

Stability of Linear Stochastic Difference Equations in Strategically Controlled Random Environments

ULRICH HORST*

Institut für Mathematik – Bereich Stochastik

Humboldt-Universität zu Berlin

Unter den Linden 6

D-10099 Berlin

E-mail:horst@mathematik.hu-berlin.de

19th March 2003

Abstract

We consider the stochastic sequence $\{Y_t\}_{t \in \mathbb{N}}$ defined recursively by the linear relation $Y_{t+1} = A_t Y_t + B_t$ in a random environment. The environment is described by the stochastic process $\{(A_t, B_t)\}_{t \in \mathbb{N}}$ and is under the simultaneous control of several agents playing a discounted stochastic game. We formulate sufficient conditions on the game which ensure the existence of Nash equilibrium in Markov strategies which has the additional property that, in equilibrium, the process $\{Y_t\}_{t \in \mathbb{N}}$ converges in distribution to a stationary regime.

AMS 2000 subject classification. 60G35, 93E15, 91A15

Key words and phrases. Stochastic difference equation, stochastic stability, stochastic games, random systems with complete connections.

*This work was initiated while I was visiting Bendheim Center for Finance, Princeton University. Grateful acknowledgement is made for hospitality. I would like to thank Hans Föllmer, Peter Imkeller, and José Scheinkman for helpful comments and valuable discussions. Financial support through the German Academic Exchange Service (DAAD), the DFG research center “Mathematics for key technologies” (FZT 86) and through the NSF grant No. 0001647 to Princeton University is gratefully acknowledged.

1 Introduction

This paper considers the stochastic sequence $\{Y_t\}_{t \in \mathbb{N}}$ defined recursively by the linear relation

$$Y_{t+1} = A_t Y_t + B_t \quad (t \in \mathbb{N}) \quad (1)$$

in the random environment $\{(A_t, B_t)\}_{t \in \mathbb{N}}$. The dynamics of the environment is under the simultaneous control of several agents who play a discounted stochastic game. We formulate sufficient conditions on the game which guarantee the existence of Nash equilibria in Markov strategies which have the additional property that, in equilibrium, the solution to (1) converges in distribution as $t \rightarrow \infty$.

Stochastic sequences of the form (1) have been extensively investigated under a mean contraction condition and under the assumption that the driving sequence $\{(A_t, B_t)\}_{t \in \mathbb{N}}$ defined on some probability space $(\Omega, \mathcal{F}, \mathbb{P})$ is stationary under the *exogenous* measure \mathbb{P} . For example, Vervaat (1979) considers the case where the environment consists of i.i.d. random variables. Brandt (1986) assumes that the driving sequence is stationary and ergodic under the law \mathbb{P} ; see also Borovkov (1998). Horst (2001) studies the case where the environment is asymptotically stationary in the sense that the process $\{(A_t, B_t)\}_{t \in \mathbb{N}}$ is stationary under some law \mathbb{Q} on (Ω, \mathcal{F}) which coincides with \mathbb{P} on the tail-field generated by $\{(A_t, B_t)\}_{t \in \mathbb{N}}$.

In view of many applications in economics it seems natural to extend the analysis of stochastic difference equations to situations in which the environment is under the simultaneous control of several agents. In such a situation the measure \mathbb{P} is no longer exogenous, but is derived through a game-theoretic solution concept. For example, the process $\{Y_t\}_{t \in \mathbb{N}}$ could be sequence of temporary equilibrium prices of a risky asset generated by the microeconomic interaction of investors who are active on a financial market. In such a model, the sequence $\{(A_t, B_t)\}_{t \in \mathbb{N}}$ may be driven by the changes in the behavioral characteristics of the agents. If the investors are ‘large’, i.e., if their behavior influences the dynamics of the random environment, then it seems natural to assume that the agents anticipate their impact on the formation of stock prices and, therefore, interact in a strategic manner. In this context, the process $\{Y_t\}_{t \in \mathbb{N}}$ may be viewed as the state sequence associated to a stochastic game, and so the probabilistic structure of the random environment specified by the measure \mathbb{P} is no longer exogenous. Instead, it is defined through an equilibrium strategy implemented by the individual investors. In order to analyze the dynamics of the price process in equilibrium, it is now desirable to have sufficient conditions which guarantee the existence of Nash equilibria

which ensure that the sequence $\{Y_t\}_{t \in \mathbb{N}}$ settles down in the long run.

In this paper we consider discounted stochastic games with weakly interacting players in which the conditional distribution of the random variable (A_t, B_t) only depends on the average action taken by the players. We formulate conditions on the game which guarantee the existence of a Nash equilibrium in Markov strategies τ such that under the induced measure \mathbb{P}^τ on (Ω, \mathcal{F}) the solution to (1) converges as $t \rightarrow \infty$. In a first step we show that stochastic games in which the state sequence follows a linear dynamics and in which the interaction between different agents is sufficiently weak have stationary equilibria in Markov strategies that depend in a Lipschitz continuous manner on the current state. This part of the paper is inspired by the work of Curtat (1996) and uses a perturbation of a *Moderate Social Influence* condition introduced in Horst and Scheinkman (2002). Under a suitable mean contraction condition on the random environment we then prove that the sequence $\{(A_t, B_t)\}_{t \in \mathbb{N}}$ defined on $(\Omega, \mathcal{F}, \mathbb{P}^\tau)$ has a nice tail structure in the sense of Horst (2001). This allows us to show that the shifted sequence $\{Y_{t+T}\}_{t \in \mathbb{N}}$ converges in law to a uniquely determined stationary process as $T \rightarrow \infty$.

The remainder of this paper is organized as follows. In Section 2 we formulate our main results. Section 3 proves the existence of a Lipschitz continuous equilibrium in Markov strategies. The convergence result for the solution to (1) is proved in Section 4.

2 Assumptions and the main results

Let $\psi := \{(A_t, B_t)\}_{t \in \mathbb{N}}$ be a sequence of \mathbb{R}^2 -valued random variables defined on some probability space $(\Omega, \mathcal{F}, \mathbb{P})$, and let $\{Y_t\}_{t \in \mathbb{N}}$ be the sequence in (1) driven by the “input” ψ . In this section we specify a probabilistic framework which allows us to analyze the asymptotic behavior of the solution to the linear stochastic difference equation (1) in a situation in which the evolution of the random environment is controlled by several strategically interacting agents.

The long run behavior of the sequence $\{Y_t\}_{t \in \mathbb{N}}$ has been intensively investigated under a mean contraction condition and under the assumption that ψ is stationary under some *exogenous* measure \mathbb{P} ; see, e.g., Brandt (1986) or Vervaat (1979). Horst (2001) assumes that the environment is asymptotically stationary and that it has a nice tail structure in the sense of the following definition.

Definition 2.1 (Horst (2001)) Let $\hat{\mathcal{F}}_t := \sigma(\{(A_s, B_s)\}_{s \geq t})$ and let

$$\mathcal{I}_\psi := \bigcap_{t \in \mathbb{N}} \hat{\mathcal{F}}_t \quad (2)$$

be the tail- σ -algebra generated by ψ . A driving sequence ψ is called nice with respect to a probability measure \mathbb{Q} on (Ω, \mathcal{F}) if the following properties are satisfied:

(i) ψ is stationary and ergodic under \mathbb{Q} and satisfies

$$\mathbb{E}_{\mathbb{Q}} \ln |A_0| < 0 \quad \text{and} \quad \mathbb{E}_{\mathbb{Q}} (\ln |B_0|)^+ < \infty \quad (3)$$

where $\mathbb{E}_{\mathbb{Q}}$ denotes the expectation with respect to the measure \mathbb{Q} .

(ii) The asymptotic behavior of ψ is the same under \mathbb{P} and \mathbb{Q} , i.e.,

$$\mathbb{P} = \mathbb{Q} \quad \text{on} \quad \mathcal{I}_\psi. \quad (4)$$

Remark 2.2 We denote by $\|\cdot\|_{\mathcal{E}}$ the total variation of a signed measure on a measurable space (E, \mathcal{E}) . Since

$$\lim_{t \rightarrow \infty} \|\mathbb{P} - \mathbb{Q}\|_{\hat{\mathcal{F}}_t} = \|\mathbb{P} - \mathbb{Q}\|_{\mathcal{I}_\psi}, \quad (5)$$

a driving sequence ψ satisfies (4) if and only if it converges to a stationary regime. Horst (2001) shows that (5) holds if, for example, ψ is driven by an underlying Markov chain that converges in total variation norm to a stationary distribution or if ψ coupling converges to a stationary sequence in the sense of Borovkov (1998). Note that (5) is equivalent to the existence of a sequence $\{c_t\}_{t \in \mathbb{N}}$ satisfying $\lim_{t \rightarrow \infty} c_t = 0$ and

$$\sup_{l \geq t} \|\mathbb{P} - \mathbb{Q}\|_{\hat{\mathcal{F}}_{t,l}} \leq c_t, \quad (6)$$

where $\hat{\mathcal{F}}_{t,l} := \sigma(\{(A_s, B_s)\}_{t \leq s \leq l})$. Here, both (5) and (6) follow from the continuity of the total variation distance along increasing and decreasing σ -algebras.

In the sequel it will be convenient to denote by $\text{Law}(Y, \mathbb{P})$ the law of a random variable Y on $(\Omega, \mathcal{F}, \mathbb{P})$ and to write \xrightarrow{w} for weak convergence of probability measures.

Let us turn to the solution $\{Y_t\}_{t \in \mathbb{N}}$ of (1). For any initial value $Y_0 = y \in \mathbb{R}$, we have the explicit representation

$$Y_t = y_t(y, \psi) := \sum_{j=0}^{t-1} \left(\prod_{i=t-j}^{t-1} A_i \right) B_{t-j-1} + \left(\prod_{i=0}^{t-1} A_i \right) y \quad (t \in \mathbb{N}). \quad (7)$$

In the stationary setting $\mathbb{P} = \mathbb{Q}$ analyzed by Brandt (1986), we may as well assume that the driving sequence is defined for all $t \in \mathbb{Z}$, due to Kolmogorov's extension theorem. Under the *mean contraction condition* (3), there exists a unique stationary solution to (1) under \mathbb{Q} driven by ψ . That is, there is a unique stationary process $\{Y_t^*\}_{t \in \mathbb{Z}}$ which satisfies the recursive relation (1) for all $t \in \mathbb{Z}$. The random variable Y_t^* is \mathbb{Q} -a.s. finite, takes the form

$$Y_t^* = \sum_{j=0}^{\infty} \left(\prod_{i=t-j}^{t-1} A_i \right) B_{t-j-1} \quad (t \in \mathbb{Z}), \quad (8)$$

and, for any initial value $y \in \mathbb{R}$ the solution $\{y_t(y, \psi)\}_{t \in \mathbb{N}}$ to (1) converges almost surely to the stationary solution in the sense that

$$\lim_{t \rightarrow \infty} |y_t(y, \psi) - Y_t^*| = 0 \quad \mathbb{Q}\text{-a.s.}$$

In the non-stationary case $\mathbb{P} \neq \mathbb{Q}$ studied in Horst (2001) the shifted sequence $\{y_{t+T}(y, \psi)\}_{t \in \mathbb{N}}$ driven by a nice input ψ converges in distribution to the unique stationary solution $\{Y_t^*\}_{t \in \mathbb{N}}$ to (1) under \mathbb{Q} :

$$\text{Law}(\{y_{t+T}(y, \psi)\}_{t \in \mathbb{N}}, \mathbb{P}) \xrightarrow{w} \text{Law}(\{Y_t^*\}, \mathbb{Q}) \quad (T \rightarrow \infty). \quad (9)$$

So far, the asymptotics of the sequence $\{Y_t\}_{t \in \mathbb{N}}$ have only been analyzed in situations where the probabilistic structure of the random environment ψ is described by an *exogenous* measure. Our aim is study to dynamics of the solution to the linear stochastic difference equation (1) in a situation in which the evolution of ψ is controlled by strategically interacting agents who play a stochastic game.

2.1 Stock price fluctuations in the presence of big players

In this section we discuss an extension of the temporary equilibrium model for asset price fluctuations by Föllmer and Schweizer (1993) where the class of stochastic games which we analyze in this paper arises rather naturally. We consider a financial market with a finite set $\mathbb{A} = \{a_1, \dots, a_N\}$ of *small* investors trading a single risky stock. In reaction to a proposed price p in period $t \in \mathbb{N}$ the agent $a \in \mathbb{A}$ forms a random excess demand $e_t^a(p, \omega)$ and the actual stock price P_t is determined implicitly by the market clearing condition, $\sum_{a \in \mathbb{A}} e_t^a(P_t, \omega) = 0$, of zero total excess demand. Following Föllmer and Schweizer (1993) we assume that the agents excess demand of the agent $a \in \mathbb{A}$ takes the log-linear form

$$e_t^a(p, \omega) = \log \hat{S}_{t+1}^a(\omega) - \log p + \eta_t^a(\omega) \quad (10)$$

where $\hat{S}_{t+1}^a(\omega)$ denotes his reference level for the following period $t + 1$ and where η_t^a may be viewed as his liquidity demand.

To simplify our analysis we assume that all the agents are fundamentalists. This means that an agents' price expectation for the following period $t + 1$ is entirely based on the idea that the stock price will move closer to his current subjective perception F_t^a of the stocks' fundamental value, and so each agent chooses a reference level of the form

$$\log \hat{S}_{t+1}^a = \log P_t + \beta_t(F_t^a - \log P_t)$$

with independent and identically distributed random coefficients $\beta_t \in (0, 1)$; see, e.g., Föllmer and Schweizer (1993) or Kirman (1998) for further details. Specifically, we assume that the agents are either optimistic or pessimistic. This means that the agents either associate a low ($F^l = 0$) or a high ($F^h = 1$) fundamental value to the stock: $F_t^a \in \{0, 1\}$.

Let $\pi_t := \frac{1}{|\mathbb{A}|} \sum_{a \in \mathbb{A}} \mathbf{1}_{\{F_t^a=1\}}$ be the fraction of optimistic agents at time t . With $\eta_t := \frac{1}{N} \sum_{a \in \mathbb{A}} \eta_t^a$, the logarithmic equilibrium price $S_t := \log P_t$ is determined via (10) as an average,

$$S_{t+1} = (1 - \beta_t) \log S_t + \beta_t \pi_t + \eta_t =: A_t S_t + B_t, \quad (11)$$

of individual price assessments and liquidity demands. We assume that the perceptions of fundamental values are independent across agents, that the average liquidity demand is $N(0, 1)$ distributed and that both the assessments of fundamentals and liquidity demands evolve independently across time. In this case $\text{Law}(S_t, \mathbb{P}) \xrightarrow{w} \mu$ for some law μ on \mathbb{R} , because $\beta_t \in (0, 1)$. Hence, the asymptotic distribution of stock prices is uniquely determined.

Let us now consider an extension of the model in which the agents choose their current benchmarks in reaction to the actions $\tau_t = (\tau_t^i)_{i=1}^M \in [0, 1]^M$ taken by some “big players”:

$$F_t^a \sim Z(\tau_t; \cdot) \quad \text{for a stochastic kernel } Z \text{ from } [0, 1]^M \text{ to } \{0, 1\}.$$

One may, for example, think of a central bank that tries to keep the “mood of the market” from becoming too optimistic and, if necessary, warns the market participants of emerging bubbles. One may also think of financial experts whose buy or sell recommendations tempt the agents into buying or selling the stock. These market participants influence the stock price process through their impact on the behavior of small investors and it seems natural to assume that these big players anticipate their impact on stock prices and thus interact in a strategic manner.

If the big players use homogeneous Markov strategies, i.e., if the actual action $\tau_t^i = \tau^i(S_t)$ of the player $i \in \{1, 2, \dots, M\}$ depends only on the current stock price, then the evolution of the of temporary logarithmic price equilibria can be described by a recursive relation of the form

$$\log S_{t+1} = A_t S_t + B_t \quad \text{where} \quad (A_t, B_t) \sim Q(\tau(S_t); \cdot), \quad (12)$$

where Q is a suitable stochastic kernel from \mathbb{R} to $(0, 1) \times \mathbb{R}$. Note that the random environment for the evolution of asset prices is longer described by an i.i.d. sequence of random variables. Instead, the mood of the market and thus the dynamics of the stock price process is under the simultaneous control of several big players with possibly conflicting interests. It is now desirable to have sufficient conditions which guarantee the existence of an equilibrium τ^* which preserves asymptotic stability of stock prices in the sense that

$$\text{Law}(S_t, \mathbb{P}^{\tau^*}) \xrightarrow{w} \mu^{\tau^*}.$$

Observe, however, that the asymptotics of the stock price process depend on the implemented equilibrium strategy and are thus no longer uniquely determined. In this sense, the presence of strategically interacting players on a financial market can be an additional source of uncertainty.

2.2 The stochastic game

The infinite-horizon discounted stochastic games $\Sigma = (I, X, (U^i), \beta, Q, y)$ which we consider in this paper are defined in terms of the following objects:

- $I = \{1, 2, \dots, M\}$ is a finite set of *players*.
- $X \subset \mathbb{R}$ is a common compact and convex *action space* for the players.
- $U^i : \mathbb{R} \times \prod_{i \in I} X \rightarrow \mathbb{R}$ is the *utility function* for player $i \in I$.
- $\beta \in (0, 1)$ is a common discount factor.
- Q is a stochastic kernel from X to \mathbb{R}^2 .
- $y \in \mathbb{R}$ is the starting point of the state sequence $\{Y_t\}_{t \in \mathbb{N}}$.

A typical action of player $i \in I$ is denoted $x^i \in X$. The actions taken by his competitors are denoted $x^{-i} \in X^{-i} := \{x^{-i} = (x^j)_{j \in I \setminus \{i\}}\}$, and $\bar{X} := \{\bar{x} = (x^i)_{i \in I} : x^i \in X\}$ is the compact set of all *action profiles*. To each action profile $\bar{x} \in \bar{X}$, we associate the average action $x := \frac{1}{M} \sum_{i \in I} x^i$.

Remark 2.3 *The assumption that all the players share a common discount factor is made merely to ease notational complexity. This condition can be dropped without altering the models' qualitative features.*

At each time $t \in \mathbb{N}$, the players observe the current position Y_t of the state sequence $\{Y_t\}_{t \in \mathbb{N}}$. They take their actions $x_t^i = \tau^i(Y_t)$ independently of each other according to a *stationary Markov strategy* $\tau^i : \mathbb{R} \rightarrow X$ and the selected action profile $\bar{x}_t = (x_t^i)_{i \in I} = (\tau^i(Y_t))_{i \in I}$ along with the present state Y_t yields the instantaneous payoff $U^i(Y_t, \bar{x}_t) = U^i(Y_t, x_t^i, x_t^{-i})$ to the agent $i \in I$. We assume that the law of motion only depends on the average action taken by the individual players. More precisely,

$$Y_{t+1} = A_t Y_t + B_t \text{ with } (A_t, B_t) \sim Q(cx_t; \cdot) \text{ for some } c > 0, \text{ where } x_t := \frac{1}{M} \sum_{i=1}^M x_t^i. \quad (13)$$

In this sense we assume that the indirect interaction between different agents, i.e., the interaction through the state sequence $\{Y_t\}_{t \in \mathbb{N}}$, is *global*. Thus, in games with many players or in games with a small $c > 0$, the impact of an individual agent on the dynamics of the random environment $\{(A_t, B_t)\}_{t \in \mathbb{N}}$ for the evolution of the state sequence $\{Y_t\}_{t \in \mathbb{N}}$ is weak.

Remark 2.4 *For our subsequent analysis it will be essential that the impact of an individual player on the law of the random environment is sufficiently weak. Assuming that the conditional distribution of the random variable (A_t, B_t) depends on the current action profile \bar{x}_t only through $cx_t = \frac{c}{M} \sum_{i \in I} x_t^i$ simplifies the formulation of an appropriate weak interaction condition. Of course, weak dependence assumptions can also be formulated differently.*

A stationary Markov strategy $\tau = (\tau^i)_{i \in I}$ along with an initial distribution μ for the starting point of the state sequence and together with the law of motion Q induces a probability measure \mathbb{P}_μ^τ on (Ω, \mathcal{F}) in the canonical way. Under the measure \mathbb{P}_μ^τ the state sequence $\{Y_t\}_{t \in \mathbb{N}}$ is a Markov chain on the state space \mathbb{R} and the *expected discounted reward* to player $i \in I$ is given by

$$J^i(y, \tau) = J^i(y, \tau^i, \tau^{-i}) := \mathbb{E}_\mu^\tau \left[\sum_{t=0}^{\infty} \beta^t U^i(Y_t, \bar{x}_t) \right]. \quad (14)$$

Here the expectation is taken with respect to the measure \mathbb{P}_μ^τ and $\tau^{-i} = (\tau^j)_{j \neq i}$. In what follows we shall write \mathbb{P}_y^τ for $\mathbb{P}_{\delta_y}^\tau$.

Definition 2.5 *A stationary Markov strategy profile τ is a Nash equilibrium for Σ if no player can increase his payoff by unilateral deviation from τ , i.e., if*

$$J^i(y, \tau) \geq J^i(y, \sigma^i, \tau^{-i}) \quad (15)$$

for all Markov strategies $\sigma^i : \mathbb{R} \rightarrow X$ and each $i \in I$.

Our objective is to formulate conditions which guarantee the existence of a Nash equilibrium in Markov strategies τ such that the Markov chain $(\{Y_t\}, \mathbb{P}_y^\tau)$ converges in law to a unique limiting distribution. To this end, we need to assume strong concavity of an agents' utility function with respect to his own action, and we have to place a quantitative bound on the dependence of the instantaneous utility for player $i \in I$ on the actions taken by his competitors.

Assumption 2.6 (i) *Uniformly in $i \in I$, the utility functions $U^i : \mathbb{R} \times \overline{X} \rightarrow \mathbb{R}$ are bounded, Lipschitz continuous and twice continuously differentiable.*

(ii) *There exists constants $L^{i,i}(y) > 0$ and $L > 0$ such that*

$$\frac{\partial^2}{\partial (x^i)^2} U^i(y, x^i, x^{-i}) \leq -L^{i,i}(y) \leq -L < 0. \quad (16)$$

In particular, the function $U^i(y, \cdot, x^{-i})$ is strongly concave on X .

(iii) *The law $Q(x; \cdot)$ has a density $q(x, \cdot)$ with respect to the Lebesgue measure λ^2 on \mathbb{R}^2 . The maps $x \mapsto q(x, \eta)$ are uniformly bounded, twice continuously differentiable on an open set containing X and uniformly Lipschitz continuous:*

$$|q(x_1, \eta) - q(x_2, \eta)| \leq \bar{L}|x_1 - x_2| \quad (\bar{L} < \infty).$$

Moreover, $|\frac{\partial}{\partial x} q(x, \eta)| \leq \varrho_1(\eta)$ and $|\frac{\partial^2}{\partial x^2} q(x, \eta)| \leq \varrho_2(\eta)$ for some functions $\varrho_1, \varrho_2 : \mathbb{R}^2 \rightarrow \mathbb{R}$ with are integrable with respect to λ^2 .

The Lipschitz continuity condition on the conditional densities $q(x; \cdot)$ translates into a norm-continuity assumption on the transition probabilities $Q(x; \cdot)$:

$$\|Q(x_1; \cdot) - Q(x_2; \cdot)\|_{\mathcal{B}} \leq \bar{L}|x_1 - x_2| \quad \text{and so} \quad \|Q(x_n; \cdot) - Q(x; \cdot)\|_{\mathcal{B}} \xrightarrow{n \rightarrow \infty} 0 \quad \text{if} \quad x_n \xrightarrow{n \rightarrow \infty} x,$$

where \mathcal{B} denotes the Borel- σ -field on \mathbb{R}^2 . Norm-continuity conditions have also been imposed by, e.g., Nowak (1985) and Duffie, Geanakoplos, MasColell, and McLennan (1994).

Let us consider an example where our assumptions on the densities $q(x; \cdot)$ can indeed be verified.

Example 2.7 *Let $f : X \rightarrow \mathbb{R}^2$ be a twice continuously differentiable function and let $\varphi_m(\cdot)$ be the density of the two-dimensional standard normal distribution with mean $m = (m_1, m_2)$ with respect to λ^2 . It is easily seen that $q(x, \eta) := \varphi_{f(x)}(\eta)$ satisfies Assumption 2.6 (iii).*

In order to establish the existence of Lipschitz continuous Nash equilibria in the stochastic game under consideration, we have to control the strength of interactions between different players. To this end, we introduce the constants

$$L^{i,j}(y) = \sup_{\bar{x} \in \bar{X}} \left| \frac{\partial^2}{\partial x_i \partial x_j} U^i(y, \bar{x}) \right| \quad (i \neq j) \quad \text{and} \quad L^i = \sup_{y \in \mathbb{R}, \bar{x} \in \bar{X}} \left| \frac{\partial^2}{\partial x_i \partial y} U^i(y, \bar{x}) \right|. \quad (17)$$

The quantity $L^{i,j}(y)$ may be viewed as a measure for the dependence of agent i 's instantaneous utility on the choice of player j , given the current state y . By analogy, L^i measures the dependence of his one-period utility on the current position of the state sequence. We put

$$\hat{L}(M) := \sup_{x \in X} \frac{c^2}{M^2} \left\| \frac{\partial^2}{\partial x^2} q(x, \cdot) \right\|_{L^1} \quad (18)$$

where $\|\cdot\|_{L^1}$ denotes the L^1 -norm with respect to λ^2 . In order to guarantee the existence of Lipschitz continuous equilibria in the discounted stochastic game Σ we need to assume that the interaction between different agents is not too strong. Since the players interact both through their instantaneous utility functions and through their individual impacts on the evolution of the state sequence, we have to control both the dependence of an action of player j on the instantaneous utility of the agent i and the dependence of the law of motion on the actions taken by an individual player. We formulate this condition in terms of a perturbation of the *Moderate Social Influence* condition in Horst and Scheinkman (2002).

Assumption 2.8 *There exists $\gamma < 1$ such that the following holds for all $i \in I$ and $x \in X$:*

$$\sum_{j \neq i} L^{i,j}(y) + \frac{\beta c^2 \|U^i\|_\infty}{M(1-\beta)} \left\| \frac{\partial^2}{\partial x^2} q(x, \cdot) \right\|_{L^1} \leq \gamma L^{i,i}(y). \quad (19)$$

Let us consider a simple case study where our Assumption 2.8 can easily be verified.

Example 2.9 Consider a law of motion Q which is of the form

$$Q\left(\frac{c \sum_{i \in I} x_i}{M}; \cdot\right) = \frac{c \sum_{i \in I} x_i}{M} Q_1(\cdot) + \left(1 - \frac{c \sum_{i \in I} x_i}{M}\right) Q_2(\cdot). \quad (20)$$

If $Q_i(\cdot)$ has a density $q_i(\cdot)$ with respect to λ^2 , then our Moderate Social Influence condition translates into an assumption on the marginal rates of substitution. Indeed, if the stochastic kernel Q takes the linear form (20), then $\frac{\partial^2}{\partial x^2} q^i(x; \cdot) = 0$, and (19) is equivalent to

$$\sum_{j \neq i} \left| \frac{\partial^2}{\partial x_i \partial x_j} U^i(y, \bar{x}) \right| \leq \gamma \left| \frac{\partial^2}{\partial x_i^2} U^i(y, \bar{x}) \right| \quad (21)$$

for all $i \in I$ and for some $\gamma < 1$. In addition, (16) requires $\inf_y \left| \frac{\partial^2}{\partial x_i^2} U^i(y, \bar{x}) \right| > 0$.

Observe that (21) is necessary for our Moderate Social Influence condition. On the other hand, if (21) holds, then we can always choose a small enough $c > 0$ or a large enough M such that Assumption 2.8 holds. In this sense, for games in which the impact of an individual agent on the dynamics of the state sequence is sufficiently weak, our condition (19) reduces to the Moderate Social Influence assumption $\sum_{j \neq i} L^{i,j}(y) \leq \gamma L^{i,i}(y)$ introduced in Horst and Scheinkman (2002).

Let us now formulate sufficient conditions for the existence of Lipschitz continuous equilibria in the stochastic game Σ . The following result will be proved in Section 3 below.

Theorem 2.10 Assume that the stochastic game Σ satisfies Assumptions 2.6 and 2.8 and

$$\hat{q} := \sup_{x \in X} \int |A| Q(x; dA, dB) < \infty, \quad q := \sup_{x \in X} \int |A| \left| \frac{\partial^2}{\partial x^2} q(x, A, B) \right| \lambda(dA, dB) < \infty. \quad (22)$$

Then there exists $C^* > 0$ such that, for all $c < C^*$ the game Σ has a stationary equilibrium in Markov strategies τ which is Lipschitz continuous. That is, there exists $L^* < \infty$ such that

$$|\tau^i(y_1) - \tau^i(y_2)| \leq L^* |y_1 - y_2| \quad (i \in I).$$

Remark 2.11 The proof of Theorem 2.10 will show that stochastic games in which the law of motion takes the linear form (20) have Lipschitz continuous equilibria for all $c > 0$.

2.3 Convergence of the state sequence

Let us now return to the solution to the stochastic difference equation (1). For a given stationary Nash equilibrium in Markov strategies τ , we denote by \mathbb{P}_y^τ the law on (Ω, \mathcal{F}) induced by τ and the starting point y on (Ω, \mathcal{F}) . In order to guarantee asymptotic stability of the process $\{Y_t\}_{t \in \mathbb{N}}$ under \mathbb{P}_y^τ , we need to assume that the following condition is satisfied.

Assumption 2.12 (i) The set $M_x := \{\eta \in \mathbb{R}^2 : q(x, \eta) > 0\}$ is convex.

(ii) There exists a constant $r < 1$ such that

$$\sup_x \int |A|Q(x; dA, dB) \leq r \quad \text{and} \quad \sup_x \int |B|Q(x; dA, dB) < \infty. \quad (23)$$

(iii) There is a measure $\tilde{\nu}$ on \mathbb{R}^2 and a constant $\varrho > 0$ such that

$$Q(x; \cdot) \geq \varrho \tilde{\nu}(\cdot) \quad \text{and} \quad \int |A| \tilde{\nu}(dA, dB) \leq r.$$

We are now ready to formulate the main results of this paper. The proofs will be given in Sections 3.2 and 4 below.

Theorem 2.13 Let τ be a stationary Lipschitz continuous equilibrium in Markov strategies for the stochastic game Σ . If Assumption 2.12 is satisfied, then the following holds:

(i) There is a probability measure \mathbb{Q}^τ on (Ω, \mathcal{F}) such that the driving sequence ψ defined on $(\Omega, \mathcal{F}, \mathbb{P}_y^\tau)$ is nice with respect to \mathbb{Q}^τ .

(ii) Under \mathbb{Q}^τ the process $\{Y_t\}_{t \in \mathbb{N}}$ defined by (1) is stationary and ergodic and

$$\text{Law}(\{Y_{t+T}\}_{t \in \mathbb{N}}, \mathbb{P}_y^\tau) \xrightarrow{w} \text{Law}(\{Y_t\}_{t \in \mathbb{N}}, \mathbb{Q}^\tau) \quad \text{for all } y \in \mathbb{R} \text{ as } T \rightarrow \infty.$$

The proof of Theorem 2.13 turns out to be an immediate consequence of the following stability result for Markov chains. Its proof is given in Section 4 below.

Theorem 2.14 Let \tilde{Q} be a stochastic kernel from \mathbb{R} to \mathbb{R}^2 and assume that $\tilde{Q}(y; \cdot)$ has a density $\tilde{q}(y, \cdot)$ with respect to λ^2 that satisfies

$$|\tilde{q}(y_1, \eta) - \tilde{q}(y_2, \eta)| \leq \tilde{L}|y_1 - y_2| \quad (\tilde{L} < \infty). \quad (24)$$

If \tilde{Q} satisfies Assumption 2.12, then the Markov chain $\{\tilde{Y}_t\}_{t \in \mathbb{N}}$ defined by the recursive relation

$$\tilde{Y}_{t+1} = A_t \tilde{Y}_t + B_t \quad \text{where} \quad (A_t, B_t) \sim \tilde{Q}(\tilde{Y}_t; \cdot)$$

has a unique stationary distribution μ^* . Moreover, with $\mathbb{P}^*(\cdot) := \int \mathbb{P}_x(\cdot) \mu^*(dx)$, we have that

$$\text{Law}(\{\tilde{Y}_{t+T}\}_{t \in \mathbb{N}}, \mathbb{P}_y) \xrightarrow{w} \text{Law}(\{\tilde{Y}_t\}_{t \in \mathbb{N}}, \mathbb{P}^*) \quad \text{for all } y \in \mathbb{R} \text{ as } T \rightarrow \infty.$$

Remark 2.15 *Our Theorem 2.14 may be viewed as a version of Theorem 2.1 in Barnsley, Demko, Elton, and Geronimo (1988). Using Theorem 2.1 in Barnsley et al. (1988) one can show that Theorem 2.10 carries over to situations in which the random variables (A_t, B_t) take only finitely many values. In fact, this case is much simpler to analyze and we leave this modification to the reader.*

Example 2.16 *Let us return to the financial market model described in Section 2.2 and assume that only two “big players” are active on the market. Suppose also that the kernel Q in (12) takes a linear form as in (20). If the small investors are optimistic (pessimistic) at least with probability $\delta > 0$, under the assumption that law of the random variable β_0 is equivalent to λ on the convex set $(0, 1)$ and under suitable regularity conditions on the big players’ instantaneous utility functions, Theorem 2.13 yields an equilibrium τ^* such that under \mathbb{P}^{τ^*} the stock price process settles down in the long run.*

We close with section with two more case studies where the assumptions on Q in Theorem 2.13 can indeed be verified.

Example 2.17 (i) *Let us return to the law of motion Q introduced in Example 2.9 below.*

In this case, (22) is clearly satisfied. Assumption 2.12 holds if, for instance, $Q_i(\cdot)$ has a strictly positive density q_i which a convex domain, if $Q_1(\cdot)$ and $Q_2(\cdot)$ satisfy (23) and if each Q_i is of the form $Q_i(\cdot) = (1 - \rho)\widehat{Q}_i(\cdot) + \rho\tilde{\nu}(\cdot)$ for a suitable measure $\tilde{\nu}$ on \mathbb{R}^2 .

(ii) *Consider now the densities $q(x, \eta) = \varphi_{f(x)}(\eta)$ where $\varphi_{f(x)}(\eta)$ is introduced in Example 2.7. We assume $f(x) = (f_1(x), x)$ for some $f_1 : X \rightarrow \mathbb{R}$. Since $\eta \mapsto \varphi_m(\eta)$ is twice continuously differentiable and because $X \subset \mathbb{R}$ is compact, it is easy to show that (22) holds. Moreover, there exists $r_1 > 0$, such that the mean-contraction condition (23) holds whenever $f_1(x) \leq r_1$.*

3 Existence of Lipschitz continuous Nash equilibria

This section is devoted to the proof of Theorem 2.10. Since the state space of our stochastic game Σ is not compact, we cannot prove the theorem directly. Instead, we shall first analyze ‘truncated’ games

$$\Sigma^N = (I, X, (U^i), \beta, Q, y, [-N, N]) \quad (N \in \mathbb{N})$$

with compact state spaces $[-N, N]$. To this end, we fix a family $\{G^N\}_{N \geq 1}$ of *truncating functions*. That is, a family of *strictly increasing* functions $G^N : \mathbb{R} \rightarrow [-N, N]$ which are Lipschitz continuous with Lipschitz constant one that satisfy $G^N(x) = x$ for $x \in [-N + \epsilon, N - \epsilon]$ ($\epsilon > 0$) and consider stochastic games Σ^N in which the conditional dynamics of the new state Y_{t+1} , given Y_t and the action profile \bar{x}_t is of the form

$$Y_{t+1} = v^N(Y_t, A_t, B_t) := G^N(A_t Y_t + B_t) \quad \text{and} \quad (A_t, B_t) \sim Q(c x_t; \cdot).$$

Using our *Moderate Social Influence* assumption we show that the games Σ^N have Lipschitz continuous Nash equilibria in Markov strategies τ^N whose Lipschitz constants do not depend on $N \in \mathbb{N}$. A Lipschitz continuous equilibrium for Σ will then be determined as a suitable accumulation point of the sequence $\{\tau^N\}_{N \in \mathbb{N}}$.

3.1 Lipschitz continuous equilibria in truncated games

The aim of this section is to establish the existence of Lipschitz continuous Nash equilibria in the truncated games Σ^N .

Proposition 3.1 *Under the assumptions of Theorem 2.10 there exists $C^* > 0$ such that for all $c \leq C^*$ the following holds:*

- (i) *For all $N \in \mathbb{N}$, the truncated game Σ^N has a Lipschitz continuous Nash equilibrium τ^N in Markov strategies.*
- (ii) *The Lipschitz constant of τ^N ($N \in \mathbb{N}$) does not depend on $N \in \mathbb{N}$.*

The proof of Proposition 3.1 needs some preparation. We put $u := \max_i \|U^i\|_\infty$ and denote by $\mathcal{L}_{L,u}([-N, N], \mathbb{R}^M)$ the class of all Lipschitz continuous functions $f : [-N, N] \rightarrow \mathbb{R}^M$ with Lipschitz constant L that satisfy $\|f\|_\infty \leq u$. To each such *average continuation function* $f = (f_i)_{i \in I}$ we associate the *reduced one-shot game* $\Sigma_f^N := (I, X, (U_f^{i,N}), \beta, y, [-N, N])$ with individual payoff functions

$$U_f^{i,N}(y, \bar{x}) = (1 - \beta)U^i(y, \bar{x}) + \beta \int f_i \circ v^N(y, A, B)q(x, A, B)\lambda(dA, dB) \quad (25)$$

viewed as functions from $[-N, N] \times \bar{X}$ to \mathbb{R} .

Remark 3.2 *Observe that $U_f^{i,N}(y, \bar{x})$ is the payoff to player $i \in I$ in the discounted stochastic game Σ^N if the game terminates after the first round, if the players receive rewards according*

to the payoff functions f_i in the second period, and if first period payoffs are discounted at the rate $1 - \beta$.

The following Lemma shows that the conditional best reply $\tau_f^{i,N}(y, x^{-i})$ of player $i \in I$ in the game Σ_f^N is uniquely determined and that Σ_f^N has a unique equilibrium $\tau_f^N(y)$. The map $\tau_f^N : [-N, N] \rightarrow \overline{X}$ is Lipschitz continuous. Its Lipschitz constant can be specified in terms of the Lipschitz constant of f , the discount factor, the bounds for the utility functions and the quantities $L^{i,j}(y)$, L^i and \hat{L} introduced in (17) and (18), respectively. In particular, it can be chosen *independently* of the specific average continuation function f and independently of N . This turns out to be the key to the proof of Proposition 3.1.

Lemma 3.3 *For $N \in \mathbb{N}$, let Σ_f^N be the reduced game with average continuation function $f \in \mathcal{L}_{L,u}([-N, N], \mathbb{R}^M)$. Under the assumptions of Theorem 2.10 the following holds:*

- (i) *The conditional best reply $\tau_f^{i,N}(y, x^{-i})$ of player $i \in I$ is uniquely determined and depends in a Lipschitz continuous manner on the actions of his competitors. More precisely,*

$$|\tau_f^{i,N}(y, x_1^{-i}) - \tau_f^{i,N}(y, x_2^{-i})| \leq \frac{(1 - \beta)L^{i,j}(y) + \beta u \hat{L}}{(1 - \beta)L^{i,i}(y) - \beta u \hat{L}} |x_1^j - x_2^j| \quad (26)$$

if $x_1^k = x_2^k$ for all $k \neq j$. Moreover,

$$|\tau_f^{i,N}(y_1, x^{-i}) - \tau_f^i(y_2, x^{-i})| \leq \tilde{L} |y_1 - y_2| \quad (27)$$

for all $y_1, y_2 \in [-N, N]$ and each $x^{-i} \in \overline{X}^{-i}$. Here

$$\tilde{L} := \sup_{y,i} \frac{(1 - \beta)L^i + \beta L(f)q}{(1 - \beta)L^{i,i}(y) - \beta u \hat{L}},$$

$L(f)$ denotes the Lipschitz constant of f and the quantity q is defined in (22).

- (ii) *The reduced game Σ_f^N has a unique equilibrium $\tau_f^N(y) = \{\tau_f^{i,N}(y)\}_{i \in I} \in \overline{X}$.*
- (iii) *The map $y \mapsto \tau_f^{i,N}(y)$ is Lipschitz continuous and the Lipschitz constant can be chosen independently of the average continuation function f and independently of N .*
- (iv) *The map $f \mapsto \tau_f^{i,N}(\cdot)$ is continuous.*

PROOF:

- (i) Let us fix an average continuation function $f \in \mathcal{L}_{L,u}([-N, N], \mathbb{R}^M)$, an action profile $x^{-i} \in \overline{X}^{-i}$ and a state $y \in [-N, N]$. By Assumption 2.6, the map $U_f^{i,N}(y, \cdot, x^{-i})$ is two times continuously differentiable on an open set containing X and our *Moderate Social Influence* condition yields

$$\frac{\partial^2}{\partial (x^i)^2} U_f^{i,N}(y, x^i, x^{-i}) \leq -(1 - \beta)L^{i,i}(y) + \beta u \hat{L} < 0.$$

Thus, an agent's utility function is strongly concave with respect to his own action, and so his conditional best reply given the choices of his competitors is uniquely determined. To establish the quantitative bound (26) on the dependence of player i 's best reply on the behavior of all the other agents, we fix a player $j \neq i$ and action profiles x_1^{-i} and x_2^{-i} which differ only at the j -th coordinate. Under the assumptions of Theorem 2.10 we have

$$\left| \frac{\partial}{\partial x^i} U_f^{i,N}(y, x^i, x_1^{-i}) - \frac{\partial}{\partial x^i} U_f^{i,N}(y, x^i, x_2^{-i}) \right| \leq \left\{ (1 - \beta)L^{i,j}(y) + \beta u \hat{L} \right\} |x_1^j - x_2^j|. \quad (28)$$

Thus, (26) follows from Theorem A.1. In view of (28) and because

$$\left| \frac{\partial}{\partial x^i} U_f^{i,N}(y_1, x^i, x^{-i}) - \frac{\partial}{\partial x^i} U_f^{i,N}(y_2, x^i, x^{-i}) \right| \leq \left\{ (1 - \beta)L^i + \beta L(f)q \right\} |x_1^j - x_2^j| \quad (29)$$

our estimate (27) also follows from Theorem A.1.

- (ii) The existence of an equilibrium in pure strategies for the static game Σ_f^N follows from strict concavity of the utility functions $U_f^{i,N}$ with respect to the player's own actions along with compactness of the action spaces using standard fixed points arguments. In order to prove uniqueness, we proceed as in the proof of Proposition 4.21 in Horst and Scheinkman (2002). In view of (16), our *Moderate Social Influence* condition yields

$$\bar{L} := \sup_{i,y} \sum_{j \neq i} \frac{(1 - \beta)L^{i,j}(y) + \beta u \hat{L}}{(1 - \beta)L^{i,i}(y) - \beta u \hat{L}} < 1.$$

Thus, given action profiles x_1^{-i} and x_2^{-i} of player i 's competitors, (26) shows that

$$|\tau_f^{i,N}(y, x_1^{-i}) - \tau_f^{i,N}(y, x_2^{-i})| \leq \bar{L} \max_j |x_1^j - x_2^j|.$$

For $x_1 \neq x_2$, we therefore obtain

$$\max_i |\tau_f^{i,N}(y, x_1^{-i}) - \tau_f^{i,N}(y, x_2^{-i})| < \max_i |x_1^i - x_2^i|.$$

Thus, the map $x \mapsto (\tau_f^{i,N}(y, x^{-i}))_{i \in I}$ has at most one fixed point. This proves uniqueness of equilibria in Σ_f^N .

(iii) Let $\tau_f^N(y)$ be an equilibrium. Then $\tau_f^{i,N}(y) = \tau_f^{i,N}(y, \{\tau_f^{j,N}(y)\}_{j \neq i})$, and so

$$\begin{aligned} |\tau_f^{i,N}(y_1) - \tau_f^{i,N}(y_2)| &\leq |\tau_f^{i,N}(y_1, \{\tau_f^{j,N}(y_1)\}_{j \neq i}) - \tau_f^{i,N}(y_1, \{\tau_f^{j,N}(y_2)\}_{j \neq i})| \\ &\quad + |\tau_f^{i,N}(y_1, \{\tau_f^{j,N}(y_2)\}_{j \neq i}) - \tau_f^{i,N}(y_2, \{\tau_f^{j,N}(y_2)\}_{j \neq i})| \\ &\leq \bar{L} \sum_{j \neq i} |\tau_f^{j,N}(y_1) - \tau_f^{j,N}(y_2)| + \tilde{L}|y_1 - y_2|. \end{aligned}$$

This yields

$$|\tau_f^{i,N}(y_1) - \tau_f^{i,N}(y_2)| \leq \frac{\tilde{L}}{(1 - \bar{L})} |y_1 - y_2|. \quad (30)$$

Hence the equilibrium mapping $\tau_f^N : [-N, N] \rightarrow \bar{X}$ is Lipschitz continuous with a constant that does not depend on $N \in \mathbb{N}$ and which depends on the average continuation function f only through its Lipschitz constant $L(f)$.

(iv) We fix $y \in [-N, N]$ and $x^{-i} \in \bar{X}^{-i}$ and apply Theorem A.1 to the map

$$(x^i, f) \mapsto U_f^{i,N}(y, x^i, x^{-i}).$$

Due to Assumption 2.6 (ii) there exists a constant $q^* < \infty$ such that

$$\left| \frac{\partial}{\partial x^i} U_{f_1}^{i,N}(y, x^i, x^{-i}) - \frac{\partial}{\partial x^i} U_{f_2}^{i,N}(y, x^i, x^{-i}) \right| \leq q^* \|f_1 - f_2\|_\infty,$$

for all $f, g \in \mathcal{L}_{L,u}([-N, N], \mathbb{R}^M)$, and so Theorem A.1 shows that there is $L_0 < \infty$ such that

$$|\tau_{f_1}^{i,N}(y, x^{-i}) - \tau_{f_2}^{i,N}(y, x^{-i})| \leq L_0 \|f_1 - f_2\|_\infty.$$

Thus, similar arguments as in the proof of (iii) yield the assertion. □

Our *Moderate Social Influence* conditions appears to be rather strong. However, for the proof of Theorem 2.10 it will be essential that the maps $\tau_f^N : [-N, N] \rightarrow \bar{X}$ are Lipschitz continuous with a constant that depends on the average continuation function f only through its Lipschitz constant. For that, we need uniqueness of equilibria in the reduced one-shot games Σ_f^N . We guarantee this by assuming that the utility functions $U_f^{i,N}$ inherit enough concavity in the player's own actions from the original reward functions U^i .

We are now ready to prove Proposition 3.1.

PROOF OF PROPOSITION 3.1: Let $B_u([-N, N], \mathbb{R}^M)$ be the class of all measurable functions $f : [-N, N] \rightarrow \mathbb{R}^M$ with $\|f\|_\infty \leq u$. We introduce an operator T from $\mathcal{L}_{L,u}([-N, N], \mathbb{R}^M)$ to $B_u([-N, N], \mathbb{R}^M)$ by

$$(Tf)_i(y) = (1 - \beta)U^i(y, \tau_f^N(y)) + \beta \int f_i \circ v^N(y, A, B)q \left(\frac{c}{M} \sum_{i \in I} \tau_f^{i,N}(y), A, B \right) \lambda(dA, dB). \quad (31)$$

A standard argument in discounted dynamic programming shows that for any fixed point F^N of T , the action profile $\tau_{F^N}^N(y)$ is an equilibrium in the non-zero sum stochastic game Σ^N . The equilibrium payoff to player $i \in I$ is given by $\frac{F^N(y)}{1-\beta}$, and the map $\tau_{F^N}^N : [-N, N] \rightarrow \overline{X}$ is Lipschitz continuous, due to Lemma 3.3. In order to establish our assertion it is therefore enough find $L < \infty$ and $C^* > 0$ such that the operator T has a fixed point in $\mathcal{L}_{L,u}([-N, N], \mathbb{R}^M)$ for all $c < C^*$. For this, we proceed in three steps:

- (i) Let $f \in \mathcal{L}_{L,u}([-N, N], \mathbb{R}^M)$. Lipschitz continuity of the utility functions and the densities along with (22) yields a constant $c_1 < \infty$ such that

$$|Tf_i(y_1) - Tf_i(y_2)| \leq c_1 (|y_1 - y_2| + \|\tau_f^N(y_1) - \tau_f^N(y_2)\|_\infty). \quad (32)$$

Thus, Lipschitz continuity of the map $\tau_f^N(\cdot)$ yields Lipschitz continuity of $Tf(\cdot)$.

- (ii) For any sequence $\{f^n\}_{n \in \mathbb{N}}$, $f^n \in \mathcal{L}_{L,u}([-N, N], \mathbb{R}^M)$, that converges to f in the topology of uniform convergence, Lemma 3.3 (v) yields $\lim_{n \rightarrow \infty} \|\tau_{f^n}^N - \tau_f^N\|_\infty = 0$. Thus, Lipschitz continuity of the reward functions and the densities gives us constants c_2 and c_3 such that

$$|Tf_i^n(y) - Tf_i(y)| \leq c_2 \|\tau_{f^n}^N - \tau_f^N\|_\infty + c_3 \{ \|f_i^n - f_i\|_\infty + \|\tau_{f^n}^N - \tau_f^N\|_\infty \},$$

and so $\lim_{n \rightarrow \infty} \|Tf^n - Tf\|_\infty = 0$. Thus, T is continuous.

- (iii) Lemma 3.3 and (30) show that there are constants d_1, d_2 and d_3 such that, for all $N \in \mathbb{N}$ and each $f \in \mathcal{L}_{L,u}([-N, N], \mathbb{R}^M)$, the equilibrium mappings $\tau_f^{i,N} : [-N, N] \rightarrow \overline{X}$ are Lipschitz continuous and that the Lipschitz constant $L^*(L)$ takes the form

$$L^* = \frac{d_1 + c \frac{d_2}{M} L}{d_3}.$$

In particular, (32) yields

$$|Tf_i(y_1) - Tf_i(y_2)| \leq c_1 \left[1 + \frac{d_1 + c \frac{d_2}{M} L}{d_3} \right] |y_1 - y_2|.$$

If we choose $C^* \leq \frac{d_3 M}{c_1 d_2}$, then for all sufficiently large $L \in \mathbb{N}$ and for each $c < C^*$, the operator T maps the set $\mathcal{L}_{L,u}([-N, N], \mathbb{R}^M)$ continuously into itself. Since $\mathcal{L}_{L,u}([-N, N], \mathbb{R}^M)$ is a compact set with respect to the topology of uniform convergence, T has a fixed point by Kakutani's theorem.

This finishes the proof □

3.2 Lipschitz continuous equilibria in discounted stochastic games and stability of the equilibrium process

We are now prepared to prove the existence of Lipschitz continuous Nash equilibria in Markov strategies for discounted stochastic games with affine state sequences.

PROOF OF THEOREM 2.10: For $N \in \mathbb{N}$, let τ^N and F^N be a Lipschitz continuous Nash equilibrium and the associated payoff function for the truncated game Σ^N . In view of Proposition 3.1 we may assume that these maps are Lipschitz continuous with a common Lipschitz constant, and so the sequences $\{\tau^N\}_{N \in \mathbb{N}}$ and $\{F^N\}_{N \in \mathbb{N}}$ are equicontinuous. Thus, by the theorem of Ascoli and Arzela, there exists a subsequence (N_k) and Lipschitz continuous functions $\tau : \mathbb{R} \rightarrow \overline{X}$ and $F : \mathbb{R} \rightarrow \mathbb{R}$ such that

$$\tau_{F^{N_k}}^{N_k}(y) \rightarrow \tau(y) \quad \text{and} \quad F^{N_k}(y) \rightarrow F(y) \quad \text{uniformly on compact sets as } k \rightarrow \infty.$$

Let us now fix $y \in \mathbb{R}$. For any compact set $K \subset \mathbb{R}^2$ we have

$$\lim_{n \rightarrow \infty} \sup_{(A,B) \in K} |F^N \circ v^N(y, A, B) - F(Ay + B)| = 0.$$

We also have $\lim_{N \rightarrow \infty} \|Q(\frac{c}{M} \sum_{i \in I} \tau^{i,N}(y); \cdot) - Q(\frac{c}{M} \sum_{i \in I} \tau^i(y); \cdot)\|_{\mathcal{B}} = 0$. Since the sequence $\{F^N \circ v^N\}_{N \in \mathbb{N}}$ is uniformly bounded, we obtain

$$F_i(y) = (1 - \beta)U^i(y, \tau(y)) + \beta \int F_i(Ay + B)q\left(\frac{c}{M} \sum_{i \in I} \tau^i(y), A, B\right) \lambda(dA, dB).$$

Now, it is easily seen that $\tau : \mathbb{R} \rightarrow \overline{X}$ is a stationary Nash equilibrium in Markov strategies for the game Σ . By construction, τ is Lipschitz continuous. □

We are now ready to prove Theorem 2.13.

PROOF OF THEOREM 2.10: For any Lipschitz continuous equilibrium τ , the stochastic kernel Q^τ from \mathbb{R} to \mathbb{R}^2 defined by $Q^\tau(y; \cdot) = Q\left(\frac{c}{M} \sum_{i=1}^M \tau^i(y); \cdot\right)$ satisfies the assumptions of Theorem 2.14. Thus, the assertion follows from Theorem 2.14. □

4 A Stability Result for a Class of Markov Chains

This section is devoted to the proof of Theorem 2.14. Throughout, we denote by K the transition kernel of the Markov chain $\{\tilde{Y}_t\}_{t \in \mathbb{N}}$:

$$K(y; \cdot) = \tilde{Q}(y; \{A, B\} : Ay + B \in \cdot).$$

Since the densities $\tilde{q}(y, \cdot)$ satisfy (24), it is easy to show that K has the Feller property.

The proof of Theorem 2.14 needs some preparation. In a first step we prove that the Feller process $\{\tilde{Y}_t\}_{t \in \mathbb{N}}$ is tight if Assumption 2.12 holds. Thus under Assumption 2.12, the Markov chain $\{\tilde{Y}_t\}_{t \in \mathbb{N}}$ has a stationary distribution. In a second step, we use a convergence result for Markov chains with *compact* state spaces to show that K has at most one invariant measure.

Proposition 4.1 *Under the assumptions of Theorem 2.14, the Markov chain $\{\tilde{Y}_t\}_{t \in \mathbb{N}}$ on \mathbb{R} has a stationary distribution.*

PROOF: Let μ_t^y be the law of the random variable \tilde{Y}_t under \mathbb{P}_y and assume that the Feller process $\{\tilde{Y}_t\}_{t \in \mathbb{N}}$ is tight. In this case the sequence

$$\{\nu_t\}_{t \in \mathbb{N}} \quad \text{where} \quad \nu_t := \frac{1}{t} \sum_{i=1}^t \mu_i^y$$

is tight as well. Thus, by Prohorov's theorem it has an accumulation point ν^* . A modification of standard arguments for Markov chains on compact state spaces given in, e.g., Liggett (1985), Proposition 1.8, shows that ν^* is a stationary measure for the Feller process $\{\tilde{Y}_t\}_{t \in \mathbb{N}}$. It is thus enough to show that the sequence $\{\tilde{Y}_t\}_{t \in \mathbb{N}}$ is tight. For this, it suffices to show that

$$\lim_{c \rightarrow \infty} \sup \left\{ \mathbb{P}_y[|\tilde{Y}_t| \geq c], t \in \mathbb{N} \right\} = 0. \quad (33)$$

By Tchebychev's inequality, (33) follows from

$$\sup \left\{ \mathbb{E}_y[|\tilde{Y}_t|], t \in \mathbb{N} \right\} < \infty. \quad (34)$$

In order to prove (34), we put

$$a := \sup_y \int |A| \tilde{Q}(y; dA, B) < 1 \quad \text{and} \quad b := \sup_y \int |B| \tilde{Q}(y; dA, dB) < \infty$$

and consider the deterministic sequence $\{\bar{Y}_t^y\}_{t \in \mathbb{N}}$ defined by the recursive relation

$$\bar{Y}_{t+1}^y = a\bar{Y}_t^y + b \quad (t \in \mathbb{N}, Y_0 = y).$$

Since $a < 1$, we have $\lim_{t \rightarrow \infty} \bar{Y}_t^y = \frac{b}{1-a} < \infty$. In particular, $\sup_t \bar{Y}_t^y < \infty$, and so it suffices to show that

$$\mathbb{E}_y|\tilde{Y}_t| \leq \bar{Y}_t^y \quad \text{for all } y \geq 0. \quad (35)$$

Clearly, (35) holds for $t = 0$. We proceed by induction and assume that (35) holds for all $t \leq T$. For the induction step, note that

$$\begin{aligned} \mathbb{E}_y|\tilde{Y}_{T+1}| &\leq \mathbb{E}_y[|A_T \tilde{Y}_T|] + \mathbb{E}_y[|B_T|] \\ &\leq \mathbb{E}_y \left[|\tilde{Y}_T| \mathbb{E}_y[|A_T| \mid |\tilde{Y}_T|] \right] + b \\ &\leq a\mathbb{E}_y|\tilde{Y}_T| + b \\ &\leq a\bar{Y}_T + b = \bar{Y}_{T+1}. \end{aligned}$$

This finishes the proof. □

In a second and more involved step, we are now going to show that the Feller process $\{\tilde{Y}_t\}_{t \in \mathbb{N}}$ has at most one stationary distribution. To this end, we need to study the transition kernel K . We denote by λ the Lebesgue measure on \mathbb{R} .

Lemma 4.2 *Under the assumptions of Theorem 2.14 the following holds:*

(i) *The support $I_y \subset \mathbb{R}$ of $K(y; \cdot)$ is an interval and $K(y; \cdot)$ is equivalent to λ on I_y :*

$$K(y; \cdot) = g_y d\lambda \quad \text{where } g_y(\cdot) > 0 \text{ on } I_y.$$

(ii) *There exists $m < \infty$ such that*

$$I_y \cap (-\infty, y) \neq \emptyset \text{ for } y \geq m \quad \text{and} \quad I_y \cap (y, \infty) \neq \emptyset \text{ for } y \leq -m.$$

PROOF:

(i) Let us fix $y \in \mathbb{R}$. The support I_y of the measure $K(y; \cdot)$ is given by the topological closure of the set $\{\hat{y} \in \mathbb{R} : \hat{y} = Ay + B \text{ for some } (A, B) \in M_y\}$, and convexity of M_y yields convexity of I_y . In order to prove equivalence of $K(y; \cdot)$ and λ on I_y , we introduce a bijection T_y on M_y by $T_y(A, B) := (A, Ay + B)$. By the transformation formula for

Lebesgue integrals the law of $T_y(A, B)$ has the density $\tilde{q}(y, \cdot) \circ T^{-1}$ with respect to λ^2 . Thus, since $\tilde{q}(y; \cdot) > 0$ on M_y , the distribution of $Ay + B$ has a strictly positive density g_y with respect to λ on I_y .

(ii) It is enough to establish the existence of a constant $m < \infty$ that satisfies

$$\tilde{Q}(y; \{(A, B) : |Ay + B| \leq |y|\}) > 0 \quad (|y| \geq m).$$

To this end, we assume to the contrary that, for all $m < \infty$ there is $|y_m| \geq m$ such that

$$\tilde{Q}(y_m; \{(A, B) : |Ay_m + B| \leq |y_m|\}) = 0. \quad (36)$$

However, by Assumption 2.12 (ii) we have $\tilde{Q}(y_m; \{(A, B) : |Ay_m| \leq r|y_m|\}) > 0$, and so (36) yields

$$\tilde{Q}(y_m; \{(A, B) : |B| \geq (1-r)|y_m|\}) = 1$$

for all $m < \infty$. Hence

$$(1-r)m \leq \sup_{|y| \geq m} \int |B| \tilde{Q}(y; dA, dB) \leq \sup_y \int |B| \tilde{Q}(y; dA, dB).$$

For $m \rightarrow \infty$, this contradicts Assumption 2.12 (i).

□

In order to prove that the Markov chain $\{\tilde{Y}_t\}_{t \in \mathbb{N}}$ has at most one invariant distribution we recall the notion of a truncation function G^N , and introduce Feller processes $\{\tilde{Y}_t^N\}_{t \in \mathbb{N}}$ with transition kernels K^N on the respective compact states spaces $[-N, N]$ by

$$\tilde{Y}_{t+1}^N = G^N(A_t \tilde{Y}_t^N + B_t) \quad \text{where} \quad (A_t, B_t) \sim \tilde{Q}(\tilde{Y}_t^N; \cdot).$$

Lemma 4.3 (i) For all $N \in \mathbb{N}$, the Markov chain $\{\tilde{Y}_t^N\}_{t \in \mathbb{N}}$ converges in distribution to a unique limiting measure μ^N .

(ii) For all sufficiently large $N \in \mathbb{N}$ and for each $y \in [-N, N]$, the measure $K^N(y; \cdot)$ is absolutely continuous with respect to $K(y; \cdot)$.

PROOF:

(i) The stochastic kernel \tilde{Q} satisfies Assumptions 2.12 (ii) and (iii) and, due to (24) also the Lipschitz continuity condition $\|\tilde{Q}(y_1; \cdot) - \tilde{Q}(y_2; \cdot)\| \leq \tilde{L}|y_1 - y_2|$. Thus, since the state space of the process $\{\tilde{Y}_t\}_{t \in \mathbb{N}}$ is compact, (i) follows from Theorem 4.1.1 in Norman (1972).

(ii) Let us assume that, for all sufficiently large N and for each $y \in [-N, N]$ we have $G^N(I_y) \subset I_y$. In this case, because the truncation functions G^N are strictly increasing, it follows as in the proof of Lemma 4.2 (i) that

$$K^N(y; \cdot) \approx \lambda \text{ on } G^N(I_y) \quad \text{and so} \quad K^N(y; \cdot) \ll K(y; \cdot).$$

In order to prove $G^N(I_y) \subset I_y$, we assume to the contrary that for all $N \in \mathbb{N}$ there exists $y_N \in (-N, N)$ such that $G^N(I_{y_N}) \not\subset I_{y_N}$. With $I_{y_N} = [a_N, b_N]$ it then follows from the definition of the truncation function G^N that $[a_N, b_N] \cap (-N, N) = \emptyset$. Thus, we may with no loss of generality assume that $a_N \rightarrow \infty$; the case $b_N \rightarrow -\infty$ along some subsequence can be analyzed by similar means. If the sequence $\{y_N\}_{N \in \mathbb{N}}$ is bounded, it has an accumulation point y^* and we let $I_{y^*} = [a^*, b^*]$. Continuity of the densities $\tilde{q}(y; \cdot)$ yields $a_N \rightarrow a^*$ and $b_N \rightarrow b^*$, and so $I_{y^*} = \emptyset$ because $I_{y_N} \cap (-N, N) = \emptyset$ for all $N \in \mathbb{N}$. Hence the sequence $\{y_N\}_{N \in \mathbb{N}}$ is unbounded, and we may with no loss of generality assume that $y_N \rightarrow \infty$. On the other hand, for all sufficiently large $N \in \mathbb{N}$, Lemma 4.2 (ii) yields $I_{y_N} \cap (-\infty, y_N) \neq \emptyset$. This, however, contradicts $I_{y_N} \cap (-N, N) = \emptyset$.

□

Now we are ready to prove that the Feller process $\{\tilde{Y}_t\}_{t \in \mathbb{N}}$ has at most one invariant distribution.

Lemma 4.4 *Under the assumption of Theorem 2.14, the Markov chain $\{\tilde{Y}_t\}_{t \in \mathbb{N}}$ has at most one invariant distribution.*

PROOF: Let us assume to the contrary that $|\{\mu : \mu K = \mu\}| > 1$. In this case there exist at least two mutually singular ergodic invariant measures μ_1, μ_2 for K . Thus, there are two disjoint sets C_1, C_2 which are stochastically closed with respect to K , i.e., which satisfy

$$K(y; C_i) = \delta_{i,j} \quad \text{for all } y \in C_j.$$

In order to prove that this contradicts Assumption 2.12, we proceed in three steps.

(i) We put $C_i^N = C_i \cap [-N, N]$. Choosing $N \in \mathbb{N}$ sufficiently large, $K^N(y; \cdot) \ll K(y; \cdot)$ for each $y \in [-N, N]$, due to Lemma 4.3 (ii). We may also assume that $\mu_i([-N, N]) > 0$. In particular, the conditional distribution μ_i^N of μ_i on $[-N, N]$ is well defined. Using absolute continuity of $K^N(y; \cdot)$ and $K(y; \cdot)$, it is easy to show that $\int K^N(y; C_i^N) \mu_i^N(dy) = \delta_{i,j}$, i.e., that the sets C_i^N are disjoint and stochastically closed with respect to K^N .

- (ii) In a second step we are going to show that the *topologically* closed sets \bar{C}_i and \bar{C}_i^N are *stochastically* closed with respect to K and K^N , respectively. To exclude trivial situations, we assume $\bar{C}_i \neq \mathbb{R}$. Let $y_n \in C_i$ and assume $\lim_{n \rightarrow \infty} y_n = y \in \bar{C}_i$ and $K(y; \bar{C}_i) < 1$. Then there exists $a < b \in \mathbb{R}$ such that $(a, b) \subset I_y \setminus \bar{C}_i$. Otherwise $I_y \cap \bar{C}_i$ would be a dense subset of I_y and thus $I_y \subset \bar{C}_i$. In particular, we can define a continuous function f on \mathbb{R} that vanishes on $[a, b]^c$ and that is strictly positive on (a, b) . Since $K(y; \cdot) \approx \lambda$ on I_y ,

$$\int f(\hat{y})K(y; d\hat{y}) > 0.$$

On the other hand, $I_{y_n} \subset \bar{C}_i$ yields

$$\int f(\hat{y})K(y_n; d\hat{y}) = 0 \quad \text{for all } n \in \mathbb{N}.$$

This, however, contradicts the Feller property of K . Using the same arguments as in (i) we see that $K(y; \bar{C}_i) = 1$ yields $K^N(y; \bar{C}_i^N) = 1$.

- (iii) Let us now assume that $\bar{C}_1 \cap \bar{C}_2 = \emptyset$. In this case it follows from Steps 1 and 2 above that \bar{C}_1^N and \bar{C}_2^N are two disjoint topologically closed subsets of $[-N, N]$ that are stochastically closed with respect to K^N . In view of Assumptions 2.8 (ii) and (iii), this contradicts Lemma 3.4.2 and Theorem 4.1.1 in Norman (1972). Thus, $\bar{C}_1 \cap \bar{C}_2 \neq \emptyset$. Let $y_n \rightarrow y$. By the Feller property of K we have $\lambda(I_{y_n} \cap I_y) > 0$ for all sufficiently large $n \in \mathbb{N}$. Since $C_i \subset \bar{C}_i$ dense we can therefore pick $y_i \in C_i$ such that

$$\lambda(I_{y_1} \cap I_{y_2}) > 0.$$

Since $K(y_i; \cdot) \approx \lambda$ on I_{y_i} we obtain $K(y_i; I_{y_1} \cap I_{y_2}) > 0$. This, however, contradicts $K(y_1; \cdot) \perp K(y_2; \cdot)$. Thus, $C_1 = C_2$, and so $\mu_1 = \mu_2$.

□

Corollary 4.5 *For any initial value $y \in \mathbb{R}$, the Markov chain $\{\tilde{Y}_t\}_{t \in \mathbb{N}}$ converges in distribution to a unique stationary measure μ^* .*

PROOF: The assertion follows from Lemma 4.2 and Lemma 4.4.

□

To prove the stability result stated in Theorem 2.14 it remains to show that the driving sequence $\{(A_t, B_t)\}_{t \in \mathbb{N}}$ is nice with respect to \mathbb{P}^* .

Proposition 4.6 *The driving sequence $\{(A_t, B_t)\}_{t \in \mathbb{N}}$ is nice with respect to \mathbb{P}^* .*

PROOF: Since the Markov chain $\{\tilde{Y}_t\}_{t \in \mathbb{N}}$ converges in law to the unique limiting distribution μ^* , the sequence $\{\tilde{Y}_t\}_{t \in \mathbb{N}}$ is stationary and ergodic under \mathbb{P}^* . Thus, we need to show that the asymptotic distribution of the random environment $\{(A_t, B_t)\}_{t \in \mathbb{N}}$ is the same under \mathbb{P}_y and under \mathbb{P}^* . To this end, let $L < \infty$ be the Lipschitz of the stochastic kernel \tilde{Q} :

$$\left\| \tilde{Q}(y_1; \cdot) - \tilde{Q}(y_2; \cdot) \right\|_{\mathcal{B}} \leq L|y_1 - y_2|. \quad (37)$$

Let $\psi_t := \{(A_1, B_1), \dots, (A_t, B_t)\}$ and let $B_t \in \mathcal{B}_t$ be the Borel- σ -field on \mathbb{R}^{2t} . Our aim is to show that there is a constant $L^* < \infty$ such that

$$|\mathbb{P}_{y_1}[\psi_t \in B_t] - \mathbb{P}_{y_2}[\psi_t \in B_t]| \leq L^*|y_1 - y_2| \quad \text{for all } y_1, y_2 \in \mathbb{R}. \quad (38)$$

For this, we proceed by induction. For $t = 1$, (38) follows from (37). Let us therefore assume that there is α_t such that

$$|\mathbb{P}_{y_1}[\psi_t \in B_t] - \mathbb{P}_{y_2}[\psi_t \in B_t]| \leq \alpha_t|y_1 - y_2| \quad \text{for all } y_1, y_2 \in \mathbb{R}.$$

For the induction step observe first that it suffices to prove (38) for sets B_t of the form $B_t = B_t^1 \times \dots \times B_t^t$ with $B_t^l \in \mathcal{B}_1$. The induction hypothesis along with our mean contraction condition yields

$$\begin{aligned} & |\mathbb{P}_{y_1}[\psi_{t+1} \in B_{t+1}] - \mathbb{P}_{y_2}[\psi_{t+1} \in B_{t+1}]| \\ &= \left| \mathbf{1}_{B_{t+1}^1}(\{A, B\}) \mathbb{P}_{Ay_1+B}[\psi_t \in B_t] \tilde{Q}(y_1; dA, dB) \right. \\ & \quad \left. - \mathbf{1}_{B_{t+1}^1}(\{A, B\}) \mathbb{P}_{Ay_2+B}[\psi_t \in B_t] \tilde{Q}(y_2; dA, dB) \right| \\ &\leq \int |\mathbb{P}_{Ay_1+B}[\psi_t \in B_t] - \mathbb{P}_{Ay_2+B}[\psi_t \in B_t]| \tilde{Q}(y_1; dA, dB) \\ & \quad + \sup_{B_1} |\mathbb{P}_{y_1}[\psi_1 \in B_1] - \mathbb{P}_{y_2}[\psi_1 \in B_1]| \\ &\leq L|y_1 - y_2| + \alpha_t \int |Ay_1 - Ay_2| \tilde{Q}(y_1; dA, dB) \\ &\leq (L + r\alpha_t) |y_1 - y_2| \end{aligned}$$

where $r < 1$ is defined in (23). Thus, $\alpha_{t+1} \leq C + r\alpha_t$, and this yields (38) with $L^* := \frac{C}{1-r}$.

Since the maps $y \mapsto \mathbb{P}_y[\psi_t \in B_t]$ are Lipschitz continuous, and because the Lipschitz constants do not depend on $t \in \mathbb{N}$ nor on $B_t \in \mathbb{B}_t$,

$$\|\mathbb{P}_y - \mathbb{Q}\|_{\hat{F}_{t,l}} \leq L^* d(\mu_t^y, \mu^*) \quad \text{for all } l \in \mathbb{N}.$$

Here μ_t^y denotes the law of \tilde{Y}_t under \mathbb{P}_y and $d(\cdot, \cdot)$ is the Vasserstein metric that induces the weak topology on the set of probability measure on \mathbb{R} . In view of Remark 2.2, this proves that ψ is nice with respect to \mathbb{P}^* . \square

We are now ready to finish the proof of Theorem 2.13.

PROOF OF THEOREM 2.14: By Proposition 4.6 the random environment ψ defined on $(\Omega, \mathcal{F}, \mathbb{P}_y)$ is nice with respect to the measure \mathbb{P}^* on (Ω, \mathcal{F}) . Thus, the assertion follows from Theorem 2.4 in Horst (2001). \square

A Parameter dependent optimization problems

The proof of Lemma 3.3 uses a result on Lipschitz continuous dependence of solutions to parameter dependent optimization problem. It is a special case of Theorem 3.1 in Montrucchio (1987):

Theorem A.1 *Let $X \subset \mathbb{R}$ be closed and convex and let $(Y, \|\cdot\|)$ be a normed space. Let $F : X \times Y \rightarrow \mathbb{R}$ be a finite function which satisfies the following conditions:¹*

(i) *For all $y \in Y$, the map $x \mapsto F(x, y)$ is concave and two times continuously differentiable on X . Moreover, there exists $\alpha > 0$ such that $\frac{\partial^2}{\partial x^2} F(x, y) \leq -\alpha$.*

(ii) *The derivative $\frac{\partial}{\partial x} F(x, y)$ of F at (x, y) satisfies the Lipschitz continuity condition*

$$\left| \frac{\partial}{\partial x} F(x, y_1) - \frac{\partial}{\partial x} F(x, y_2) \right| \leq L \|y_1 - y_2\|$$

for all $y_1, y_2 \in Y$ and all $x \in X$.

Under the above assumptions there exists a unique map $\theta : Y \rightarrow X$ that satisfies $\sup_{x \in X} F(x, y) = F(\theta(y), y)$. Moreover, θ is Lipschitz continuous and

$$|\theta(y_1) - \theta(y_2)| \leq \frac{L}{\alpha} \|y_1 - y_2\| \quad \text{for all } y_1, y_2 \in Y.$$

¹Montrucchio (1987) formulates this theorem under the additional assumption of Y being a closed and convex subset of a Hilbert space. His proof, however, shows that this assumptions is redundant.

References

- BARNESLEY, M. F., S. G. DEMKO, J. H. ELTON, AND J. S. GERONIMO (1988): “Invariant Measures for Markov Processes Arising from Iterated Function Systems with Place-Dependent Probabilities,” *Ann. Inst. Henri Poincaré*, 24, 367–394.
- BOROVKOV, A. A. (1998): *Ergodicity and Stability of Stochastic Processes*. John Wiley & Sons, New York.
- BRANDT, A. (1986): “The Stochastic Equation $Y_{n+1} = A_n Y_n + B_n$ with Stationary Coefficients,” *Adv. Appl. Probab.*, 18, 211–220.
- CURTAT, L. (1996): “Markov Equilibria in Stochastic Games with Complementaries,” *Games and Economic Behavior*, 17, 177–199.
- DUFFIE, D., J. GEANAKOPOLOS, A. MASCOLELL, AND A. MCLENNAN (1994): “Stationary Markov Equilibria,” *Econometrica*, 62 (4), 745–781.
- FÖLLMER, H., AND M. SCHWEIZER (1993): “A Microeconomic Approach to Diffusion Models for Stock Prices,” *Mathematical Finance*, 3, 1–23.
- HORST, U. (2001): “The Stochastic Equation $Y_{t+1} = A_t Y_t + B_t$ with Non-Stationary Coefficients,” *J. Appl. Prob.*, 38, 80–95.
- (2001): *Asymptotics of locally and globally interacting Markov chains arising in microstructure models for financial markets*, Shaker Verlag, Aachen.
- HORST, U. AND, J. SCHEINKMAN (2002): “Equilibria in Systems of Social Interactions,” *Working paper*. Princeton University.
- KIRMAN, A. (1998): “On the transitory nature of gurus,” Working paper, EHESS and Université de Marseille III.
- LIGGETT, T. (1985): *Interacting Particle Systems*, Springer Verlag, Berlin.
- MONTRUCCHIO, L. (1987): “Lipschitz Continuous Policy Functions for Strongly Concave Optimization Problems,” *Journal of Mathematical Economics*, 16, 259–273.
- NORMAN, F. M. (1972): *Markov Processes and Learning Models*. Academic Press, New York.

NOWAK, A. (1985): “Existence of Equilibrium Stationary Strategies in Discounted Non-Cooperative Stochastic Games with Uncountable State Space,” *J. Optim. Theory Appl.*, 45, 591–602.

VERVAAT, W. (1979): “On a Stochastic Difference Equation and a Representation of Non-Negative Infinitely Divisible Random Variables,” *Adv. Appl. Probab.*, 11, 750–783.