



Continuous-Time Mean Field Markov Decision Models

Nicole Bäuerle¹  · Sebastian Höfer¹

Accepted: 3 June 2024 / Published online: 22 June 2024
© The Author(s) 2024

Abstract

We consider a finite number of N statistically equal agents, each moving on a finite set of states according to a continuous-time Markov Decision Process (MDP). Transition intensities of the agents and generated rewards depend not only on the state and action of the agent itself, but also on the states of the other agents as well as the chosen action. Interactions like this are typical for a wide range of models in e.g. biology, epidemics, finance, social science and queueing systems among others. The aim is to maximize the expected discounted reward of the system, i.e. the agents have to cooperate as a team. Computationally this is a difficult task when N is large. Thus, we consider the limit for $N \rightarrow \infty$. In contrast to other papers we treat this problem from an MDP perspective. This has the advantage that we need less regularity assumptions in order to construct asymptotically optimal strategies than using viscosity solutions of HJB equations. The convergence rate is $1/\sqrt{N}$. We show how to apply our results using two examples: a machine replacement problem and a problem from epidemics. We also show that optimal feedback policies from the limiting problem are not necessarily asymptotically optimal.

Keywords Continuous-time Markov decision process · Mean field problem · Process limits · Pontryagin's maximum principle

Mathematics Subject Classification 90C40 · 60G55 · 60F17

1 Introduction

We consider a finite number of N statistically equal agents, each moving on a finite set of states according to a continuous-time Markov Decision Process. Transition intensities of the agents and generated rewards can be controlled and depend not only on the state and action of the agent itself, but also on the states of the other agents.

✉ Nicole Bäuerle
nicole.baeuerle@kit.edu

Sebastian Höfer
sebastian.hoefer@kit.edu

¹ Department of Mathematics, Karlsruhe Institute of Technology (KIT), Karlsruhe, Germany

Interactions like this are typical for a wide range of models in e.g. biology, epidemics, finance, social science and queueing systems among others. The aim is to maximize the expected discounted reward of the system, i.e. the agents have to cooperate as a team. This can be implemented by a central controller who is able to observe the whole system and assigns actions to the agents. Though this system itself can be formulated as a continuous-time Markov Decision Process, the established solution procedures are not really practical since the state space of the system is complicated and of high cardinality. Thus, we consider the limit $N \rightarrow \infty$ when the number of agents tends to infinity and analyze the connection between the limiting optimization problem, which is a deterministic control problem, and the N agents problem.

Investigations like this are well-known under the name *Mean-field approximation*, because the mean dynamics of the agents can be approximated by differential equations for a measure-valued state process. This is inspired by statistical mechanics and can be done for different classes of stochastic processes for the agents. In our paper we restrict our investigation to continuous-time Markov chains (CTMC). Earlier, more practical studies in this spirit with CTMC, but *without control* are e.g. [1, 2] which consider illustrating examples to discuss how the mean-field method is used in different application areas. The convergence proof there is based on the law of large numbers for centred Poisson processes, see also [3]. The authors of [4] look at so-called reaction networks which are chemical systems involving multiple reactions and chemical species. They take approximations of multiscale nature into account and show that 'slow' components can be approximated by a deterministic equation. Reference [5] formulates some simple conditions under which a CTMC may be approximated by the solution to a differential equation, with quantifiable error probabilities. They give different applications. Reference [6] explores models proposed for the analysis of BitTorrent P2P systems and provide the arguments to justify the passage from the stochastic process, under adequate scaling, to a fluid approximation driven by a differential equation. A more recent application is given in [7] where a multi-type analogue of Kingman's coalescent as a death chain is considered. The aim is to characterize the behaviour of the replicator coalescent as it is started from an initial population that is arbitrarily large. This leads to a differential equation called the replicator equation. A similar control model as ours is considered in [8]. However, there the author uses a finite time horizon and solves the problem with HJB equations. This requires a considerable technical overhead like viscosity solutions and more assumptions on the model data like Lipschitz properties which we do not need here.

A related topic are fluid models. Fluid models have been introduced in queueing network theory since there is a close connection between the stability of the stochastic network and the corresponding fluid model, [9]. They appear under 'fluid scaling' where time in the CTMC for the stochastic queueing network is accelerated by a factor N and the state is compressed by factor $1/N$. Fluid models have also been used to approximate the optimal control in these networks, see e.g. [10–14]. In [15] different scales of time are treated for the approximation and some components may be replaced by differential equations. But there is no mean-field interaction in all of these fluid models.

There are also investigations about controlled mean-field Markov decision processes and their limits in discrete time. An early paper is [16] where the mean-field

limit for increasing number of agents is considered in a model where only the central controller is allowed to choose one action. However, in order to get a continuous limit the authors have to interpolate and rescale the original discrete-time processes. This implies the necessity for assumptions on the transition probabilities. The authors show the convergence of the scaled value functions and derive asymptotically optimal strategies. The recent papers [17–20] discuss the convergence of value functions and asymptotically optimal policies in discrete time. In contrast to our paper they allow a common noise. The limit problem is then a controlled stochastic process in discrete-time.

Another strand of literature considers continuous-time mean-field *games* on a finite number of states [21–25]. These papers among others considers the construction of asymptotically optimal Nash-equilibria from a limiting equation. The exception is [24] where it is shown that any solution of the limiting game can be approximated by ϵ_N -Nash equilibria in the N player game. However, all these papers deal with the convergence of the HJB equations which appear in the N player game to a limiting equation, called the Master equation ([26]) which is a deterministic PDE for the value function. This approach needs sufficient regularity of the value functions and many assumptions. Reference [25] considers the problem with common noise and reduces the mean field equilibrium to a system of forward-backward systems of (random) ordinary differential equations.

The *contribution of our paper* is first to establish and investigate the limit of the controlled continuous-time Markov decision processes. In contrast to previous literature which works with the HJB equation this point of view requires less assumptions e.g. we do not need Lipschitz conditions on the model data. Second, we are also able to construct an asymptotically optimal strategy for the N agents model. Our model is general, has only a few, easy to check assumptions and allows for various applications. The advantage of our limiting optimization problem is that we can apply Pontryagin's maximum principle easily which is often more practical than deterministic dynamic programming. Further, we show that an optimal feedback policy in the deterministic problem does not necessarily imply an asymptotically optimal policy for the N agents problems. Third, we obtain a convergence rate in a straightforward way. Fourth, we can consider finite and infinite time horizon at the same time. There is essentially no difference. We restrict the presentation mainly to the infinite time horizon.

Our paper is organized as follows: In the next section we introduce our N agents continuous-time Markov decision process. The aim is to maximize the expected discounted reward of the system. In Sect. 3 we introduce a measure-valued simplification which is due to the symmetry properties of the problem and which reduces the cardinality of the state space. The convergence theorem if the number of agents tends to infinity can be found in Sect. 4. It is essentially based on martingale convergence arguments. In Sect. 5 we construct a sequence of asymptotically optimal strategies from the limiting model for the N agents model. We also show that different implementations may be possible and that the rate of convergence is at most $1/\sqrt{N}$. Finally in Sect. 6 we discuss three applications. The first one is a machine replacement problem when we have many machines, see e.g. [27]. The second one is the spreading of malware which is based on the classical SIR model for spreading infections, [16, 28]. The last example shows that one has to be careful with feedback policies.

2 The N Agents Continuous-Time Markov Decision Process

We consider a finite number of N statistically equal agents, each moving on a finite set of states S according to a continuous-time Markov Decision Process. The vector $\mathbf{x}_t = (x_t^1, \dots, x_t^N) \in S^N$ describes the state of the system at time $t \in [0, \infty)$, where x_t^k is the state of agent $k = 1, \dots, N$. The action space of one agent is a compact Borel set A . The action space of the system is accordingly A^N . We denote an action of the system by $\mathbf{a} = (a^1, \dots, a^N) \in A^N$ where a^k is the action chosen by agent $k = 1, \dots, N$.

Let $D(i) \subset A$ be the set of actions available for an agent in state $i \in S$ which we again assume to be compact. Then the set of admissible actions for the system in state $\mathbf{x} \in S^N$ is given by $\mathbf{D}(\mathbf{x}) := D(x^1) \times \dots \times D(x^N) \subset A^N$. The set of admissible state-action combinations for one agent is denoted by $D := \{(i, a) \in S \times A \mid a \in D(i), \forall i \in S\}$.

For the construction of the system state process we follow the notation of [29]. The state process of the system is defined on the measurable space $(\Omega, \mathcal{F}) := ((S^N \times \mathbb{R}_+)^{\infty}, \mathcal{B}((S^N \times \mathbb{R}_+)^{\infty}))$. We denote an element of Ω by $\omega = (\mathbf{x}_0, t_1, \mathbf{x}_1, t_2, \dots)$. Now define

$$\begin{aligned}\tilde{\mathbf{X}}_n &: \Omega \rightarrow S^N, & \tilde{\mathbf{X}}_n(\omega) &= \mathbf{x}_n, & n \in \mathbb{N}_0, \\ \tau_n &: \Omega \rightarrow \mathbb{R}_+, & \tau_n(\omega) &= t_n, & n \in \mathbb{N}, \\ T_n &:= \sum_{k=1}^n \tau_k, & T_0 &:= 0.\end{aligned}$$

The controlled state process of the system is then given by

$$\mathbf{X}_t := \sum_{n \in \mathbb{N}_0} \mathbb{1}_{\{T_n \leq t < T_{n+1}\}} \tilde{\mathbf{X}}_n, \quad t \in [0, \infty).$$

The construction of the process can be interpreted as follows: The random variables τ_n describe the sojourn times of the system in states $\tilde{\mathbf{X}}_{n-1}$. Based on the sojourn times, T_n describes the time of the n -th jump of the process and $\tilde{\mathbf{X}}_n$ the state of the process on the interval $[T_n, T_{n+1})$. By construction the continuous-time state process (\mathbf{X}_t) has piecewise constant càdlàg-paths and the embedded discrete-time process is $(\tilde{\mathbf{X}}_n)$.

The system is controlled by policies. W.l.o.g. we restrict here to Markovian stationary policies. Further, we allow for randomized decisions, i.e. each agent can choose a probability distribution on A as its action. Hence a policy for the system is given by a collection of N stochastic kernels $\pi(d\mathbf{a} \mid \mathbf{x}) = (\pi^k(da \mid \mathbf{x}))_{k=1, \dots, N}$, where

$$\pi^k : S^N \times \mathcal{B}(A) \rightarrow [0, 1], \quad (\mathbf{x}, \mathcal{A}) \mapsto \pi^k(\mathcal{A} \mid \mathbf{x}) \quad (\text{kernel for agent } k).$$

$\pi^k(\mathcal{A} \mid \mathbf{x})$ is the stochastic kernel (it is here considered as a relaxed control) with which agent k chooses an action, given the state \mathbf{x} of the system. Naturally, it should hold that the kernel is concentrated on admissible actions, i.e. $\pi^k(D(x^k) \mid \mathbf{x}) = 1$ for all agents $k = 1, \dots, N$.

The action process is thus defined by

$$\pi_t := \sum_{n \in \mathbb{N}_0} \mathbb{1}_{\{T_n < t \leq T_{n+1}\}} \pi(\cdot \mid \tilde{\mathbf{X}}_n), \quad t \in [0, \infty).$$

In contrast to the state process, the action process has piecewise constant càglàd-paths. This means that a new decision can only be taken after a change of state has already occurred. The general theory on continuous-time Markov decision processes states that the optimal policy can be found among the piecewise constant, deterministic, stationary policies. In particular, varying the action continuously on the interval $[T_n, T_{n+1})$ does not increase the value of the problem. Also randomization does not increase the value, but in view of the sections to come, we already allowed for randomization (relaxation) here.

To prepare the description of the transition mechanism in our model, we define the empirical distribution of the agents over the states, i.e.

$$\mu[\mathbf{x}] := \frac{1}{N} \sum_{k=1}^N \delta_{x^k}.$$

where δ_{x^k} is the Dirac measure in point x^k . The transition intensities for one agent are given by a signed kernel

$$q : S \times A \times \mathbb{P}(S) \times \mathcal{P}(S) \rightarrow \mathbb{R}, \quad (i, a, \mu, \Gamma) \mapsto q(\Gamma \mid i, a, \mu) = \sum_{j \in \Gamma} q(\{j\} \mid i, a, \mu).$$

Here $\mathbb{P}(S)$ is the set of all probability distributions on S and $\mathcal{P}(S)$ is the power set of S . Note that the transition of an agent depends not only on its own state and action, but also on the empirical distribution of all agents over the states.

We make the following assumptions on q :

- (Q1) $q(\{j\} \mid i, a, \mu) \geq 0$ for all $i, j \in S, j \neq i, a \in D(i), \mu \in \mathbb{P}(S)$.
- (Q2) $\sum_{j \in S} q(\{j\} \mid i, a, \mu) = 0$ for all $(i, a) \in D, \mu \in \mathbb{P}(S)$.
- (Q3) $\sup_{i,a,j,\mu} |q(\{j\} \mid i, a, \mu)| =: q_{max} < \infty$.
- (Q4) $\mu \mapsto q(\{j\} \mid i, a, \mu)$ is continuous w.r.t. weak convergence for all $i, j \in S, a \in D(i)$.
- (Q5) $a \mapsto q(\{j\} \mid i, a, \mu)$ is continuous for all $i, j \in S, \mu \in \mathbb{P}(S)$.

Note that (Q3) follows from (Q4) and (Q5), but since it is important we list it here. Based on the transition intensities for one agent, the transition intensities of the system are given by

$$q(\{(x^1, \dots, x^{k-1}, j, x^{k+1}, \dots, x^N)\} \mid \mathbf{x}, \mathbf{a}) := q(\{j\} \mid x^k, a^k, \mu[\mathbf{x}]) \tag{2.1}$$

for all $(\mathbf{x}, \mathbf{a}) \in \mathbf{D}(\mathbf{x})$, $j \in S$, $j \neq x^k$ and

$$q(\{j\}|\mathbf{x}, \mathbf{a}) := \sum_{k=1}^N q(\{x^k\}|x^k, a^k, \mu[\mathbf{x}]).$$

All other intensities are zero. The intensity in Eq. (2.1) describes the transition of agent k from state $x^k \in S$ to state $j \in S$, while all other agents stay in their current state. Since only one agent can change its state at a time, this definition is sufficient to describe the transition mechanism of the system.

Further we set (in a relaxed sense) for a decision rule $\pi^k(da|\mathbf{x})$

$$q(\{x^1, \dots, x^{k-1}, j, x^{k+1}, \dots, x^N\}|\mathbf{x}, \pi) = \int_A q(\{j\}|x^k, a, \mu[\mathbf{x}])\pi^k(da|\mathbf{x}).$$

Note that in a certain sense there is abuse of notation here since we use the letter q both for the agent transition intensity and for the system transition intensity. It should always be clear from the context which one is meant.

The probability measure of the N agent process is now given by the following transition kernels

$$\mathbb{P}^\pi(\tau_n \leq t, \tilde{\mathbf{X}}_n \in B|\tilde{\mathbf{X}}_{n-1}) = \int_0^t q(B|\tilde{\mathbf{X}}_{n-1}, \pi)e^{s \cdot q(\{\tilde{\mathbf{X}}_{n-1}\}|\tilde{\mathbf{X}}_{n-1}, \pi)} ds$$

for all $t \geq 0$ and $B \in \mathcal{P}(S^N)$. In particular, the sojourn times τ_n are exponentially distributed with parameter $-q(\{\tilde{\mathbf{X}}_{n-1}\}|\tilde{\mathbf{X}}_{n-1}, \pi)$ respectively. Note that by using this construction, the probability measure depends on the chosen policy. This construction is more convenient when the transition intensities are given. In case the system is described by transition functions and external noise it is easier to use a common probability space which does not depend on the policy. Of course these two points of view are equivalent.

Returning to the model’s control mechanism, keep in mind that the policy of an agent $\pi^k(da | \mathbf{x})$ is allowed to depend on the state of the whole system, i.e. we assume that each agent has information about the position of all other agents. Therefore, we can interpret our model as a centralized control problem, where all information is collected and shared by a central controller.

The goal of the central controller is to maximize the social reward of the system. In order to implement this, we introduce the (stationary) reward function for one agent as

$$r : D \times \mathbb{P}(S) \rightarrow \mathbb{R}, \quad (i, a, \mu) \mapsto r(i, a, \mu),$$

which does not only depend on the state and action of the agent, but also on the empirical distribution of the system. We make the following assumptions on the reward function:

- (R1) For all $(i, a) \in D$ the function $\mu \mapsto \mu(i)r(i, a, \mu)$ is continuous w.r.t. weak convergence.

(R2) For all $i \in S$ and $\mu \in \mathbb{P}(S)$ the function $a \mapsto r(i, a, \mu)$ is continuous.

Since the set of admissible actions $D(i)$ is compact, (R1) and (R2) imply that the following expression is bounded:

$$\sup_{(i,a) \in D, \mu \in \mathbb{P}(S)} |\mu(i)r(i, a, \mu)| < \infty. \tag{2.2}$$

The (social) reward of the system is the average of the agents' rewards

$$r(\mathbf{x}, \mathbf{a}) := \frac{1}{N} \sum_{k=1}^N r(x^k, a^k, \mu[\mathbf{x}]), \tag{2.3}$$

or, in a relaxed sense for a decision rule $\pi^k(da | \mathbf{x})$

$$r(\mathbf{x}, \pi) := \frac{1}{N} \sum_{k=1}^N \int_A r(x^k, a, \mu[\mathbf{x}])\pi^k(da | \mathbf{x}).$$

The aim is now to find the *social optimum*, i.e. to maximize the joint expected discounted reward of the system over an infinite time horizon. For a policy π , a discount rate $\beta > 0$ and an initial configuration $\mathbf{x} \in S^N$ define the value function

$$\begin{aligned} V_\pi(\mathbf{x}) &= \mathbb{E}_\pi^\mathbf{x} \left[\int_0^\infty e^{-\beta t} r(\mathbf{X}_t, \pi_t) dt \right] \\ V(\mathbf{x}) &= \sup_\pi V_\pi(\mathbf{x}). \end{aligned} \tag{2.4}$$

We are not discussing solution procedures for this optimization problem here since we simplify it in the next section and present asymptotically optimal solution methods in Sect. 5.

3 The Measure-Valued Continuous-Time Markov Decision Process

As N is getting larger, so does the state space S^N , which could make the model increasingly complex and impractical to solve. Therefore, we seek for some simplifications. An obvious approach which is common for these kind of models, is to exploit the symmetry of the system by capturing not the state of every single agent, but the relative or empirical distribution of the agents across the $|S|$ states.

Thus, let $\mu_i^N := \mu[\mathbf{X}_t]$ and define as new state space the set of all distributions which are empirical measures of N atoms

$$\mathbb{P}_N(S) := \{ \mu \in \mathbb{P}(S) \mid \mu = \mu[\mathbf{x}], \text{ for } \mathbf{x} \in S^N \}.$$

It holds that the new state process μ_t^N is the same as

$$\mu_t^N = \sum_{n \in \mathbb{N}_0} \mathbb{1}_{\{T_n \leq t < T_{n+1}\}} \mu[\tilde{\mathbf{X}}_n], \quad t \in [0, \infty).$$

As action space take the $|S|$ -fold Cartesian product $\mathbb{P}(A)^{|S|}$ of $\mathbb{P}(A)$. Hence, an action is given by $|S|$ probability measures $\alpha(da) = (\alpha^i(da))_{i \in S}$ with $\alpha^i(D(i)) = 1$. Hereby the i -th component indicates the distribution of the agents' actions in state $i \in S$. The set of admissible state-action combinations of the new model is given by $\hat{D} := \mathbb{P}_N(S) \times \mathbb{P}(A)^{|S|}$.

For the policies we restrict again to Markovian, stationary policies given by a collection of $|S|$ stochastic kernels $\hat{\pi}(da|\mu) = (\hat{\pi}^i(da|\mu))_{i \in S}$, where

$$\hat{\pi}^i : \mathbb{P}_N(S) \times \mathcal{B}(A) \rightarrow [0, 1], \quad (\mu, \mathcal{A}) \mapsto \hat{\pi}^i(\mathcal{A} | \mu) \quad (\text{kernel for state } i).$$

where $\hat{\pi}^i(D(i) | \mu) = 1$. In what follows we denote $\tilde{\mu}_n^N := \mu[\tilde{\mathbf{X}}_n]$. Then we can express the action process by setting

$$\hat{\pi}_t := \sum_{n \in \mathbb{N}_0} \mathbb{1}_{\{T_n < t \leq T_{n+1}\}} \hat{\pi}(\cdot | \tilde{\mu}_n^N), \quad t \in [0, \infty). \tag{3.1}$$

The transition intensities of the process $(\mu_t^N)_{t \geq 0}$ are given by

$$q(\{\mu^{i \rightarrow j}\} | \mu, \alpha) = N \mu(i) \int_A q(\{j\} | i, a, \mu) \alpha^i(da), \quad \mu \in \mathbb{P}_N(S), \alpha \in \mathbb{P}(A)^{|S|}, \tag{3.2}$$

with $\mu^{i \rightarrow j} := \mu - \frac{1}{N} \delta_i + \frac{1}{N} \delta_j$ for all $i, j \in S, i \neq j$ if $\mu(i) > 0$. This intensity describes the transition of one arbitrary agent in state $i \in S$ to state $j \in S$, while all other agents stay in their current state. Note that the intensity follows from the usual calculations for continuous-time Markov chains, in particular from the fact that if X, Y are independent random variables with $X \sim Exp(\lambda), Y \sim Exp(\nu)$, then $X \wedge Y \sim Exp(\lambda + \nu)$. In the situation in Eq. (3.2) we have $N \mu(i)$ agents in state i . Further we set for all $\mu \in \mathbb{P}_N(S)$ and $\alpha \in \mathbb{P}(A)^{|S|}$

$$q(\{\mu\} | \mu, \alpha) := - \sum_{i, \mu(i) > 0} \sum_{j \neq i} q(\{\mu^{i \rightarrow j}\} | \mu, \alpha).$$

All other intensities are zero, since again only one agent can change its state at a time.

The probability distribution of the measure-valued process under a fixed policy $\hat{\pi}$ is now given by the following transition kernels

$$\mathbb{P}^{\hat{\pi}}(\tau_n \leq t, \tilde{\mu}_n^N \in B | \tilde{\mu}_{n-1}^N) = \int_0^t q(B | \tilde{\mu}_{n-1}^N, \hat{\pi}) e^{s \cdot q(\{\tilde{\mu}_{n-1}^N\} | \tilde{\mu}_{n-1}^N, \hat{\pi})} ds$$

for all $t \geq 0$ and $B \subset \mathbb{P}_N(S)$ measurable, where the random variables (τ_n) are the same as before.

The reward function of the system is derived from the reward for one agent:

$$r(\mu, \alpha) := \sum_{i \in S} \int r(i, a, \mu) \alpha^i(da) \mu(i).$$

In view of Eq. (2.2) $r(\mu, \alpha)$ is bounded. The aim in this model is again to maximize the joint expected discounted reward of the system over an infinite time horizon. For a policy $\hat{\pi}$, a discount rate $\beta > 0$ and an initial configuration $\mu \in \mathbb{P}_N(S)$ define the value function

$$\begin{aligned} V_{\hat{\pi}}^N(\mu) &= \mathbb{E}_{\hat{\pi}}^{\mu} \left[\int_0^{\infty} e^{-\beta t} r(\mu_t^N, \hat{\pi}_t) dt \right] \\ V^N(\mu) &= \sup_{\hat{\pi}} V_{\hat{\pi}}^N(\mu). \end{aligned} \tag{3.3}$$

We can now show that both formulations in Eqs. (2.4) and (3.3) are equivalent in the sense that the optimal values are the same. Of course, an optimal policy in the measure-valued setting can directly be implemented in the original problem. The advantage of the measure-valued formulation is the reduction of the cardinality of the state space. Suppose for example that $S = \{0, 1\}$, i.e. all agents are either in state 0 or state 1. Then $|S^N| = 2^N$ in the original formulation whereas $|\mathbb{P}_N(S)| = N + 1$ in the second formulation. A proof of the next theorem can be found in the appendix.

Theorem 3.1 *It holds that $V(\mathbf{x}) = V^N(\mu)$ for $\mu = \mu[\mathbf{x}]$ for all $\mathbf{x} \in S^N$.*

Remark 3.2 It is possible to extend the previous result to a situation where reward and transition intensity both also depend on the empirical distribution of actions, see e.g. [18]. However, due to the definition of the Young topology which we use later it is not possible to transfer the convergence results to this setting.

The problem we have introduced is a classical continuous-time Markov Decision Process and can be solved with the established theory accordingly. Thus, we obtain:

Theorem 3.3 *There exists a continuous function $v : \mathbb{P}_N(S) \rightarrow \mathbb{R}$ satisfying*

$$\beta v(\mu) = \sup_{\alpha \in \mathbb{P}(A)^{|S|}} \left\{ r(\mu, \alpha) + \int v(v)q(dv|\mu, \alpha) \right\}$$

for all $\mu \in \mathbb{P}_N(S)$ and there exists a maximizer $\hat{\pi}(\cdot|\mu)$ of the r.h.s. such that $v = V^N$ and $\hat{\pi}$ determines the optimal policy by Eq. (3.1).

The theorem follows from Theorem 4.6, Lemma 4.4 in [30] or Theorem 3.1.2 in [29]. Theorem 3.3 implies a solution method for problem (3.3). It can e.g. be solved by value or policy iteration. However, as already discussed, even in this simplified setting, the computation may be inefficient if N is large, since this leads to a large state space.

4 Convergence of the State Process

In this section we discuss the behaviour of the system when the number of agents tends to infinity. In this case we obtain a deterministic limit control model which serves as an asymptotic upper bound for our optimization problem with N agents. Moreover, an optimal control of the limit model can be used to establish a sequence of asymptotically optimal policies for the N agents model.

In what follows we consider (μ_t^N) as a stochastic element of $D_{\mathbb{P}_N(S)}[0, \infty)$, the space of càdlàg paths with values in $\mathbb{P}_N(S)$ equipped with the Skorokhod J_1 -topology and metric d_{J_1} . On $\mathbb{P}_N(S)$ we choose the total variation metric $\|\cdot\|_{TV}$.

Further, we consider $\hat{\pi}^i$ as a stochastic element in $\mathcal{R} := \{\rho : \mathbb{R}_+ \rightarrow \mathbb{P}(A) \mid \rho \text{ measurable}\}$ endowed with the Young topology (cf. [31]). It is possible to show that \mathcal{R} is compact and metrizable. Measurability and convergence in \mathcal{R} can be characterized as in Lemma 4.1. These statements follow directly from the fact that the Young topology is the coarsest topology such that the mappings

$$\rho \mapsto \int_0^\infty \int_A \psi(t, a) \rho_t(da) dt$$

are continuous for all real functions ψ on $\mathbb{R}_+ \times A$ where ψ is a Carathéodory function, i.e. ψ is continuous in a and measurable in t where ψ is integrable in the sense that $\int_0^\infty \sup_a |\psi(t, a)| dt < \infty$.

Lemma 4.1 (a) $\rho : \mathbb{R}_+ \rightarrow \mathbb{P}(A)$ is measurable if and only if ρ is a transition probability from \mathbb{R}_+ into A .

(b) Let $\rho^n, \rho \in \mathcal{R}$. $\rho^n \rightarrow \rho$ for $n \rightarrow \infty$ if and only if

$$\int_0^\infty \int_A \psi(t, a) \rho_t^n(da) dt \rightarrow \int_0^\infty \int_A \psi(t, a) \rho_t(da) dt$$

for all measurable functions $\psi : \mathbb{R}_+ \times A \rightarrow \mathbb{R}$ such that $a \mapsto \psi(t, a)$ is continuous for all $t \geq 0$ and $\int_0^\infty \sup_a |\psi(t, a)| dt < \infty$.

In a first step we define for $N \in \mathbb{N}$, a fixed policy $\hat{\pi}^N$ and arbitrary $j \in S$, the one-dimensional process

$$M_t^N(j) := \mu_t^N(j) - \mu_0^N(j) - \int_0^t \sum_{\nu \in \mathbb{P}_N(S)} (\nu(j) - \mu_s^N(j)) q(\{\nu\} | \mu_s^N, \hat{\pi}_s^N) ds.$$

Then $(M_t^N(j))$ are martingales w.r.t. the filtration $\mathcal{F}_t^N = \sigma(\mu_s^N, s \leq t)$. This follows from the Dynkin formula, see e.g. [31], Proposition 14.13. Next we can express the process $(M_t^N(j))$ a bit more explicitly. Note that the difference $\nu(j) - \mu_s^N(j)$ can either be $-1/N$ if an agent changes from state j to a state $k \neq j$ or it could be $1/N$ if an agent changes from state $i \neq j$ to state j . Since by (Q2)

$$\sum_{k \neq j} \int q(\{k\} | j, a, \mu_s^N) \hat{\pi}_s^{N,j}(da) = - \int q(\{j\} | j, a, \mu_s^N) \hat{\pi}_s^{N,j}(da) \quad (4.1)$$

we obtain by inserting the intensity in Eq. (3.2) and by using Eq. (4.1)

$$\begin{aligned}
 M_t^N(j) &= \mu_t^N(j) - \mu_0^N(j) - \int_0^t \sum_{k \neq j} -\frac{1}{N} N \mu_s^N(j) \int q(\{k\}|j, a, \mu_s^N) \hat{\pi}_s^{N,j}(da) ds \\
 &\quad - \int_0^t \sum_{i \neq j} \frac{1}{N} N \mu_s^N(i) \int q(\{j\}|i, a, \mu_s^N) \hat{\pi}_s^{N,i}(da) ds \\
 &= \mu_t^N(j) - \mu_0^N(j) - \int_0^t \sum_{i \in S} \mu_s^N(i) \int q(\{j\}|i, a, \mu_s^N) \hat{\pi}_s^{N,i}(da) ds. \quad (4.2)
 \end{aligned}$$

With this representation we can prove that the sequence of stochastic processes $(M^N(j))$ converges weakly (denoted by \Rightarrow) in the Skorokhod J_1 -topology to the zero process. The proof of this lemma together with the proof of the next theorem can be found in the appendix.

Lemma 4.2 *We have for all $j \in S$ that*

$$(M_t^N(j))_{t \geq 0} \Rightarrow 0, \quad N \rightarrow \infty.$$

Next we show that an arbitrary state-action process sequence is relatively compact which implies the existence of converging subsequences.

Theorem 4.3 *A sequence of arbitrary state-action processes $(\mu^N, \hat{\pi}^N)_N$ is relatively compact. Thus, there exists a subsequence (N_k) which converges weakly*

$$(\mu^{N_k}, \hat{\pi}^{N_k}) \Rightarrow (\mu, \hat{\pi}), \text{ for } k \rightarrow \infty.$$

Moreover, the limit $(\mu, \hat{\pi})$ satisfies

- (a) (μ_t) has a.s. continuous paths,
- (b) and for each component j we have

$$\mu_t(j) = \mu_0(j) + \int_0^t \sum_{i \in S} \mu_s(i) \int q(\{j\}|i, a, \mu_s) \hat{\pi}_s^i(da) ds.$$

5 The Deterministic Limit Model

Consider the following deterministic optimization problem:

$$\begin{aligned}
 (F) \quad & \sup_{\hat{\pi}} \int_0^\infty e^{-\beta t} r(\mu_t, \hat{\pi}_t) dt, \\
 & \text{s.t. } \mu_0 \in \mathbb{P}(S), \hat{\pi}_t^i \in \mathbb{P}(A), \hat{\pi}_t^i(D(i)) = 1,
 \end{aligned}$$

$$\mu_t(j) = \mu_0(j) + \int_0^t \sum_{i \in S} \mu_s(i) \int q(\{j\}|i, a, \mu_s) \hat{\pi}_s^i(da) ds, \quad \forall t \geq 0, j = 1, \dots, |S|.$$

Note that the theory of continuous-time Markov processes implies that μ_t is automatically a distribution. Hence one of the $|S|$ differential equations in (F) may be skipped. Also note that when the transition intensity and the reward are linear in the action, relaxation of the control is unnecessary. We denote the maximal value of this problem by $V^F(\mu_0)$. We show next, that this value provides an asymptotic upper bound to the value of problem (3.3).

Theorem 5.1 *For all $(\mu_0^N) \subset \mathbb{P}_N(S)$, $\mu_0 \in \mathbb{P}(S)$ with $\mu_0^N \Rightarrow \mu_0$ and for all sequences of policies $(\hat{\pi}_t^N)$ we have*

$$\limsup_{N \rightarrow \infty} V_{\hat{\pi}_t^N}^N(\mu_0^N) \leq V^F(\mu_0).$$

Proof According to Theorem 4.3 we can choose a subsequence (N_k) of corresponding state and action processes such that

$$(\mu^{N_k}, \hat{\pi}^{N_k}) \Rightarrow (\mu, \hat{\pi}), \text{ for } k \rightarrow \infty.$$

For convenience we still denote this sequence by (N) . We show that

$$\begin{aligned} \lim_{N \rightarrow \infty} V_{\hat{\pi}_t^N}^N(\mu_0^N) &= \lim_{N \rightarrow \infty} \mathbb{E} \left[\int_0^\infty e^{-\beta t} r(\mu_t^N, \hat{\pi}_t^N) dt \right] \\ &= \mathbb{E} \left[\int_0^\infty e^{-\beta t} r(\mu_t, \hat{\pi}_t) dt \right] \leq V^F(\mu_0). \end{aligned}$$

The last inequality is true due to the fact that by Theorem 4.3 the limit process $(\mu, \hat{\pi})$ satisfies the constraints of problem (F).

Let us show the second equality. We obtain by bounded convergence (r is bounded)

$$\lim_{N \rightarrow \infty} \mathbb{E} \left[\int_0^\infty e^{-\beta t} r(\mu_t^N, \hat{\pi}_t^N) dt \right] = \mathbb{E} \left[\int_0^\infty e^{-\beta t} \lim_{N \rightarrow \infty} r(\mu_t^N, \hat{\pi}_t^N) dt \right].$$

Further we have

$$\begin{aligned} &\left| \int_0^\infty e^{-\beta t} \sum_{i \in S} \int_A r(i, a, \mu_t^N) \hat{\pi}_t^{N,i}(da) \mu_t^N(i) dt \right. \\ &\quad \left. - \int_0^\infty e^{-\beta t} \sum_{i \in S} \int_A r(i, a, \mu_t) \hat{\pi}_t^i(da) \mu_t(i) dt \right| \end{aligned}$$

$$\begin{aligned} &\leq \left| \int_0^\infty e^{-\beta t} \sum_{i \in S} \int_A r(i, a, \mu_t^N) \hat{\pi}_t^{N,i}(da) \mu_t^N(i) dt \right. \\ &\quad \left. - \int_0^\infty e^{-\beta t} \sum_{i \in S} \int_A r(i, a, \mu_t) \hat{\pi}_t^{N,i}(da) \mu_t(i) dt \right| \\ &+ \left| \int_0^\infty e^{-\beta t} \sum_{i \in S} \int_A r(i, a, \mu_t) \hat{\pi}_t^{N,i}(da) \mu_t(i) dt \right. \\ &\quad \left. - \int_0^\infty e^{-\beta t} \sum_{i \in S} \int_A r(i, a, \mu_t) \hat{\pi}_t^i(da) \mu_t(i) dt \right|. \end{aligned}$$

The second expression tends to zero for $N \rightarrow \infty$ due to the definition of the Young topology and the fact that $a \mapsto r(i, a, \mu)$ is continuous by (R2). The first expression can be bounded from above by

$$\begin{aligned} &\int_0^\infty e^{-\beta t} \sum_{i \in S} \int_A \left| r(i, a, \mu_t^N) \mu_t^N(i) - r(i, a, \mu_t) \mu_t(i) \right| \hat{\pi}_t^{N,i}(da) dt \\ &\leq \int_0^\infty e^{-\beta t} \sum_{i \in S} \sup_{a \in D(i)} \left| r(i, a, \mu_t^N) \mu_t^N(i) - r(i, a, \mu_t) \mu_t(i) \right| dt \end{aligned}$$

which also tends to zero for $N \rightarrow \infty$ due to (R1), (R2), Lemma 7.1 and dominated convergence. Thus, the statement follows. \square

On the other hand we are now able to construct a strategy which is asymptotically optimal in the sense that the upper bound in the previous theorem is attained in the limit. Suppose that $(\mu^*, \hat{\pi}^*)$ is an optimal state-action trajectory for problem (F). Then we can consider for the N agents problem the strategy

$$\hat{\pi}_t^{N,i} := \hat{\pi}_t^{*,i}$$

which applies at time t the kernel $\hat{\pi}_t^{*,i}$ irrespective of the state μ_t^N the process is in. More precisely, the considered strategy is deterministic and not a feedback policy.

Theorem 5.2 *Suppose $\hat{\pi}^*$ is an optimal strategy for (F) where the corresponding differential equation in (F) has a unique solution and let $(\mu_0^N) \subset \mathbb{P}_N(S)$ be such that $\mu_0^N \Rightarrow \mu_0 \in \mathbb{P}(S)$. Then if we use strategy $\hat{\pi}^*$ for problem (3.3) for any N we obtain*

$$\lim_{N \rightarrow \infty} V_{\hat{\pi}^*}^N(\mu_0^N) = V^F(\mu_0).$$

Thus, we call $\hat{\pi}^*$ asymptotically optimal.

Proof First note that $\hat{\pi}^*$ is an admissible policy for any N . Further let (μ_t^N) be the corresponding state process when N agents are present. Since the corresponding differential equation in (F) has a unique solution, every subsequence (N_k) is such that

$$\mu^{N_k} \Rightarrow \mu^*, \text{ for } k \rightarrow \infty$$

holds (Theorem 4.3). Using the same arguments as in the last proof we obtain

$$\lim_{N \rightarrow \infty} \mathbb{E} \left[\int_0^\infty e^{-\beta t} r(\mu_t^N, \hat{\pi}_t^*) dt \right] = \mathbb{E} \left[\int_0^\infty e^{-\beta t} r(\mu_t^*, \hat{\pi}_t^*) dt \right] = V^F(\mu_0).$$

Together with the previous theorem, the statement is shown. □

Remark 5.3 (a) In order to guarantee the unique solvability, it is sufficient to assume Lipschitz continuity for $\mu \mapsto q(\{j\}|i, a, \mu)$. More precisely, instead of (Q4) we have to assume (Q4') which is given below. The proof follows from the Theorem of Picard-Lindelöf. Example 5.4 shows what may happen if the differential equation for (μ_t) in (F) has multiple solutions.

(b) Note that the construction of asymptotically optimal policies which we present here, works in the same way when we consider control problems with finite time horizon. I.e. instead of Eq. (3.3) we consider

$$\sup_{\hat{\pi}} \mathbb{E}_{\mu}^{\hat{\pi}} \left[\int_0^T e^{-\beta t} r(\mu_t^N, \hat{\pi}_t) dt + g(\mu_T^N) \right] \tag{5.1}$$

with possibly a terminal reward $g(\cdot)$ for the final state. In this case (F) is given with a finite time horizon

$$\begin{aligned} & \sup_{\hat{\pi}} \int_0^T e^{-\beta t} r(\mu_t, \hat{\pi}_t) dt + g(\mu_T) \\ & \text{s.t. } \mu_0 \in \mathbb{P}(S), \hat{\pi}_t^i \in \mathbb{P}(A), \hat{\pi}_t^i(D(i)) = 1, \\ & \mu_t(j) = \mu_0(j) + \int_0^t \sum_{i \in S} \mu_s(i) \int q(\{j\}|i, a, \mu_s) \hat{\pi}_s^i(da) ds, \\ & \forall t \in [0, T], j = 1, \dots, |S|. \end{aligned} \tag{5.2}$$

Theorem 5.2 holds accordingly.

(c) General statements about the existence of optimal controls in (F) can only be made under additional assumptions. A classical result is the Theorem of Filippov-Cesari (see [32] Theorem 8 in Chapter II.8 for the finite time horizon problem and Theorem 15 in Chapter III.7 for the infinite horizon problem). It states the existence of an optimal control (for the finite horizon problem) under the following assumptions:

- (i) There exist admissible pairs $(\hat{\pi}, \mu)$, (for example by assuming Lipschitz continuity like in (a))
- (ii) A is closed and bounded (which we assume here)
- (iii) μ is bounded for all controls (which we have here)
- (iv) For fixed μ the set $\{(r(\mu, \alpha) + \gamma, f_1(\mu, \alpha)), \gamma \leq 0, \alpha \in A\}$ is convex where f_1 is the r.h.s. of the differential equation in (F).

(d) Suppose we obtain for problem (F) an optimal feedback rule $\hat{\pi}_t(\cdot) = \hat{\pi}(\cdot|\mu_t)$. If $\mu \mapsto \hat{\pi}(\cdot|\mu)$ is continuous, this feedback rule is also asymptotically optimal

for problem Eq. (3.3). The proof can be done in the same way as before. If the mapping is not continuous, the convergence may not hold (see application 6.3).

- (e) Natural extensions of our model that we have not included in the presentation are resource constraints. For example the total sum of fractions of a certain action may be limited, i.e. we restrict the set $\mathbb{P}(A)^{|S|}$ by requiring that $\sum_{i \in S} \hat{\pi}_i^i(\{a^0\}|\mu) \leq c < |S|$ for a certain action $a^0 \in A$. As long as the constraint yields a compact subset of $\mathbb{P}(A)^{|S|}$ our analysis also covers this case.

Example 5.4 In this example we discuss what may happen if the differential equation for (μ_t) in (F) has multiple solutions. Suppose the state space is $S = \{1, 2\}$ and the system is uncontrolled. State 1 is absorbing, i.e. $q(\{1\}|1, \mu) = q(\{2\}|1, \mu) = 0$ (since the system is uncontrolled we skip the action from the notation). So agents can only change from state 2 to 1. The intensity of such a change is

$$q(\{1\}|2, \mu) = \begin{cases} \frac{(\mu_t(1))^{\frac{1}{3}}}{1-\mu_t(1)}, & \text{if } \mu_t(1) \leq 0.99 \\ \frac{0.99^{\frac{1}{3}}}{0.01} & \text{if } \mu_t(1) \geq 0.99. \end{cases}$$

Intensities are bounded and continuous. Since the two probabilities satisfy $\mu_t(1) + \mu_t(2) = 1$ we can concentrate on $\mu_t(1)$. The differential equation for $\mu_t(1)$ in (F) is

$$\mu_t'(1) = \mu_t(1)q(\{1\}|1, \mu_t) + (1 - \mu_t(1))q(\{1\}|2, \mu) = (\mu_t(1))^{\frac{1}{3}}$$

as long as $\mu_t(1) \leq 0.99$. If $\mu_0(1) = 0$, there are two solutions of this initial value problem: $\mu_t(1) \equiv 0$ and $\mu_t(1) = (\frac{2}{3}t)^{\frac{3}{2}}$ for $\mu_t(1) \leq 0.99$. Now consider the following sequence (μ_0^N) : For N even we set $\mu_0^N = (0, 1)$ (all N agents start in state 2), for N odd we set $\mu_0^N = (1/N, N - 1/N)$ (exactly one agent starts in state 1). Obviously $(\mu_0^N) \Rightarrow (0, 1)$. However, when we consider the even subsequence we obtain $\mu_t^N(1) \equiv 0$ since the intensity to change from 2 to 1 remains 0. The uneven subsequence converges against the second solution $\mu_t(1) = (\frac{2}{3}t)^{\frac{3}{2}}$ as long as $\mu_t(1)$ is below 0.99. Thus, when we skip the assumption of a unique solution in Theorem 5.2 we only obtain $\limsup_{N \rightarrow \infty} V_{\hat{\pi}^*}^N(\mu_0^N) \leq V^F(\mu_0)$, see Theorem 5.1.

Under stricter assumptions it is possible to prove that the rate of convergence in the finite horizon problem Eq. (5.1) is $1/\sqrt{N}$. In order to obtain this rate we need Lipschitz conditions on the reward function and the intensity functions. More precisely assume

(R1') For all $(i, a) \in D$ there exists a uniform constant $L_1 > 0$ s.t.

$$|r(i, a, \mu) - r(i, a, \nu)| \leq L_1 \|\mu - \nu\|_{TV}, \quad |g(\mu) - g(\nu)| \leq L_1 \|\mu - \nu\|_{TV}$$

for all $\mu, \nu \in \mathbb{P}(S)$.

(Q4') For all $(i, a) \in D, j \in S$ there exists a uniform constant $L_2 > 0$ s.t.

$$|q(\{j\}|i, a, \mu) - q(\{j\}|i, a, \nu)| \leq L_2 \|\mu - \nu\|_{TV}$$

for all $\mu, \nu \in \mathbb{P}(S)$.




Fig. 1 Colourful lines: State trajectories $\mu_t^N(1)$ for $N = 100$ (red) and $N = 10000$ (green) agents in Example 5.4 when one agent starts in state 1. Black line: Deterministic limit process $\mu_t(1) = (\frac{2}{3}t)^{\frac{3}{2}}$ (Color figure online)

Denote by $\hat{\pi}^*$ the optimal control of the limiting problem (5.2), $V^{F,T}(\mu_0)$ the corresponding value and let

$$V_{\hat{\pi}^*}^{N,T}(\mu_0^N) := \mathbb{E}_{\mu_0^N}^{\hat{\pi}^*} \left[\int_0^T e^{-\beta t} r(\mu_t^N, \hat{\pi}_t^*) dt + g(\mu_T^N) \right].$$

Then we can state the following convergence rate

Theorem 5.5 *In the finite horizon setting under assumption (Q1)–(Q5) with (Q4) replaced by (Q4') and (R1'), (R2), suppose that $\mathbb{E}[\|\mu_0^N - \mu_0\|_{TV}] \leq \frac{L_0}{\sqrt{N}}$ for a constant $L_0 > 0$. Then*

$$\left| V_{\hat{\pi}^*}^{N,T}(\mu_0^N) - V^{F,T}(\mu_0) \right| \leq \frac{\tilde{L}}{\sqrt{N}}$$

for a constant $\tilde{L} > 0$ which is independent of N , but depends on T .

The statement about the convergence rate can be extended to the infinite horizon problem when the discount factor is large enough. Also note that $\mathbb{E}[\|\mu_0^N - \mu_0\|_{TV}] \leq \frac{L_0}{\sqrt{N}}$ is satisfied if e.g. the states of the N agents are sampled i.i.d. from μ_0 .

A direct implementation of policy $\hat{\pi}^*$ in the problem Eq. (3.3) might make it necessary to update the policy continuously. This can be avoided by using the following policy instead. We assume here that $t \mapsto \hat{\pi}_t^*$ is piecewise continuous. Thus, let $(t_n)_{n \in \mathbb{N}}$

be the discontinuity points in time of $\hat{\pi}^*$ and define the set

$$\{T_n^N, n \in \mathbb{N}\} \cup \{t_n, n \in \mathbb{N}\} =: \{\tilde{T}_1^N < \tilde{T}_2^N < \dots\}$$

where T_n^N describes the time of the n -th jump of the N agents process. Then (\tilde{T}_n^N) is the ordered sequence of the time points in this set. Define

$$\pi_t^{N,*} := \sum_{n=0}^{\infty} \hat{\pi}_{\tilde{T}_n}^* \mathbb{1}_{[\tilde{T}_n^N, \tilde{T}_{n+1}^N)}(t). \tag{5.3}$$

The idea of the action process $\pi_t^{N,*}$ is to adapt it to $\hat{\pi}^*$ only when an agent changes its state or when $\hat{\pi}^*$ has a jump, and to keep it constant otherwise. It can be shown that this sequence of policies is also asymptotically optimal.

Theorem 5.6 *Suppose $\hat{\pi}^*$ is a piecewise continuous optimal strategy for (F) where the corresponding differential equation in (F) has a unique solution and let $(\mu_0^N) \subset \mathbb{P}_N(S)$ be such that $\mu_0^N \Rightarrow \mu_0 \in \mathbb{P}(S)$. Then if we use the strategy $(\pi_t^{N,*})$ of Eq. (5.3) for problem (3.3) for any N we obtain*

$$\lim_{N \rightarrow \infty} V_{\hat{\pi}^{N,*}}^N(\mu_0^N) = V^F(\mu_0).$$

Proof In light of the proof of Theorem 5.2 it is enough to show that $\pi^{N,*} \Rightarrow \pi^*$. Indeed, the convergence can be shown \mathbb{P} -a.s. Now $(\pi^{N,*})$ converges in J_1 -topology against π^* on $[0, \infty)$ if and only if $(\pi^{N,*})|_{[0, T]}$, the restriction to $[0, T]$, converges in the finite J_1 -topology to the restriction $\pi_{[0, T]}^*$ for all T which are continuity points of the limit function (see [33] Sect. 16, Lemma 1). Since $\hat{\pi}^*$ is piecewise continuous we can consider the convergence on each compact interval of the partition separately. Indeed we have if $t \in [\tilde{T}_n^N, \tilde{T}_{n+1}^N]$

$$\|\pi_t^{N,*} - \hat{\pi}_t^*\|_{TV} \leq \sup_{s \in [\tilde{T}_n^N, \tilde{T}_{n+1}^N]} \|\hat{\pi}_s^* - \hat{\pi}_t^*\|_{TV}.$$

Since $t \mapsto \hat{\pi}_t^*$ is continuous on this interval and since all $|\tilde{T}_{n+1}^N - \tilde{T}_n^N|$ converge to zero for $N \rightarrow \infty$ uniformly (the jump intensity increases with N) we have that the right hand side converges to zero for $N \rightarrow \infty$ uniformly in t which implies the statement. □

Remark 5.7 Let us briefly discuss the main differences to [8] where a similar model is considered. In [8] the author considers a finite horizon problem where model data is not necessarily stationary, i.e. reward and transition intensities may depend on time. Moreover, he solves the corresponding optimization problems (N -agents and limit problem) via HJB equations. This requires the notion of viscosity solutions and more regularity assumptions in terms of Lipschitz continuity of reward and transition intensities. Using the MDP perspective, we can state our solution theorem for the N agents

problem (in form of a Bellman equation) and the convergence result under weaker continuity conditions. For the convergence to hold we use randomized policies whereas in [8] the author sticks to deterministic policies throughout. The obtained convergence rates under Lipschitz assumptions are the same whereas our proof is simpler and more direct. In [8] the problem is further discussed under stronger assumption. In contrast we present some applications next in order to show how to use the results of the previous sections.

6 Applications

In this section we discuss two applications of the previously derived theorems and one example which shows that state processes under feedback policies do not necessarily have to converge. More precisely we construct in two applications asymptotically optimal strategies for stochastic N agents systems from the deterministic limit problem (F). The advantage of our problem (F) in contrast to the master equation is that it can be solved with the help of *Pontryagin's maximum principle* which gives necessary conditions for an optimal control and is in many cases easier to apply than dynamic programming. For examples see [10, 11, 13, 34] and for the theory see e.g. [32, 35].

6.1 Machine Replacement

The following application is a simplified version of the deterministic control problem in [27]. A mean-field application can be found in [36]. Suppose a company has N statistically equal machines. Each machine can either be in state 0='working' or in state 1='broken', thus $S = \{0, 1\}$. Two actions are available: 0='do nothing' or 1='repair', thus $A = \{0, 1\}$. A working machine does not need repair, so $D(0) = \{0\}$. The transition rates are as follows: A working machine breaks down with fixed rate $\lambda_{wb} > 0$. A broken machine which gets repaired changes to the state 'working' with rate $\lambda_{bw} > 0$. Thus, we can summarize the transition rates of one machine by

$$q(\{1\}|0, 0, \mu_t^N) = \lambda_{wb}, \quad q(\{0\}|1, a_t, \mu_t^N) = \lambda_{bw}\delta_{\{a_t=1\}}.$$

The diagonal elements of the intensity matrix are given by

$$q(\{0\}|0, 0, \mu_t^N) = -\lambda_{wb}, \quad q(\{1\}|1, a_t, \mu_t^N) = -\lambda_{bw}\delta_{\{a_t=1\}},$$

and all other intensities are zero. Obviously (Q1)–(Q5) are satisfied. The initial state of the system is $\mu_0^N = (1, 0)$, i.e. all machines are working in the beginning. Each working machine produces a reward rate $g > 0$ whereas we have to pay a fixed cost of $C > 0$ when we have to call the service for repair, i.e.

$$r(i, a, \mu_t^N) = g\delta_{\{i=0\}} - C\delta_{\{a=1\}}\delta_{\{i=1\}} \frac{1}{1 - \mu_t^N(0)}.$$

Hence we obtain an interaction of the agents in the reward. Note that (R1), (R2) are satisfied. This yields the reward rate for the system

$$r(\mu_t^N, \hat{\pi}_t) = g\mu_t^N(0) - C(1 - \hat{\pi}_t^1(\{0\}|\mu_t^N)).$$

Thus, problem (F) in this setting is given by (we denote the limit by $(\mu_t(0), \mu_t(1)) =: (\mu_t^0, 1 - \mu_t^0)$ and let $\alpha_t^0 := \hat{\pi}_t^1(\{0\}|\mu_t)$):

$$\begin{aligned} (F) \quad & \sup_{(\alpha_t)} \int_0^T g \cdot \mu_t^0 - C \cdot (1 - \alpha_t^0) dt, \\ & s.t. \text{ for all } t \in [0, T] \\ & \mu_t^0 = 1 + \int_0^t \lambda_{bw}(1 - \mu_s^0)(1 - \alpha_s^0) - \lambda_{wb}\mu_s^0 ds. \end{aligned}$$

We briefly explain how to solve this problem using Pontryagin’s maximum principle. The Hamiltonian function to (F) is given by

$$\begin{aligned} H(\mu_t^0, \alpha_t^0, p_t, t) &= g\mu_t^0 - C(1 - \alpha_t^0) + p_t(\lambda_{bw}(1 - \mu_t^0)(1 - \alpha_t^0) - \lambda_{wb}\mu_t^0) \\ &= (1 - \alpha_t^0)(\lambda_{bw}p_t(1 - \mu_t^0) - C) + g\mu_t^0 - \lambda_{wb}p_t\mu_t^0 \end{aligned}$$

where (p_t) is the adjoint function. Pontryagin’s maximum principle yields the following sufficient conditions for optimality ([32, 35]):

Lemma 6.1 *The control $(\alpha_t^{0,*})$ with the associated trajectory $(\mu_t^{0,*})$ is optimal for (F) if there exists a continuous and piecewise continuously differentiable function (p_t) such that for all $t > 0$:*

- (i) $\alpha_t^{0,*}$ maximizes $\alpha \mapsto H(\mu_t^0, \alpha, p_t, t)$ for $\alpha \in [0, 1]$,
- (ii) $\dot{p}_t = -g + p_t(\lambda_{wb} + \lambda_{bw}(1 - \alpha_t^0))$ at those points where p_t is differentiable,
- (iii) $p(T) = 0$.

Inspecting the Hamiltonian it is immediately clear from (i) that the optimal control is essentially ‘bang-bang’. For a numerical illustration we solved (F) for the parameters $C = 1, g = 2, \lambda_{wb} = 1, \lambda_{bw} = 2$ and $T = 4$. Here it is optimal to do nothing until time point $t^* = \ln 2$. Then it is optimal to repair the fraction $\alpha^{0,*} = 1/2$ of the broken machines which keeps the number of working machines at $1/2$. Finally, $\ln 2$ time units before the end, we do again nothing and wait until the end of the time horizon. A numerical illustration of the optimal trajectory $\mu_t^{0,*}$ of the deterministic problem together with simulated paths under this policy for different number of N can be found in Fig. 2, left. A number of different simulations for $N = 1000$ are shown in Fig. 2, right. The simulated paths are quite close to the deterministic trajectory.

The optimal value in the deterministic model is $V^F(1, 0) = \frac{9}{2} - \frac{3}{2} \ln(2) \approx 3.4603$. If we simulate ten times the trajectory of the state process for $N = 1000$ machines while following the asymptotically optimal policy and take the average of the respective values, we obtain a mean of 3.43612 which is slightly less than the value for (F), cp. Theorem 5.1.

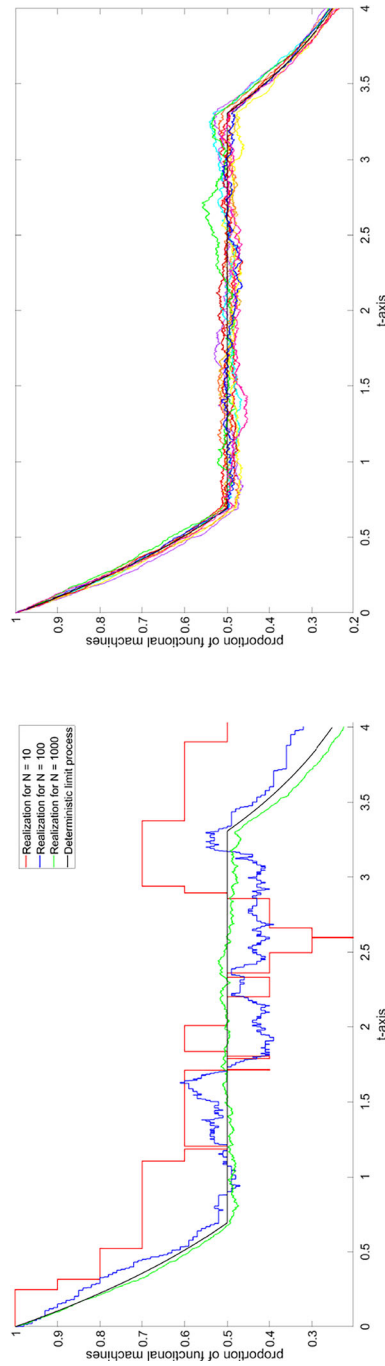


Fig. 2 Left: State trajectories for different numbers N of machines executing the optimal control for (F) . Right: Ten state trajectories for $N = 1000$ machines executing the asymptotically optimal control (Color figure online)




Fig. 3 Transition intensities of one device between the possible states

6.2 Spreading Malware

This example is based on the deterministic control model considered in [28], see resp. [16], and treats the propagation of a virus in a mobile wireless network. It is based on the classical SIR model by Kermack-McKendrick, [37]. Suppose there are N devices in the network. A device can be in one of the following states: *Susceptible* (S), *Infective* (I), *Dead* (D) or *Recovered* (R). A device is in the susceptible state if it is not contaminated yet, but prone to infection. A device is infective if it is contaminated by the virus. It is dead if the virus has destroyed the software and recovered if the device has already a security patch which makes it immune to the virus. The states D and R are absorbing. The joint process $\mu_t^N = (S_t^N, I_t^N, D_t^N, R_t^N)$ is a controlled continuous-time Markov chain where X_t^N represents the fraction of devices in state $X \in \{S, I, D, R\}$. The control is a strategy of the virus which chooses the rate $a(t) \in [0, \bar{a}]$, at which infected devices are destroyed. In this model we have $S_t^N + I_t^N + D_t^N + R_t^N = 1$ and $S_t^N, I_t^N, D_t^N, R_t^N \geq 0$. The transition rates of one device are as follows: A susceptible device gets infected with rate $\lambda_{SI}I_t$ with $\lambda_{SI} > 0$. The rate is proportional to the number of infected devices and we thus have an interaction of one agent with the empirical distribution of the others. And it gets recovered with rate $\lambda_{SR} > 0$ which is the rate the security patch is distributed. An infected device gets killed by the virus with rate $a(t) \in [0, \bar{a}]$ chosen by the attacker and gets recovered at rate $\lambda_{IR} > 0$. The rates are shown in Fig. 3.

The intensities of one device at time t are summarized by

$$q(\{I\}|S, \cdot, \mu_t^N) = \lambda_{SI}I_t^N, \quad q(\{R\}|S, \cdot, \mu_t^N) = \lambda_{SR},$$

$$q(\{D\}|I, a_t, \mu_t^N) = a_t, \quad q(\{R\}|I, \cdot, \mu_t^N) = \lambda_{IR}.$$

Thus, the diagonal elements of the intensity matrix are given by

$$q(\{S\}|S, \cdot, \mu_t^N) = -\lambda_{SI}I_t^N - \lambda_{SR}, \quad q(\{I\}|I, a_t, \mu_t^N) = -a_t - \lambda_{IR},$$

$$q(\{D\}|D, \cdot, \mu_t^N) = q(\{R\}|R, \cdot, \mu_t^N) = 0$$

and all other intensities are zero. Note that (Q1)–(Q5) are satisfied and that since the intensities are linear in a , there is no need for a relaxed control. The initial state of the

network is $\mu_0^N = (S_0^N, I_0^N, D_0^N, R_0^N) = (1 - I_0, I_0, 0, 0)$ with $0 < I_0 < 1$. The aim of the virus is to produce as much damage as possible over the time interval $[0, T]$, evaluated by

$$\mathbb{E} \left[D_T^N + \frac{1}{T} \int_0^T (I_t^N)^2 dt \right]$$

which is given when we choose $r(i, a, \mu) = \frac{1}{T}(\mu(2))^2$ (the second component of μ squared) and an appropriate terminal reward. (R1) and (R2) are satisfied. Thus, problem (F) in this setting is given by (we denote the limit by $\mu_t = (S_t, I_t, D_t, R_t)$)

$$\begin{aligned} (F) \quad & \sup_{(a_t)} D_T + \frac{1}{T} \int_0^T I_t^2 dt, \\ & \text{s.t. } a_t \in [0, \bar{a}], \text{ and for all } t \in [0, T] \\ & S_t = 1 - I_0 + \int_0^t -\lambda_{SI} I_s S_s - \lambda_{SR} S_s ds, \\ & I_t = I_0 + \int_0^t \lambda_{SI} I_s S_s - \lambda_{IR} I_s - a_t I_s ds, \\ & D_t = \int_0^t a_t I_s ds. \end{aligned}$$

A solution of this deterministic control problem can be found in [28]. It is shown there that a critical time point $t_1 \in [0, T]$ exists such that $a_t = 0$ on $t \in [0, t_1]$ and $a_t = \bar{a}$ on $t \in (t_1, T]$. Thus, the attacker is not destroying devices from the beginning because this lowers the number of devices which can get infected. Instead, she first waits to get more infected devices before setting the kill rate to a maximum.

A numerical illustration can be found in Fig. 4. There we can see the trajectories of the optimal state distribution in (F) and simulated paths for $N = 1000$ devices for $\lambda_{SI} = 0.6, \lambda_{SR} = \lambda_{IR} = 0.2, \bar{a} = 1, T = 10$. The optimal time point for setting a_t to the maximum is here 4.9. The simulated paths are almost indistinguishable from the deterministic trajectories.

6.3 Resource Competition

This example shows that feedback policies in the deterministic problem are not necessarily asymptotically optimal when implemented in the N agents problem. The infinite horizon problem (F) could also be solved using an HJB equation which would provide (under sufficient regularity) a feedback control $\hat{\pi}(\cdot|\mu)$. I.e. we obtain the optimal control by $\hat{\pi}_t = \hat{\pi}(\cdot|\mu_t)$. This feedback function could also be used in the N agents model. However, in this case convergence of the N agents model to the deterministic model like in Theorem 5.2 is not guaranteed. Convergence may fail when discontinuities in the feedback function are present. The example is an adaption of the queuing network considered in [38, 39] to our setting. Suppose the state space is given by $S = \{1, 2, 3, 4, 5, 6, 7, 8\}$. Agents starting in state 1 change to state 2, then 3 and

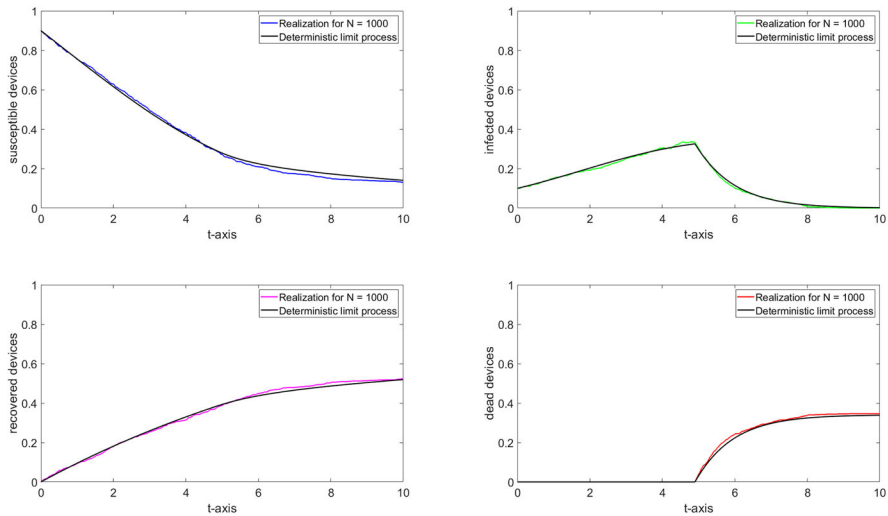


Fig. 4 State trajectories for $N = 1000$ devices under optimal control for $\lambda_{SI} = 0.6, \lambda_{SR} = \lambda_{IR} = 0.2, \bar{a} = 1, T = 10$. (Color figure online)

are finally absorbed in state 4. Agents starting in state 5 change to state 6, then 7 and are finally absorbed in state 8. The aim is to get the agents in the absorbing states as quickly as possible by activating the intensities in states 2,3,6 and 7. The intensity for leaving states 1 and 5 is $\lambda_1 = \lambda_5 = 1$, the full intensity for leaving states 2 and 6 is $\lambda_2 = \lambda_6 = 6$ and finally the full intensity for leaving states 3 and 7 is $\lambda_3 = \lambda_7 = 1.5$. The action space is $A = \{0, 1\}$ where actions have to be taken in states 2, 3, 6 and 7 and determine the activation of the transition intensity. Action $a = 0$ means that the intensity is deactivated and $a = 1$ that it is fully activated. There is a resource constraint such that the sum of the activation probabilities in states 2 and 7 as well as the sum of the activation probabilities in states 3 and 6 are constraint by 1 (see remark on p.13). When we denote the randomized control by $\hat{\pi}_t^2 = a_t, \hat{\pi}_t^7 = 1 - a_t, \hat{\pi}_t^6 = b_t, \hat{\pi}_t^3 = 1 - b_t, a_t, b_t \in [0, 1]$ then the intensities are given by

$$q(\{3\}|2, a_t, \mu_t^N) = a_t \lambda_2, \quad q(\{4\}|3, 1 - b_t, \mu_t^N) = (1 - b_t) \lambda_3,$$

$$q(\{7\}|6, b_t, \mu_t^N) = b_t \lambda_6, \quad q(\{8\}|7, 1 - a_t, \mu_t^N) = (1 - a_t) \lambda_7.$$

An illustration of this model can be seen in Fig. 5.

The initial state distribution is given by $\mu_0 = (\frac{5}{14}, \frac{1}{14}, \frac{1}{14}, 0, \frac{5}{14}, \frac{1}{14}, \frac{1}{14}, 0)$ where we assume for the simulation that we have $N = 1400$ agents. Now suppose further that agents in the absorbing states 4 and 8 produce no cost whereas agents in state 3 and 7 are the most expensive as soon as there are at least 0.01% of the population present. This optimization criterion leads to a priority rule where agents in state 3 receive priority (and thus full capacity) over those in state 6 (as long as there are at least 0.01% present) and agents in state 7 receive priority (and thus full capacity) over those in state 2 (as long as there are at least 0.01% present). In the deterministic problem the priority rule can be implemented such that once the number of agents in state 3 and 7 fall to




Fig. 5 Transition intensities of one agent for the resource constraint problem

the threshold of 0.01% of the population it is possible to keep this level. This is not possible in the N agents problem. The priority switch leads to blocking the agents in the other line, see Fig. 6. The blue line shows the state trajectories in the deterministic model. The red line is a realization of the system for $N = 1400$ agents where we use the deterministic open-loop control of Theorem 5.2. We see that the state processes converge. Finally the green line is a realization of the $N = 1400$ agents model under the priority rule. We can see that here state processes do not converge.

7 Appendix

7.1 Auxiliary Result

Lemma 7.1 *Let X be a separable metric space, Y be compact metric and $f : X \times Y \rightarrow \mathbb{R}$ continuous. Then $x_n \rightarrow x$ for $n \rightarrow \infty$ implies*

$$\lim_{n \rightarrow \infty} \sup_{y \in Y} |f(x_n, y) - f(x, y)| = 0.$$

For a proof see e.g. Lemma B.12, [40].

7.2 Proof of Theorem 3.1

First of all observe that the reward function r in Eq. (2.3) in the N agents problem is symmetric, i.e. $r(\mathbf{x}, \mathbf{a}) = r(s(\mathbf{x}), s(\mathbf{a}))$ for any permutation $s(\cdot)$ of the vectors. Moreover, the agent transition intensities $q(\cdot|i, a, \mu[\mathbf{X}_t])$ depend only on the own state of the agent and on $\mu[\mathbf{X}_t]$. Thus, the optimal policy in the N agents problem at time t only depends on $\mu[\mathbf{X}_t]$. Now for a decision rule π for the N agents problem define for all states $i \in S$:

$$\hat{\pi}^i(da|\mu) := \frac{1}{N\mu(i)} \sum_{k=1}^N \pi^k(da|\mathbf{x}) \mathbb{1}_{\{x^k=i\}}$$




Fig. 6 State trajectories for $N = 1400$ agents. Deterministic trajectory (blue), realization under deterministic open loop (red), realization under feedback priority rule (green) (Color figure online)

where $\mu = \mu[\mathbf{x}]$. On the right-hand side we consider all agents in state i and take a convex combination of their action distributions as the action distribution in state i . If π depends only on $\mu[\mathbf{x}]$, then this is also true for $\hat{\pi}$.

Choosing $\hat{\pi}$ in the measure-valued MDP yields the reward (again $\mu = \mu[\mathbf{x}]$)

$$\begin{aligned} r(\mu, \hat{\pi}) &= \sum_{i \in S} \int r(i, a, \mu) \frac{1}{N\mu(i)} \sum_{k=1}^N \pi^k(da|\mathbf{x}) \mathbb{1}_{\{x^k=i\}} \mu(i) \\ &= \frac{1}{N} \sum_{k=1}^N \sum_{i \in S} \mathbb{1}_{\{x^k=i\}} \int r(i, a, \mu) \pi^k(da|\mathbf{x}) = r(\mathbf{x}, \pi). \end{aligned}$$

Thus, the reward in both formulations is the same. Finally the transition intensity in the N agents model that one agent changes its state from i to j is given by (again $\mu = \mu[\mathbf{x}]$)

$$\begin{aligned} &\sum_{k=1}^N \mathbb{1}_{\{x^k=i\}} \int q(\{j\}|i, a, \mu) \pi^k(da|\mathbf{x}) \\ &= N\mu(i) \int_A q(\{j\}|i, a, \mu) \frac{1}{N\mu(i)} \sum_{k=1}^N \pi^k(da|\mathbf{x}) \mathbb{1}_{\{x^k=i\}} \\ &= N\mu(i) \int_A q(\{j\}|i, a, \mu) \hat{\pi}^i(da|\mu) = q(\{i \rightarrow j\}|\mu, \hat{\pi}). \end{aligned}$$

Thus, the empirical measure process of the N agents problem is statistically equal to the measure-valued MDP process and they produce the same expected reward under measure-dependent policies which implies the result. A formal proof has to be done by induction like in [20] Thm. 3.3.

7.3 Proof of Lemma 4.2

First we show that $M_t^N(j)$ is bounded for fixed t :

$$\begin{aligned} |M_t^N(j)| &= \left| \mu_t^N(j) - \mu_0^N(j) - \int_0^t \sum_{i \in S} \mu_s^N(i) \int q(\{j\}|i, a, \mu_s^N) \hat{\pi}_s^{N,i}(da) ds \right| \\ &\leq |\mu_t^N(j) - \mu_0^N(j)| + \int_0^t \sum_{i \in S} \mu_s^N(i) \int |q(\{j\}|i, a, \mu_s^N)| \hat{\pi}_s^{N,i}(da) ds \\ &\leq 1 + q_{max} \cdot t < \infty \end{aligned}$$

Therefore $(M_t^N(j))_{t \geq 0}$ are square-integrable martingales. Now we take advantage of the fact that there are only jumps of height $\frac{1}{N}$ in our model, since no two agents change

their state simultaneously. With the quadratic variation of the process we obtain

$$\begin{aligned} \mathbb{E}[(M_t^N(j))^2] &= \mathbb{E}[\langle M_t^N(j) \rangle] \leq \frac{1}{N^2} \mathbb{E}[\#\text{jumps in } [0, t]] \\ &\leq \frac{1}{N^2} \cdot N \cdot q_{\max} \cdot t = \frac{1}{N} \cdot q_{\max} \cdot t \xrightarrow{N \rightarrow \infty} 0. \end{aligned}$$

Doob’s L^p -inequality provides on $[0, t]$

$$\mathbb{E}[(\sup_{s \in [0, t]} M_s^N(j))^2] \leq 4 \cdot \mathbb{E}[(M_t^N(j))^2] \xrightarrow{N \rightarrow \infty} 0.$$

Thus for the sequence $(\sup_{s \in [0, t]} M_s^N(j))_{N \in \mathbb{N}}$ it holds that

$$\sup_{s \in [0, t]} M_s^N(j) \xrightarrow{L^2} 0.$$

Now we can find a suitable probability space $(\Omega, \mathcal{F}, \mathbb{P})$, such that for \mathbb{P} -almost all $\omega \in \Omega$ the sequence of functions $((M_s^N(j)(\omega))_{s \in [0, t]})_{N \in \mathbb{N}}$ converges uniformly to the zero-function.

The finite-dimensional distributions with arbitrary time-points $t_1, \dots, t_k \in [0, t]$ then obviously fulfill

$$(M_{t_1}^N(j), \dots, M_{t_n}^N(j)) \xrightarrow{\text{a.s.}} (0, \dots, 0)$$

and therefore in particular

$$(M_{t_1}^N(j), \dots, M_{t_n}^N(j)) \Rightarrow (0, \dots, 0).$$

Here \Rightarrow is the usual weak convergence of random vectors in \mathbb{R}^n . To apply Theorem VI.16 in [41] we check Aldous’ condition. Let (δ_N) be a sequence of positive numbers with $\delta_N \rightarrow 0$ and (σ_N) a sequence of stopping times w.r.t. (\mathcal{F}_t^N) with values in $[0, t]$. Then we have

$$\mathbb{E}[(M_{\sigma_N}^N(j))^2] \leq \mathbb{E}[(\sup_{s \in [0, t]} M_s^N(j))^2] \leq 4 \cdot \mathbb{E}[(M_t^N(j))^2] \xrightarrow{N \rightarrow \infty} 0.$$

Further, for N sufficiently large it holds that

$$\mathbb{E}[(M_{\sigma_N + \delta_N}^N(j))^2] \leq \mathbb{E}[(\sup_{s \in [0, 2t]} M_s^N(j))^2] \leq 4 \cdot \mathbb{E}[(M_{2t}^N(j))^2] \xrightarrow{N \rightarrow \infty} 0.$$

Therefore $M_{\sigma_N}^N(j)$ and $M_{\sigma_N + \delta_N}^N(j)$ converge in L^2 to 0 (and thus their difference). Hence, the conditions of Theorem VI.16 in [41] are fulfilled, and the sequence $(M_t^N(j))$ converges weakly on $[0, \infty)$ towards 0 in the sense of the Skorokhod J_1 -metric.

7.4 Proof of Theorem 4.3

We start by showing the relative compactness of a sequence $(\mu^N)_N$. We use Theorem 2.7 in [42]. The sequence $(\mu^N)_N$ has paths in $D_{\mathbb{P}(S)}[0, \infty)$, where $\mathbb{P}(S)$ is complete and separable with respect to the total variation distance.

In what follows let σ be an arbitrary (\mathcal{F}_t^N) -stopping time with $\sigma \leq T$ a.s.

For every $\varepsilon > 0$ and rational $t \geq 0$ choose the compact set $\Gamma_{t,\varepsilon} \equiv \mathbb{P}(S)$. Then we obtain by construction of the model

$$\mathbb{P}(\mu_t^N \in \Gamma_{t,\varepsilon}) = 1.$$

Moreover, for every $T > 0$ it holds that

$$\begin{aligned} & \lim_{\delta \rightarrow 0} \limsup_{N \rightarrow \infty} \sup_{\sigma} \mathbb{E}[\min\{1, \|\mu_{\sigma}^N - \mu_{\sigma+\delta}^N\|_{TV}\}] \\ & \leq \lim_{\delta \rightarrow 0} \limsup_{N \rightarrow \infty} \sup_{\sigma} \mathbb{E}[\|\mu_{\sigma}^N - \mu_{\sigma+\delta}^N\|_{TV}] \\ & \leq \lim_{\delta \rightarrow 0} \limsup_{N \rightarrow \infty} \sup_{\sigma} \mathbb{E}[\# \text{ state changes in } [\sigma, \sigma + \delta]] \cdot \frac{1}{N} \\ & \leq \lim_{\delta \rightarrow 0} \limsup_{N \rightarrow \infty} \sup_{\sigma} N \cdot q_{max} \cdot \delta \cdot \frac{1}{N} = 0. \end{aligned}$$

The second inequality holds because $\|\mu_s^N - \mu_t^N\|_{TV} = \frac{1}{N}$, provided that in $[s, t]$ only one state change occurs, i.e. one agent changes its state. Theorem 2.7 in [42] now states that $(\mu^N)_N$ is relatively compact.

Since \mathcal{R} is compact, so is $\mathcal{R}^{|\mathcal{S}|}$ and we obtain directly the relative compactness of $(\hat{\pi}^N)_N$. The relative compactness of the sequence of state-action-processes $(\mu^N, \hat{\pi}^N)_N$ then follows by Proposition 3.2.4 in [43]. Thus, a converging subsequence exists. To ease the notation we will still denote it by (N) .

To prove the continuity of the limit state process define for arbitrary $\mu \in D_{\mathbb{P}(S)}[0, \infty)$

$$\begin{aligned} J(\mu, u) &= \sup_{0 \leq t \leq u} \|\mu_t - \mu_{t-}\|_{TV}. \\ J(\mu) &= \int_0^\infty e^{-u} J(\mu, u) du. \end{aligned}$$

For the sequence of state processes $(\mu^N)_N$ we get

$$\lim_{N \rightarrow \infty} J(\mu^N) = \lim_{N \rightarrow \infty} \int_0^\infty e^{-u} \sup_{0 \leq t \leq u} \|\mu_t^N - \mu_{t-}^N\|_{TV} du \leq \lim_{N \rightarrow \infty} \frac{1}{N} = 0.$$

We exploit the fact that there can be at most jumps of height $\frac{1}{N}$ in the state processes with N agents. Theorem 3.10.2 a) in [43] then implies the a.s. continuity of the limit state process $(\mu_t^*)_{t \geq 0}$.

In particular, due to the Skorokhod representation theorem we find a probability space such that convergence of $\mu^N \Rightarrow \mu^*$ holds almost surely in J_1 and is uniformly on compact sets such as $[0, t]$ since μ^* is a.s. continuous (see p. 383 in [44]). Thus, component-wise for almost all ω in the probability space above we obtain:

$$\lim_{N \rightarrow \infty} \sup_{0 \leq s \leq t} \|\mu_s^N(\omega) - \mu_s^*(\omega)\|_{TV} = 0$$

for every $t \in [0, \infty)$.

Finally we have to take the limit $N \rightarrow \infty$ in Eq. (4.2). By the previous Lemma 4.2 we know that the martingale on the left-hand side converges to zero and that $\mu_t^N(\omega) \rightarrow \mu_t^*$. Now consider the integral on the right-hand side:

$$\begin{aligned} & \left| \int_0^t \sum_{i \in S} \mu_s^N(i) \int q(\{j\} | i, a, \mu_s^N) \hat{\pi}_s^{N,i}(da) ds - \int_0^t \sum_{i \in S} \mu_s^*(i) \int q(\{j\} | i, a, \mu_s^*) \hat{\pi}_s^{*,i}(da) ds \right| \\ & \leq \left| \int_0^t \sum_{i \in S} \mu_s^N(i) \int q(\{j\} | i, a, \mu_s^N) \hat{\pi}_s^{N,i}(da) ds - \int_0^t \sum_{i \in S} \mu_s^*(i) \int q(\{j\} | i, a, \mu_s^*) \hat{\pi}_s^{N,i}(da) ds \right| \\ & \quad + \left| \int_0^t \sum_{i \in S} \mu_s^*(i) \int q(\{j\} | i, a, \mu_s^*) \hat{\pi}_s^{N,i}(da) ds - \int_0^t \sum_{i \in S} \mu_s^*(i) \int q(\{j\} | i, a, \mu_s^*) \hat{\pi}_s^{*,i}(da) ds \right|. \end{aligned}$$

The second expression tends to 0 for $N \rightarrow \infty$ due to the definition of the Young topology and the fact that $a \mapsto q(\{j\} | i, a, \mu_s^*)$ is continuous by assumption. The first expression can be bounded by

$$\begin{aligned} & \int_0^t \sum_{i \in S} \int \left| \mu_s^N(i) q(\{j\} | i, a, \mu_s^N) - \mu_s^*(i) q(\{j\} | i, a, \mu_s^*) \right| \hat{\pi}_s^{N,i}(da) ds \\ & \leq \int_0^t \sum_{i \in S} \sup_{a \in D(i)} \left| \mu_s^N(i) q(\{j\} | i, a, \mu_s^N) - \mu_s^*(i) q(\{j\} | i, a, \mu_s^*) \right| ds \end{aligned}$$

which also tends to zero due to dominated convergence, (Q4),(Q5) and Lemma 7.1. Now putting things together, Eq. (4.2) implies that the limit satisfies the stated differential equation.

7.5 Proof of Theorem 5.5

Fix $\hat{\pi}^*$ and let (μ_t) be the unique solution of

$$\mu_t(j) = \mu_0(j) + \int_0^t \sum_{i \in S} \mu_s(i) \int q(\{j\} | i, a, \mu_s) \hat{\pi}_s^{*,i}(da) ds, \quad \forall t \in [0, T], j = 1, \dots, |S|.$$

Further recall from Eq. (4.2)

$$\begin{aligned} \mu_t^N(j) &= \mu_0^N(j) + \int_0^t \sum_{i \in S} \mu_s^N(i) \int q(\{j\}|i, a, \mu_s^N) \hat{\pi}_s^{*,i}(da) ds + M_t^N(j), \\ \forall t \in [0, T], j &= 1, \dots, |S|. \end{aligned}$$

Now we obtain

$$\begin{aligned} & \left| V_{\pi^*}^{N,T}(\mu_0^N) - V^{F,T}(\mu_0) \right| \\ & \leq \mathbb{E}_{\mathbf{x}}^{\hat{\pi}^*} \left[\int_0^T e^{-\beta t} |r(\mu_t^N, \hat{\pi}_t^*) - r(\mu_t, \hat{\pi}_t^*)| dt + |g(\mu_T^N) - g(\mu_T)| \right] \\ & \leq \mathbb{E}_{\mathbf{x}}^{\hat{\pi}^*} \left[\int_0^T e^{-\beta t} L_3 \|\mu_t^N - \mu_t\|_{TV} dt + L_1 \|\mu_T^N - \mu_T\|_{TV} \right]. \end{aligned}$$

Note that (R1') implies the Lipschitz continuity of $\mu \mapsto r(\mu, \pi)$ with a constant $L_3 > 0$ since r is bounded. Next consider

$$\begin{aligned} & |\mu_t^N(j) - \mu_t(j)| \\ & \leq |\mu_0^N(j) - \mu_0(j)| + \int_0^t \sum_{i \in S} \mu_s^N(i) \int |q(\{j\}|i, a, \mu_s^N) - q(\{j\}|i, a, \mu_s)| \hat{\pi}_s^{*,i}(da) ds \\ & \quad + \int_0^t \sum_{i \in S} |\mu_s^N(i) - \mu_s(i)| \int q(\{j\}|i, a, \mu_s) \hat{\pi}_s^{*,i}(da) ds + |M_t^N(j)| \\ & \leq |\mu_0^N(j) - \mu_0(j)| + (L_2 + 2q_{max}) \int_0^t \|\mu_s^N - \mu_s\|_{TV} ds + |M_t^N(j)|. \end{aligned}$$

For the last term we have already shown in Lemma 4.2 that

$$\mathbb{E}_{\mathbf{x}}^{\hat{\pi}^*} [|M_t^N(j)|] \leq \sqrt{\mathbb{E}_{\mathbf{x}}^{\hat{\pi}^*} [(M_t^N(j))^2]} \leq \frac{\sqrt{q_{max} t}}{\sqrt{N}}.$$

Thus, from the two previous inequalities we get with $L_4 := |S| \sqrt{q_{max} T} / 2$ that

$$\begin{aligned} \mathbb{E}_{\mathbf{x}}^{\hat{\pi}^*} [\|\mu_t^N - \mu_t\|_{TV}] & \leq \mathbb{E}_{\mathbf{x}}^{\hat{\pi}^*} [\|\mu_0^N - \mu_0\|_{TV}] + \frac{L_4}{\sqrt{N}} + L_5 \int_0^t \mathbb{E}_{\mathbf{x}}^{\hat{\pi}^*} [\|\mu_s^N - \mu_s\|_{TV}] ds \\ & \leq \frac{L_6}{\sqrt{N}} + L_5 \int_0^t \mathbb{E}_{\mathbf{x}}^{\hat{\pi}^*} [\|\mu_s^N - \mu_s\|_{TV}] ds. \end{aligned}$$

Finally, Gronwall's inequality implies that for all $t \in [0, T]$

$$\mathbb{E}_{\mathbf{x}}^{\hat{\pi}^*} [\|\mu_t^N - \mu_t\|_{TV}] \leq \frac{L_7}{\sqrt{N}} e^{L_5 T}$$

which in turn implies the statement.

Acknowledgements The authors would like to thank two anonymous referees for their comments which helped to improve the paper.

Funding Open Access funding enabled and organized by Projekt DEAL.

Declarations

Competing interests The authors have no relevant financial or non-financial interests to disclose.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Bortolussi, L., Hillston, J., Latella, D., Massink, M.: Continuous approximation of collective system behaviour: A tutorial. *Perform. Eval.* **70**(5), 317–349 (2013)
2. Kolesnichenko, A., Senni, V., Pourranjbar, A., Remke, A.: Applying mean-field approximation to continuous time Markov chains. *Stochastic Model Checking. Rigorous Dependability Analysis Using Model Checking Techniques for Stochastic Systems: International Autumn School, ROCKS 2012, Vahrn, Italy, October 22–26, 2012, Advanced Lectures*, pp. 242–280 (2014)
3. Kurtz, T.G.: Solutions of ordinary differential equations as limits of pure jump Markov processes. *J. Appl. Probab.* **7**(1), 49–58 (1970)
4. Ball, K., Kurtz, T.G., Popovic, L., Rempala, G.: Asymptotic analysis of multiscale approximations to reaction networks. *Ann. Appl. Probab.* **16**(4), 1925–1961 (2006)
5. Darling, R.W., Norris, J.R.: Differential equation approximations for Markov chains. *Probab. Surv.* **5**, 37–79 (2008)
6. Aspirot, L., Mordecki, E., Rubino, G.: Fluid limits applied to peer to peer network analysis. In: 2011 Eighth International Conference on Quantitative Evaluation of Systems, pp. 13–20. IEEE (2011)
7. Kyrianiou, A., Peñaloza, L., Rogers, T.: The replicator coalescent. Preprint at <http://arxiv.org/abs/2207.00998> (2022)
8. Cecchin, A.: Finite state n-agent and mean field control problems. *ESAIM* **27**, 31 (2021)
9. Meyn, S.: Stability and optimization of queueing networks and their fluid models. *Lect. Appl. Math. Am. Math. Soc.* **33**, 175–200 (1997)
10. Avram, F., Bertsimas, D., Ricard, M.: Fluid models of sequencing problems in open queueing networks; an optimal control approach. *Inst. Math. Appl.* **71**, 199 (1995)
11. Weiss, G.: Optimal draining of fluid re-entrant lines: some solved examples. *Stoch. Netw. Theory Appl.* **4**, 19–34 (1996)
12. Bäuerle, N.: Asymptotic optimality of tracking policies in stochastic networks. *Ann. Appl. Probab.* **10**(4), 1065–1083 (2000)
13. Bäuerle, N.: Optimal control of queueing networks: an approach via fluid models. *Adv. Appl. Probab.* **34**(2), 313–328 (2002)
14. Čudina, M., Ramanan, K.: Asymptotically optimal controls for time-inhomogeneous networks. *SIAM J. Control Optim.* **49**(2), 611–645 (2011)
15. Yin, G.G., Zhang, Q.: *Continuous-Time Markov Chains and Applications: A Singular Perturbation Approach*, vol. 37. Springer, Cham (2012)
16. Gast, N., Gaujal, B., Le Boudec, J.-Y.: Mean field for Markov decision processes: from discrete to continuous optimization. *IEEE Trans. Autom. Control* **57**(9), 2266–2280 (2012)
17. Carmona, R., Laurière, M., Tan, Z.: Model-free mean-field reinforcement learning: mean-field MDP and mean-field q-learning. Preprint at <http://arxiv.org/abs/1910.12802> (2019)

18. Motte, M., Pham, H.: Mean-field Markov decision processes with common noise and open-loop controls. *Ann. Appl. Probab.* **32**(2), 1421–1458 (2022)
19. Motte, M., Pham, H.: Quantitative propagation of chaos for mean field Markov decision process with common noise. *Electron. J. Probab.* **28**, 1–24 (2023)
20. Bäuerle, N.: Mean field Markov decision processes. *Appl. Math. Optim.* **88**(1), 12 (2023)
21. Gomes, D.A., Mohr, J., Souza, R.R.: Continuous time finite state mean field games. *Appl. Math. Optim.* **68**(1), 99–143 (2013)
22. Basna, R., Hilbert, A., Kolokoltsov, V.N.: A $1/n$ Nash equilibrium for non-linear Markov games of mean-field-type on finite state space. Preprint at <http://arxiv.org/abs/1403.0426> (2014)
23. Bayraktar, E., Cohen, A.: Analysis of a finite state many player game using its master equation. *SIAM J. Control Optim.* **56**(5), 3538–3568 (2018)
24. Cecchin, A., Fischer, M.: Probabilistic approach to finite state mean field games. *Appl. Math. Optim.* **81**(2), 253–300 (2020)
25. Belak, C., Hoffmann, D., Seifried, F.T.: Continuous-time mean field games with finite state space and common noise. *Appl. Math. Optim.* **84**, 3173–3216 (2021)
26. Cardaliaguet, P., Delarue, F., Lasry, J.-M., Lions, P.-L.: *The Master Equation and the Convergence Problem in Mean Field Games* (ams-201). Princeton University Press, Princeton (2019)
27. Thompson, G.L.: Optimal maintenance policy and sale date of a machine. *Manag. Sci.* **14**(9), 543–550 (1968)
28. Khouzani, M., Sarkar, S., Altman, E.: Maximum damage malware attack in mobile wireless networks. *IEEE/ACM Trans. Netw.* **20**(5), 1347–1360 (2012)
29. Pionovskiy, A., Zhang, Y.: Continuous-time Markov decision processes. *Probab. Theory Stoch. Model.* **2020**, 10 (2020)
30. Guo, X., Hernández-Lerma, O.: *Continuous-Time Markov Decision Processes*. Springer, Cham (2009)
31. Davis, M.H.: *Markov Models and Optimization*. Routledge (2018)
32. Seierstad, S.: *Optimal Control Theory with Economic Applications*. Elsevier, North-Holland (1987)
33. Billingsley, P.: *Convergence of Probability Measures*. Wiley, Hoboken (2013)
34. Bäuerle, N., Rieder, U.: Optimal control of single-server fluid networks. *Queueing Syst.* **35**, 185–200 (2000)
35. Zabczyk, J.: *Mathematical Control Theory*. Springer, Cham (2020)
36. Huang, M., Ma, Y.: Mean field stochastic games: Monotone costs and threshold policies. In: 2016 IEEE 55th Conference on Decision and Control (CDC), pp. 7105–7110. IEEE (2016)
37. Daley, D.J., Gani, J.: *Epidemic Modelling: An Introduction*. Cambridge University Press, Cambridge (2001)
38. Kumar, P., Seidman, T.I.: Dynamic instabilities and stabilization methods in distributed real-time scheduling of manufacturing systems. In: Proceedings of the 28th IEEE Conference on Decision and Control, pp. 2028–2031. IEEE (1989)
39. Rybko, A.N., Stolyar, A.L.: Ergodicity of stochastic processes describing the operation of open queueing networks. *Problems Inform. Transmission* **28**(3), 199–220 (1992)
40. Lange, D.K.: *Cost Optimal Control of Piecewise Deterministic Markov Processes Under Partial Observation*. Karlsruhe Institut für Technologie (KIT), Karlsruhe (2017). <https://doi.org/10.5445/IR/1000069448>
41. Pollard, D.: *Convergence of Stochastic Processes*. Springer, Cham (1984)
42. Kurtz, T.G.: Approximation of Population Processes. Volume 36 of CBMS-NSF Regional Conf. Series in Appl. Math. SIAM (1981)
43. Ethier, S.N., Kurtz, T.G.: *Markov Processes*. Wiley, New York (1986)
44. Whitt, W.: *Stochastic-Process Limits: An Introduction to Stochastic-Process Limits and Their Application to Queues*. Springer, Cham (2002)