{s(i)} = 0$. Therefore (using (5.8) and (5.10)):

$$\text{If } \omega \in G_{j,i}(\varepsilon) \text{ then } P_\omega[\tilde{\omega} \notin G_{j,i}(\varepsilon)] \leq n\gamma_2 \frac{1}{m} + n\alpha_2\mu. \quad (5.11)$$

As $a^{i,s(i)} = b^i$, we have:

$$\text{If } R^{i,s(i)} = 1 \text{ then } b^i \in S^i. \quad (5.12)$$




Figure 2: The Markov Chain over Time

5.2 The Proof of the Main Theorem

To prove the Main Theorem, we estimate transition times (using Proposition A.3) and thereby obtain the invariant distribution of the system (Proposition A.5).

Specifically (assume for simplicity that $|M(i)| = m$ for all $i \in N$), the proof is based on the following two properties of our dynamic system, which hold for any μ , m , and $Q \in \Xi(m, \mu)$ we choose:

P1 The expected time to reach BI_ϵ is at most $C_1(m + 1/\mu) \exp(C_1 m \mu)$, for some constant C_1 .

P2 The expected time to leave BI_ϵ is at least $C_2 m / \mu$, for some constant C_2 .

If we assume that $m\mu \leq D$ for some $D < \infty$ (the case $m\mu \geq D$ is treated in Hart [2]), then the Markov chain looks like Figure 2. Therefore, on average, every $C_1(m + 1/\mu) + C_2 m / \mu$ periods, the system is in BI_ϵ at least $C_2 m / \mu$ periods, and that ratio is the invariant probability of BI_ϵ , and we obtain (3.2).

Let the Game Γ , constants γ_1 , γ_2 , σ , α_1 , and α_2 be fixed. The following results are true for all $\mu < 1/(2n\alpha_2)$, \mathbf{m} such that $m > 2n\gamma_2$, and $Q \in \Xi(\mathbf{m}, \mu)$.

We start by showing that **P1** holds. The proof of **P1** is by backward induction, and for each node $i \in N$, we first estimate the time required until the node is reached (Proposition 5.1), and once the node is reached, there

is selection there towards the backward induction strategy, and we estimate the time needed until most individuals play that strategy (Proposition 5.2).

In order for node i to be reached, we need that in all nodes on the path from the root to i , at least one individual will play the strategy that leads toward i . This can always occur due to mutation, and as there is no selection in nodes that are not reached, we need sequential rather than simultaneous mutations, and the expected time until the node is reached is of the order of $1/\mu$.

By the induction assumption, we have $L^i = 1$, and therefore, once the node is reached, the backward induction strategy is the unique best reply. Therefore, there is selection towards that strategy, and after $m_i \leq \bar{m}$ periods, most individuals in i play the backward induction strategy.

Proposition 5.1. *For any $\varepsilon < \lambda$ there exists a constant $C_{1,1} = C_{1,1}(\Gamma, \varepsilon)$ such that $U(\Omega_{\mathbf{m}}, G_{j,i}(\varepsilon)) \leq C_{1,1}(U(\Omega_{\mathbf{m}}, G_j(\varepsilon)) + 1/\mu)$ for all $i \in N$ and $i < j \leq n + 1$.*

Proof. Without loss of generality assume that $s(i)$ is a node (if $s(i)$ is a terminal vertex, the only difference is that $s(i)$ is not numbered), $s(i) = i + 1$, and that the nodes along the path from the root to i are $1, 2, \dots, i$.

Let¹² $H_k = \{\omega \in G_j(\varepsilon) : R^{k,k+1} = 0, R^{k+1,s(i)} = 1\}$ for $k = 1, \dots, i$; i.e., H_k is the set of all states in $G_j(\varepsilon)$ such that $s(i)$ is reached from $k + 1$, but in node k no one plays the strategy towards $k + 1$. Let $H_0 = G_{j,i}(\varepsilon)$ and let $H_{i+1} = \Omega_m \setminus G_j(\varepsilon)$. If $s(i)$ is not reached, then there is a node where no one plays the strategy towards $s(i)$, and there is a unique maximum node of this

¹²Where $R^{i,i}(\omega) = 1$ for all ω .

kind. Therefore $\{H_k\}_{k=0}^{i+1}$ is a partition of Ω_m .

Let $1 \leq k \leq i$ and $\omega \in H_k$. Then $\omega \in G_j(\varepsilon)$, $R^{k,k+1} = 0$ and $R^{k+1,s(i)} = 1$.

Therefore, for all $k < l \leq i$ we have $R^l = 0$ and $R^{l,l+1} = 1$.

To get¹³ $\tilde{\omega} \in B_k$ we need either $\tilde{\omega} \in H_{i+1}$, and by (5.10), we have $P_\omega[\tilde{\omega} \in H_{i+1}] = P_\omega[\tilde{\omega} \notin G_j(\varepsilon)] = O(\mu)$, or $\tilde{\omega} \in H_l$ for some $k < l \leq i$, and, therefore, there must be $R^{l,l+1}(\tilde{\omega}) = 0$, which can happen only by mutation (by (5.2)). Therefore there is a constant g_k (which depends on the game and not on m or μ) such that

$$P_\omega[\tilde{\omega} \in B_k] \leq g_k \mu. \quad (5.13)$$

To get $\tilde{\omega} \in A_k$, we need $\tilde{\omega} \in G_j(\varepsilon)$, $R^{k,k+1}(\tilde{\omega}) = 1$, and $R^{l,l+1}(\tilde{\omega}) = 1$ for all $k < l \leq i$. As each one of those events depends on different nodes (the first on nodes $l \geq j \geq i+1$, the second on node $l = k$, and the third on nodes $k < l \leq i$), we can use the conditional independence (2.1) to calculate the probabilities of those events.

By (5.10), we have $P_\omega[\tilde{\omega} \in G_j(\varepsilon)] \geq (1 - n\alpha_2\mu)$. To get $R^{k,k+1}(\tilde{\omega}) = 1$, we need $q(k)$ to choose by mutation the strategy $a^{k,k+1}$, and by (5.1) we get $P_\omega[R^{k,k+1}(\tilde{\omega}) = 1] \geq \alpha_1\mu$. Finally, by (5.2), we have $P_\omega[R^{l,l+1}(\tilde{\omega}) = 1] \geq (1 - \alpha_2\mu)$ for all $k < l \leq i$. Therefore, there is a constant $f_k > 0$ such that

$$P_\omega[\tilde{\omega} \in A_k] \geq (1 - n\alpha_2\mu)\alpha_1\mu(1 - \alpha_2\mu)^{i-k} \geq f_k\mu. \quad (5.14)$$

¹³Recall that for a partition $\{H_k\}_{k=0}^K$, we use $A_k = \bigcup_{l=0}^{k-1} H_l$ and $B_k = \bigcup_{l=k+1}^K H_l$.

Using Corollary A.4 on $\{H_k\}_{k=0}^{i+1}$ with (5.13) and (5.14) we get

$$\begin{aligned} U(\Omega_{\mathbf{m}}, G_{j,i}(\varepsilon)) &= U(\Omega_{\mathbf{m}}, H_0) \leq \left(U(\Omega_{\mathbf{m}}, \Omega_{\mathbf{m}} \setminus H_{i+1}) + \sum_{k=1}^i \frac{1}{f_k \mu} \right) \prod_{k=1}^i \left(1 + \frac{g_k \mu}{f_k \mu} \right) \\ &\leq C_{1,1} \left(U(\Omega_{\mathbf{m}}, G_j(\varepsilon)) + \frac{1}{\mu} \right), \end{aligned}$$

for $C_{1,1} = (1 + \sum_{k=1}^i 1/f_k) \prod_{k=1}^i (1 + g_k/f_k)$. \square

Proposition 5.2. *For any $\varepsilon < \lambda$ there exists a constant $C_{1,2} = C_{1,2}(\Gamma, \varepsilon)$ such that $U(\Omega_{\mathbf{m}}, G_i(\varepsilon)) \leq C_{1,2} (U(\Omega_{\mathbf{m}}, G_{i+1,i}(\varepsilon)) + \overline{m}) \exp(C_{1,2}(\overline{m}\mu + \overline{m}/m))$ for all $i \in N$.*

Proof. Without loss of generality assume that εm_i is an integer. Put $G = G_{i+1,i}(\varepsilon)$ and $K = m_i - \varepsilon m_i$. Let $H_K = \Omega_{\mathbf{m}} \setminus G$, let $H_k = \{\omega \in G : Y^i(\omega) = (k + \varepsilon m_i)/m_i\}$ for $k = 1, \dots, K-1$, and let $H_0 = \{\omega \in G : Y^i(\omega) \leq \varepsilon\} \subset G_i(\varepsilon)$. Then $\{H_k\}_{k=0}^{K-1}$ is a partition of G according to the value of Y^i , and $\{H_k\}_{k=0}^K$ is a partition of $\Omega_{\mathbf{m}}$.

Let $1 \leq k < K$ and $\omega \in H_k$. Then, by (5.9) and (5.12), we have $L^i = 1$, $R^i = 1$, and $b^i \in S^i$.

To get $\tilde{\omega} \in B_k$ we need either $\tilde{\omega} \in H_l$ for $k < l < K$, and therefore we have $Y^i(\tilde{\omega}) > Y^i(\omega)$, or $\tilde{\omega} \in H_K$, and therefore $\tilde{\omega} \notin G$. By (5.5) and (5.11), there is a constant g (which does not depend on k) such that

$$P_{\omega}[\tilde{\omega} \in B_k] \leq g\left(\mu + \frac{1}{m}\right). \quad (5.15)$$

To get $\tilde{\omega} \in H_{k-1} \subset A_k$, we need $\tilde{\omega} \in G$ and $Y^i(\tilde{\omega}) = Y^i(\omega) - 1/m_i$. Notice that if $Y^i(\tilde{\omega}) = Y^i(\omega) - 1/m_i < 1$, then at least one individual plays

$b^i = a^{i,s(i)}$, and therefore $R^{i,s(i)}(\tilde{\omega}) = 1$. Therefore, to get $\tilde{\omega} \in H_{k-1}$ we need $\tilde{\omega} \in G_{i+1}(\varepsilon)$, $R^i(\tilde{\omega}) = 1$, and $Y^i(\tilde{\omega}) = Y^i(\omega) - 1/m_i$. Again, those events depend on different nodes, and we use the conditional independence (2.1).

By (5.10), we have $P_\omega[\tilde{\omega} \in G_{i+1}(\varepsilon)] \geq (1 - n\alpha_2\mu)$. By (5.8), we have $P_\omega[R^i(\tilde{\omega}) = 1] \geq (1 - n\gamma_2/m)$. To get $Y^i(\tilde{\omega}) = Y^i(\omega) - 1/m_i$, we need $q(i)$ to be chosen from the set $E = \{q \in M(i) : w_q^i \neq b^i\}$ and change his strategy to b^i . As $|E|/m_i = Y^i(\omega) = (k + \varepsilon m_i)/m_i \geq \varepsilon$, we have (using (4.2) and (5.3)) $P_\omega[Y^i(\tilde{\omega}) = Y^i(\omega) - 1/m_i] \geq \gamma_1\varepsilon\sigma$. Therefore, there is a constant $f > 0$ such that

$$P_\omega[\tilde{\omega} \in A_k] \geq P_\omega[\tilde{\omega} \in H_{k-1}] \geq (1 - n\alpha_2\mu)(1 - \frac{n\gamma_2}{m})\gamma_1\varepsilon\sigma \geq f\varepsilon. \quad (5.16)$$

Using Corollary A.4 on $\{H_k\}_{k=0}^K$ with (5.15) and (5.16) we get

$$\begin{aligned} U(\Omega_{\mathbf{m}}, G_i(\varepsilon)) &\leq U(\Omega_{\mathbf{m}}, H_0) \leq \left(U(\Omega_{\mathbf{m}}, \Omega_{\mathbf{m}} \setminus H_K) + \sum_{k=1}^{K-1} \frac{1}{f\varepsilon} \right) \prod_{k=1}^{K-1} \left(1 + \frac{g(\mu + \frac{1}{m})}{f\varepsilon} \right) \\ &\leq (U(\Omega_{\mathbf{m}}, G_{i+1,i}(\varepsilon)) + \frac{m_i}{f\varepsilon}) \left(1 + \frac{g}{f\varepsilon} (\mu + \frac{1}{m}) \right)^{m_i} \\ &\leq \left(1 + \frac{1}{f\varepsilon} \right) \cdot (U(\Omega_{\mathbf{m}}, G_{i+1,i}(\varepsilon)) + \overline{m}) \cdot \exp\left(\frac{g}{f\varepsilon} (\overline{m}\mu + \frac{\overline{m}}{m}) \right) \\ &\leq C_{1,2} \cdot (U(\Omega_{\mathbf{m}}, G_{i+1,i}(\varepsilon)) + \overline{m}) \exp(C_{1,2}(\overline{m}\mu + \overline{m}/m)), \end{aligned}$$

for $C_{1,2} = 1 + (g + 1)/(f\varepsilon)$. □

We now use Propositions 5.1 and 5.2 to estimate the expected time to reach BI_ε ,¹⁴ and hence show that **P1** holds.

¹⁴When $m_i = m$ for all $i \in N$, this becomes Proposition 3.6.

Proposition 5.3. *For any $\varepsilon > 0$ there exists a constant $C_1 = C_1(\Gamma, \varepsilon)$ such that $U(\Omega_{\mathbf{m}}, G_1(\varepsilon)) \leq C_1(1/\mu + \overline{m}) \exp(C_1(\overline{m}\mu + \overline{m}/m))$.*

Proof. If $\varepsilon_1 \geq \varepsilon_2 > 0$ then $G_1(\varepsilon_2) \subset G_1(\varepsilon_1)$ and therefore $U(\Omega_{\mathbf{m}}, G_1(\varepsilon_1)) \leq U(\Omega_{\mathbf{m}}, G_1(\varepsilon_2))$. Thus, assume that $\varepsilon < \lambda$. Using Propositions 5.1 and 5.2 we get, for all¹⁵ $i \in N$,

$$\begin{aligned} U(\Omega_{\mathbf{m}}, G_i(\varepsilon)) &\leq C_{1,2} \cdot (U(\Omega_{\mathbf{m}}, G_{i+1,i}(\varepsilon)) + \overline{m}) \exp(C_{1,2}(\overline{m}\mu + \overline{m}/m)) \\ &\leq C_{1,2} \cdot (C_{1,1} \cdot (U(\Omega_{\mathbf{m}}, G_{i+1}(\varepsilon)) + \frac{1}{\mu}) + \overline{m}) \exp(C_{1,2}(\overline{m}\mu + \overline{m}/m)) \\ &\leq C_{1,1}C_{1,2}(U(\Omega_{\mathbf{m}}, G_{i+1}(\varepsilon)) + \frac{1}{\mu} + \overline{m}) \exp(C_{1,2}(\overline{m}\mu + \overline{m}/m)). \end{aligned}$$

Using induction and the equality $U(\Omega_{\mathbf{m}}, G_{n+1}(\varepsilon)) = 0$, we get,

$$U(\Omega_{\mathbf{m}}, G_1(\varepsilon)) \leq C_1(\frac{1}{\mu} + \overline{m}) \exp(C_1(\overline{m}\mu + \overline{m}/m)),$$

for $C_1 = n(C_{1,1}C_{1,2})^n$. □

We now estimate the expected time to leave BI_ε , and show that **P2** holds. Notice that if $\varepsilon < \lambda$, then in any state in BI_ε , and for any node i , either node i is not reached, or node i is reached and b^i is the unique best reply of i . Therefore, the proportion of b^i will decrease only by mutation.

Proposition 5.4. *Let $0 < \varepsilon_1 < \varepsilon_2 < \lambda$. Then there exists a constant $C_2 = C_2(\Gamma, \varepsilon_2 - \varepsilon_1) > 0$ such that $U(\omega, \Omega_{\mathbf{m}} \setminus G_1(\varepsilon_2)) \geq C_2(m/\mu)$ for every $\omega \in G_1(\varepsilon_1)$.*

¹⁵We may assume without loss of generality that $C_{1,1} \geq 1$ and $C_{1,2} \geq 1$.

Proof. For every $\omega \in G_1(\varepsilon_2)$, and every $i \in N$, we have $P_\omega[Y^i(\tilde{\omega}) > Y^i(\omega) + 1/m] \leq P_\omega[Y^i(\tilde{\omega}) > Y^i(\omega) + 1/m_i] = 0$, and $P_\omega[Y^i(\tilde{\omega}) > Y^i(\omega)] \leq \alpha_2\mu$ (by (5.7), (5.9), and (5.6)). Let $Z(\omega) = \max_{i \in N} Y^i(\omega)$; then $P_\omega[Z(\tilde{\omega}) > Z(\omega) + 1/m] = 0$ and $P_\omega[Z(\tilde{\omega}) > Z(\omega)] \leq n\alpha_2\mu$. Therefore, for every $z \leq \lambda$, a move from $Z(\omega) = z$ to $z < Z(\omega) \leq z + 1/m$ takes, on average, at least $1/(n\alpha_2\mu)$ periods.

As $Z(\omega) \leq \varepsilon_1$ for every $\omega \in G_1(\varepsilon_1)$ and $Z(\omega) > \varepsilon_2$ for every $\omega \notin G_1(\varepsilon_2)$, a move from $G_1(\varepsilon_1)$ to $\Omega_{\mathbf{m}} \setminus G_1(\varepsilon_2)$ takes on average at least $(\varepsilon_2 - \varepsilon_1)m/(n\alpha_2\mu)$ periods, or $U(\omega, \Omega_{\mathbf{m}} \setminus G_1(\varepsilon_2)) \geq C_2(m/\mu)$ for every $\omega \in G_1(\varepsilon_1)$. \square

Now we can prove Theorem 4.1:

Proof of Theorem 4.1. Without loss of generality assume that $\varepsilon < \lambda$. Using Proposition 5.3, there is a constant C_1 such that $U(\Omega_{\mathbf{m}}, G_1(\varepsilon/2)) \leq C_1(1/\mu + \bar{m}) \exp(C_1(\bar{m}\mu + \bar{m}/m))$. Using Proposition 5.4, there is a constant $C_2 > 0$ such that $U(\omega, \Omega_{\mathbf{m}} \setminus G_1(\varepsilon)) \geq C_2(m/\mu)$ for every $\omega \in G_1(\varepsilon/2)$. By Proposition A.5 we have

$$\begin{aligned} \hat{\pi}_Q[X \setminus BI_\varepsilon] &= \pi_Q[\Omega_{\mathbf{m}} \setminus G_1(\varepsilon)] \\ &\leq \pi_Q[G_1(\varepsilon)] \frac{C_1(1/\mu + \bar{m}) \exp(C_1(\bar{m}\mu + \bar{m}/m))}{C_2(m/\mu)} \\ &\leq \frac{C_1}{C_2} \left(\frac{1}{m} + \mu \frac{\bar{m}}{m} \right) \exp(C_1(\bar{m}\mu + \frac{\bar{m}}{m})) \\ &\leq C \left(\left(\frac{1}{m} + \mu \frac{\bar{m}}{m} \right) \exp(C\mu\bar{m}) \exp(C\frac{\bar{m}}{m}) \right), \end{aligned}$$

for $C = C_1(1 + 1/C_2)$. \square

A Appendix: Some General Results on Markov Chains

In this appendix, we prove a few general propositions on Markov chains.

Lemma A.1. *For every $\Omega_1, \Omega_2 \subset \Omega$, and $\omega \in \Omega$, we have $U(\omega, \Omega_2) \leq U(\omega, \Omega_1) + U(\Omega_1, \Omega_2)$.*

Proof. Let $T_1 = T(\Omega_1)$ and define $T_2 = \min \{t \geq T_1 : X_t \in \Omega_2\}$; then $T_2 \geq T(\Omega_2)$. By the strong Markov property (see Revuz [5, Chapter 1, Theorem 3.5]), we have for any t , $P[T_2 = T_1 + t \mid X_0 = \omega] = \sum_{\omega' \in \Omega_1} P[X_{T_1} = \omega' \mid X_0 = \omega] \cdot P[T(\Omega_2) = t \mid X_0 = \omega']$. Therefore, $U(\omega, \Omega_2) \leq E[T_2 \mid X_0 = \omega] = E[T_1 \mid X_0 = \omega] + \sum_{\omega' \in \Omega_1} P[X_{T_1} = \omega' \mid X_0 = \omega] \cdot E[T(\Omega_2) \mid X_0 = \omega] \leq U(\omega, \Omega_1) + U(\Omega_1, \Omega_2)$. \square

Let Ω be a finite state space, and let $\{H_k\}_{k=0}^K$ be a partition of Ω . Using the notations introduced in Section 5.1, we can now prove the propositions we need about transition times.

Lemma A.2. *For all $1 \leq k \leq l \leq K$ we have $U(H_l, A_k) \leq \sum_{i=k}^l U(H_i, A_i)$.*

Proof. Let $1 \leq l \leq K$; we will use induction on k . For $k = l$, we have $U(H_l, A_k) = U(H_l, A_l)$. Let $k < l$. Since $U(A_k, A_k) = 0$ and $U(H_k, A_k) \geq 0$, we get $U(A_k \cup H_k, A_k) = \max \{U(A_k, A_k), U(H_k, A_k)\} = U(H_k, A_k)$. Therefore, using Lemma A.1 and induction, we get

$$\begin{aligned} U(H_l, A_k) &\leq U(H_l, A_k \cup H_k) + U(A_k \cup H_k, A_k) \\ &= U(H_l, A_{k+1}) + U(H_k, A_k) \leq \sum_{i=k}^l U(H_i, A_i). \end{aligned} \quad \square$$

Proposition A.3. Assume that for all $1 \leq k \leq K$ there are constants $f_k > 0$ and g_k , such that $Q(\omega, A_k) \geq f_k$ and $Q(\omega, B_k) \leq g_k$ for all $\omega \in H_k$. Then

$$U(\Omega, H_0) \leq \sum_{k=1}^K \left[\frac{1}{f_k} \prod_{l=1}^{k-1} \left(1 + \frac{g_l}{f_l} \right) \right] \leq \left(\sum_{k=1}^K \frac{1}{f_k} \right) \left[\prod_{l=1}^{K-1} \left(1 + \frac{g_l}{f_l} \right) \right].$$

Proof. For every $1 \leq k \leq K$, let $u_k = U(H_k, A_k)$ and $v_k = u_k + \dots + u_K$, and let $v_{K+1} = 0$. By Lemma A.2, for all $1 \leq k \leq K$ we have $U(B_{k-1}, A_k) = \max_{l=k, \dots, K} U(H_l, A_k) \leq v_k$.

Let $\omega_k \in H_k$ such that $U(H_k, A_k) = u_k$ is attained at ω_k , i.e., $u_k = U(\omega_k, A_k)$. As $\omega_k \notin A_k$, we have, first for $k = K$,

$$\begin{aligned} u_K &= 1 + \sum_{\omega \in \Omega} q_{\omega_K, \omega} U(\omega, A_K) \\ &= 1 + \sum_{\omega \in H_K} q_{\omega_K, \omega} U(\omega, A_K) + \sum_{\omega \in A_K} q_{\omega_K, \omega} U(\omega, A_K) \\ &\leq 1 + U(H_K, A_K) Q(\omega_K, H_K) + 0 \\ &\leq 1 + u_K(1 - f_K) = 1 + v_{K+1} g_K + u_K(1 - f_K), \end{aligned}$$

and for $1 \leq k < K$, we have

$$\begin{aligned} u_k &= 1 + \sum_{\omega \in B_k} q_{\omega_k, \omega} U(\omega, A_k) + \sum_{\omega \in H_k} q_{\omega_k, \omega} U(\omega, A_k) + \sum_{\omega \in A_k} q_{\omega_k, \omega} U(\omega, A_k) \\ &\leq 1 + U(B_k, A_k) Q(\omega_k, B_k) + U(H_k, A_k) Q(\omega_k, H_k) + 0 \\ &\leq 1 + U(B_k, A_k) Q(\omega_k, B_k) + U(H_k, A_k) (1 - Q(\omega_k, A_k)) \\ &\leq 1 + v_{k+1} g_k + u_k(1 - f_k). \end{aligned}$$

After rearrangement we have $u_k \leq 1/f_k + (g_k/f_k)v_{k+1}$, or $v_k = u_k +$

$v_{k+1} \leq 1/f_k + (1 + g_k/f_k)v_{k+1}$. Applying backward induction on k , starting with $v_{K+1} = 0$, we get $v_1 \leq \sum_{k=1}^K \left[1/f_k \prod_{l=1}^{k-1} (1 + g_l/f_l) \right]$. As $U(\Omega, H_0) = U(B_0, A_1) \leq v_1$, and $\prod_{l=1}^{k-1} (1 + g_l/f_l) \leq \prod_{l=1}^{K-1} (1 + g_l/f_l)$ for all $1 \leq k \leq K$, we complete the proof. \square

Remark A.1. In the proof of Proposition A.3 we use the inequality $v_K \leq 1/f_K$, and of all the assumptions on the one-step transition probabilities on H_K we only need the assumption that v_K is finite.

Corollary A.4. *Using the notations in Proposition A.3, with the assumptions on the one-step transition probabilities only on H_k for $1 \leq k < K$, we have*

$$U(\Omega, H_0) \leq \left(U(\Omega, \Omega \setminus H_K) + \sum_{k=1}^{K-1} \frac{1}{f_k} \right) \left[\prod_{l=1}^{K-1} \left(1 + \frac{g_l}{f_l} \right) \right].$$

Proof. If $v_K = U(\Omega, \Omega \setminus H_K) = \infty$, the claim is trivial. Otherwise, let $f_K = 1/v_K$. Using Proposition A.3 and Remark A.1 we complete the proof. \square

Remark A.2. For $H_K = \phi$, we have $U(\Omega, \Omega \setminus H_K) = U(\Omega, \Omega) = 0$, and, therefore, using the partition $\{H_k\}_{k=0}^{K-1}$, we get

$$U(\Omega, H_0) \leq \left(\sum_{k=1}^{K-1} \frac{1}{f_k} \right) \prod_{l=1}^{K-2} \left(1 + \frac{g_l}{f_l} \right) \leq \left(U(\Omega, \Omega \setminus H_K) + \sum_{k=1}^{K-1} \frac{1}{f_k} \right) \left[\prod_{l=1}^{K-1} \left(1 + \frac{g_l}{f_l} \right) \right].$$

Proposition A.5. *Assume that the Markov chain is irreducible and aperiodic with an invariant distribution π . Let $\Omega_1 \subseteq \Omega_2 \subseteq \Omega$ and let C_1 and $C_2 > 0$ be constants such that $U(\Omega, \Omega_1) \leq C_1$, and $U(\omega, \Omega \setminus \Omega_2) \geq C_2$ for every $\omega \in \Omega_1$. Then $\pi[\Omega \setminus \Omega_2] \leq \pi[\Omega_2] \cdot C_1/C_2$.*

Proof. Let $X = (X_t)_{t=0}^\infty \in \Omega^\mathbb{N}$. Let $T_0 = T_0(X) = 0$ and define successively for $n \geq 1$: $T_{2n-1} = T_{2n-1}(X) = \min \{t > T_{2n-2} : X_t \in \Omega_1\}$ and

$$T_{2n} = T_{2n}(X) = \min \{t > T_{2n-1} : X_t \in \Omega \setminus \Omega_2\}.$$

For every $t \geq 1$ define: $H_t = H_t(X) = \max \{n : T_{2n} < t\}$, $Q_t = Q_t(X) = (1/t) \sum_{n=0}^{t-1} 1_{\{X_n \in \Omega_2\}}$ and $P_t = P_t(X) = (1/t) \sum_{n=0}^{t-1} 1_{\{X_n \in \Omega \setminus \Omega_2\}}$.

Then for every t

$$\begin{aligned} P_t &= \frac{1}{t} \sum_{n=0}^{t-1} 1_{\{X_n \in \Omega \setminus \Omega_2\}} = \frac{1}{t} \left[\sum_{n=0}^{T_1-1} 1_{\{X_n \in \Omega \setminus \Omega_2\}} + \cdots + \sum_{n=T_{2H_t}}^{t-1} 1_{\{X_n \in \Omega \setminus \Omega_2\}} \right] \\ &\leq \frac{1}{t} [T_1 + \cdots + (T_{2H_t-1} - T_{2H_t-2}) + (T_{2H_t+1} - T_{2H_t})] = \frac{\sum_{n=1}^{H_t+1} (T_{2n-1} - T_{2n-2})}{t}, \\ Q_t &= \frac{1}{t} \sum_{n=0}^{t-1} 1_{\{X_n \in \Omega_2\}} = \frac{1}{t} \left[\sum_{n=0}^{T_1-1} 1_{\{X_n \in \Omega_2\}} + \cdots + \sum_{n=T_{2H_t}}^{t-1} 1_{\{X_n \in \Omega_2\}} \right] \\ &\geq \frac{1}{t} [(T_2 - T_1) + \cdots + (T_{2H_t} - T_{2H_t-1})] = \frac{\sum_{n=1}^{H_t} (T_{2n} - T_{2n-1})}{t}. \end{aligned}$$

If $\{X_t\}_{t=0}^\infty$ is distributed according to Q , then for every X_0 we have $\lim_{t \rightarrow \infty} E[P_t \mid X_0] = \pi[\Omega \setminus \Omega_2]$ and $\lim_{t \rightarrow \infty} E[Q_t \mid X_0] = \pi[\Omega_2]$. By the strong Markov property, we have $E[T_{2n-1} - T_{2n-2}] \leq U(\Omega, \Omega_1) \leq C_1$ and $E[T_{2n} - T_{2n-1}] \geq C_2$. Therefore,

$$\begin{aligned} \pi[\Omega \setminus \Omega_2] &= \lim_{t \rightarrow \infty} E[P_t] \leq \overline{\lim}_{t \rightarrow \infty} E \left[\frac{\sum_{n=1}^{H_t+1} (T_{2n-1} - T_{2n-2})}{t} \right] \\ &\leq \overline{\lim}_{t \rightarrow \infty} \frac{E[H_t] + 1}{t} \cdot C_1 = C_1 \overline{\lim}_{t \rightarrow \infty} \frac{E[H_t]}{t} = \frac{C_1}{C_2} \overline{\lim}_{t \rightarrow \infty} \frac{E[H_t] C_2}{t} \\ &\leq \frac{C_1}{C_2} \overline{\lim}_{t \rightarrow \infty} E \left[\frac{\sum_{n=1}^{H_t} (T_{2n} - T_{2n-1})}{t} \right] \leq \frac{C_1}{C_2} \overline{\lim}_{t \rightarrow \infty} E[Q_t] = \frac{C_1}{C_2} \pi[\Omega_2]. \quad \square \end{aligned}$$

Acknowledgments

The author thanks Sergiu Hart for his guidance and insight, and the referees and the associate editor for helping to improve the paper. This research was

supported by the Institute of Mathematics and by the Center for the Study of Rationality, The Hebrew University of Jerusalem.

References

- [1] W. Feller. *An Introduction to Probability Theory and Its Applications*, Volume I. Wiley, New York, 1950.
- [2] S. Hart. Evolutionary dynamics and backward induction. *Games and Economic Behavior*, 41:227–264, 2002.
- [3] M. Kandori, G. Mailath, and R. Rob. Learning, mutation, and long-run equilibrium in games. *Econometrica*, 61:29–56, 1993.
- [4] H.W. Kuhn. Extensive games and the problem of information. In A.W. Kuhn, H.W. Tucker, editor, *Contributions to the Theory of Games II*, volume 28 of *Annals of Mathematics Studies*, pages 193–216. Princeton University Press, 1953.
- [5] D. Revuz. *Markov Chains*. North-Holland, 1975.
- [6] H.P. Young. The evolution of conventions. *Econometrica*, 61:57–84, 1993.