Journal of Mathematical Economics 58 (2015) 46–60

Contents lists available at ScienceDirect

Journal of Mathematical Economics

journal homepage: www.elsevier.com/locate/jmateco

Stochastic stability on general state spaces☆

Jonathan Newton

School of Economics, University of Sydney, Australia

ARTICLE INFO

Article history:
Received 3 October 2013
Received in revised form
30 November 2014
Accepted 25 March 2015
Available online 7 April 2015

Keywords:
Learning
Stochastic stability
General state space

ABSTRACT

This paper studies stochastic stability methods applied to processes on general state spaces. This includes settings in which agents repeatedly interact and choose from an uncountable set of strategies. Dynamics exist for which the stochastically stable states differ from those of any reasonable finite discretization. When there are a finite number of rest points of the unperturbed dynamic, sufficient conditions for analogues of results from the finite state space literature are derived and studied. Illustrative examples are given.

1. Introduction

The occurrence of social learning and the convergence of agents' behavior via processes of adaptive behavior is well-documented within economics (e.g. Chong et al., 2006; Selten and Apesteguia, 2005). The possibility of multiple resting points for such processes naturally leads one to question which of these stable states is more plausible from an economic perspective. Strongly influenced by evolutionary game theory (Smith and Price, 1973), a literature has grown that analyses the robustness of stable states of social learning dynamics to random errors made by players in their choice of action (Kandori et al., 1993; Young, 1993a). These ideas have been applied to a variety of economic situations, including bargaining (Binmore et al., 2003; Naidu et al., 2010), Nash demand games (Young, 1993b; Agastya, 1999), exchange economies (Serrano and Volij, 2008), local interaction on networks and the persistence of altruistic behavior (Eshel et al., 1998).

A common approach when assessing the robustness of stable states of social learning dynamics has been that pioneered by Kandori et al. (1993) and Young (1993a), building on the work of Freidlin and Wentzell (1984). Agents are assumed to make errors independently and when they do make an error are assumed to play a strategy chosen at random from a distribution with full support on a finite set of strategies. This imposes a mathematical structure on the process that leads to clear and appealing characterization results.

Unfortunately, such results cannot be straightforwardly applied when agents have non-finite sets of strategies.[1] Even assuming the convergence of the underlying social learning dynamic, the addition of random errors can lead to behavior which hinders efforts to obtain a clear cut characterization of the long run pattern of play. This paper takes up the task of analyzing the problems and intricacies which arise and, when there are a finite number of rest points of the underlying dynamic, determines a set of sufficient conditions which enable existing results to be applied to models with continuous state spaces. These conditions include a continuity requirement on error distributions and players' responses as a function of the current state, an asymptotic stability condition and a condition which ensures a specific type of discontinuity does not occur at rest points of the underlying dynamic. Examples are given showing how no subset of the conditions is sufficient on its own.

Fortunately, all of these conditions are satisfied for many common models found in economics. Typical error distributions of the kind described above coupled with the continuous best responses found in many models of industrial organization will often satisfy all of the conditions. This study applies the theory to linear quadratic games and to population models in the style of Kandori et al. (1993).

A related paper is that of Feinberg (2006), which also looks at discrete time, continuous state space processes. However, the paper in question imposes the strong condition that the perturbed

☆ I wish to thank Heather Battey, Yossi Feinberg, Johannes Hörner, George Mailath and the audience at the Econometric Society Australasian Meetings for helpful comments. The author is the recipient of a Discovery Early Career Researcher Award funded by the Australian Research Council (Grant Number: DE130101768).

E-mail address: jonathan.newton@sydney.edu.au.

[1] An early paper in the literature (Foster and Young, 1990) has an infinite state space and a continuous time dynamic in which perturbations are modeled as a Wiener process. However, it differs markedly from the majority of the literature, in which the error distributions are irrelevant to the stability results as long as they have full support.

http://dx.doi.org/10.1016/j.jmateco.2015.03.007
0304-4068/

process be governed by transition probabilities that are continuous functions of the current state of the process. The bulk of the analysis in the current paper concerns situations where this is not the case. Moreover, Feinberg considers a particular unperturbed dynamic and state space, whereas the current paper is more general in its scope. Schenk-Hoppé (2000) adapts the results of Freidlin and Wentzell (1984) and Ellison (2000) for use in finding stochastically stable states in a continuous strategy oligopoly model equipped with an imitation dynamic.

By considering finite state spaces, Young (1993a) dispenses with the need for regularity assumptions found in treatments of perturbed dynamics by Freidlin and Wentzell (1984), Kifer (1988, 1990). Specifically, all finite state spaces are compact, continuity requirements become unnecessary, and the probability of non-convergence to some stable state in given finite time need no longer be bounded by a function of error probabilities. The treatment of the current paper incorporates some finiteness in that the set of orders of magnitude of one step transition probabilities is taken to be finite. This allows us to use weaker continuity requirements on transition probabilities. We also dispense with compactness assumptions on the state space. From an economics perspective this enables, for example, the use of the Cartesian plane as the state space and the use of error probabilities which are independent across players.

The paper is organized as follows. Section 2 introduces the ideas of the paper via two motivating examples. Section 3 describes the processes of interest, gives convergence results, looks at transition probabilities between stable states, and defines a useful regularity property, showing how this property allows the problems associated with infinite state spaces to be circumvented. Section 4 gives sufficient conditions for this property to hold and discusses each of the conditions, giving examples of the problems which arise if any condition fails to hold. Section 5 gives examples. Section 6 solves an example from Section 2 for which our regularity condition fails to hold. Section 7 concludes. Formal proofs are relegated to the Appendix.

2. Motivating examples

This paper focuses on situations where agents follow some rule when deciding how to behave. The rule can be deterministic or random, cautious or hasty, imitative or best responding: any kind of behavioral bias or irregularity can be represented. Usually the rule is adaptive in the sense that an agent's behavior is intended to improve his lot. What really matters is that the rule has the Markov property: the past per se does not affect the future, although features of the present shaped by the past, including memories, are allowed to do so. We analyze situations where behavior over time will converge towards one of a number of stable states. As long as there is some probability of convergence to more than one stable state, this is predictively awkward. The possibility of random errors or idiosyncratic play justifies the introduction of perturbed versions of the process which help in obtaining long run predictions. There is a well-developed literature which deals with these problems for finite state spaces,[2] so the first question that must be addressed is whether there is benefit to be had from dealing directly with processes on general state spaces, rather than with finite discrete approximations.

2.1. Discretization can fail to represent the original process accurately

There is not always a suitable finite discretization of a process available. To illustrate, we present the following example. Consider a Markov process with state space $X = [0, 1] \subset \mathbb{R}$ endowed with the Euclidean distance metric. Let the process be governed by the Markov kernel $P(., .)$. The Markov kernel is a generalized analogue of transition probabilities on Markov chains. $P(x, A)$ gives the probability with which the process moves from state x to any state within a set of states A. For notational ease, for $y \in X$, we identify $P(., y) := P(., \{y\})$. Let $P(x, x^2) = 1$. This process has a set of stable states $\Lambda = \{0, 1\}$: from $x^* \in \Lambda$, $P(x^*, x^*) = 1$. We examine a perturbed variant of the process in which each period, with probability $1 - \varepsilon$ the unperturbed process is followed, and with probability ε the new state is drawn from the uniform distribution $U[\![0, 1]\!]$. This perturbed process has an invariant measure π_ε which converges to a measure with all weight on $\{0\}$ as $\varepsilon \to 0$: the set of *stochastically stable*[3] states is $\{0\}$.

Any discretized state space and process should satisfy some properties in order for it to be a reasonable representation of the original process. We suggest the following as reasonable restrictions on the discretized state space $X_\Delta \subseteq X$ and the discretized unperturbed process $P_\Delta(., .)$: (a) From a state $x \in X_\Delta$, if a set $A \subseteq X$ is reached with positive probability under the original process, then the closest states to A in X_Δ (under the original metric) are reached with positive probability under the discretized process $P_\Delta(., .)$; (b) If, from a state $x \in X_\Delta$, under the original process the set of states in X which are closer to $z \in X_\Delta$ than to any other point in X_Δ is never reached with positive probability, then z is never reached with positive probability under the discretized process; (c) Stable states of the original process are states of the discretized process and therefore stable states of the discretized process by (b).

We take as a discretization of the perturbation (the uniform distribution on X) any distribution on X_Δ that places positive probability on all states in X_Δ. Now, for any finite discretization satisfying our conditions, as $\varepsilon \to 0$, the limit of π_ε places positive probability on all states in $\{0, 1\}$: discretizing the process has given us one additional stochastically stable state.[4]

Finding the stochastically stable states of the original process in this section turns out to be simple. The reason for this is that far enough along any convergent path to a stable state, the probability under the perturbed process of moving to the basin of attraction of another given stable state is of constant order of ε. For example, from any convergent path to 0 under the unperturbed process $P(., .)$, at any given future period t the probability under the perturbed process of being in the basin of attraction of state is 1 is of order $\varepsilon^\infty = 0$. There do not exist convergent paths to 0 with escape probabilities of different orders of ε. We shall define Property C as the absence of multiple paths which converge to the same stable state and have different orders of escape probability. When Property C holds, we show that variants of results used heavily in the finite state space stochastic stability literature can be used. An important part of the current paper gives sufficient conditions under which Property C holds.

2.2. Multiple convergent paths

The next example can be considered as a model in which there are two possible focal points for a social norm. There are $n \geq 2$

[2] See also Bergin and Lipman (1996), van Damme and Weibull (2002), Beggs (2005).

[3] The use of the term 'stochastic stability' in the economics literature refers almost exclusively to states with positive weight under some limiting measure. Other uses of the term appear in the literature on dynamic processes. This paper follows the economic usage.

[4] It may be remarked that, for this example, there exist sequences of finite discretizations such that the limit (of the sequence of discretizations) of the limits (as $\varepsilon \to 0$) of π_ε converges to the stochastically stable states of the original process. Such a sequence does not always exist, as is apparent from the example in Section 2.2.

agents who contribute some real amount of effort towards the provision of a public good. If at least some threshold number of agents contribute at least some focal amount (which we take to be 1), then agents converge towards that level of contribution. Otherwise they converge towards a zero contribution. Consider a state space $X = [0, x_{\max}]^n \subset \mathbb{R}^n_+, x_{\max} > 1, n \in \mathbb{N}_+$. Let $x = (x_1, \ldots, x_n) \in X$ denote a representative element. Define $I(x)$ as the set of players who contribute at least 1 in a given state x:

$$I(x) = \{i \in \{1, \ldots, n\} : x_i \geq 1\}.$$

For some $k < n, k \in \mathbb{N}_+$, we define $P(., .)$ as follows:

$$\text{If } |I(x)| < k \text{ then } P\left(x, \frac{x}{2}\right) = 1$$

$$\text{If } |I(x)| \geq k \text{ then } P\left(x, \frac{x + 1^n}{2}\right) = 1.$$

The process has stable states $\Lambda = \{0^n, 1^n\}$. We examine a perturbed variant of the process in which each period, with probability $1 - \sum_{i=1}^n \varepsilon^i$ the unperturbed process is followed, and with probability ε^i the new state is drawn from the distribution:

$$G_i(x, .) \sim U[\{\bar{x} \in X : (|\{r : x_r = \bar{x}_r\}| = n - i)\}].$$

That is, with probability ε^i, exactly i agents randomly choose their contribution from a uniform distribution on $[0, x_{\max}]$. There exist convergent paths to 1^n with $|I(x)| \in \{k, k + 1, \ldots, n\}$. From a convergent path to 1^n for which $|I(x)| = k$, a move to the basin of attraction of 0^n in a single period is an event with probability of order ε. From a convergent path to 1^n for which $|I(x)| = n$, a move to the basin of attraction of 0^n in a single period is an event with probability of order ε^{n-k+1}. Property C does not hold. The only stochastically stable state of this process turns out to be 0^n. Showing that this is the case is complicated by Property C not holding, so recourse to more general methods is necessary (see Section 6). For completeness, we note that if $k < n - k + 1$, then any reasonable finite discretization of the process leads to 1^n being selected as the unique stochastically stable state.

3. A general model of perturbed adaptive behavior

A very general model is presented: the unperturbed dynamic can be any Markov process on any separable metric space, the only assumption being the nonemptiness of, and convergence of the process to, a set of stable states. The perturbed model then allows a broad class of perturbations which includes independent random errors such as are found in the traditional stochastic stability literature, but also allows correlated errors and any type of state dependent behavior.[5]

3.1. Quantitative characterization

The first step is to model an unperturbed dynamic which gives the behavior of agents in the absence of random perturbations. Let Φ be a Markov process on a separable metric space X with kernel $P(x, A), x \in X, A \in \mathcal{B}(X)$, where $\mathcal{B}(X)$ is the Borel σ-algebra.

Definition 1. The set of *stable* states is defined:

$$\Lambda := \{x \in X : P(x, x) = 1\}.$$

Assumption 1.

$$\Lambda \neq \emptyset, \quad \Lambda \text{ is closed.}$$

Definition 2. The *basin of attraction* W_i of $x_i^* \in \Lambda$ is:

$$W_i := \Big\{x \in X : \text{ for every open } V \supseteq \{x_i^*\}, \\ P^t(x, V) \to 1 \text{ as } t \to \infty\Big\}.$$

Define $W := \bigcup_i W_i$. Let $W_i^\delta := W_i \cap B_\delta(x_i^*)$ and $W^\delta := \bigcup_i W_i^\delta$, where $B_\delta(x)$ is the open ball of radius δ centered at x. We now introduce an assumption which guarantees that wherever you start in the state space you end up arbitrarily close to some element of Λ—the unperturbed process Φ converges to a stable state. This assumption is necessary to the purpose of this paper, which is to give tools by which to select from several stable states. If convergence were not assumed, then equilibrium selection would become a secondary issue.

Assumption 2 (*Convergence*).

$$\forall\, \delta > 0,\ x \in X,\ \exists T_{\delta x} \text{ such that } P^t(x, W^\delta) = 1 \quad \text{for all } t \geq T_{\delta x}.$$

Note, from any $x \in X$, it is assumed that convergence occurs in bounded time. Boundedness is unnecessary when state spaces are finite as analysis in such a case can proceed without a distinction being drawn between stable and unstable states.[6]

Φ can be taken to represent some unperturbed dynamic which describes the evolution of the strategies of players in a game, with each entry in a vector $x \in X$ describing the strategy chosen by a player in some previous time period. In this context the kernel $P(x, A)$ can be taken to be some sort of (not necessarily continuous) best response or imitation dynamic. Let $\{\Phi_\varepsilon\}_\varepsilon, \varepsilon \in (0, \frac{1}{M}), M \in \mathbb{N}_+$, be a family of Markov processes on the state space X with kernels $P_\varepsilon(x, A)$. Define:

$$P_\varepsilon(x, A) = \left(1 - \sum_{i=1}^{M} \varepsilon^i\right) P(x, A) \\ + \sum_{\{i \in \mathbb{N}_+ : i < M\}} \varepsilon^i G_i(x, A) + \varepsilon^M G_M(A)$$

where $G_i(x, .)$, $G_M(.)$ are probability measures on $\mathcal{B}(X)$. For $A \in \mathcal{B}(X)$, $G_i(., A)$ are non-negative $\mathcal{B}(X)$-measurable functions on X. As a sum of $\mathcal{B}(X)$-measurable functions, $P_\varepsilon(., A)$ is $\mathcal{B}(X)$-measurable. Note that $P_\varepsilon(x, X) = 1$ is satisfied. $\{\Phi_\varepsilon\}_\varepsilon$ is a subset of the class of Puiseux Markov processes. It is not necessary for the powers of ε to be integers, but here they are assumed so for expositional ease. Note that perturbations according to G_i occur with probabilities that approach zero at rate ε^i as ε is taken to zero.[7]

As all of our measures have $\mathcal{B}(X)$ as their domain and all our functions are $\mathcal{B}(X)$-measurable, we can use the property that an integral over a sum of measures is equal to the sum of integrals over those measures[8] to show (Lemma 1 in the Appendix) that any positive transition probability over a finite number of periods has an order of magnitude given by ε to some integer. That is, for any $x \in X, A \in \mathcal{B}(X), T \in \mathbb{N}_+$, if $P_\varepsilon^T(x, A) > 0$, then

$$l\varepsilon^r < P_\varepsilon^T(x, A) \leq (M + 1)^T \varepsilon^r$$

for some positive real l and non-negative integer r. This fact is used heavily in proving the results of the paper.

[5] For example, perturbations could include coalitional behavior such as that found in Newton (2012), depend on relative payoffs as in Blume (1993), or show some degree of intentionality as in Naidu et al. (2010).

[6] Inertia can still be modeled under processes satisfying Assumption 2, but it will always be a finite amount of inertia. For example a state $(x, n) \notin \Lambda$ could proceed with positive probability to a state $(x, n + 1) \notin \Lambda$ up to $(x, \bar{n})$ for some, possibly very large, $\bar{n} \in \mathbb{N}_+$.

[7] Setting $\varepsilon = e^{\frac{-1}{\eta}}$, the processes in this paper satisfy: $\lim_{\eta \to 0} \eta \log P_\eta(x, U) = -\inf_{y \in U} \rho(x, y)$ for open sets U and a function $\rho(., .) : X \times X \to \mathbb{R}$. $\rho(., .)$ is not necessarily continuous and so does not necessarily satisfy the conditions of Kifer (1990). For example, independent error models do not satisfy continuity of $\rho(., .)$.

[8] See Fremlin (2001), 234H(c).

In the most common models of stochastic stability, which we refer to as *independent error* models (Young, 1993a; Kandori et al., 1993), the state is composed of strategy profiles, each player has an independent probability of making an error with probability ε and players who make errors play strategies chosen from a given distribution with full support. Such a process satisfies the definition above.[9] We now show that for any given value of ε, Φ_ε has a unique invariant measure. This measure is predictive in the sense that it gives the frequencies with which given sets of states will be observed in the long run.

Proposition 1. *Φ_ε has a unique invariant probability measure π_ε.*

Proposition 2. *For all $x \in X$,*

$$\sup_{A \in \mathcal{B}(X)} |P_\varepsilon^t(x, A) - \pi_\varepsilon(A)| \to 0 \quad \text{as } t \to \infty.$$

The independence of $G_M(.)$ from x makes $G_M(.)$ an irreducibility measure. That is, from any $x \in X$, any $A \in \mathcal{B}(X)$ such that $G_M(A) > 0$ will eventually be reached by the process. The presence of the $G_M(.)$ term in $P_\varepsilon(.,.)$ is sufficient, but not necessary, for the existence of a unique invariant measure and ergodicity. In fact, for several examples later in the paper, the analysis is independent of $G_M(.)$.

3.2. Invariant measures when perturbations are rare

In the preceding subsection it was shown that Φ_ε has a unique invariant measure π_ε which depends on ε. In order to predict long term behavior under small perturbations it helps to analyze the limit as $\varepsilon \to 0$. As stochastic stability is primarily used as a tool of equilibrium selection, we desire that, as ε gets small, π_ε place arbitrarily small probability mass on sets of states which are not close to the stable states of the unperturbed process. Assumption 2 does not guarantee this (see counterexample in Appendix B.1). Hence we impose the following assumption, that restricts the perturbed dynamic so that, from some point in the state space, the process far enough in the future is likely to be close to a stable state.

Assumption 3 (*Perturbed Stability*)**.**

$$\forall \delta > 0,\ \exists x \in X,\ \varepsilon_\delta, r_\delta > 0 \text{ s.t. } \lim_{t\to\infty} P_\varepsilon^t(x, W^\delta) > 1 - \varepsilon^{r_\delta}$$
for all $\varepsilon < \varepsilon_\delta$.

Ergodicity implies that the condition of Assumption 3 holding for some x implies that it holds for all x. Note that if the $T_{\delta x}$ in Assumption 2 are independent of x, then Assumption 3 is automatically satisfied.[10,11]

Proposition 3. *For any $\eta \in (0, 1]$, $A \in \mathcal{B}(X)$ such that $\Lambda \cap cl(A) = \emptyset$ there exists $\hat{\varepsilon}$ such that $\pi_\varepsilon(A) < \eta$ for all $\varepsilon < \hat{\varepsilon}$.*

Corollary 1. *If π is a limiting measure (in the sense of weak convergence of measures) of π_ε as $\varepsilon \to 0$, then π is an invariant measure of the unperturbed process Φ. Specifically, $\pi(\Lambda) = 1$.*

So, the addition of perturbations to the model can be seen as a way of selecting between the alternative invariant measures of the unperturbed process. We now move to find conditions under which processes on a general state space can be analyzed using similar tools to those used in the finite state space literature. The following assumption is made:

Assumption 4.

$|\Lambda| < \infty$.

Cases for which $|\Lambda| = \infty$ can sometimes be analyzed using a careful application of the results of Freidlin and Wentzell (1984).[12] Note that $|\Lambda| < \infty$ implies that $\Lambda = \{x_1^*, \ldots, x_\nu^*\}$ for some $\nu \in \mathbb{N}_+$. Limiting invariant measures in many examples turn out to place all of the probability mass on a single stable state, predicting that in the long run the process should be observed to be at or near that state almost all of the time.

Definition 3. Stable states $x^* \in \Lambda$ with $\pi(x^*) > 0$ are called *stochastically stable*.

The rest of the paper devotes itself to the question of how to find stochastically stable states and the analysis of intricacies that can arise due to having an infinite state space.

3.3. Transition probabilities between stable states

In order to find stochastically stable states it will be necessary to determine the magnitudes of the transition probabilities between the basins of attraction of different stable states. These magnitudes are given as powers of ε. The following Bachmann–Landau asymptotic notation expresses the idea of f being bounded below by g.

$$f(\varepsilon) \in \Omega(g(\varepsilon)) \Leftrightarrow \exists k > 0,\ \bar{\varepsilon} \text{ s.t. } \quad \forall \varepsilon < \bar{\varepsilon},\ kg(\varepsilon) \le |f(\varepsilon)|.$$

Define:

$$V(x_k^*, x_j^*) := \inf\left\{i : \ \forall x \in W_k,\ \left(\exists t : \ P_\varepsilon^t(x, W_j) \in \Omega(\varepsilon^i)\right)\right\} \wedge \infty$$
$$V^-(x_k^*, x_j^*) := \inf\Big\{i : \ \forall \delta > 0,\ \exists x \in W_k^\delta : \left(\exists t : \ P_\varepsilon^t(x, W_j) \in \Omega(\varepsilon^i)\right)\Big\} \wedge \infty$$

which can be interpreted as resistances measuring the difficulty of moving from the basin of attraction of x_k^* to the basin of attraction of x_j^*. $V(x_k^*, x_j^*)$ and $V^-(x_k^*, x_j^*)$ measure respectively the most unlikely and easiest way in which a move from close to x_k^* to W_j could occur. The values of $V(x_k^*, x_j^*)$ and $V^-(x_k^*, x_j^*)$ will depend on which error distributions G_i are used along paths between neighborhoods of the two stable states. As paths occur with probabilities of order ε^r for positive integers r, it follows that V and V^-, when finite, must take integer values.

Before proceeding further, some notation is required. $F_\delta(t)$ is the set of states which under the unperturbed dynamic *do not* converge to W^δ within t periods. $\bar{V}_i$ is simply the maximum finite value from the $V^-(x_i^*, .)$ functions.

Definition 4.

$$F_\delta(t) := \{x \in X : \exists\, t' \ge t \text{ s.t. } P^{t'}(x, W^\delta) < 1\},$$
$$\bar{V}_i := \max_{j : V^-(x_i^*, x_j^*) < \infty} V^-(x_i^*, x_j^*).$$

The definition of $F_\delta(t)$ together with Assumption 2 implies that for any $x \in X$, for large enough t, $x \notin F_\delta(t)$. For very large t, $F_\delta(t)$ can be understood as the set of states with very large convergence times $T_{\delta x}$.

[9] Independent error models do not have $P_\varepsilon(x, A)$ continuous in x for any given $A \in \mathcal{B}(X)$ and so do not satisfy the assumptions of Feinberg (2006).

[10] As for $t \ge T_{\delta x}$, small enough ε, we then have $P_\varepsilon^t(x, W^\delta) \ge (1 - \sum_{i=1}^M \varepsilon^i)^{T_{\delta x}} P^{T_{\delta x}}(x, W^\delta) = (1 - \sum_{i=1}^M \varepsilon^i)^{T_{\delta x}} > (1 - M\varepsilon)^{T_{\delta x}} > 1 - 2T_{\delta x}M\varepsilon > 1 - \varepsilon^{\frac{1}{2}}$.

[11] Note that the proof of Proposition 3 is the only place in this paper where Assumption 3 is used.

[12] See Section 6, the Appendix, and Schenk-Hoppé (2000) for an introduction to the ideas involved.

The following assumption dictates that from starting points close to any stable state x_k^*, the process cannot be more likely to transit to states from which convergence times are arbitrarily large than it is to transit to the basin of attraction of any other stable state x_j^* for which $V^-(x_k^*, x_j^*) < \infty$. It further prohibits unboundedly large convergence times for states arbitrarily close to stable states. This and Assumption 2 are stronger than convergence assumptions made in the finite state space literature. However, these assumptions are not redundant: the extension of 'mistake counting' methods of determining stochastic stability to general state spaces (Proposition 4) does not hold without Assumptions 2 and 5. Counterexamples given in Appendix B demonstrate that even when state spaces are countable (but not finite), results can fail when these assumptions are dropped.

Assumption 5 (*Fast Convergence*)**.** There exists $\hat{\delta}$ such that for all $\delta < \hat{\delta}$,

$$\exists T_{\delta+} : \ \forall x_i^* \in \Lambda, \ x \in W_i^\delta, \ \nexists t : \ P_\varepsilon^t(x, F_\delta(T_{\delta+})) \in \Omega(\varepsilon^{\bar{V}_i - 1}),$$

and $W^{\hat{\delta}} \cap F_\delta(T_{\delta+}) = \emptyset$.

Note that if there is a uniform bound on convergence times, that is the $T_{\delta x}$ in Assumption 2 do not depend on x, then setting $T_{\delta+} = T_{\delta x}$, we have that $F_\delta(T_{\delta+}) = \emptyset$, so Assumption 5 is satisfied. This will always be the case when the state space is finite.

A regularity property is now defined that allows a single magnitude of transition probability to characterize transitions from states within the basin of attraction of, and close to, one stable state to the basin of attraction of another stable state. We shall further require that the appropriate orders of transition probabilities can occur within a bounded number of periods T,[13] and that there exists some uniform lower bound on transition probabilities for given ε.[14]

Definition 5 (*Property C*)**.** *Property C* is said to hold when

(C1) $V(x_k^*, x_j^*) = V^-(x_k^*, x_j^*)$ for all $x_k^*, x_j^* \in \Lambda$, and

(C2) For $V(x_k^*, x_j^*) \neq \infty$, there exists $\delta_{kj} > 0$ such that for all $\tilde{\delta} > 0$, there exist $T_{kj}(\tilde{\delta}) \in \mathbb{N}_+$, $l > 0$ such that $\forall x \in W_k^{\delta_{kj}}$, $P_\varepsilon^{T_{kj}(\tilde{\delta})}(x, W_j^{\tilde{\delta}}) > l\varepsilon^{V(x_k^*, x_j^*)}$.

The convergence of the process to stable states (Assumptions 2 and 3) ensures that, for small ε, the process spends almost all of the time close to stable states. Property C allows a single value to characterize the order of magnitude of transition probabilities between these small neighborhoods of stable states. A strengthening of convergence assumptions (Assumption 5) enables us to equate these probabilities to the probabilities which govern the related process which is only observed when the state is close to a stable state.[15] As there are a finite number of stable states, the problem of determining the stochastically stable states is reduced to a discrete problem: the sets W_k^δ for which $\pi_\varepsilon(W_k^\delta) \nrightarrow 0$ as $\varepsilon \to 0$ are determined solely by the values of $V(.,.)$. Analogues of results from the finite state space stochastic stability literature can then be used.

3.4. Characterization restated for infinite state spaces

Let $L = \{1, \ldots, \nu\}$ index the states in Λ. A graph on L is an *i-graph* if each $j \neq i$ has a single exiting directed edge, and the graph has no cycles. Let $\mathcal{G}(i)$ denote the set of all i-graphs. Letting $k \to j$ denote a directed edge from k to j, define:

$$\mathcal{V}(i) = \min_{g \in \mathcal{G}(i)} \sum_{(k \to j) \in g} V(x_k^*, x_j^*),$$

$$L_{\min} = \{i \in L : \mathcal{V}(i) = \min_{j \in L} \mathcal{V}(j)\}.$$

Note that it is possible that $\mathcal{V}(i) = \infty$ for some, but not all, $i \in L$. An analogue of the key result of Young (1993a) and Kandori et al. (1993) can now be stated.

Proposition 4. *If Assumptions 1–5 and Property* C *hold,* $\pi_\varepsilon \Rightarrow \pi$ *as* $\varepsilon \to 0$*, then*

$$\pi(x_i^*) > 0 \Leftrightarrow i \in L_{\min}.$$

The proof analyzes the process restricted to small neighborhoods of stable states. The invariant probability measure of the restricted process is a restriction and scaling of the invariant measure of the original process. For fixed ε, from any state, any of these neighborhoods which have positive invariant measure can be reached with positive probability in bounded time. This allows the construction of a finite state space Markov chain with invariant measure equal to the invariant measure of the restricted process. Property C gives that the order of the transition probabilities of this chain are precisely $\varepsilon^{V(.,.)}$, and the finite problem can be solved.[16]

4. Sufficient conditions for Property C

Given the usefulness of Property C in allowing the use of an analogous characterization to that in the finite state space literature, a natural question to ask is under what conditions it holds and whether or not these conditions are plausible and commonly satisfied. The next proposition concerns itself with finding sufficient conditions for Property C.

Proposition 5 gives conditions for Property C that can be satisfied by independent error models of stochastic stability. This is important as such models are commonly found in the literature. Firstly, a continuity requirement is placed on the unperturbed dynamic and the error distributions. For the unperturbed dynamic, this requirement is implied by the weak Feller property $P(y, .) \Rightarrow P(x, .)$ as $y \to x$, corresponding in a game theoretic context to continuity of the response correspondence of the underlying game. For the error distributions, the requirement is satisfied by independent error models. Secondly, asymptotic stability is imposed in the neighborhoods of some stable states. Thirdly, a condition is given which restricts the behavior of the process at a stable state according to the behavior of the process at nearby states. Following the statement of the proposition, a series of examples illustrates the role of each condition.

First, define an attainability property for each $x^* \in \Lambda$ which holds when there is a positive probability of ending up spending time in the basin of attraction of x^* when the initial state is distributed according to $G_M(.)$.

Definition 6 (*Attainable Stable States*)**.** For a given stable state x_i^*, let

$$B_i = \{x \in X : \exists t \ s.t. \ P_\varepsilon^t(x, W_i) > 0\}.$$

[13] This will eliminate the possibility of the existence of an infinite sequence $\{x(i)\}_{i\in\mathbb{N}}$ converging to x_k^* such that $\lim_{i\to\infty} \min\{t \in \mathbb{N}_+ : P_\varepsilon^t(x(i), W_j) \in \Omega(\varepsilon^{V(x_k^*, x_j^*)})\} \to \infty$. If such a sequence existed, there could exist a distribution on W_k such that on entering W_k according to this distribution the process would have infinite expected waiting time until an escape probability of order $\Omega(\varepsilon^{V(x_k^*, x_j^*)})$ is possible.

[14] This eliminates the possibility of the existence of an infinite sequence $\{x(i)\}_{i\in\mathbb{N}}$ converging to x_k^* such that $\lim_{i\to\infty} \sup\{l \in \mathbb{R}_{++} : P_\varepsilon^T(x(i), W_j) \geq l\varepsilon^{V(x_k^*, x_j^*)}\} \to 0$.

[15] This step is formalized in Lemma 6 in Appendix A and counterexamples showing the importance of Assumptions 2 and 5 to Lemma 6 are given in Appendix B.

[16] As a consequence of Proposition 4, "radius–coradius" results (Ellison, 2000, citing a no longer available paper of Evans, 1993), which follow immediately from the i-graph characterization of stochastically stable states, will also hold when appropriately restated.

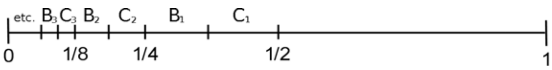

Fig. 1. Example: (i) does not hold for $P(.,.)$.

Note that B_i is independent of ε, as any transition which has positive probability for some $\varepsilon > 0$, has positive probability for all $\varepsilon > 0$. Define the set of attainable stable states $\mathcal{A} = \{x_i^* \in \Lambda : G_M(B_i) > 0\}$. Note that $x_i^* \in \Lambda \setminus \mathcal{A}$ implies $\pi_\varepsilon(W_i) = 0$.

Proposition 5. *Under Assumptions 1–5, if the following conditions hold then Property C holds.*

(i) *For all $A \in \mathcal{B}(X)$ open, $x_1 \in X$, there exist $\delta_{Ax_1} > 0$, $\xi_{Ax_1} > 0$, such that if $x_2 \in X$ satisfies $d(x_1, x_2) < \delta_{Ax_1}$ then:*

$$P(x_1, A) > 0 \Rightarrow P(x_2, A) > \xi_{Ax_1}$$
$$G_q(x_1, A) > 0 \Rightarrow G_q(x_2, A) > \xi_{Ax_1}, \quad q = 1, \ldots, M-1.$$

(ii) *For all $x_i^* \in \mathcal{A}$, there exists $\tilde{\delta} > 0$ such that $W_i^{\tilde{\delta}}$ is open.*

(iii) *For all $x_j^*, x_k^* \in \Lambda$, $\exists\, t$:*

$$P_\varepsilon^t(x_k^*, W_j) \in \Omega(\varepsilon^{V^-(x_k^*, x_j^*)}).$$

A sketch of the proof is as follows. Take a finite path from x_k^* to W_j. Condition (iii) ensures that no easier path to W_j exists from any points in W_k which are close enough to x_k^*. Condition (i), used iteratively (Lemma 8 in the Appendix), shows that from points close to x_k^*, similar paths can be followed with similar probabilities. Condition (ii) ensures that paths which are similar enough to the original path must enter W_j. The rest of this section consists of stylized examples designed to demonstrate how Property C can fail if any of the conditions of Proposition 5 do not hold.

4.1. Example: (i) does not hold for the unperturbed dynamic

Define $X = [0, 1]^3$. Let the unperturbed process Φ represent a game played repeatedly by 3 players with strategy spaces $[0, 1]$. Each period each player plays a best response to the actions of the other two players in the previous period. For $i \in \mathbb{N}_+$, define $C_i = \left(\frac{3}{4}\frac{1}{2^i}, \frac{1}{2^i}\right]$. Define $C = \bigcup_{i=1}^\infty C_i$. Let $B_i = \left(\frac{1}{2}\frac{1}{2^i}, \frac{3}{4}\frac{1}{2^i}\right]$, $B = \bigcup_{i=1}^\infty B_i$ (see Fig. 1). Let best response correspondences be symmetric and anonymous:

$$BR(a, b) = \begin{cases} C_{i+1}, & \text{if } a \in C_i, b \notin C \quad \text{for some } i \\ C_{i+1}, & \text{if } a \in C_i, b \in C_j \quad \text{for some } i, j, i \geq j \\ 1, & \text{if } a \in \left(\frac{1}{2}, 1\right], b \notin C \\ B_{i+1}, & \text{if } a \in B_i, b \in B_j \quad \text{for some } i, j, i \geq j \\ 0, & \text{if } a = 0, b \in B \cup \{0\}. \end{cases}$$

Note that $x_0^* = (0, 0, 0)$ and $x_1^* = (1, 1, 1)$ are the only stable states of the unperturbed dynamic. Let the perturbed process Φ_ε be such that each player independently with probability ε plays an action chosen uniformly at random from $[0, 1]$ instead of playing a best response. $V(x_0^*, x_1^*) = 3$. $V^-(x_0^*, x_1^*) = 1$. There is in effect one convergent path to x_0^* via states in $B \times B \times B$ from which it is easy to escape and another convergent path via states in $C \times C \times C$ from which it is difficult to escape. As $\varepsilon \to 0$, $\pi_\varepsilon \Rightarrow \pi$ where $\pi((0, 0, 0)) = 1$.

Intuitively, (i) not holding for $P(.,.)$ carries the implication that even when two states are extremely close to one another there is no guarantee that similar random shocks will lead to similar responses by the players. It is always possible to choose two states x_1, x_2 arbitrarily close to one another such that from x_1 a single player choosing a given random action would lead to completely different short run behavior of the process to that which would occur if from x_2 exactly the same player chose exactly the same random action.

4.2. Example: (i) does not hold for error distributions

Let $X = [0, 1]$. Let there be two stable states $x_k^* = 0$, $x_j^* = 1$, let $W_i \supset B_\delta(x_i^*)$ for some δ, $i = j, k$, so (ii) is satisfied. For $x \in W_k$, let $P(x, {}^x/_{10}) = 1$. The error process guarantees that (iii) is satisfied:

$G_1(x, x) = 1$, if the first non-zero digit in the decimal expansion of x is 1.

$G_1(x, W_j) = 1$, otherwise

$$G_2(W_j) = G_2(W_k) = \frac{1}{2}.$$

Then $V^-(x_k^*, x_j^*) = 1$ and $V(x_k^*, x_j^*) = 2$ no matter whether or not (i) is satisfied by $P(.,.)$.

4.3. Example: (ii) does not hold

Define $X = [0, 1]^2$. Let (x, y) describe an element of the state space (Fig. 2). Let:

$$P\left((x, y), \left(\frac{x}{2}, y + \min\left\{\frac{1}{2}, y\right\}(1 - y)\right)\right) = 1$$

then there are 2 stable states, $x_k^* = (0, 1)$ and $x_j^* = (0, 0)$. Let:

$$A = \{(x, y) \in X : y = 0\}$$
$$G_1(x, A) = \frac{1}{2} \quad \text{for } x \in U := \{(x, y) \in X : x + 2y \geq 2\},$$

and let $G_1(x,.)$ be uniform on X otherwise. Let $G_2(A) = 1$. Then (i) and (iii) hold but $V^-(x_k^*, x_j^*) = 1$ and $V(x_k^*, x_j^*) = 2$. Intuitively, although (i) implies that from any point close to x_k^* but not in U it is possible to reach any point arbitrarily close to x_j^* with a probability of order ε, (ii) not holding means that this is not sufficient for convergence to x_j^* and the process ends up reconverging towards x_k^*.

4.4. Example: (iii) does not hold

Define $X = [0, 1] \times [-1, 1]$. Let (x, y) describe an element of the state space. Let:

$$P\left((x, y), \left(\frac{x}{2}, \frac{y}{2}\right)\right) = \frac{\max\left\{0, \frac{3}{5} - x, \frac{1}{2} - y\right\}}{\max\left\{0, \frac{3}{5} - x, \frac{1}{2} - y\right\} + \max\left\{0, \min\left\{x - \frac{2}{5}, y\right\}\right\}}$$

$$P\left((x, y), \left(\frac{x+1}{2}, \frac{y+1}{2}\right)\right) = \frac{\max\left\{0, \min\left\{x - \frac{2}{5}, y\right\}\right\}}{\max\left\{0, \frac{3}{5} - x, \frac{1}{2} - y\right\} + \max\left\{0, \min\left\{x - \frac{2}{5}, y\right\}\right\}}$$

then there are 2 stable states, $x_k^* = (0, 0)$ and $x_j^* = (1, 1)$. Note that under the unperturbed dynamic the process will in each period move in a straight line towards one of the stable states. Let:

$$G_1((x, y),.) \sim U[\{(x', y') : x' = x \text{ or } y' = y\}]$$
$$G_2(.) \sim U[X].$$

Then (i) and (ii) hold but $V^-(x_k^*, x_j^*) = 1$ and $V(x_k^*, x_j^*) = 2$. Although there exists a convergent path to x_k^* from which W_j can be reached with a probability of order ε (such as those in Area A in Fig. 3), the limit x_k^* does not have this property and so we cannot rule out the existence of convergent paths with lower escape probabilities (such as those in Area B in the figure).

5. Economic examples

5.1. Linear quadratic games

We apply the theory to two player games with strategies $y_i \in \mathbb{R}_+$ and the following payoff functions:

$$u_i(y_i, y_j) = a_i y_i^2 + b_i y_i y_j + c_i y_i + d_i y_j,$$
$$a_i, b_i < 0;\ c_i > 0;\ a_i, b_i, c_i, d_i \in \mathbb{R}.$$

This class of games includes public goods problems with strategic substitutes and Cournot duopolies with linear demand and quadratic costs. Oechssler and Riedel (2001) showed that under the replicator dynamic with symmetric payoff functions such games converge to the interior equilibrium in which players play $y_i^{int} = \frac{b_i c_j - 2c_i a_j}{4a_i a_j - b_i b_j}$.[17] We assume a multiplicity of equilibria: $y_i^{int} > 0$ so that $x_I = (y_1^{int}, y_2^{int})$ is a Nash equilibrium; $4a_i a_j < b_i b_j$ so that corner equilibria exist: $x_{c1} = (y_1^{cnr}, 0), x_{c2} = (0, y_2^{cnr}), y_i^{cnr} = \frac{-c_i}{2a_i}$. Let the unperturbed dynamic be a Markov process on $X = \mathbb{R}_+^2$ in which each period one player best responds to the current action of the other player, following which the other player best responds to the new action of the first player. Let the metric on X be Euclidean distance. The best response of player i to y_j is

$$BR_i(y_j) = \max\left\{\frac{-b_i y_j - c_i}{2a_i}, 0\right\}$$

Φ has kernel:

$$P\left((y_i, y_j), \left(BR_i(y_j), BR_j(BR_i(y_j))\right)\right) = \frac{1}{2}, \quad i = 1, 2.$$

This gives:

$$W_I = \{x_I\}, \quad W_{ci} = \{(y_i, y_j) : y_i > y_i^{int}, y_j < y_j^{int}\}.$$

Note that $P(.,.)$ satisfies condition (i) of Proposition 5. We analyze two possible perturbed dynamics.

5.1.1. Uniform local perturbations

For some small $\varsigma > 0$, define:

$$G_1((y_1, y_2), .) \sim U[B_\varsigma((y_1, y_2))]$$
$$G_n((y_1, y_2), .) = \int_X G_1((y_1, y_2), dx) G_{n-1}(x, .), \quad n = 2, \ldots, M-1$$
$$G_M((0, 0)) = 1.$$

These $G_n(.,.)$, $G_M(.)$ satisfy condition (i) of Proposition 5. $G_M(.)$ is not necessary for the results of this section, although it is of interest to note that any Markovian dynamic with some small, bounded below, probability of an 'Armageddon' event will satisfy the conditions for ergodicity. $\mathcal{A} = \{x_{c1}, x_{c2}\}$ and it is clear that for some $\delta > 0$, W_{ci}^δ are open so condition (ii) is also satisfied. For almost all values of ς, for any escape path from a state close to x_{ci} to W_{cj}^δ there is a similar path from x_{ci} itself to W_{cj}^δ, so condition (iii) is generically satisfied.[18,19] Now, for large enough M, the order of perturbations required to move from near x_{ci} to W_{cj} is given by:

$$V(x_{ci}, x_{cj}) = \left\lceil \min\left\{\frac{y_j^{int}}{\varsigma}, \frac{y_i^{cnr} - y_i^{int}}{\varsigma}\right\}\right\rceil$$

and applying Proposition 4 we obtain:

Proposition 6.

$$\min\left\{y_j^{int}, y_i^{cnr} - y_i^{int}\right\} \geq \min\left\{y_i^{int}, y_j^{cnr} - y_j^{int}\right\}$$
$$\iff \exists\, \hat{\varsigma} \text{ such that } \forall\, \varsigma < \hat{\varsigma},\ \exists\, \hat{M} \text{ such that } \forall\, M > \hat{M},\ \pi(x_{ci}) > 0.$$

5.1.2. Proportional perturbations

For some $k \in (0, 1)$, define:

$$G_1((y_1, y_2), .) \sim U[[ky_1, y_1] \times [ky_2, y_2]]$$
$$G_n((y_1, y_2), .) = \int_X G_1((y_1, y_2), dx) G_{n-1}(x, .), \quad n = 2, \ldots, M-1$$
$$G_M((0, 0)) = 1.$$

Similarly to above, conditions (i), (iii), (iv) are satisfied. Now, for large M,[20]

$$V(x_{ci}, x_{cj}) = \left\lceil \frac{\log\left(\frac{y_i^{int}}{y_i^{cnr}}\right)}{\log k} \right\rceil$$

and applying Proposition 4 we obtain:

Proposition 7.

$$\frac{y_i^{cnr}}{y_i^{int}} \geq \frac{y_j^{cnr}}{y_j^{int}}$$
$$\iff \exists \hat{k} :\ \forall\, k > \hat{k},\ \exists\, \hat{M} \text{ such that } \forall\, M > \hat{M},\ \pi(x_{ci}) > 0.$$

5.2. Sampling a population

Take a two player symmetric matrix game Γ in which a player has a set N of possible actions, $|N| = n$. Let there be a continuum of agents on the unit interval. The state space is defined as the proportions in which each action is played at a point in time: X is the unit $(n-1)$-simplex. In period t, independently of his previous actions, with probability $1-\alpha$ any given agent plays the same action as at time $t-1$. With probability α he randomly and uniformly samples a finite number k of the actions of players in period $t-1$ before playing a best response to the mixed strategy σ, which has each action being played with a probability equal to its proportion in his sample. If multiple best responses exist we assume that they are chosen with equal probability. Denote the distribution of such best responses to an action profile x by $BR(x)$. Then:

$$P(x, \tilde{x}) = 1, \quad \tilde{x} := (1-\alpha)x + \alpha BR(x).$$

Note that as the probability of drawing any given sample is continuous in x, $P(x, .)$ is itself continuous and does not violate condition (i) of Proposition 5. We restrict attention to games for which this process satisfies Assumption 2. It is trivial to construct games which do not satisfy Assumption 2 under this process, for example the 2 by 2 matrix of zeros.

Any stable state x^* is close to a Nash equilibrium of Γ in the following sense:

17 $y_i^{int} = \frac{-c}{2a+b}$ under symmetry.

18 The exception to this is when an expression for $V(x_{ci}, x_{cj})$ would be an integer even before the ceiling function is applied.

19 Thus Assumption 5 is satisfied as for any δ s.t. $W^\delta \cap B_\delta(x_I) = \{x_I\}$, $\gamma > 0$, for large enough t, $F_\delta(t) \subseteq \{(y_1, y_2) : \min_{i\in\{1,2\}} d(y_i, y_i^{int}) < \gamma\} \setminus \{x_I\}$. So for small enough δ and γ, correspondingly large $\tilde{t}$, $x \in W_i^\delta$, we have that $P_\varepsilon^t(x, F_\delta(\tilde{t}))$ is of the order of $\varepsilon^{V(x_i^*, x_j^*)}$ for all t.

20 $V(x_{ci}, x_{cj})$ is the lowest integer V such that $k^V y_i^{cnr} < y_i^{int}$.

Proposition 8. *For any $\xi > 0$, there exists $\bar{k} \in \mathbb{N}_+$ such that if $k > \bar{k}$, for any $x^* \in \Lambda$ there exists a symmetric Nash equilibrium x^{NE} of Γ such that $|x^* - x^{NE}| < \xi$.*

The perturbations are defined as follows. For some small $\varsigma > 0$, define:

$$G_1(x, .) \sim U[B_\varsigma(x)]$$

$$G_n(x, .) = \int_X G_1(x, dy)G_{n-1}(y, .), \quad n = 2, \dots, M-1$$

$$G_M(.) \sim U[X].$$

These $G_n(., .)$, $G_M(.)$ satisfy condition (i) of Proposition 5. $G_M(.)$ is not necessary for the results of this section.

Proposition 9. *If x^{NE} is a strict symmetric pure Nash equilibrium, then there exists $\bar{k} \in \mathbb{N}_+$ such that if $k > \bar{k}$, then $x^{NE} \in \mathcal{A}$ and x^{NE} is asymptotically stable.*

In a game such as that of Fig. 4 in which there are two strict symmetric Nash equilibria, by Proposition 9, for large enough k both of these Nash equilibria correspond to stable states in $\mathcal{A}$ such that the basin of attraction includes some open ball centered on the Nash equilibrium. Close to the mixed Nash equilibrium, for large enough k, there is one stable state $x_{int} = (p, 1-p)$ which is not in $\mathcal{A}$. Hence condition (ii) of Proposition 5 is satisfied. Then, as long as neither p nor $1-p$ is an integer multiple of ς, condition (iii) is also satisfied.[21,22] x_{int} converges to the mixed Nash equilibrium as $k \to \infty$. Identical arguments to those in Young (1998) then give:

Proposition 10. *In a 2 by 2 game with two symmetric strict pure Nash equilibria, one of which is risk dominant, there exists $\bar{k} \in \mathbb{N}_+$ such that if $k > \bar{k}$, then the risk dominant equilibrium is uniquely stochastically stable for small enough ς for large enough M.*

6. What if Property C does not hold?

Analysis is less straightforward if Property C does not hold. In some cases, however, a simple argument can still be used to find stochastically stable states. In the example in Section 2.2, Conditions (i), (ii) and (iii) of Proposition 5 all fail at state 1^n. Property C does not hold as $V^-(1^n, 0^n) = 1$ whereas $V(1^n, 0^n) = n - k + 1$. However, the state space can still be partitioned into a finite collection of disjoint sets such that, under $P_\varepsilon(., .)$, any of the sets can be reached from any $x \in X$ with a probability bounded below by a strictly positive number. Such a partition is:

$$X = X_0 \cup X_k \cup X_{k+1} \dots \cup X_n$$

where for $i = k, \dots, n$:

$$X_i = \{x \in X : |I(x)| = i\}$$

and

$$X_0 = \{x \in X : |I(x)| < k\}.$$

Bounds can be found for the transition probabilities between these sets. Freidlin and Wentzell (1984) style tree arguments can then be used. The 'flow' of probability mass from X_i to X_{i-1} for $i = k+1, \dots, n$, and from X_k to X_0 can be shown to be of order ε, so a tree of order ε^{n-k+1} rooted at X_0 can be constructed. Any tree not rooted at X_0 must include an edge leaving X_0 of order ε^k. Taking limits of the invariant measures then gives the conclusion that all probability mass in the limit is concentrated in X_0 and therefore, by Proposition 3, on 0^n.

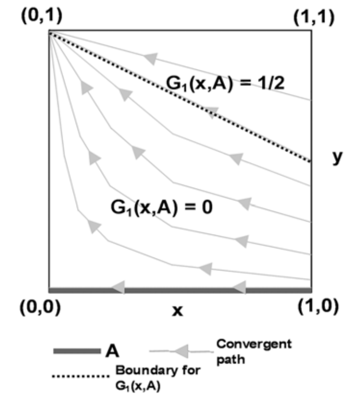

Fig. 2. Example: (ii) does not hold.

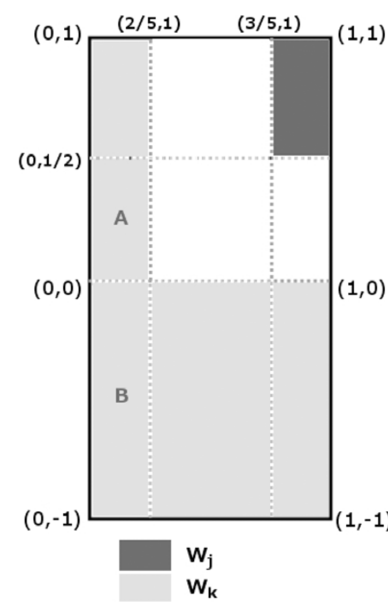

Fig. 3. Example: (iii) does not hold. Although convergent paths to $(0, 0)$ within A only require an order of ε transition to reach W_j, convergent paths within B require an order of ε^2 transition. So W_j cannot always be reached with $\Omega(\varepsilon)$ probabilities.

	L	R
L	a, a	b, c
R	c, b	d, d

Fig. 4. A two player strategic game. (L, L) and (R, R) assumed to be strict Nash equilibria.

Proposition 11. *In the example of Section 2.2 the unique stochastically stable state is 0^n.*

Given the simple nature of the counterexamples in Section 4, a similar analysis to the above can be carried out. In Example 4.1, W_0 can be partitioned into states that converge within $C \times C \times C$, those that converge within $B \times B \times B$, and those that attain x_0^*. In Example 4.2, W_k can be partitioned into states whose first digit is 1 and states whose first digit is not 1. In Example 4.3, W_k can be partitioned into the two areas on either side of the dotted line in Fig. 2. In Example 4.4, W_k can be partitioned into states with $y > 0$ and those with $y \leq 0$.

7. Concluding remarks

This paper has demonstrated how commonly used stochastic stability methods can be applied to settings with infinite state

[21] If $p = 2\varsigma$, from $(0, 1)$ it takes an order ε^3 event to move to the basin of attraction of $(1, 0)$, but from $(\xi, 1-\xi)$, $\xi \in (0, \varsigma)$, an order ε^2 event is all that is required.

[22] Assumption 5 can then be seen to be satisfied by a similar argument to that in Section 5.1 (see Footnote 19).

spaces, corresponding to situations in which economic agents choose from infinite strategy sets. It includes sufficient conditions for the straightforward application of existing results in the literature to such settings. Moreover, the analysis of the complications that can arise with general state spaces aids understanding of the problems of which one should be aware when applying ideas of robustness to random perturbations to processes which do not satisfy all of our conditions.

Another approach when seeking to find stochastically stable states for games with infinite strategy sets is to discretize the strategy space and the transition kernel. This is not always simple and can lead to problems such as the absence of simple closed form best response functions and nonexistence of equilibrium. Difficulties can be met when passing to the limit of any discretization as it becomes fine. Moreover, examples in this paper show that any discretization satisfying plausible criteria can lead to the selection of different equilibria to those selected when the analysis is carried out directly on the original state space and process. Sometimes it may be better to analyze stochastic stability whilst remaining in a non-finite world. This paper gives tools with which to aid that endeavor.

Appendix A. Proofs

Definition 7. A measure φ on $\mathcal{B}(X)$ is an *irreducibility measure* and Φ_ε is *φ-irreducible* if for all $x \in X$, whenever $\varphi(A) > 0$, there exists some $t > 0$, possibly depending on both A and x, such that $P_\varepsilon^t(x, A) > 0$.

Definition 8. A set $A \in \mathcal{B}(X)$ is *petite* if there exist a nontrivial[23] measure ν and a probability distribution a on $\mathbb{Z}_+$ such that $\forall x \in A$,

$$\sum_t a(t)P_\varepsilon^t(x, .) \geq \nu(.)$$

$A \in \mathcal{B}(X)$ is *$\nu_{\tilde{t}}$-small* if there exist such ν, a, with $a(\tilde{t}) = 1$ for some $\tilde{t} \in \mathbb{Z}_+$.

Definition 9. The process is said to be *strongly aperiodic* if there exists a ν_1-small set $A \in \mathcal{B}(X)$ with $\nu_1(A) > 0$.

Definition 10. For $A \in \mathcal{B}(X)$, $\tau_A := \min\{t \geq 1 : \Phi_\varepsilon^t \in A\} \wedge \infty$.

Proof of Proposition 1. $G_M(.)$ is an irreducibility measure on $\mathcal{B}(X)$ as for any $A \in \mathcal{B}(X)$ with $G_M(A) > 0$ we have $P_\varepsilon(x, A) \geq \varepsilon^M G_M(A) > 0$ for all $x \in X$. Letting $a(1) = 1$, $\nu(.) = \varepsilon^M G_M(.)$ we see that the set X is petite as for all $x \in X$, $P_\varepsilon(x, .) \geq \varepsilon^M G_M(.) = \nu(.)$. As X is the entire state space, $\tau_X \equiv 1$. Combined with irreducibility and petiteness of X this implies the existence of a unique invariant probability measure π_ε for Φ_ε.[24] □

Proof of Proposition 2. Take some $A \in \mathcal{B}(X)$ with $G_M(A) > 0$. For all $x \in A$, $P_\varepsilon(x, A) \geq \varepsilon^M G_M(A) > 0$. Letting $\nu_1(.) = \varepsilon^M G_M(.)$, we see that A is ν_1-small. Therefore the process is strongly aperiodic. This and the uniqueness and finiteness of π_ε imply the result.[25] □

Proof of Proposition 3. As $\Lambda \cap cl(A) = \emptyset$, there exists δ such that $W^\delta \cap cl(A) = \emptyset$. By Assumption 3 there exist x, ε_δ, r_δ such that

$$\lim_{t\to\infty} P_\varepsilon^t(x, W^\delta) > 1 - \varepsilon^{r_\delta} \quad \text{for all } \varepsilon < \varepsilon_\delta.$$

By Proposition 2, $|P_\varepsilon^t(x, W^\delta) - \pi_\varepsilon(W^\delta)| \to 0$ as $t \to \infty$, so $\pi_\varepsilon(W^\delta) > 1 - \varepsilon^{r_\delta}$. So, for small enough ε, $\pi_\varepsilon(W^\delta) > 1 - \eta$, and we have

$$\pi_\varepsilon(A) \leq \pi_\varepsilon(cl(A)) \leq 1 - \pi_\varepsilon(W^\delta) < 1 - (1 - \eta) = \eta. \quad \square$$

Proof of Corollary 1. For $n \in \mathbb{N}$, let

$$S_n = \{x \in X : d(x, \Lambda) > 2^{-n}\},$$
$$\bar{S}_n = \{x \in X : d(x, \Lambda) \geq 2^{-n}\}.$$

For all $n \in \mathbb{N}$, $\bar{S}_n$ is closed, so $\bar{S}_n = cl(\bar{S}_n)$. As $\bar{S}_n \cap \Lambda = \emptyset$, Proposition 3 implies $\pi_\varepsilon(\bar{S}_n) \to 0$ as $\varepsilon \to 0$. So, as $S_n \subseteq \bar{S}_n$, $\pi_\varepsilon(S_n) \to 0$ as $\varepsilon \to 0$. S_n is open, so by the definition of weak convergence, it must be that $\pi(S_n) = 0$. As Λ is closed, we have that

$$X \setminus \Lambda = \bigcup_{n=1}^{\infty} S_n, \quad \text{and therefore}$$

$$\pi(X \setminus \Lambda) = \pi\left(\bigcup_{n=1}^{\infty} S_n\right) \leq \sum_{n=1}^{\infty} \pi(S_n) = 0.$$

So the only states in X which can have positive probability under π are in Λ. □

To aid conciseness, denote $G_0(x, .) := P(x, .)$ and $G_M(x, .) := G_M(.)$. Denote, for $q_t \in \{0, 1, \ldots, M\}$, $t = 1, \ldots, T$,

$$G_{(q_1,\ldots,q_T)}(x, .) := \int_X G_{q_1}(x, dy_1) \int_X G_{q_2}(y_1, dy_2) \ldots \times \int_X G_{q_{T-1}}(y_{T-2}, dy_{T-1}) G_{q_T}(y_{T-1}, .). \tag{A.1}$$

Observe that for small ε, $1/2 < (1 - \varepsilon - \cdots - \varepsilon^M) < 1$. For the rest of this section, we assume that this holds. We have

$$P_\varepsilon(x, .) \leq \sum_{q=0}^{M} \varepsilon^q G_q(x, .) \leq 2P_\varepsilon(x, .)$$

with the second inequality strict for any $P_\varepsilon(x, A) > 0$. Similarly,

$$\begin{aligned} P_\varepsilon^T(x, .) &= \int_X P_\varepsilon(x, dy_1) \int_X P_\varepsilon(y_1, dy_2) \ldots \\ &\quad \times \int_X P_\varepsilon(y_{T-2}, dy_{T-1}) P_\varepsilon(y_{T-1}, .) \\ &\leq \sum_{q_1,\ldots,q_T} \int_X \varepsilon^{q_1} G_{q_1}(x, dy_1) \int_X \varepsilon^{q_2} G_{q_2}(y_1, dy_2) \ldots \\ &\quad \times \int_X \varepsilon^{q_{T-1}} G_{q_{T-1}}(y_{T-2}, dy_{T-1}) \varepsilon^{q_T} G_{q_T}(y_{T-1}, .) \\ &= \sum_{q_1,\ldots,q_T} \varepsilon^{q_1+\cdots+q_T} G_{(q_1,\ldots,q_T)}(x, .) \leq 2^T P_\varepsilon^T(x, .) \end{aligned} \tag{A.2}$$

with the second inequality strict for any $P_\varepsilon^T(x, A) > 0$.

Lemma 1. *For $x \in X$, $A \in \mathcal{B}(X)$, $T \in \mathbb{N}_+$, $P_\varepsilon^T(x, A) > 0$, let $r = \min\{q_1 + \cdots + q_T \mid G_{(q_1,\ldots,q_T)}(x, A) > 0\}$. Then, for some $l > 0$ independent of ε,*

$$l\varepsilon^r < P_\varepsilon^T(x, A) \leq (M+1)^T \varepsilon^r.$$

Proof. Considering a term in (A.2) such that $\tilde{q}_1 + \cdots + \tilde{q}_T \leq r$, $G_{(\tilde{q}_1,\ldots,\tilde{q}_T)}(x, A) > 0$, we have

$$P_\varepsilon^T(x, A) > \frac{1}{2^T} G_{(\tilde{q}_1,\ldots,\tilde{q}_T)}(x, A)\, \varepsilon^{\tilde{q}_1+\cdots+\tilde{q}_T} \geq l\,\varepsilon^r,$$
$$\text{for } l = \frac{1}{2^T} G_{(\tilde{q}_1,\ldots,\tilde{q}_T)}(x, A).$$

If $q_1 + \cdots + q_T \geq r$ for all strictly positive terms in (A.2), then as there are at most $(M+1)^T$ such terms,

$$P_\varepsilon^T(x, A) \leq \sum_{q_1,\ldots,q_T} \varepsilon^{q_1+\cdots+q_T} \leq (M+1)^T \varepsilon^r. \quad \square$$

[23] A measure which is not everywhere zero.

[24] Meyn and Tweedie (1994, Theorem 3.2).

[25] Meyn and Tweedie (2009, Theorem 13.0.1).

Lemma 2. (i) $x_j^* \in \mathcal{A} \Longleftrightarrow \forall x_k^* \in \Lambda,\ V(x_k^*, x_j^*) < \infty$.
(ii) $x_j^* \notin \mathcal{A} \Longrightarrow \forall x_k^* \in \mathcal{A},\ V(x_k^*, x_j^*) = \infty$.
(iii) $x_j^* \in \mathcal{A} \Longrightarrow \forall \delta,\ x \in X,\ Pr_x(\tau_{W_j^\delta} < \infty) = 1$.

Proof. Let B_j be as in Definition 6. Let $x_j^* \in \mathcal{A}$. Then $G_M(B_j) > 0$. Fix ε. Let

$B_j^{T\delta n} := \left\{x \in B_j :\ \forall t \geq T,\ P_\varepsilon^t(x, W_j^\delta) > 1/2^t n\right\}$.

Fix $\delta > 0$. By Definition 6, $x \in B_j$ implies that $P_\varepsilon^{t_1}(x, W_j) > 0$ for some t_1. This implies that for large enough t_2, n,

$$\begin{aligned}
P_\varepsilon^{t_2}(x, W_j^\delta) &= \int_X P_\varepsilon^{t_1}(x, dy) P_\varepsilon^{t_2-t_1}(y, W_j^\delta) \\
&\geq \int_{W_j} P_\varepsilon^{t_1}(x, dy) \underbrace{P_\varepsilon^{t_2-t_1}(y, W_j^\delta)}_{\substack{\geq (1/2)^{t_2-t_1} P^{t_2-t_1}(y, W_j^\delta) \\ \text{by (A.2)}}} \\
&\geq \left(\frac{1}{2}\right)^{t_2-t_1} \int_{W_j} P_\varepsilon^{t_1}(x, dy) P^{t_2-t_1}(y, W_j^\delta) \\
\underbrace{\geq}_{\substack{\text{by bounded} \\ \text{convergence}}} &\left(\frac{1}{2}\right)^{t_2} \int_{W_j} P_\varepsilon^{t_1}(x, dy) \\
&= \left(\frac{1}{2}\right)^{t_2} P_\varepsilon^{t_1}(x, W_j) > \frac{1}{2^{t_2}} \frac{1}{n},
\end{aligned}$$

where the third inequality comes from bounded convergence as Assumption 2 implies that for $y \in W_j$, $P^{t_2-t_1}(., W_j^\delta) \to 1$ pointwise as $t_2 \to \infty$. So $x \in B_j^{t_2 \delta n}$.

The above implies $B_j = \cup_{T \in \mathbb{N}_+} \cup_{n \in \mathbb{N}_+} B_j^{T\delta n}$. This union is countable, so $G_M(B_j) > 0$ implies that $G_M(B_j^{T\delta n}) > 0$ for some T, n. So, for all $x \in X$, $P_\varepsilon(x, B_j^{T\delta n}) \geq \varepsilon^M G_M(B_j^{T\delta n}) > 0$. For $x \in B_j^{T\delta n}$, $P_\varepsilon^T(x, W_j^\delta) > 1/2^T n$. Combining, for all $x \in X$,

$$\begin{aligned}
P_\varepsilon^{T+1}(x, W_j^\delta) &= \int_X P_\varepsilon(x, dy) P_\varepsilon^T(y, W_j^\delta) \\
&\geq \int_{B_j^{T\delta n}} P_\varepsilon(x, dy) \underbrace{P_\varepsilon^T(y, W_j^\delta)}_{> 1/2^T n} \\
&\geq \frac{1}{2^T n} \int_{B_j^{T\delta n}} P_\varepsilon(x, dy) \\
&= \frac{1}{2^T n} \underbrace{P_\varepsilon(x, B_j^{T\delta n})}_{\geq \varepsilon^M G_M(B_j^{T\delta n})} \geq \frac{1}{2^T n} \varepsilon^M G_M(B_j^{T\delta n}) > 0.
\end{aligned}$$

Hence it follows, by the definition of $V(.,.)$, that for all $x_k^* \in \Lambda$, $V(x_k^*, x_j^*) \leq M(T+1) < \infty$. The uniform lower bound on $P_\varepsilon^{T+1}(., W_j^\delta)$ implies that $\tau_{W_j^\delta} < \infty$ with probability 1.

Let x_i^* be such that $V(x_k^*, x_i^*) < \infty$ for all $x_k^* \in \Lambda$. Let $\delta = \delta_{ji}$ from the definition of Property C. For $x \in B_j^{T\delta_{ji} n}$, $P_\varepsilon^T(x, W_j^{\delta_{ji}}) > 1/2^T n$. By Property (C2) and $V(x_j^*, x_i^*) < \infty$, there exist $t, l > 0$, such that for all $x \in W_j^{\delta_{ji}}$, $P_\varepsilon^t(x, W_i) > l\varepsilon^{V(x_j^*, x_i^*)}$. Hence for all $x \in B_j^{T\delta_{ji} n}$,

$$\begin{aligned}
P_\varepsilon^{T+t}(x, W_i) &= \int_X P_\varepsilon^T(x, dy) P_\varepsilon^t(y, W_i) \\
&\geq \int_{W_j^{\delta_{ji}}} P_\varepsilon^T(x, dy) \underbrace{P_\varepsilon^t(y, W_i)}_{> l\varepsilon^{V(x_j^*, x_i^*)}} \\
&\geq l\varepsilon^{V(x_j^*, x_i^*)} \int_{W_j^{\delta_{ji}}} P_\varepsilon^T(x, dy) \\
&= l\varepsilon^{V(x_j^*, x_i^*)} \underbrace{P_\varepsilon^T(x, W_j^{\delta_{ji}})}_{> 1/2^T n} > \frac{1}{2^T n} l\varepsilon^{V(x_j^*, x_i^*)} > 0.
\end{aligned}$$

So $x \in B_j^{T\delta_{ji} n}$ implies $x \in B_i$. That is, $B_j^{T\delta_{ji} n} \subseteq B_i$ for all T, n. From above we know that $G_M(B_j^{T\delta_{ji} n}) > 0$ for some T, n, so it must be that $G_M(B_i) > 0$ and $x_i^* \in \mathcal{A}$. By contraposition, if $x_i^* \notin \mathcal{A}$, it must be that $V(x_j^*, x_i^*) = \infty$. $\square$

Definition 11. Fix $\bar\delta$, $T_{\bar\delta+}$, so that for all $x_k^*, x_j^* \in \mathcal{A}$, $x \in W_k^{\bar\delta}$,

(i) $\nexists t :\ P_\varepsilon^t(x, W_j) \in \Omega(\varepsilon^{V^-(x_k^*, x_j^*)-1})$, (ii) $\bar\delta < \delta_{kj}$,
(iii) $\bar\delta$, $T_{\bar\delta+}$ satisfy Assumption 5.

To see that Definition 11 makes sense, note that (i) can be satisfied due to the definition of V^-. For (ii), δ_{kj} is as in the definition of Property C. Further define:

$\underline{T} := \max_{k,j:\, x_k^*, x_j^* \in \mathcal{A}} T_{kj}(\bar\delta) + T_{\bar\delta+}$.

Lemma 3. *For all $t \geq \underline{T}$, $x_k^*, x_j^* \in \mathcal{A}$, there exist $l > 0$, $\bar\varepsilon$ such that for all $x \in W_k^{\bar\delta}$, $\varepsilon < \bar\varepsilon$,*

$$l\varepsilon^{V(x_k^*, x_j^*)} < P_\varepsilon^t(x, W_j^{\bar\delta}) \leq (M+1)^t \varepsilon^{V(x_k^*, x_j^*)}. \tag{A.3}$$

Proof. By Definition 11(ii), $\bar\delta < \delta_{kj}$, so Property (C2) implies that there exists $\hat l > 0$ such that for all $x \in W_k^{\bar\delta}$,

$$P_\varepsilon^{T_{kj}(\bar\delta)}(x, W_j^{\bar\delta}) > \hat l\, \varepsilon^{V(x_k^*, x_j^*)}. \tag{A.4}$$

By Definition 11(iii), $\bar\delta$, $T_{\bar\delta+}$ satisfy Assumption 5, hence $W_j^{\bar\delta} \cap F_{\bar\delta}(T_{\bar\delta+}) = \emptyset$. Then, as $t \geq \underline{T}$ implies $t - T_{kj}(\bar\delta) \geq T_{\bar\delta+}$, we have, for all $x \in W_j^{\bar\delta}$,

$$\begin{aligned}
P_\varepsilon^{t-T_{kj}(\bar\delta)}(x, W_j^{\bar\delta}) &\underbrace{\geq}_{\text{by (A.2)}} (1/2)^{t-T_{kj}(\bar\delta)} \underbrace{G_{(0,\ldots,0)}}_{t-T_{kj}(\bar\delta)\text{zeros}} (x, W_j^{\bar\delta}) \\
&= (1/2)^{t-T_{kj}(\bar\delta)} \underbrace{P^{t-T_{kj}(\bar\delta)}(x, W_j^{\bar\delta})}_{=1 \text{ as } t-T_{kj}(\bar\delta) \geq T_{\bar\delta+}} = (1/2)^{t-T_{kj}(\bar\delta)}.
\end{aligned} \tag{A.5}$$

Combining (A.4) and (A.5), for all $x \in W_k^{\bar\delta}$,

$$\begin{aligned}
P_\varepsilon^t(x, W_j^{\bar\delta}) &= \int_X P_\varepsilon^{T_{kj}(\bar\delta)}(x, dy) P_\varepsilon^{t-T_{kj}(\bar\delta)}(y, W_j^{\bar\delta}) \\
&\geq \int_{W_j^{\bar\delta}} P_\varepsilon^{T_{kj}(\bar\delta)}(x, dy) \underbrace{P_\varepsilon^{t-T_{kj}(\bar\delta)}(y, W_j^{\bar\delta})}_{\geq (1/2)^{t-T_{kj}(\bar\delta)} \text{ by(A.5)}} \\
&\geq (1/2)^{t-T_{kj}(\bar\delta)} \int_{W_j^{\bar\delta}} P_\varepsilon^{T_{kj}(\bar\delta)}(x, dy) \\
&= (1/2)^{t-T_{kj}(\bar\delta)} \underbrace{P_\varepsilon^{T_{kj}(\bar\delta)}(x, W_j^{\bar\delta})}_{> \hat l \varepsilon^{V(x_k^*, x_j^*)} \text{ by (A.4)}} \\
&> \underbrace{(1/2)^{t-T_{kj}(\bar\delta)} \hat l}_{=:l}\, \varepsilon^{V(x_k^*, x_j^*)} = l\varepsilon^{V(x_k^*, x_j^*)}
\end{aligned}$$

and we have the first inequality in (A.3). For $x \in W_k^{\bar\delta}$, consider the expansion of $P_\varepsilon^t(x, W_j^{\bar\delta})$ in the form (A.2). If there exists a strictly positive term with $q_1 + \cdots + q_t \le V(x_k^*, x_j^*) - 1$, then by Lemma 1 we have $P_\varepsilon^t(x, W_j^{\bar\delta}) > l\varepsilon^{V(x_k^*,x_j^*)-1}$, so $P_\varepsilon^t(x, W_j^{\bar\delta}) \in \Omega(\varepsilon^{V(x_k^*,x_j^*)-1})$. Together with Property (C1) this implies $P_\varepsilon^t(x, W_j^{\bar\delta}) \in \Omega(\varepsilon^{V^-(x_k^*,x_j^*)-1})$, contradicting Definition 11(i). So all strictly positive terms have $q_1+\cdots+q_t \ge V(x_k^*, x_j^*)$ and by Lemma 1 we have $P_\varepsilon^t(x, W_j^{\bar\delta}) \le (M+1)^t \varepsilon^{V(x_k^*,x_j^*)}$. $\square$

Define $\hat\Phi_\varepsilon$ as the Markov process with kernel $\hat P_\varepsilon(.,.) = P_\varepsilon^{\underline{T}}(.,.)$. This process also has invariant measure π_ε. Define:

$$\hat\tau_A(k) := \min\{t > \hat\tau_A(k-1) : \hat\Phi_\varepsilon^t \in A\}; \quad \hat\tau_A(0) = 0.$$

Let $\bar W^\delta = \bigcup_{i: x_i^* \in \mathcal{A}} W_i^\delta$ and define $\tilde\Phi_\varepsilon$ as the process $\hat\Phi_\varepsilon$ only observed when it lies in $\bar W^{\bar\delta}$, $\tilde\Phi_\varepsilon^t = \hat\Phi_\varepsilon^{\hat\tau_{\bar W^{\bar\delta}}(t)}$. It follows from Lemma 2(iii) that $Pr_x(\hat\tau_{\bar W^{\bar\delta}}(t) < \infty) = 1$. Then the invariant measure of $\tilde\Phi_\varepsilon$ is given by:

$$\tilde\pi_\varepsilon(.) = \frac{\pi_\varepsilon(.)}{\pi_\varepsilon(\bar W^{\bar\delta})}.$$

Denote:

$$_A\hat P_\varepsilon^t(x, B) := Pr_x(\hat\Phi_\varepsilon^t \in B, \hat\tau_A(1) \ge t); \quad x \in X;\ A, B \in \mathcal{B}(X).$$

The kernel of $\tilde\Phi_\varepsilon$ is given by:

$$\tilde P_\varepsilon(x, A) = \sum_{t=1}^{\infty} {}_{\bar W^{\bar\delta}}\hat P_\varepsilon^t(x, A), \quad A \in \mathcal{B}(\bar W^{\bar\delta}).$$

Lemma 4. *There exists $\bar\varepsilon$ such that for all $\varepsilon < \bar\varepsilon$, $x \notin F_{\bar\delta}(T_{\bar\delta+})$,*

$$Pr_x(\hat\tau_{W^{\bar\delta}} > t | \tau_{F_{\bar\delta}(T_{\bar\delta+})} > t\underline{T}) < \varepsilon^{\frac{t}{2}}. \tag{A.6}$$

Proof. For conciseness, write $F = F_{\bar\delta}(T_{\bar\delta+})$. Note that

$$\begin{aligned}
&\int_{X\setminus F} P_\varepsilon(., dy_1) \int_{X\setminus F} P_\varepsilon(y_1, dy_2) \ldots \\
&\quad \times \int_{X\setminus F} P_\varepsilon(y_{\underline{T}-2}, dy_{\underline{T}-1}) P_\varepsilon(y_{\underline{T}-1}, X \setminus (F \cup W^{\bar\delta})) \\
&\le \int_X P_\varepsilon(., dy_1) \int_X P_\varepsilon(y_1, dy_2) \ldots \\
&\quad \times \int_X P_\varepsilon(y_{\underline{T}-2}, dy_{\underline{T}-1}) P_\varepsilon(y_{\underline{T}-1}, X \setminus (F \cup W^{\bar\delta})) \\
&= P_\varepsilon^{\underline{T}}(., X \setminus (F \cup W^{\bar\delta})) \le P_\varepsilon^{\underline{T}}(., X \setminus W^{\bar\delta})
\end{aligned} \tag{A.7}$$

and that for $x \notin F$,

$$\begin{aligned}
P_\varepsilon^{\underline{T}}(x, X \setminus W^{\bar\delta}) &= 1 - P_\varepsilon^{\underline{T}}(x, W^{\bar\delta}) \\
&\le 1 - (\underbrace{1 - \varepsilon - \cdots - \varepsilon^M}_{>1-M\varepsilon})^{\underline{T}} \underbrace{P^{\underline{T}}(x, W^{\bar\delta})}_{=1} \\
&\le 1 - (1 - M\varepsilon)^{\underline{T}} < 1 - (1 - 2\underline{T}M\varepsilon) \\
&= (2\underline{T}M\varepsilon) < \varepsilon^{\frac{3}{4}}.
\end{aligned} \tag{A.8}$$

Now see that, for $x \notin F$,

$$\begin{aligned}
&Pr_x(\hat\Phi_\varepsilon^n \notin W^{\bar\delta}, n = 1, \ldots, t;\ \Phi_\varepsilon^m \notin F, m = 1, \ldots, t\underline{T}) \\
&= \int_{X\setminus F} P_\varepsilon(x, dy_1) \int_{X\setminus F} P_\varepsilon(y_1, dy_2) \ldots \\
&\times \int_{X\setminus F} P_\varepsilon(y_{\underline{T}-2}, dy_{\underline{T}-1}) \int_{X\setminus(F\cup W^{\bar\delta})} P_\varepsilon(y_{\underline{T}-1}, dy_{\underline{T}}) \\
&\times \int_{X\setminus F} \cdots \int_{X\setminus(F\cup W^{\bar\delta})} P_\varepsilon(y_{(t-1)\underline{T}-1}, dy_{(t-1)\underline{T}}) \\
&\times \underbrace{\iint_{X\setminus F} P_\varepsilon(y_{(t-1)\underline{T}}, dy_{(t-1)\underline{T}+1}) \ldots \int_{X\setminus F} P_\varepsilon(y_{t\underline{T}-2}, dy_{t\underline{T}-1}) P_\varepsilon(y_{t\underline{T}-1}, X \setminus (F \cup W^{\bar\delta}))}_{<\varepsilon^{\frac{3}{4}} \text{ by (A.7) and (A.8)}} \\
&< \varepsilon^{\frac{3}{4}} \int_{X\setminus F} P_\varepsilon(x, dy_1) \int_{X\setminus F} P_\varepsilon(y_1, dy_2) \ldots \int_{X\setminus F} P_\varepsilon(y_{\underline{T}-2}, dy_{\underline{T}-1}) \\
&\times \int_{X\setminus(F\cup W^{\bar\delta})} P_\varepsilon(y_{\underline{T}-1}, dy_{\underline{T}}) \int_{X\setminus F} \cdots \int_{X\setminus(F\cup W^{\bar\delta})} P_\varepsilon(y_{(t-1)\underline{T}-1}, dy_{(t-1)\underline{T}}) \\
&= \varepsilon^{\frac{3}{4}} \int_{X\setminus F} P_\varepsilon(x, dy_1) \int_{X\setminus F} P_\varepsilon(y_1, dy_2) \ldots \int_{X\setminus F} P_\varepsilon(y_{\underline{T}-2}, dy_{\underline{T}-1}) \\
&\times \int_{X\setminus(F\cup W^{\bar\delta})} P_\varepsilon(y_{\underline{T}-1}, dy_{\underline{T}}) \int_{X\setminus F} \ldots P_\varepsilon(y_{(t-1)\underline{T}-1}, X \setminus (F \cup W^{\bar\delta})) \\
&< \cdots < \varepsilon^{\frac{3}{4}t}
\end{aligned} \tag{A.9}$$

where the strict inequalities result from the repeated use of (A.7) and (A.8). Observe that for $x \notin F$,

$$P(x, F) = 0 \quad \text{and therefore } P(x, X \setminus F) = 1. \tag{A.10}$$

To see this is so, note that by definition, $F_{\bar\delta}(T_{\bar\delta+}) = \cup_{t \ge T_{\bar\delta+}} \{y \in X : P^t(y, W^{\bar\delta}) < 1\}$, a countable union of measurable sets. Therefore, if $P(x, F) > 0$, then there exists $t \ge T_{\bar\delta+}$ such that $P(x, \{y \in X : P^t(y, W^{\bar\delta}) < 1\}) > 0$. This implies that $P^{t+1}(x, W^{\bar\delta}) < 1$, implying $x \in F$, a contradiction.

Now, using (A.10), for $x \notin F$,

$$\begin{aligned}
&Pr_x(\Phi_\varepsilon^m \notin F, m = 1, \ldots, t\underline{T}) \\
&= \int_{X\setminus F} P_\varepsilon(x, dy_1) \int_{X\setminus F} P_\varepsilon(y_1, dy_2) \ldots \\
&\quad \times \int_{X\setminus F} P_\varepsilon(y_{t\underline{T}-2}, dy_{t\underline{T}-1}) P_\varepsilon(y_{t\underline{T}-1}, X \setminus F) \\
&\ge \int_{X\setminus F} \frac{1}{2} P(x, dy_1) \int_{X\setminus F} \frac{1}{2} P(y_1, dy_2) \ldots \\
&\quad \times \int_{X\setminus F} \frac{1}{2} P(y_{t\underline{T}-2}, dy_{t\underline{T}-1}) \frac{1}{2} \underbrace{P(y_{t\underline{T}-1}, X \setminus F)}_{=1 \text{ by (A.10)}} \\
&= \left(\frac{1}{2}\right)^{t\underline{T}} \int_{X\setminus F} P(x, dy_1) \int_{X\setminus F} P(y_1, dy_2) \ldots \\
&\quad \times \int_{X\setminus F} P(y_{t\underline{T}-2}, dy_{t\underline{T}-1}) \\
&= \left(\frac{1}{2}\right)^{t\underline{T}} \int_{X\setminus F} P(x, dy_1) \int_{X\setminus F} P(y_1, dy_2) \ldots \\
&\quad \times \underbrace{P(y_{t\underline{T}-2}, X \setminus F)}_{=1 \text{ by (A.10)}} \\
&= \cdots = \left(\frac{1}{2}\right)^{t\underline{T}}.
\end{aligned} \tag{A.11}$$

Using the definition of conditional probability, (A.9) and (A.11), for $x \notin F$,

$$\begin{aligned}
&Pr_x(\hat\tau_{W^{\bar\delta}} > t | \tau_F > t\underline{T}) \\
&= Pr_x(\hat\Phi_\varepsilon^n \notin W^{\bar\delta}, n = 1, \ldots, t | \Phi_\varepsilon^m \notin F, m = 1, \ldots, t\underline{T}) \\
&= \frac{Pr_x(\hat\Phi_\varepsilon^n \notin W^{\bar\delta}, n = 1, \ldots, t;\ \Phi_\varepsilon^m \notin F, m = 1, \ldots, t\underline{T})}{Pr_x(\Phi_\varepsilon^m \notin F, m = 1, \ldots, t\underline{T})} \\
&< \frac{\varepsilon^{\frac{3}{4}t}}{\left(\frac{1}{2}\right)^{t\underline{T}}} = \left(\frac{\varepsilon^{\frac{3}{4}}}{\left(\frac{1}{2}\right)^{\underline{T}}}\right)^t < \varepsilon^{\frac{t}{2}}. \quad \square
\end{aligned}$$

Lemma 5. *For all $x_i^* \in \Lambda$, there exists $\bar\varepsilon$ such that for all $\varepsilon < \bar\varepsilon$, $x \in W_i^{\bar\delta}$,*

$$Pr_x(\tau_{F_{\bar\delta}(T_{\bar\delta+})} \le \tilde t) \le \tilde t(M+1)^{\tilde t}\varepsilon^{\bar V_i}. \tag{A.12}$$

Proof.

$$\begin{aligned} Pr_x(\tau_{F_{\bar\delta}(T_{\bar\delta+})} \le \tilde t) &\le \sum_{t=1}^{\tilde t} P_\varepsilon^t(x, F_{\bar\delta}(T_{\bar\delta+})) \\ &\le \sum_{t=1}^{\tilde t}(M+1)^t\varepsilon^{\bar V_i} \le \tilde t(M+1)^{\tilde t}\varepsilon^{\bar V_i} \end{aligned}$$

where the second inequality follows from Lemma 1 and Assumption 5. □

Lemma 6. *Let $\tilde t > 2\max_{x_k^*, x_j^* \in \mathcal{A}} V(x_k^*, x_j^*)$. Then, for all $x_k^*, x_j^* \in \mathcal{A}$, there exist $l > 0$, $\bar\varepsilon$ such that for all $x \in W_k^{\bar\delta}$, $x \notin \cup_{l: x_l^* \notin \mathcal{A}} B_l$, $\varepsilon < \bar\varepsilon$,*

$$\begin{aligned} l\varepsilon^{V(x_k^*,x_j^*)} < \tilde P_\varepsilon(x, W_j^{\bar\delta}) &< (\tilde t(M+1)^{\tilde t\underline{T}} + 1 \\ &\quad + \tilde t\underline{T}(M+1)^{\tilde t\underline{T}})\varepsilon^{V(x_k^*,x_j^*)}. \end{aligned} \tag{A.13}$$

Proof.

$$\begin{aligned} \tilde P_\varepsilon(x, W_j^{\bar\delta}) &= \sum_{t=1}^{\infty} {}_{\bar W^{\bar\delta}}\hat P_\varepsilon^t(x, W_j^{\bar\delta}) \ge \hat P_\varepsilon(x, W_j^{\bar\delta}) \\ &= P_\varepsilon^{\underline{T}}(x, W_j^{\bar\delta}) > l\varepsilon^{V(x_k^*,x_j^*)} \end{aligned}$$

where the final inequality uses Lemma 3. Note that $x \in W_k^{\bar\delta}$, $x \notin \cup_{l: x_l^* \notin \mathcal{A}} B_l$, $x_i^* \notin \mathcal{A}$ implies $P_\varepsilon^t(x, W_i^{\bar\delta}) = 0$ for all t, hence $Pr_x(\hat\tau_{\bar W^{\bar\delta}} > \tilde t) = Pr_x(\hat\tau_{W^{\bar\delta}} > \tilde t)$. Furthermore, using the law of total probability, Lemmas 3–5, and Property C, we have

$$\begin{aligned} \tilde P_\varepsilon(x, W_j^{\bar\delta}) &= \sum_{t=1}^{\infty} {}_{\bar W^{\bar\delta}}\hat P_\varepsilon^t(x, W_j^{\bar\delta}) \\ &\le \sum_{t=1}^{\tilde t} {}_{\bar W^{\bar\delta}}\hat P_\varepsilon^t(x, W_j^{\bar\delta}) + Pr_x(\hat\tau_{\bar W^{\bar\delta}} > \tilde t) \\ &\le \sum_{t=1}^{\tilde t} \hat P_\varepsilon^t(x, W_j^{\bar\delta}) + Pr_x(\hat\tau_{W^{\bar\delta}} > \tilde t \,|\, \tau_{F_{\bar\delta}(T_{\bar\delta+})} > \tilde t\underline{T}) \\ &\quad \times Pr_x(\tau_{F_{\bar\delta}(T_{\bar\delta+})} > \tilde t\underline{T}) + Pr_x(\hat\tau_{W^{\bar\delta}} > \tilde t \,|\, \tau_{F_{\bar\delta}(T_{\bar\delta+})} \le \tilde t\underline{T}) \\ &\quad \times Pr_x(\tau_{F_{\bar\delta}(T_{\bar\delta+})} \le \tilde t\underline{T}) \quad \text{[by law of total probability]} \\ &\le \sum_{t=1}^{\tilde t} \hat P_\varepsilon^t(x, W_j^{\bar\delta}) + Pr_x(\hat\tau_{W^{\bar\delta}} > \tilde t \,|\, \tau_{F_{\bar\delta}(T_{\bar\delta+})} > \tilde t\underline{T}) \\ &\quad + Pr_x(\tau_{F_{\bar\delta}(T_{\bar\delta+})} \le \tilde t\underline{T}) \\ &< \tilde t(M+1)^{\tilde t\underline{T}}\varepsilon^{V(x_k^*,x_j^*)} + \varepsilon^{\frac{\tilde t}{2}} + \tilde t\underline{T}(M+1)^{\tilde t\underline{T}}\varepsilon^{\bar V_k} \\ &\quad \text{[by Lemmas 3–5 respectively]} \\ &\le \tilde t(M+1)^{\tilde t\underline{T}}\varepsilon^{V(x_k^*,x_j^*)} + \underbrace{\varepsilon^{V(x_k^*,x_j^*)}}_{\substack{\text{as } \tilde t > 2V(x_k^*,x_j^*) \\ \text{by statement} \\ \text{of lemma.}}} \\ &\quad + \tilde t\underline{T}(M+1)^{\tilde t\underline{T}} \underbrace{\varepsilon^{V(x_k^*,x_j^*)}}_{\substack{\text{as } \bar V_k \ge V^-(x_k^*,x_j^*) \\ = V(x_k^*,x_j^*) \\ \text{by defn of } \bar V_k \\ \text{and (C1)}}}. \quad \square \end{aligned}$$

Lemma 7 (Freidlin and Wentzell, 1984, Lemmas 3.1, 3.2)**.** *Assume there exists a partition of X into finitely many disjoint sets $\{X_i | i \in L\}$, $|L| = \nu$, such that, for all $i, j \in L$,*

$$\exists c_{ij} > 0 \quad s.t. \inf_{x \in X_i} P_\varepsilon(x, X_j) \ge c_{ij}.$$

For given invariant measure $\pi_\varepsilon(.)$, let:

$$p_{ij} := \frac{1}{\pi_\varepsilon(X_i)} \int_{X_i} P_\varepsilon(x, X_j)\pi_\varepsilon(dx).$$

For $g \in \mathcal{G}(i)$, define:

$$vol(g) := \prod_{(j\to k)\in g} p_{jk}; \qquad Q_i := \sum_{g \in \mathcal{G}(i)} vol(g)$$

then:

$$\pi_\varepsilon(X_i) = \frac{Q_i}{\sum_{j=1}^{\nu} Q_j}.$$

Subsequent proofs use the following additional items of Bachmann–Landau asymptotic notation, which express the ideas of f being bounded by g above, and both above and below respectively:

$$\begin{aligned} &f(\varepsilon) \in O(g(\varepsilon)) \Leftrightarrow \exists k > 0,\ \bar\varepsilon \text{ s.t. } \forall \varepsilon < \bar\varepsilon,\ |f(\varepsilon)| \le kg(\varepsilon) \\ &f(\varepsilon) \in \Theta(g(\varepsilon)) \Leftrightarrow \exists k_1, k_2 > 0,\ \bar\varepsilon \text{ s.t. } \forall \varepsilon < \bar\varepsilon, \\ &\quad k_1 g(\varepsilon) \le |f(\varepsilon)| \le k_2 g(\varepsilon). \end{aligned}$$

Proof of Proposition 4. Let $\tilde L := \{j \in L : x_j^* \in \mathcal{A}\}$. Let $\tilde{\mathcal{G}}(.)$, $\tilde{\mathcal{V}}(.)$, $\tilde L_{\min}$ be as $\mathcal{G}(.)$, $\mathcal{V}(.)$, $L_{\min}$, only defined on $\tilde L$ instead of on L. Note that for $x_j^* \in \mathcal{A}$, $x_k^* \notin \mathcal{A}$, by Lemma 2(ii) we have $V(x_j^*, x_k^*) = \infty$. This implies that for $x_k^* \notin \mathcal{A}$, $\mathcal{V}(k) = \infty$. Lemma 2(i) implies that for $x_j^* \in \mathcal{A}$, $\mathcal{V}(j) < \infty$. Therefore, $x_k^* \notin \mathcal{A}$ implies $k \notin L_{\min}$. A further implication is that, for $i \in L_{\min}$, $x_j^* \in \mathcal{A}$, $x_k^* \notin \mathcal{A}$, the edge $(j \to k)$ will not be in any graph g that solves the minimization problem in the definition of $\mathcal{V}(i)$. Therefore, such a graph $g \in \mathcal{G}(i)$ must comprise the union of

(a) a graph $\tilde g \in \tilde{\mathcal{G}}(i)$ that is induced from g by the vertex set $\tilde L$, and
(b) a set g^- of directed edges exiting vertices in $L \setminus \tilde L$, such that there is a path from each $k \in L \setminus \tilde L$ to some $j \in \tilde L$.

Note that any set of edges that satisfies (b) does so independently of i, so the choice of g^- in the solution of the minimization problem in the definition of $\mathcal{V}(i)$, $i \in \tilde L$, is independent of i. Therefore, $L_{\min}$ is solely determined by solving the problem on $\tilde L$. That is,

$$i \in L_{\min} \iff i \in \tilde L_{\min}. \tag{A.14}$$

Partition $\bar W^{\bar\delta}$ into $\{W_i^{\bar\delta}\}_{i \in \tilde L}$ and define, for $i, j \in \tilde L$:

$$\tilde p_{\varepsilon,ij} := \frac{1}{\tilde\pi_\varepsilon(W_i^{\bar\delta})} \int_{W_i^{\bar\delta}} \tilde P_\varepsilon(x, W_j^{\bar\delta})\tilde\pi_\varepsilon(dx).$$

It follows from Definition 6, that $\tilde\pi_\varepsilon(\cup_{l: x_l^* \notin \mathcal{A}} B_l) = \pi_\varepsilon(\cup_{l: x_l^* \notin \mathcal{A}} B_l) = 0$, therefore for $x_i^*, x_j^* \in \mathcal{A}$, the bounds on $\tilde P_\varepsilon(x, W_j^{\bar\delta})$ in Lemma 6 imply that $\tilde p_{\varepsilon,ij} \in \Theta(\varepsilon^{V(x_i^*,x_j^*)})$. Define $vol(.)$ and Q_i as in Lemma 7 with $p_{ij} = \tilde p_{\varepsilon,ij}$. Now, by definition of $\Theta(.)$,

$$vol(g) = \prod_{(j\to k)\in g} \underbrace{\tilde p_{\varepsilon,jk}}_{\substack{\in \Theta(\varepsilon^{V(x_j^*,x_k^*)}) \\ \text{by Lemma 6}}} \in \Theta(\varepsilon^{\sum_{(j\to k)\in g} V(x_j^*,x_k^*)})$$

and

$$Q_i = \sum_{g \in \tilde{\mathcal{G}}(i)} vol(g) \in \Theta(\varepsilon^{\min_{g \in \tilde{\mathcal{G}}(i)} \sum_{(j\to k)\in g} V(x_j^*,x_k^*)}) = \Theta(\varepsilon^{\tilde{\mathcal{V}}(i)}).$$

So, by Lemma 7 and the definitions of $\Theta(.)$, $\tilde{L}_{\min}$, we have

$$\tilde{\pi}_\varepsilon(W_i^{\tilde{\delta}}) = \frac{Q_i}{\sum_{j\in\tilde{L}} Q_j} \in \Theta(1) \iff \sum_{j\in\tilde{L}} Q_j \in \Theta(\varepsilon^{\tilde{\mathcal{V}}(i)})$$
$$\iff \tilde{\mathcal{V}}(i) \le \tilde{\mathcal{V}}(j) \quad \forall j \in L \iff i \in \tilde{L}_{\min}. \tag{A.15}$$

Note that for $x_j^* \notin \mathcal{A}$, by the definition of $\mathcal{A}$, $\pi_\varepsilon(W_j^{\bar{\delta}}) = 0$, so by Proposition 3, $\lim_{\varepsilon\to 0} \pi_\varepsilon(\bar{W}^{\bar{\delta}}) = 1$. Finally, see that

$$\tilde{\pi}_\varepsilon(W_i^{\bar{\delta}}) \in \Theta(1)$$
$$\iff \lim_{\varepsilon\to 0} \tilde{\pi}_\varepsilon(W_i^{\bar{\delta}}) = \lim_{\varepsilon\to 0} \frac{\pi_\varepsilon(W_i^{\bar{\delta}})}{\pi_\varepsilon(\bar{W}^{\bar{\delta}})} > 0$$
$$\underbrace{\iff}_{\substack{\text{as by Proposition 3,}\\ \lim_{\varepsilon\to 0}\pi_\varepsilon(\bar{W}^{\bar{\delta}})=1}} \lim_{\varepsilon\to 0} \pi_\varepsilon(W_i^{\bar{\delta}})$$
$$= \pi(W_i^{\bar{\delta}}) > 0 \underbrace{\iff}_{\text{by Corollary 1}} \pi(x_i^*) > 0. \tag{A.16}$$

Combining (A.14), (A.15), (A.16) gives $i \in L_{\min} \iff \pi(x_i^*) > 0$. □

Lemma 8. *Under Condition* (i) *of Proposition 5, for $x \in X$, $A \in \mathcal{B}(X)$, A open, $T \in \mathbb{N}_+$, there exist δ_{Ax}, ξ_{Ax}, such that for all $y \in B_{\delta_{Ax}}(x)$,*

$$G_{(q_1,\ldots,q_T)}(x, A) > 0 \Longrightarrow G_{(q_1,\ldots,q_T)}(y, A) > \xi_{Ax}.$$

Proof. For $T = 1$, the statement of the lemma is simply Condition (i). We shall prove the lemma by induction on T. Assume it is true for $T = t - 1$. Now, as X is a separable metric space and hence strongly Lindelöf, the support of $G_{q_1}(x, .)$ is a Borel set with measure 1 under $G_{q_1}(x, .)$. Thus we can integrate over the support of $G_{q_1}(x, .)$ rather than over the entire space.

$$G_{(q_1,\ldots,q_t)}(x, A) > 0$$
$$\Longrightarrow \int_{\operatorname{supp} G_{q_1}(x,.)} G_{q_1}(x, dy) G_{(q_2,\ldots,q_t)}(y, A) > 0$$

which implies there exists $y_1 \in \operatorname{supp} G_{q_1}(x, .)$ such that $G_{(q_2,\ldots,q_t)}(y_1, A) > 0$. Then, by the inductive hypothesis, there exist δ_{Ay_1}, ξ_{Ay_1}, such that $y \in E_1 := B_{\delta_{Ay_1}}(y_1)$ implies that $G_{(q_2,\ldots,q_t)}(y, A) > \xi_{Ay_1}$.

Also, by the definition of support, $y_1 \in \operatorname{supp} G_{q_1}(x, .)$ implies that $G_{q_1}(x, E_1) > 0$. Hence, by Condition (i), there exist δ_{E_1x}, ξ_{E_1x} such that for all $w \in E_0 := B_{\delta_{E_1x}}(x)$, $G_{q_1}(w, E_1) > \xi_{E_1x}$.

So, for all $w \in E_0$,

$$\begin{aligned} G_{(q_1,\ldots,q_t)}(w, A) &\ge \int_{E_1} G_{q_1}(w, dy) G_{(q_2,\ldots,q_t)}(y, A) \\ &\ge \xi_{Ay_1} \int_{E_1} G_{q_1}(w, dy) \\ &= \xi_{Ay_1} G_{q_1}(w, E_1) > \xi_{Ay_1}\xi_{E_1x} =: \xi_{Ax}. \quad \square \end{aligned}$$

Proof of Proposition 5. We know from Condition (iii) that there exists T_1 such that $P_\varepsilon^{T_1}(x_k^*, W_j) \in \Omega(\varepsilon^{V^-(x_k^*,x_j^*)})$. Let $\tilde{\delta}$ satisfy Condition (ii). By Lemma 1 there exist $q_1 + \cdots + q_{T_1} \le V^-(x_k^*, x_j^*)$ such that the term $G_{(q_1,\ldots,q_{T_1})}(x_k^*, W_j)$ in the expansion of $P_\varepsilon^{T_1}(x_k^*, W_j)$ given by expression (A.2) is strictly positive.

Expanding $G_{(q_1,\ldots,q_{T_1})}(x_k^*, W_j)$ as per expression (A.1), we see that the similar expansion of $G_{(q_1,\ldots,q_{T_1})}(x_k^*, W_j \setminus F_{\tilde{\delta}}(t))$ differs only in the final term, which is $G_{q_{T_1}}(y_{T_1-1}, W_j \setminus F_{\tilde{\delta}}(t))$ rather than $G_{q_{T_1}}(y_{T_1-1}, W_j)$. Now, for all $z \in W_j$, there exists t such that $z \notin F_{\tilde{\delta}}(t)$, so $\mathbb{1}_{W_j\setminus F_{\tilde{\delta}}(t)}(z) \to \mathbb{1}_{W_j}(z)$ as $t \to \infty$. So by bounded convergence

$$\begin{aligned} &G_{q_{T_1}}(y_{T_1-1}, W_j \setminus F_{\tilde{\delta}}(t)) \\ &= \int_X \mathbb{1}_{W_j\setminus F_{\tilde{\delta}}(t)}(z) G_{q_{T_1}}(y_{T_1-1}, dz) \\ &\xrightarrow{t\to\infty} \int_X \mathbb{1}_{W_j}(z) G_{q_{T_1}}(y_{T_1-1}, dz) = G_{q_{T_1}}(y_{T_1-1}, W_j). \end{aligned} \tag{A.17}$$

This holds for all $y_{T_1-1} \in X$, so $G_{q_{T_1}}(., W_j \setminus F_{\tilde{\delta}}(t))$ converges pointwise to $G_{q_{T_1}}(., W_j)$. Using bounded convergence repeatedly on the integrals in the expansion of $G_{(q_1,\ldots,q_{T_1})}(x_k^*, W_j \setminus F_{\tilde{\delta}}(t))$ given by (A.1), we see that $G_{(q_1,\ldots,q_{T_1})}(x_k^*, W_j \setminus F_{\tilde{\delta}}(t)) \to G_{(q_1,\ldots,q_{T_1})}(x_k^*, W_j)$ as $t \to \infty$, so we can choose T_2 large enough that $G_{(q_1,\ldots,q_{T_1})}(x_k^*, W_j \setminus F_{\tilde{\delta}}(T_2)) > 0$.

By Lemma 1 we then have $P_\varepsilon^{T_1}(x_k^*, W_j \setminus F_{\tilde{\delta}}(T_2)) \in \Omega(\varepsilon^{V^-(x_k^*,x_j^*)})$. For all $z \in W_j \setminus F_{\tilde{\delta}}(T_2)$, $t \ge T_2$, $P^t(z, W_j^{\tilde{\delta}}) = 1$. So, for $T = T_1 + T_2$, $P_\varepsilon^T(x_k^*, W_j^{\tilde{\delta}}) \in \Omega(\varepsilon^{V^-(x_k^*,x_j^*)})$.

Expanding $P_\varepsilon^T(x_k^*, W_j^{\tilde{\delta}})$ as in (A.2), there must be a strictly positive term $G_{(q_1,\ldots,q_T)}(x_k^*, W_j^{\tilde{\delta}})$ in the expansion such that $q_1 + \cdots + q_T \le V^-(x_k^*, x_j^*)$ or Lemma 1 gives a contradiction to $P_\varepsilon^T(x_k^*, W_j^{\tilde{\delta}}) \in \Omega(\varepsilon^{V^-(x_k^*,x_j^*)})$. So, by Lemma 8, there exist δ_{kj}, ξ such that for $z \in B_{\delta_{kj}}(x_k^*)$, $G_{(q_1,\ldots,q_T)}(z, W_j^{\tilde{\delta}}) > \xi$ and therefore

$$P_\varepsilon^T(z, W_j^{\tilde{\delta}}) \ge \varepsilon^{q_1+\cdots+q_T} G_{(q_1,\ldots,q_T)}(z, W_j^{\tilde{\delta}}) > \varepsilon^{V^-(x_k^*,x_j^*)}\xi. \tag{A.18}$$

This shows that for some $l > 0$, for all $x \in W_k^{\delta_{kj}}$, $P_\varepsilon^T(x, W_j^{\tilde{\delta}}) > l\varepsilon^{V^-(x_k^*,x_j^*)}$. So, by definition of $V(.,.)$, we have $V(x_k^*, x_j^*) \le V^-(x_k^*, x_j^*)$. As $V(x_k^*, x_j^*) \ge V^-(x_k^*, x_j^*)$ by definition, we have $V(x_k^*, x_j^*) = V^-(x_k^*, x_j^*)$, so (C1) holds.

Now, $\tilde{\delta}$ can be chosen arbitrarily small and still satisfy Condition (ii), so we can assume $\tilde{\delta} < \hat{\delta}$, where $\hat{\delta}$ is as in Assumption 5. Now, by Assumption 5, for any $\tilde{\delta}' < \tilde{\delta}$ we can choose $T_{\tilde{\delta}'+}$ such that $F_{\tilde{\delta}'}(T_{\tilde{\delta}'+}) \cap W_j^{\tilde{\delta}} = \emptyset$. Hence, for all $x \in W_j^{\tilde{\delta}}$, $P^{T_{\tilde{\delta}'+}}(x, W_j^{\tilde{\delta}'}) = 1$. Therefore, for $T' = T + T_{\tilde{\delta}'+}$, for all $x \in W_k^{\delta_{kj}}$,

$$\begin{aligned} P_\varepsilon^{T'}(x, W_j^{\tilde{\delta}'}) &= \int_X P_\varepsilon^T(x, dy) P_\varepsilon^{T_{\tilde{\delta}'+}}(y, W_j^{\tilde{\delta}'}) \\ &\ge \int_{W_j^{\tilde{\delta}}} P_\varepsilon^T(x, dy) P_\varepsilon^{T_{\tilde{\delta}'+}}(y, W_j^{\tilde{\delta}'}) \\ &\ge \int_{W_j^{\tilde{\delta}}} P_\varepsilon^T(x, dy) \left(\frac{1}{2}\right)^{T_{\tilde{\delta}'+}} \underbrace{P^{T_{\tilde{\delta}'+}}(y, W_j^{\tilde{\delta}'})}_{=1 \text{ as } y\notin F_{\tilde{\delta}'}(T_{\tilde{\delta}'+})} \\ &= \left(\frac{1}{2}\right)^{T_{\tilde{\delta}'+}} \underbrace{P_\varepsilon^T(x, W_j^{\tilde{\delta}})}_{\substack{>\varepsilon^{V^-(x_k^*,x_j^*)}\xi \\ \text{by (A.18)}}} > \left(\frac{1}{2}\right)^{T_{\tilde{\delta}'+}} \varepsilon^{V^-(x_k^*,x_j^*)}\xi. \end{aligned}$$

This shows that for some $l' > 0$, for all $x \in W_k^{\delta_{kj}}$, $P_\varepsilon^{T'}(x, W_j^{\tilde{\delta}'}) > l'\varepsilon^{V^-(x_k^*,x_j^*)}$, so (C2) holds. □

Proof of Proposition 8. Assume the statement is incorrect for some $\xi > 0$. There must be an infinite sequence of stable states $\{x^*(k)\}_k$ corresponding to increasing values of k such that none of these stable states are within ξ of any Nash equilibrium. As the sequence is bounded it contains a convergent subsequence. Restrict attention to such a subsequence. Denote its limit by $\bar{x}^*$. As $\bar{x}^*$ is not a

Nash equilibrium at least one of the actions played is not a best response to the true distribution of play. Assume that action i is one of these actions. $\bar{x}_i^* > 0$. Then there exist $\underline{k}, \eta > 0$, such that for all $k > \underline{k}$, $x_i^*(k) > \eta$. Note that as $x^*(k)$ is a stable state, $BR_i(x^*(k)) = x_i^*(k)$.

There exists γ such that for all x such that $|x - \bar{x}^*| < \gamma$, best responses to the true distribution of x are a subset of the best responses to the true distribution of $\bar{x}^*$. Define $f_k(\sigma)$ as the probability, at $x^*(k)$, that a sample of size k gives the mixed strategy σ. As $k \to \infty$, $f_k(\sigma)$ weakly converges to the probability measure with point mass on $\bar{x}^*$. This implies:

$$\forall \eta\, \exists \bar{k} : \forall k > \bar{k}, \sum_{\sigma \notin B_\gamma(\bar{x}^*)} f_k(\sigma) < \eta,$$

which implies:

$$BR_i(x^*(k)) < \eta$$

and we have a contradiction. □

Proof of Proposition 9. Assume that $x_i^{NE} = 1$. Let x^t be such that $x_i^t = 1 - \xi$, $\xi \in (0, 1)$. There exists $s \in \mathbb{R}$ such that if i is not a unique best response to the strategy σ then:

$$\sum_{j \neq i} \sigma_j \geq s.$$

The proportion of players changing strategy who will sample such a σ is:

$$\tilde{\xi} = \sum_{j=\lceil ks \rceil}^{k} \binom{k}{j} \xi^j (1 - \xi)^{k-j} \in \Theta(\xi^{\lceil ks \rceil}).$$

That is, of the players called to update their action, a proportion at least $(1 - \tilde{\xi})$ will choose action i. Assuming k is large enough that $\lceil ks \rceil \geq 2$, there exists $\bar{\xi}$ such that:

$$\forall \xi < \bar{\xi}, \quad \tilde{\xi} < \xi.$$

So for any $x^t \in B_{\bar{\xi}}(x^{NE})$ we have that $x_i^t = 1 - \xi$ for some $\xi < \bar{\xi}$ and:

$$\begin{aligned} x_i^{t+1} &= (1 - \alpha) x_i^t + \alpha BR_i(x_i^t) \geq (1 - \alpha) x_i^t + \alpha(1 - \tilde{\xi}) \\ &= (1 - \alpha)(1 - \xi) + \alpha(1 - \tilde{\xi}) > 1 - \xi = x_i^t. \end{aligned}$$

So we have convergence to x^{NE} from an open ball centered on x^{NE}, so $x^{NE} \in \Lambda$ and x^{NE} is asymptotically stable. This open ball is reached with positive probability from anywhere in the state space so $x^{NE} \in \mathcal{A}$. □

Proof of Proposition 11. For $i = k + 1, \ldots, n$, $x \in X_i$:

$$n\varepsilon > P_\varepsilon(x, X_{i-1}) \geq \varepsilon \frac{i}{n} \frac{1}{x_{\max}},$$

therefore $p_{i,i-1} \in \Theta(\varepsilon)$. Also, for $x \in X_k$:

$$n\varepsilon > P_\varepsilon(x, X_0) \geq \varepsilon \frac{k}{n} \frac{1}{x_{\max}},$$

and $p_{k,0} \in \Theta(\varepsilon)$. So there exists a 0-graph $\tilde{g}$ with $vol(\tilde{g}) \in \Theta(\varepsilon^{n-k+1})$. Therefore $Q_0 \in \Theta(\varepsilon^{n-k+1})$. For $x \in X_0$, $i \neq 0$:

$$P_\varepsilon(x, X_i) \in O(\varepsilon^k)$$

so any i-graph g, $i \neq 0$, has $vol(g) \in O(\varepsilon^{(n-k)+k}) = O(\varepsilon^n)$. Therefore $Q_i \in O(\varepsilon^n)$. Using the formula for $\pi_\varepsilon(X_i)$ in Lemma 7, we see that $\pi_\varepsilon(X_i) \to 0$ as $\varepsilon \to 0$ for all $i \neq 0$. So $\pi_\varepsilon(X_0) \to 1$ as $\varepsilon \to 0$ and by Proposition 3, π_ε approaches the distribution with point mass on 0^n. □

Appendix B. Counterexamples

B.1. The role of Assumption 3

Here we present a counterexample when Assumption 3 does not hold. Let $X = \{x_m\}_{m \in \mathbb{N}}$ and let X be equipped with the discrete metric. Let $\Lambda = \{x_0\}$. Let

$$P(x_{m+1}, x_m) = 1, \quad m \in \mathbb{N},$$

$$G_1(x_m, \{x_{m'} : m' \geq \bar{m}\}) = \left(\frac{1}{\bar{m} - m}\right)^{\frac{1}{4}} \quad \text{for } \bar{m} > m,$$

$$G_2(x_0) = 1.$$

Note that for any $x_m \in X$, $m \in \mathbb{N}_+$, $P^m(x_m, \Lambda) = 1$ so Assumption 2 holds. Now note that

$$\text{If } t' \leq \frac{\log \frac{1}{2}}{\log(1 - \varepsilon^2)}, \quad \text{then } (1 - \varepsilon^2)^{t'} \geq \frac{1}{2},$$

implying that for $t' \leq \frac{\log \frac{1}{2}}{\log(1-\varepsilon^2)}$, $m > t'$,

$$Pr_{x_m}(\Phi_\varepsilon^t \neq x_0 \text{ for } t = 1, \ldots, t') = (1 - \varepsilon^2)^{t'} \geq \frac{1}{2}. \tag{B.1}$$

Consider ε such that $1/2\varepsilon^2$ is an integer, and note that,

$$G_1\left(x_0, \left\{x_m : m \geq \frac{1}{\varepsilon^2}\right\}\right) = \left(\frac{1}{\frac{1}{\varepsilon^2}}\right)^{\frac{1}{4}} = \varepsilon^{\frac{1}{2}}. \tag{B.2}$$

Further note that

$$\frac{1}{2\varepsilon^2} < \frac{\log \frac{1}{2}}{\log(1 - \varepsilon^2)} \quad \text{and} \quad \frac{1}{\varepsilon^2} > \frac{1}{2\varepsilon^2}, \tag{B.3}$$

so (B.1) holds for $t' = 1/2\varepsilon^2$, $m = 1/\varepsilon^2$.

Now, as $G_2(x_0) > 0$, $\pi_\varepsilon(x_0) > 0$. The invariant measure of other sets as a proportion of $\pi_\varepsilon(x_0)$ is then given by the expected number of periods the process spends in those sets between visits to x_0.[26] For the set $X \setminus \{x_0\}$, this quantity is bounded below by

$$\begin{aligned} &Pr_{x_0}(\tau_{\{x_0\}} > 1/2\varepsilon^2) \frac{1}{2\varepsilon^2} \\ &= Pr_{x_0}(\Phi^t \neq x_0,\ t = 1, \ldots, 1/2\varepsilon^2) \frac{1}{2\varepsilon^2} \\ &\geq \varepsilon\, \underbrace{G_1\left(x_0, \left\{x_m : m \geq \frac{1}{\varepsilon^2}\right\}\right)}_{=\varepsilon^{\frac{1}{2}} \text{ by (B.2).}} \underbrace{(1 - \varepsilon^2)^{\frac{1}{2\varepsilon^2}}}_{\substack{\geq 1/2 \text{by (B.1)} \\ \text{and (B.3).}}} \frac{1}{2\varepsilon^2} \\ &\geq \varepsilon\, \varepsilon^{\frac{1}{2}} \frac{1}{2} \frac{1}{2\varepsilon^2} = \frac{1}{4\varepsilon^{\frac{1}{2}}}. \end{aligned}$$

This is true for arbitrarily small ε, and $1/4\varepsilon^{\frac{1}{2}} \to \infty$ as $\varepsilon \to 0$. Therefore, $\pi_\varepsilon(X \setminus \{x_0\}) \not\to 0$ as $\varepsilon \to 0$. Proposition 3 does not hold.

B.2. The role of Assumption 2

Here we present a counterexample when Assumption 2 does not hold. Let $X = \{x_k^*, x_j^*\} \cup \{x_m\}_{m \in \mathbb{N}_+}$ and let X be equipped with the discrete metric. In accordance with previous notation, let $\Lambda = \{x_k^*, x_j^*\}$. Let

$$P(x_m, x_k^*) = \frac{1}{m + 1}, \qquad P(x_m, x_{m+1}) = \frac{m}{m + 1}, \qquad m \in \mathbb{N}_+,$$

[26] Meyn and Tweedie (2009, Theorem 10.4.9).

$$G_1(x_k^*, x_1) = 1, \qquad G_1(x_j^*, x_j^*) = 1,$$
$$G_1(x_m, x_j^*) = 1, \quad m \in \mathbb{N}_+,$$
$$G_2(x_k^*) = 1.$$

Note that for any $x \in X$, $P^t(x, \Lambda) \to 1$ as $t \to \infty$, but for any $x \notin \Lambda$, there does not exist a t such that $P^t(x, \Lambda) = 1$. Hence Assumption 2 is violated. Observe that $P_\varepsilon^t(x_k^*, x_j^*) \in \Theta(\varepsilon^2)$ for all $t \in \mathbb{N}_+$. As for small enough δ, $B_\delta(x_k^*) = \{x_k^*\}$, we have that $V(x_k^*, x_j^*) = V^-(x_k^*, x_j^*) = 2$. Similarly, $V(x_j^*, x_k^*) = V^-(x_j^*, x_k^*) = 2$, so (C1) holds. Noting that $P_\varepsilon^2(x_k^*, x_j^*) = P_\varepsilon^1(x_j^*, x_k^*) = \varepsilon^2$, we see that (C2) also holds, so Property C holds. However,

$$\begin{aligned}\tilde{P}_\varepsilon(x_k^*, x_j^*) &= \varepsilon\left(\varepsilon + (1-\varepsilon)\frac{1}{2}\varepsilon \right.\\ &\quad \left. + (1-\varepsilon)^2\frac{1}{2}\frac{2}{3}\varepsilon + (1-\varepsilon)^3\frac{1}{2}\frac{2}{3}\frac{3}{4}\varepsilon + \cdots\right)\\ &= \varepsilon^2\sum_{m=0}^{\infty}\frac{1}{m+1}(1-\varepsilon)^m = \varepsilon^2\frac{(-\log\varepsilon)}{1-\varepsilon},\end{aligned}$$

so for any $l \in \mathbb{R}_{++}$, there exists $\bar{\varepsilon}$ such that for all $\varepsilon < \bar{\varepsilon}$,

$$\tilde{P}_\varepsilon(x_k^*, x_j^*) = \varepsilon^2\frac{(-\log\varepsilon)}{1-\varepsilon} > \varepsilon^2(-\log\varepsilon) > l\varepsilon^2.$$

This contradicts Lemma 6.

B.3. The role of Assumption 5

Here we present a counterexample when Assumption 5 does not hold. Let $X = \{x_k^*, x_j^*\} \cup \{x_m\}_{m\in\mathbb{N}_+}$ and let X be equipped with the discrete metric. Let $\Lambda = \{x_k^*, x_j^*\}$. Let

$$P(x_1, x_k^*) = 1, \qquad P(x_{m+1}, x_m) = 1, \quad m \in \mathbb{N}_+,$$
$$G_1(x_j^*, x_j^*) = 1, \qquad G_1(x_m, x_j^*) = 1, \quad m \in \mathbb{N}_+,$$
$$G_1(x_k^*, \{x_m : m \geq \bar{m}\}) = \left(\frac{1}{\bar{m}}\right)^{\frac{1}{2}}$$
$$G_2(x_k^*) = 1.$$

Note that for any $x_m \in X$, $P^m(x_m, \Lambda) = 1$ so Assumption 2 holds. In a similar manner to the example in Appendix B.2 we see that $P_\varepsilon^t(x_k^*, x_j^*) \in \Theta(\varepsilon^2)$ for all $t \in \mathbb{N}_+$, $V(x_k^*, x_j^*) = V^-(x_k^*, x_j^*) = 2$ and that Property C is satisfied. For small δ, $W^\delta = \Lambda$, so

$$F_\delta(t) = \{x \in X : \exists t' \geq t \text{ s.t. } P^{t'}(x, \Lambda) < 1\} = \{x_m : m > t\}$$

and as $P_\varepsilon(x_k^*, \{x_m : m > t\}) \in \Omega(\varepsilon)$ for all t, Assumption 5 does not hold. Note that

$$\text{If } \frac{\log\frac{3}{4}}{\log(1-\varepsilon^2)} \geq m \geq \left\lceil\frac{\log\frac{1}{4}}{\log(1-\varepsilon-\varepsilon^2)}\right\rceil,$$
$$\text{then } (1-\varepsilon^2)^m \geq \frac{3}{4} \quad\text{and}\quad (1-\varepsilon-\varepsilon^2)^m \leq \frac{1}{4},$$

implying that for $m \geq \left\lceil\frac{\log\frac{1}{4}}{\log(1-\varepsilon-\varepsilon^2)}\right\rceil$,

$$\begin{aligned}&Pr_{x_m}(\Phi_\varepsilon^t = x_j^* \quad \text{for some } t \leq m;\ \Phi_\varepsilon^t \neq x_k^* \text{ for } t = 1,\ldots,m)\\ &\geq 1 - \underbrace{(1-(1-\varepsilon^2)^m)}_{\substack{\text{Prob at least}\\ \text{one } G_2(.)\text{event}}} - \underbrace{(1-\varepsilon-\varepsilon^2)^m}_{\substack{\text{Prob no } G_1(\ldots)\\ \text{or } G_2(.) \text{ events}}}\\ &\geq 1 - \left(1-\frac{3}{4}\right) - \left(\frac{1}{4}\right) = \frac{1}{2}. \qquad \text{(B.4)}\end{aligned}$$

Now,

$$\begin{aligned}&G_1\left(x_k^*, \left\{x_m : m \geq \left\lceil\frac{\log\frac{1}{4}}{\log(1-\varepsilon-\varepsilon^2)}\right\rceil\right\}\right)\\ &= \left(\frac{1}{\left\lceil\frac{\log\frac{1}{4}}{\log(1-\varepsilon-\varepsilon^2)}\right\rceil}\right)^{\frac{1}{2}} > \left(\frac{1}{\frac{\log\frac{1}{4}}{\log(1-\varepsilon-\varepsilon^2)}+1}\right)^{\frac{1}{2}}\\ &= \left(\frac{\log(1-\varepsilon-\varepsilon^2)}{\log\frac{1}{4}+\log(1-\varepsilon-\varepsilon^2)}\right)^{\frac{1}{2}} > \left(\frac{\log(1-\varepsilon-\varepsilon^2)}{2\log\frac{1}{4}}\right)^{\frac{1}{2}}\\ &> \left(\frac{-\varepsilon}{2\log\frac{1}{4}}\right)^{\frac{1}{2}} = \left(-2\log\frac{1}{4}\right)^{-\frac{1}{2}}\varepsilon^{\frac{1}{2}}. \qquad \text{(B.5)}\end{aligned}$$

So, using (B.4) and (B.5),

$$\begin{aligned}\tilde{P}_\varepsilon(x_k^*, x_j^*) &\geq \varepsilon\, G_1\left(x_k^*, \left\{x_m : m \geq \left\lceil\frac{\log\frac{1}{4}}{\log(1-\varepsilon-\varepsilon^2)}\right\rceil\right\}\right)\frac{1}{2}\\ &> \varepsilon\left(-2\log\frac{1}{4}\right)^{-\frac{1}{2}}\varepsilon^{\frac{1}{2}}\frac{1}{2} = \frac{1}{2}\left(-2\log\frac{1}{4}\right)^{-\frac{1}{2}}\varepsilon^{\frac{3}{2}}.\end{aligned}$$

This contradicts Lemma 6.

References

Agastya, M., 1999. Perturbed adaptive dynamics in coalition form games. J. Econom. Theory 89, 207–233.
Beggs, A., 2005. Waiting times and equilibrium selection. Econom. Theory 25, 599–628.
Bergin, J., Lipman, B.L., 1996. Evolution with state-dependent mutations. Econometrica 64, 943–956.
Binmore, K., Samuelson, L., Young, P., 2003. Equilibrium selection in bargaining models. Games Econom. Behav. 45, 296–328.
Blume, L.E., 1993. The statistical mechanics of strategic interaction. Games Econom. Behav. 5, 387–424.
Chong, J.K., Camerer, C.F., Ho, T.H., 2006. A learning-based model of repeated games with incomplete information. Games Econom. Behav. 55, 340–371.
Ellison, G., 2000. Basins of attraction, long-run stochastic stability, and the speed of step-by-step evolution. Rev. Econom. Stud. 67, 17–45.
Eshel, I., Samuelson, L., Shaked, A., 1998. Altruists, egoists, and hooligans in a local interaction model. Amer. Econ. Rev. 88, 157–179.
Feinberg, Y., 2006. Evolutionary dynamics and long-run selection. B. E. J. Theor. Econ. 6.
Foster, D., Young, H.P., 1990. Stochastic evolutionary game dynamics. Theor. Popul. Biol. 38, 219–232.
Freidlin, M.I., Wentzell, A.D., 1984. Random Perturbations of Dynamical Systems. Springer, ISBN: 9780387983622, p. 430, second ed. (1998).
Fremlin, D.H., 2001. Measure Theory: Broad Foundations. Torres Fremlin, ISBN: 978-0953812929.
Kandori, M., Mailath, G.J., Rob, R., 1993. Learning, mutation, and long run equilibria in games. Econometrica 61, 29–56.
Kifer, Y., 1988. Random Perturbations of Dynamical Systems. Birkhäuser, ISBN: 9780817633844, p. 294.
Kifer, Y., 1990. A discrete-time version of the Wentzell–Friedlin theory. Ann. Probab. 18, 1676–1692.
Meyn, S.P., Tweedie, R.L., 1994. State-dependent criteria for convergence of Markov chains. Ann. Appl. Probab. 4, 149–168.
Meyn, S.P., Tweedie, R.L., 2009. Markov Chains and Stochastic Stability. Springer-Verlag, London, New York.
Naidu, S., Hwang, S.H., Bowles, S., 2010. Evolutionary bargaining with intentional idiosyncratic play. Econom. Lett. 109, 31–33.
Newton, J., 2012. Coalitional stochastic stability. Games Econom. Behav. 75, 842–854.
Oechssler, J., Riedel, F., 2001. Evolutionary dynamics on infinite strategy spaces. Econom. Theory 17, 141–162.
Schenk-Hoppé, K.R., 2000. The evolution of Walrasian behavior in oligopolies. J. Math. Econom. 33, 35–55.
Selten, R., Apesteguia, J., 2005. Experimentally observed imitation and cooperation in price competition on the circle. Games Econom. Behav. 51, 171–192.
Serrano, R., Volij, O., 2008. Mistakes in cooperation: the stochastic stability of edgeworth's recontracting. Econ. J. 118, 1719–1741.
Smith, J.M., Price, G.R., 1973. The logic of animal conflict. Nature 15–18.
van Damme, E., Weibull, J.W., 2002. Evolution in games with endogenous mistake probabilities. J. Econom. Theory 106, 296–315.
Young, H.P., 1993a. The evolution of conventions. Econometrica 61, 57–84.
Young, H.P., 1993b. An evolutionary model of bargaining. J. Econom. Theory 59, 145–168.
Young, H.P., 1998. Individual Strategy and Social Structure. Princeton University Press.