Stationary Equilibria of Mean Field Games with Finite State and Action Space: Existence, Computation, Stability, and a Myopic Adjustment Process

Dissertation

zur Erlangung des Doktorgrades
an der Fakultät Mathematik, Informatik und Naturwissenschaften
Fachbereich Mathematik
der Universität Hamburg

vorgelegt von

Berenice Anne Neumann

Hamburg, 2019

Tag der Disputation: 17. Januar 2020

Folgende Gutachter empfehlen die Annahme der Dissertation:
Prof. Dr. Sören Christensen
Prof. Dr. Natalie Neumeyer
Prof. Dr. Frank Thomas Seifried
Prof. Dr. Frank Riedel

Vorsitzender der Prüfungskommission:
Prof. Dr. Jens Struckmeier

Acknowledgements

First and foremost, I sincerely wish to thank my supervisor Prof. Dr. Sören Christensen for his encouragement and support over the last three years. I would also like to express my gratitude to Prof. Dr. Natalie Neumeyer, Prof. Dr. Frank Thomas Seifried and Prof. Dr. Frank Riedel for reviewing this thesis. Furthermore, I wish to thank anonymous reviewers and seminar participants for their comments and suggestions. Moreover, I would like to thank the weekly "Dilo group" for many pleasant lunch breaks and my colleagues in the Stochastic and Statistics group for a friendly working environment. Finally, I would like to thank my parents and my friends for their support.

Contents

1. Introduction

Game theory is a rather young branch of mathematics that provides a toolbox to analyse the strategic behaviour of interacting agents in contexts where the individual outcomes are influenced by the decisions of each individual agent. It became an independent field of research with the publication of “Theory of Games and Economic Behaviour” by Von Neumann and Morgenstern (1953). Since then many developments in game theory itself, but also in applications of this theory have been made. These applications include all fields of social sciences. In economics the applications range from bargaining, auction models, oligopolies and mechanism design to network formation and general equilibrium theory.

In game theory the interaction of many agents in an intertemporal context is of particular interest. However, models of the interaction of many agents over a certain time horizon are difficult and often intractable. Therefore, several simplifying assumptions have been introduced, which are satisfactory only for small classes of examples: In general equilibrium theory one uses the reasonable assumption that “prices mediate all interactions”, but this is not possible for many other effects like externalities or public goods (Guéant et al., 2011, p.209). In the case of growth models one often relies on representative agents models, where it is assumed that all agents are identical and then the decision of a representative agent is analysed. However, for several problems, like the effect of income inequality, the differences among agents are essential and cannot be captured in a representative agent model (Gomes et al., 2015, p.5).

Mean field games, which have been introduced by Lasry and Lions (2007) and Huang et al. (2006), use a different simplification that relies on the idea of mean field approximation from particle physics: One assumes that the agents do not observe the behaviour of the other agents individually, but only the distribution of the other agents' behaviour. This is sensible in view of the applications since it is reasonable to assume that from the eyes of an individual player the behaviour of the other agents gets “lost in the crowd” when the number of players increases (Guéant et al., 2011, p.208). Moreover, it overcomes the previously mentioned tractability problems since the agents can no longer implement complex inter-individual strategies.

Although mean field game theory relies on the idea of mean field approximation, it is not a model from econophysics that describes an economic phenomenon using physical models. Instead, it works inside the “classical economic paradigm” (Guéant et al., 2011, p.208). Indeed, mean field games are differential games satisfying four hypothesis, namely rational expectations, continuum of agents, anonymous agents and interactions of the mean field type, of which the first three are common in game theory (Guéant, 2009a,b). The first hypothesis states that the "agents' predictions [...] are not systematically wrong" in equilibrium (Gomes et al., 2015, p.2). The second hypothesis, introduced by Aumann (1964), is a well-accepted approximation of "many" agents, which often yields through the use of analysis to tractable models. Moreover, it is, as Aumann (1964) argued, the mathematical appropriate way to formalize "negligibility of agents" in perfect competition. The third hypothesis states that the agents are indistinguishable. More precisely, it states that permuting the agents does not change the outcome of the game. The forth hypothesis is specific to mean field games and states that an agent cannot take into account the behaviour of each individual agent, but that he will consider summary statistics of this behaviour. We remark that under some continuity conditions a game with a large number of anonymous agents automatically has interactions of the mean field type in the limit. More precisely, Guéant (2009a, Appendix 1) shows that symmetrical functions of N variables converge to functions of measures under suitable continuity assumptions.

In general, we can describe a mean field game as follows: Every individual agent controls his state variable that is given by a controlled stochastic dynamic system where the dynamics might depend on the current population distribution. The preferences of the individual agent depend on his own state and action as well as the current population distribution. Thus, an agent plays a game not against all individuals but against the population as a whole. Therefore, also the equilibrium notion is different to classical differential games. Namely, we say that a strategy is a mean field equilibrium, if when all players adopt this strategy an individual agent has no incentive to deviate from playing this strategy. In other words, the equilibrium strategy is indeed a best response to the entire population playing the equilibrium strategy.

In the first mean field game models that have been analysed, the players controlled a diffusion and optimized their expected reward over a finite time horizon. Then, in order to find a mean field equilibrium, we have to solve an individual control problem, namely, given any flow of population distributions what is the optimal strategy for an individual agent, as well as a fixed point problem, namely, find a flow of population distributions such that the dynamic of an agent using the optimal strategy equals the flow of population distributions. Under some conditions one can show that an equilibrium is given by a solution of a system of a Hamilton-Jacobi-Bellman equation, accounting for the individual control problem, coupled with a Fokker-Planck-equation, accounting for the fixed point problem. We remark that the system is half forward and half backward, which is non-standard. In the literature in particular existence and uniqueness of solutions, numerical methods for the forward-backward systems as well as approximation of N-player-games through mean field games (which in turn justifies the continuum of agents hypothesis) are discussed.

In general there are two approaches to mean field games, namely an analytic one dating back to Lasry and Lions (2007) and a probabilistic one dating back to Huang et al. (2006) (Basna et al. (2014)). The first approach indeed yields the system of a Hamilton-Jacobi-Bellman equation coupled with a Fokker-Planck-equation and the second approach yields to forward-backward stochastic differential equations. As we will also see in this thesis, for some problems one approach is more suitable than the other. For a detailed introduction to diffusion-based mean field games consider Carmona and Delarue (2018a,b) and Bensoussan et al. (2013), who mainly utilize the probabilistic approach, and Cardaliaguet (2013), who mainly utilizes the analytic approach. Furthermore, several economic applications have been considered, for an overview consider Guéant et al. (2011), Caines et al. (2017) as well as the references therein. These applications in particular include growth models (Gomes et al., 2015; Guéant, 2009a) and Bertrand and Cournot competition (Chan and Sircar, 2015).

There is also a growing body of literature regarding models with different underlying stochastic dynamic systems as well as application of those: These include mean field games of timing with applications to bank run and R&D completion (Carmona et al., 2017; Nutz, 2018; Nutz and Zhang, 2019), static mean field games with applications to ranking effects for asset managers (Guéant, 2009a) as well as mean field games with a finite state space, where a large number of applications has been considered. These include the spread of corruption (Kolokoltsov and Malafeyev, 2017), botnet defence models (Kolokoltsov and Bensoussan, 2016), inspection games (Kolokoltsov and Yang, 2015), labour market models (Guéant, 2009a), paradigm shift models (Besancenot and Dogguy, 2015; Gomes et al., 2014b) and consumer choice models (Gomes et al., 2014b).

Most of the theoretical results for finite state models presented so far (see next section for an overview) rely on the assumption that there is a unique optimizer of the Hamiltonian. This assumption is satisfied for the applications that exhibit fully controllable transition rates (for example Kolokoltsov and Yang (2015) and Gomes et al. (2014b)). However, this assumption will in general not be satisfied for models with finite action spaces like Kolokoltsov and Bensoussan (2016) and Kolokoltsov and Malafeyev (2017). Indeed, the assumption will only hold true if one action is the optimizer of the Hamiltonian irrespective of the other parameters (in particular irrespective of the population distribution), which means that we consider a setting without a real strategic choice. However, the effect of strategic interaction is the main interest when analysing such models. Nonetheless, finite action spaces are reasonable in many settings since agents will often face finitely many choice options (stick to one's behaviour or change it, choose among finitely many providers). Thus, theoretical models capturing finite action spaces are desirable.

In this thesis, we therefore analyse a theoretical model relevant for economic applications with finite state and action space. More precisely, we investigate the existence, the computation, the stability and the explanatory power of stationary equilibria in mean field games with finite state and action space. We focus on stationary equilibria in the sense that the individual player maximizes the expected discounted reward. We remark that this formulation is mostly analogous to stationary (ergodic) equilibria in the sense that the individual player maximizes the expected average reward, which are often considered in particular in the literature regarding diffusion-based mean field game models (Kolokoltsov and Bensoussan, 2016; Kolokoltsov and Malafeyev, 2017). The reasons for our focus on stationary equilibria are diverse: First, because these are the main interest in economic applications. Second, in light of the limited tractability also of mean field games, this is the most promising approach to obtain results with respect to equilibrium computation. Third, since one can often establish convergence of dynamic equilibria to stationary equilibria (see Gomes et al. (2013) for a theoretical result, Guéant (2009a) as well as Besancenot and Dogguy (2015) for such a result in an economic example) or find a learning procedure converging towards a stationary equilibrium (Mouzouni (2018) in the setting of a diffusion-based model).

The remaining part of the introduction is organized as follows: Section 1.1 describes the contributions of the thesis. Section 1.2 describes the related literature on mean field games with finite state space and Section 1.3 discusses related economic equilibrium concepts. Finally, Section 1.4 contains an outline of the thesis.

1.1. Contributions of the Thesis

This thesis contributes to the theory of mean field games with finite state and action space in two fields. First, we investigate the existence, the computation and stability of stationary mean field equilibria. Second, we investigate in how far stationary mean field equilibria describe dynamic behaviour of partially rational agents in the corresponding games. As a byproduct, some contributions to the theory of nonlinear Markov chains are made. This section describes the contributions of the thesis in detail.

We work inside the model of Doncel et al. (2016a), although we present a new probabilistic formulation of that model. At several points in the thesis, in particular in the first part, this allows us to obtain results that could not be easily obtained in the original analytic formulation. The agents control a continuous time Markov chain, which transition rates depend on the current population distribution, and the agents maximize the expected discounted reward on an infinite time horizon. The instantaneous reward depends on the current state, the current action and the current population distribution. We are interested in dynamic as well as stationary equilibria, which in our setting are described by the (flow of) population distribution(s) as well as a strategy such that the strategy is optimal given the (flow of) population distribution(s) and, moreover, the (flow of) population distribution(s) is indeed the distribution of the individual agent using this strategy.

First, we will look at the existence of dynamic as well as stationary mean field equilibria. In the dynamic case, Doncel et al. (2016a) formulated an existence result for mixed strategy equilibria under a continuity assumption. We fill all the missing (measure-theoretical and topological) details in their sketch of the proof, where the main idea is to consider a best-response map in population distributions and not in strategies because this yields an upper semi-continuous best-response correspondence. Thereafter, we modify this approach to prove existence of stationary equilibria in mixed strategies under the same conditions. This is a rather strong result compared to the setting of Gomes et al. (2013), where existence of stationary (ergodic) equilibria could only be established under additional assumptions. It turns out that our proof relies heavily on the probabilistic formulation of the model since we use the characterization of optimal strategies that we will derive in context of equilibrium computation as well as other results regarding stationary points of continuous time Markov chains.

As a second step, we will look at the computation of stationary equilibria. As in diffusion-based models, computing mean field equilibria means to solve for each possible population distribution an individual control problem and look for each optimal strategy whether it induces a fixed point for which the strategy is again optimal. With respect to the individual control problem, we prove that the decision problem is equivalent to a continuous time Markov decision process with infinite time horizon and expected discounted reward criterion. This allows us to show that a stationary strategy is optimal if and only if it is a convex combination of optimal deterministic stationary strategies, which are in turn those that maximize the right-hand-side of the optimality equation of the equivalent Markov decision process. In order to tackle the fixed point problem for general dynamics we provide a cut criterion to reduce the set of candidates for being a fixed point given some optimal strategy. This cut criterion is basically a reformulation of the balance equation. However, it yields in many examples to a strategy-free condition every fixed point has to satisfy since often agents cannot influence the transitions between one set of the states to the rest of states. Moreover, we obtain in the case of irreducible dynamics an explicit representation of the stationary points given the continuous time Markov chain for a fixed strategy and a fixed population distribution. Using this we prove that all mean field equilibria have a distribution that is a fixed point of the set-valued map that maps every population distribution to the convex hull of the stationary points given those deterministic strategies that are optimal for that point. This yields that it is sufficient to find all fixed points of this map and thereafter solve a linear equation to determine the corresponding equilibrium strategy. In the case of constant dynamics we even obtain, after solving the individual control problem for each population distribution, a closed form expression of all distributions of stationary mean field equilibria. We illustrate the described techniques for three toy examples, for which we can now compute all stationary equilibria.

Up to the knowledge of the author, the literature covers only the computation of stationary equilibria with deterministic equilibrium strategies in particular examples (for example Kolokoltsov and Bensoussan (2016), Kolokoltsov and Malafeyev (2017)). This thesis presents the first general results regarding the computation of stationary equilibria of mean field games with finite state and action space as well as the first examples for which all equilibria have been computed.

As a next step, we enter into the question under which conditions stationary mean field equilibria are stable against slight perturbations of the model. This is a classical question in game theory since in many settings the exact model parameters are estimated and, thus, they are erroneous (see Chapter 7 for more details). For this we will generalize the notions of essential and strongly stable equilibria for mean field games with finite state and action space. For the analysis we again have to consider the two subproblems, the individual control problem and the fixed point problem, at first individually and thereafter combine the results. The question whether a strategy stays optimal/non-optimal under slight model perturbations can be easily answered with "yes" in most cases. However, the question which fixed points of the dynamics are stable against slight model perturbations, which we will call essential stationary points, does not have a simple answer. This is not surprising due to the close connection to essential fixed points, for which also no full characterization results for spaces other than $[0, 1]$ are known (Del Prete et al., 1999). Nonetheless, we obtain, analogously to the results for essential fixed points, that a unique stationary point of the dynamics is automatically essential and that the set of all those transition rate matrix functions for which only essential stationary points exist is (topologically) generic. Moreover, we illustrate that strong stability cannot be expected in general and analyse the effects of changes in the class of all dynamics with constant transition rate matrix functions.

Putting the results for the two subproblems together we obtain the following positive and negative results regarding the question whether a stationary mean field equilibrium is essential: First, we obtain that a stationary mean field equilibrium is essential whenever the equilibrium distribution is an essential stationary point of the dynamics and the equilibrium strategy is a deterministic strategy and, moreover, the unique optimal stationary strategy at the equilibrium distribution. Second, we obtain that stationary mean field equilibria with a randomized equilibrium strategy that randomizes over more than two deterministic stationary strategies are in general not essential. Third, given constant irreducible dynamics, a stationary mean field equilibrium that randomizes over exactly two deterministic stationary strategies is essential whenever for each of the two deterministic stationary strategies over which the strategy randomizes there are population distributions arbitrarily close to the equilibrium distribution such that this strategy is the unique optimal stationary strategy. In the end we show that also the set of all games with only essential stationary mean field equilibria is (topologically) generic, which means that "many" games have only essential equilibria. Up to the knowledge of the author, these stability considerations have not been undertaken for mean field games so far.

In the second part of the thesis, we focus on the question to what extend stationary equilibria explain dynamic behaviour in the mean field game model since we cannot assume that agents are able to compute dynamic equilibria: Indeed, in our setting with infinite time horizon we cannot expect an agent to compute a dynamic mean field equilibrium since the individual control problem is in general not tractable. More precisely, there are no results regarding the computation of optimal policies for Markov decision processes with infinite time horizon and non-stationary transition rates and rewards. Also in the case of finite time horizons we cannot expect that agents can indeed compute dynamic mean field equilibria. Namely, the results by Belak et al. (2019) yield that dynamic equilibria are given by a system of ordinary differential equations partially running forward, partially running backwards in time. However, these systems are non-standard and often only numerical computation of solutions is possible.

The observation that agents cannot compute equilibria is also true for classical Nash equilibria. However, it is observed in experiments as well as in practice that (at least after playing a game for several rounds) agents "learn" to play equilibrium strategies. This led to the "learning theory" in games, where decision mechanisms for partially rational agents are introduced and then their convergence or non-convergence towards Nash equilibria is studied. We introduce such a learning procedure for our mean field game model, namely a myopic adjustment process, and justify that it is indeed a sensible way of describing partially rational decision making in our model. The process we introduce will be the solution of a differential inclusion and we prove that for any initial condition trajectories of this process exist. Thereafter, we study the convergence behaviour locally as well as globally. More precisely, we show that we have local convergence towards stationary mean field equilibria with a deterministic equilibrium strategy that is the unique optimal strategy at the equilibrium distribution whenever we face constant irreducible dynamics or whenever the Jacobian matrix has a certain eigenvalue structure. In the first case we derive explicitly the radius of convergence, in the second case we only obtain an implicit characterization. Moreover, given additional assumptions we can even show (almost) global convergence towards stationary equilibria. To be more precise, we prove that the trajectories either converge towards a deterministic mean field equilibrium or remain inside a set of all those points for which at least two strategies are simultaneously optimal. The assumptions include that the underlying nonlinear Markov chains are converging towards a stationary distribution in the limit, that the sets of points for which a particular deterministic strategy is optimal are simple and that the trajectories given the different deterministic strategies are in some sense consistent. We illustrate that these convergence results can be applied in examples. We emphasize that the convergence results increase the predictive power of stationary mean field equilibria since they show that stationary equilibria will not only emerge if the population distribution equals the equilibrium distribution, but that under certain conditions they will naturally arise in the interaction of partially rational agents.

In the context of mean field games two learning procedures have been studied previously. On the one hand, fictitious play for repeated games with finite time horizon has been analysed for diffusion-based models (Cardaliaguet and Hadikhanloo, 2017; Hadikhanloo, 2017; Briani and Cardaliaguet, 2018) as well as discrete time finite state space models

(Hadikhanloo, 2018). On the other hand, Mouzouni (2018) introduced a learning procedure similar to our myopic adjustment process for diffusion-based mean field games with finite time horizon. He studies existence in a general setting as well as local convergence given a quadratic Hamiltonian and the Lasry-Lions monotonicity condition. However, the convergence result is only local and requires strong assumptions, moreover, the models and, thus, the used methods differ substantially.

Additionally, nonlinear continuous time Markov chains are considered. These are processes in which the transition probabilities do not only depend on the current state, but also on the current distribution of the process. We prove that given a Lipschitz continuous transition rate matrix function we can always define a nonlinear Markov process whose nonlinear generator is the transition rate matrix function. Thereafter, we investigate the limit behaviour of these processes. More precisely, we provide several examples that the limit behaviour of nonlinear Markov chains is more complex than for classical time-homogeneous Markov chains. In particular, we illustrate that we can face cycles and that it might happen that a nonlinear Markov chain with irreducible nonlinear generator is not strongly ergodic, but convergence towards several stationary points happens. Furthermore, we provide a sufficient criterion for a nonlinear Markov chain to have a unique stationary point as well as a sufficient criterion for a nonlinear Markov chain with two or three states to be strongly ergodic. Up to the knowledge of the author, nonlinear Markov chains with finite action space and continuous time has been introduced in Kolokoltsov (2010) and no further investigations have been made. The closest contributions in the literature are ergodicity results for discrete time nonlinear Markov chains, which can be found in Butkovsky (2012), Butkovsky (2014) and Saburov (2016).

1.2. Literature Review: Finite State Mean Field Game Models

This section reviews the previous research on mean field games with finite state space. We restrict our attention to models in continuous time although the study of finite state models started with the discrete time model of Gomes et al. (2010). We will start by reviewing the models as well as results obtained for them. Thereafter, we will discuss their relation to the model discussed in this thesis.

There are several finite state mean field game models and they mainly differ with respect to the underlying controlled stochastic dynamic system. The first models have been introduced with fully controllable transition rates (Gomes et al., 2013) and fully controllable transition rates with the restriction that the players cannot reach from a given state all other states (Guéant, 2011, 2015). Thereafter, more general stochastic dynamic systems have been considered: Carmona and Delarue (2018a, Section 7.2) consider a finite state mean field game model where the dynamics are given by a continuous time Markov chain with transition rates depending on the current action and the current population distribution. Cecchin and Fischer (2018) consider a finite state mean field game model where the individual dynamics are given by a stochastic differential equation driven by a stationary Poisson random measure where the coefficients again might depend on the current action and the current population distribution. Carmona and Wang (2018) consider an extended mean field game model with finite state space where the dynamics are given by a continuous time Markov chain with a generator that might depend on the current action and the current distribution of states as well as actions. Doncel et al. (2016a, 2019) consider a mean field game model with finite action spaces where the individual dynamics are given by a differential equation specifying the transition rates, which might depend on the current action and the current population distribution. Belak et al. (2019) again consider a mean field game model where the dynamics are given by a continuous time Markov chain with a generator that depends on the current action and the current distribution of states. Moreover, they include common noise, which are random events that influence the distribution of all players simultaneously.

With respect to these models several questions have been discussed: In all models conditions that ensure the existence of a mean field equilibrium are proposed. These conditions usually include that there is a unique maximizer of the Hamiltonian. Only Cecchin and Fischer (2018), Doncel et al. (2016a, 2019) and Belak et al. (2019) prove existence (possibly in mixed strategies) under weaker assumptions. Moreover, except for some approximation results (see Cecchin and Fischer (2018) and Doncel et al. (2016a)) and the derivation of a forward-backward system (see Belak et al. (2019)) the assumption that there is a unique maximizer of the Hamiltonian is crucial for all other results.

The question of uniqueness of equilibria has also been considered for several models and as for diffusion-based models two uniqueness regimes have been identified, namely, a small time horizon (for example Gomes et al. (2013) and Cecchin and Fischer (2018)) and the Lasry-Lions monotonicity condition (for example Gomes et al. (2013), Carmona and Delarue (2018a, Section 7.2) and Carmona and Wang (2018)). Additionally, in several settings (Gomes et al. (2013), Guéant (2011, 2015), Carmona and Delarue (2018a, Section 7.2) as well as Belak et al. (2019)) a forward-backward system of ordinary differential equations is derived, which describes the dynamic mean field equilibria of the game. We emphasize that the derivation in Belak et al. (2019) does not require that there is a unique maximizer of the Hamiltonian.

For many models the relation to N-player games has been considered, where two questions are of central interest (see Cecchin and Fischer (2018)): The first question, which is more commonly analysed, is the question whether mean field equilibria approximate Nash equilibria in N-player games such that the approximation error is arbitrarily small when N increases. The second question is whether a sequence of Nash equilibria of the corresponding N-player games converges towards a mean field equilibrium as N tends to infinity. With respect to approximation Carmona and Wang (2018), Cecchin and Fischer (2018) and Doncel et al. (2016a, 2019) constructed ϵ_N-Nash equilibria for the corresponding N-player games from the optimal strategy of the mean field equilibrium, where the last two results do not rely on the assumption that a unique optimizer of the Hamiltonian exists. The first result in the convergence direction has been proposed in Gomes et al. (2013), but only for a small time horizon. Recently, Bayraktar and Cohen (2018) as well as Cecchin and Pelino (2019) proved for arbitrary time horizons convergence under the assumption that the associated master equation admits a (unique) regular solution and that there is a unique equilibrium for the N-player system. However, the assumptions necessary to ensure the existence of a (unique) regular solution of the master equation in particular imply uniqueness of the mean field equilibrium, which is rather restrictive. Cecchin et al. (2019) analyse a first example in which there are three mean field equilibria, but still a unique Nash equilibrium for the N-player game, and convergence is observed towards one of the three mean field equilibria. Also Doncel et al. (2016a,b, 2019) considers the convergence question, but in their setting the prelimit N-player games are formulated in discrete time. They prove that convergence holds if the equilibria for the N-player games are local strategies (which only depend on the current state and time) and provide an example illustrating that convergence does not hold if one considers Markov strategies (which also depend on the current distribution of all players). The intuitive reason for this is that the "tit-for-tat"-principle cannot be applied for a continuum of (negligible) players (see Doncel et al. (2016b) for more details).

Another branch of research is whether a trend towards stationary equilibria exists. More precisely, one is interested under which conditions a dynamic equilibrium converges towards some stationary equilibrium. Gomes et al. (2013) proved convergence towards stationary (ergodic) equilibria under the Lasry-Lions monotonicity condition as well as a contractivity assumption regarding the Hamiltonian. Later, Ferreira and Gomes (2014) proved convergence towards stationary (ergodic) equilibria for potential games, a class of mean field games with additional structural properties that often allows for deeper results. More precisely, in these games the costs split into two additive term, one depending on the current state and current population distribution being, moreover, the gradient of a convex function, the other one depending only on the current state and action).

Finally, we remark that several attempts have been made to make the forward-backward systems that describe dynamic mean field equilibria at least numerically tractable. In Gomes et al. (2014b) a method that relies on the prelimit games with finitely many players is proposed for two state mean field games. In Gomes et al. (2014a) several equivalent formulations of mean field equilibria for two state mean field games are introduced and several numerical methods for these equivalent formulations are compared. Gomes and Saúde (2018) proposed for games satisfying the Lasry-Lions monotonicity condition a numerical method relying on a monotone operator that is furthermore a contraction.

We conclude by discussing the relation of the models with the model used in this thesis: Most models and results rely, as we have discussed, on the assumption that a unique optimizer of the Hamiltonian exists. However, as stated before, this assumption is not satisfied for finite action spaces. Thus, only the models introduced without this assumption are related to our model for mean field games with finite state and action space. More precisely, the models by Cecchin and Fischer (2018) and Belak et al. (2019) cover a version of the model discussed in this thesis where the time horizon is now finite and yield existence, approximation results as well as a characterization via a system of ordinary differential equations. The model of Doncel et al. (2016a, 2019) is the model used in this thesis, however, we will work with a probabilistic reformulation to utilize the corresponding theory.

1.3. Literature Review: Related Equilibrium Concepts

This section reviews related theoretical economic models with a continuum of anonymous agents as well as the related notion of oblivious equilibria in discrete time dynamic stochastic games with a large (but finite) number of players.

The research regarding games with a continuum of anonymous agents started with Aumann's contribution on perfect competition (Aumann, 1964). Thereafter, general static games with a continuum of players and dynamic games in discrete time have been discussed: Schmeidler (1973) introduced general static non-cooperative games in normal form with a continuum of players and provided, using measure-theoretic arguments, sufficient criteria for the existence of equilibria in mixed as well as in pure strategies. Mas-Colell (1984) reformulated the games in Schmeidler (1973) in terms of distributions and obtained sufficient conditions for the existence of general as well as symmetric equilibria in pure strategies.

The first dynamic (discrete time) games with a continuum of anonymous players have been introduced in Jovanovic and Rosenthal (1988) under the name sequential anonymous games. They prove existence of dynamic as well as stationary equilibria under the same conditions. In particular, in the dynamic case it is shown that the equilibrium is characterized by a system of equations that exhibit a forward-backward structure. Bergin and Bernhardt (1992) extend this model by adding aggregated uncertainty, which is in the mathematical literature for mean field games often called common noise, and prove existence of dynamic equilibria in this more complex setting. Bergin and Bernhardt (1995) extend the model to capture aggregated uncertainty described by continuous state spaces. Furthermore, they prove existence of equilibria in Markov strategies for stationary models where the evolution of the aggregated uncertainty is described by a

Markov process. Moreover, the papers describe several possible economic applications of sequential games. However, none of the papers provides any computational tools.

For discrete time dynamic games with a large, but finite number of players an equilibrium concept closely related to our notion of stationary mean field equilibria has been introduced. Namely, Weintraub et al. (2008) introduced the notion of oblivious equilibria in the context of a dynamic model of imperfect competition with a large, but finite number of firms. A central motivation for the introduction of oblivious equilibria was to overcome the following two problems of the standard solution concept Markov perfect equilibria for dynamic stochastic games (see Fudenberg and Tirole (1991, Chapter 13) for an overview): First, in large player settings Markov perfect equilibria become intractable. Second, it is not reasonable to assume in large player settings that an individual player keeps track of all other agents' states individually. An oblivious equilibrium is an equilibrium in oblivious strategies, which take only the current individual state and the long-run average population distribution into account. The intuition behind this equilibrium concept is that in large player settings simultaneous state changes average out and, thus, it is reasonable to assume that the aggregated distribution remains roughly constant over time. Indeed, it is proven that oblivious equilibria are a good approximation to Markov perfect equilibria in their context whenever the equilibrium distribution satisfies a light-tail condition.

Thereafter, Adlakha et al. (2015) discuss oblivious equilibria for more general games. They provide under suitable assumptions, which in particular include a finite action space or the existence of a unique optimizer of the Hamiltonian, general existence results and prove that stationary equilibria approximate Markov perfect equilibria. Light and Weintraub (2018) again consider games under more general assumptions, in particular capturing unbounded state spaces. They obtain an existence and uniqueness result as well as a comparative statics result for a parameter such that the transition probability kernel increases in this parameter. However, also here all results are established relying on the assumption that there is a unique maximizer of the Hamiltonian. Furthermore, Adlakha and Johari (2013) consider games with strategic complementarities, which means that the payoffs and transition probabilities satisfy certain monotonicity conditions. Relying on the monotonicity properties, they prove existence of equilibria and obtain a comparative static result for a parameter that influences the payoffs and transition rates in a monotone fashion. Moreover, they consider two learning procedures that converge towards the smallest/largest equilibrium. More precisely, they consider a best response dynamic, where in each step the individual agent computes the best response to the stationary distribution induced by the strategies used by the other agents in the preceding period, and a myopic adjustment process, where the agents compute the best response to the current population distribution.

The definition of oblivious equilibria in their setting is similar to the definition of stationary mean field games in our setting. However, the underlying models differ substantially, namely they consider a finite number of players, whereas we consider a continuum of players, and they consider discrete time, whereas we consider continuous time. In particular, the discrete decision points can be criticized since for a large group of players it is unreasonable to assume that they all take their decisions at the same time and do not exploit possible benefits for a deviation of this scheme. In addition, we remark that they cover for the setting with a finite state and action space only an existence result, but no structural results regarding computational tools, comparative statics or learning procedures.

1.4. Outline of the Thesis

In Chapter 2 we introduce time-inhomogeneous continuous time Markov chains as well as Markov decision processes, which are the main building blocks of our model. In Chapter 3 we introduce and discuss the mean field game model with finite state and action space that we consider in this thesis. In Chapter 4 we present rigorous existence proofs for dynamic as well as stationary mean field equilibria.

Chapters 5-7 cover the analysis of stationary mean field equilibria: In Chapter 5 we provide several results that help to find stationary mean field equilibria with arbitrary dynamics as well as sharper results for games satisfying an irreducibility assumption. In Chapter 6 three toy models are introduced and the application of the tools introduced in the preceding chapter is illustrated. In Chapter 7 we analyse under which conditions stationary equilibria are stable against slight model perturbations and we prove that stability is a (topological) generic property.

Chapters 8-9 are concerned with the question in how far stationary mean field equilibria explain dynamic behaviour of partially rational agents. In Chapter 8 we motivate and introduce a learning procedure, namely a myopic adjustment process, as a reasonable form of decision making of partially rational agents in the mean field games we consider. Moreover, we prove existence of trajectories as well as local convergence towards deterministic equilibria under certain conditions. In Chapter 9 we prove a global convergence result for the myopic adjustment process under several assumptions and illustrate the application of this theorem in the examples of Chapter 6. Additionally, we introduce in Chapter 9 nonlinear Markov chains and investigate their limit behaviour.

Finally, we list some further directions of research in Chapter 10. Appendix A contains several technical lemmata as well as an introduction to Caratheodory solutions, which arise at several points of the thesis.

2. Preliminaries

In this chapter we lay the foundations to introduce and analyse our model. In Section 2.1 we will introduce the notion of time-inhomogeneous continuous time Markov chains. Thereafter, we will introduce in Section 2.2 the notion of Markov decision processes with expected discounted reward criterion and infinite time horizon in continuous as well as discrete time. In the end of the chapter we will then prove the first central result of the thesis, namely, that the set of all optimal stationary strategies for a Markov decision process is the convex hull of all those deterministic stationary strategies that choose actions that maximize the right-hand side of the optimality equation.

2.1. Time-Inhomogeneous Continuous Time Markov Chains

In this section we will summarize without proofs the relevant material on time-inhomogeneous continuous time Markov chains (hereafter time-inhomogeneous CTMC) following the presentation in Guo and Hernández-Lerma (2009, Appendices B+C). The purpose of this is twofold: First, we describe the dynamics of the individual agents by time-inhomogeneous Markov chains. Second, the notion of time-inhomogeneous CTMCs is necessary for the discussion of Markov decision processes in Section 2.2.

A stochastic process $(X_t)_{t\geq 0}$ with values in a countable set $\mathcal{S}$, which we will call state space, is called a *time-inhomogeneous continuous time Markov chain* if for any finite sequence $0 \leq t_1 < t_2 < \ldots < t_n < t_{n+1}$ and choice of states $i_1, \ldots, i_{n-1}, i, j \in \mathcal{S}$ we have

$$\mathbb{P}(X_{t_{n+1}} = j | X_{t_1} = i_1, X_{t_2} = i_2, \ldots, X_{t_{n-1}} = i_{n-1}, X_{t_n} = i) = \mathbb{P}(X_{t_{n+1}} = j | X_{t_n} = i)$$

whenever $\mathbb{P}(X_{t_1} = i_1, X_{t_2} = i_2, \ldots, X_{t_{n-1}} = i_{n-1}, X_{t_n} = i) > 0$. Intuitively, this condition, called Markov property, means that the future path of the process depends on the past only through the current state.

The probabilities $p(s, i, t, j) := \mathbb{P}(X_t = j | X_s = i)$ ($i, j \in \mathcal{S}$, $0 \leq s \leq t$) are called *transition (probability) function* and satisfy for all $i, j \in \mathcal{S}$, $t \geq s \geq 0$

(a) $p(s,i,t,j) \geq 0$ and $\sum_{j\in\mathcal{S}} p(s,i,t,j) \leq 1$

(b) $p(s,i,s,j) = \delta_{ij}$

(c) and the Chapman-Kolmogorov equation

$$p(s,i,t,j) = \sum_{k\in\mathcal{S}} p(s,i,v,k)p(v,k,t,j)$$

for all $s \leq v \leq t$.

The simple time-homogeneous case arises if for all $i,j \in \mathcal{S}, t \geq s \geq 0$ we have

$$\mathbb{P}(X_t = j|X_s = i) = \mathbb{P}(X_{t+u} = j|X_{s+u} = i)$$

for all $u \geq 0$. In this case we call $p_{ij}(t) := p(0,i,t,j)$ the stationary transition function and we can apply the standard theory, which is described in Asmussen (2003). For the rest of the section we will focus on the time-inhomogeneous case.

A function $p(s,i,t,j)$ defined for $t \geq s \geq 0$ and $i,j \in \mathcal{S}$ is called transition function if it satisfies the properties $(a) - (c)$. It is called *standard*, if it further satisfies that the right-sided limit $\lim_{t\downarrow s} p(s,i,t,j) = \delta_{ij}$ uniformly in $j \in \mathcal{S}$ for each fixed $i \in \mathcal{S}$, $s \geq 0$. A standard transition function is called *regular* if $\sum_{j\in\mathcal{S}} p(s,i,t,j) = 1$ for all $i \in \mathcal{S}$ and $t \geq s \geq 0$.

One can show that for a standard transition function $p(s,i,tj)$ the right-sided limit

$$\lim_{t\downarrow s} \frac{p(s,i,t,j) - \delta_{ij}}{t-s}$$

exists in $[-\infty, 0]$ if $i = j$ and in $[0,\infty)$ if $i \neq j$. We refer to these limits as *transition rates* and denote them by $Q_{ij}(s)$. We say that a transition function is *stable* if $-Q_{ii}(s)$ is finite for every $i \in \mathcal{S}$, $s \geq 0$.

Given a stable standard transition function $p(s,i,t,j)$ we find a corresponding Markov chain $(X_t)_{t\geq 0}$ that has $p(s,i,t,j)$ as transition functions and is a right process, which means that it is right-continuous and has finite left-hand limits at every t.

In the time-homogeneous setting a CTMC is most of the time described by specifying the infinitesimal *generator*, often also called *transition rate matrix*, which collects the transition rates Q_{ij} from state i to state j and satisfies for all $i \in \mathcal{S}$ that

$$Q_{ii} \leq 0, \quad Q_{ij} \geq 0 \quad \forall j \neq i \quad \text{and} \quad \sum_{j\in\mathcal{S}} Q_{ij} = 0.$$

This matrix is the corner stone of analysis because it enables to compute the transition probability function and to derive several qualitative properties of the CTMC itself. The time-inhomogeneous analogue is the $Q(t)$-matrix:

For all $i, j \in \mathcal{S}$ let $Q_{ij}(\cdot)$ be a real-valued function defined on $[0, \infty)$. The function $Q(\cdot) := ((Q_{ij}(\cdot))_{i,j\in\mathcal{S}})_{t\geq 0}$ is said to be a *inhomogeneous* $Q(t)$*-matrix* if for every $i \in \mathcal{S}$ and $t \geq 0$ it holds that

(i) $Q_{ij}(t) \geq 0$ for all $j \neq i$ and $-\infty < Q_{ii}(t) \leq 0$

(ii) and $\sum_{j\in\mathcal{S}} Q_{ij}(t) \leq 0$.

If in addition $\sum_{j\in\mathcal{S}} Q_{ij}(t) = 0$ for all $i \in \mathcal{S}$, $t \geq 0$ we say that $Q(\cdot)$ is *conservative*.

Let $((Q_{ij}(\cdot))_{i,j\in\mathcal{S}})_{t\geq 0}$ be $Q(t)$-matrix that satisfies that $Q_{ij}(t)$ is Borel measurable in $t \geq 0$ for each $i, j \in \mathcal{S}$ and that $Q_{ii}(t)$ is Lebesgue integrable over any finite interval in $[0, \infty)$. A $Q(t)$*-function* is a transition function $p(s, i, t, j)$ such that for all $i, j \in \mathcal{S}$, $t \geq 0$ and for almost every $s \in [0, t]$ the partial derivatives $\frac{\partial p(s,i,t,j)}{\partial s}$ exist and for which, moreover, the right-sided limits

$$\lim_{t\downarrow s} \frac{p(s, i, t, j) - \delta_{ij}}{t - s}$$

exist and equal $Q_{ij}(s)$.

Given any $Q(t)$-matrix that satisfies the previously stated measurability and integrability conditions we can construct a transition function $p(s, i, t, j)$ that is a standard and stable $Q(t)$-function. It furthermore satisfies the Kolmogorov forward and backward equations

$$\begin{aligned} \frac{\partial p(s, i, t, j)}{\partial t} &= \sum_{k\in\mathcal{S}} p(s, i, t, k) Q_{kj}(t) \\ \frac{\partial p(s, i, t, j)}{\partial s} &= -\sum_{k\in\mathcal{S}} Q_{ik}(s) p(s, k, t, j) \end{aligned} \tag{2.1}$$

for all $i, j \in \mathcal{S}$ and almost all $t \geq s \geq 0$.

If additionally the $Q(t)$-matrix is conservative and we find constants $L_1, L_2 > 0$ as well as a non-negative real-valued function $w(\cdot)$ on $\mathcal{S}$ such that

$$\sum_{j\in\mathcal{S}} w(j) Q_{ij}(t) \leq L_1 w(i) \quad \text{and} \quad -Q_{ii}(t) \leq L_2 w(i) \quad \text{for all } i \in \mathcal{S} \text{ and } t \geq 0, \tag{2.2}$$

then the $Q(t)$-function is unique and regular.

2.2. Infinite Horizon Markov Decision Processes

In the context of our mean field game model, Markov decision processes are of particular importance. The reason is that given a stationary population distribution the optimization problem faced by the individual player is equivalent to a continuous time infinite horizon Markov decision process with the expected discounted reward criterion (see Section 5.1). In order to establish this result and to utilize the existing theory some standard facts on continuous time Markov decision processes, which can be found in Guo and Hernández-Lerma (2009), are reviewed in Subsection 2.2.1. For further analysis, we moreover present the concept of uniformization presented in Kakumanu (1977) and review in Subsection 2.2.2 selected results on discrete time Markov decision processes, for which Puterman (1994) is a standard reference. In Subsection 2.2.3 we then derive the first central result of the thesis. Namely, that the set of all optimal stationary strategies is given as the set of all optimal stationary deterministic strategies, which in turn are exactly those that choose actions that maximize the right hand side of the optimality equation.

2.2.1. Continuous Time Markov Decision Processes

The continuous time control model presented in Guo and Hernández-Lerma (2009) is a five-tuple

$$\{\mathcal{S}, \mathcal{A}, (A(i))_{i\in\mathcal{S}}, Q_{ija}, r_{ia}\}$$

consisting of a denumerable state space $\mathcal{S}$, an action space $\mathcal{A}$ being a Borel subset of a complete, separable metric space with Borel-σ-algebra $\mathcal{B}(\mathcal{A})$, a family of sets $A(i) \subseteq \mathcal{A}$, which contain all those actions that can be chosen when the system is in state i, as well as transition rates Q_{ija} and reward functions r_{ia}. Let

$$K := \{(i, a)|i \in \mathcal{S}, a \in A(i)\}$$

be the set of all feasible state-action pairs. The transition rates Q_{ija} are measurable in $a \in A(i)$ for each fixed $i, j \in \mathcal{S}$ and satisfy $Q_{ija} \geq 0$ for all $(i, a) \in K$ and $j \neq i$. We furthermore assume that they are conservative, which means that $\sum_{j\in\mathcal{S}} Q_{ija} = 0$ for all $(i, a) \in K$, and that they are stable, which means that $\sup_{a\in A(i)} -Q_{iia} < \infty$ for all $i \in \mathcal{S}$. The *reward function* r_{ia} is a real-valued function that is measurable in $a \in A(i)$ for each fixed $i \in \mathcal{S}$ and which can, despite its name, also admit negative values. We restrict our attention to models with a uniformly bounded reward function, that is we assume that we find an $M \in \mathbb{R}$ such that $|r_{ia}| \leq M$ for all $i \in \mathcal{S}$ and $a \in \mathcal{A}$.

The controller can choose his actions according to a *randomized Markov strategy*, which is a real valued function $(i, C, t) \mapsto \pi_{iC}(t)$ that satisfies the following conditions:

(i) For all $i \in \mathcal{S}$ and $C \in \mathcal{B}(A(i))$ the mapping $t \mapsto \pi_{iC}(t)$ is measurable on $[0, \infty)$.

(ii) For all $i \in \mathcal{S}$, $t \geq 0$ the function $C \mapsto \pi_{iC}(t)$ is a probability measure on $\mathcal{B}(\mathcal{A})$, where $\pi_{iC}(t)$ denotes the probability that an action in C is taken when the system's state at time t is i. Furthermore the function satisfies that $\pi_{iA(i)}(t) = 1$, which states that only admissible actions are selected in such a strategy.

We say that such a randomized Markov strategy $\pi_{iC}(t)$ is *stationary* if $\pi_{iC}(t) = \pi_{iC}$ for all $t \in [0, \infty)$. A *deterministic stationary strategy* is a function $d : \mathcal{S} \to \mathcal{A}$ such that $d(i) \in A(i)$ for all $i \in \mathcal{S}$, which can be identified by a randomized stationary strategy via $\pi_{iC} = \delta_{d(i)}(C)$, with $\delta_{d(i)}(\cdot)$ being the Dirac measure for the action $d(i)$.

Informally, we can describe the system's behaviour as follows: Suppose that the system is at time $t \geq 0$ in state i and that the controller chooses the action $a \in A(i)$. Then over the interval $[t, t + dt]$ the decision maker will receive an infinitesimal reward $r_{ia}dt$ and a transition will occur to state $j \neq i$ with probability $Q_{ija}dt + o(dt)$ or the system will remain in state i with probability $1 + Q_{iia}dt + o(dt)$.

To model the system's behaviour formally we define for any Markovian strategy π the associated transition rate matrix as

$$Q^{\pi}_{ij}(t) := \int_{A(i)} Q_{ija}\pi_{i\mathrm{d}a}(t) \quad \text{for } i, j \in \mathcal{S} \text{ and } t \geq 0,$$

where we integrate with respect to a. If the transition rates are uniformly bounded (i.e. $||Q|| := \sup_{i \in \mathcal{S}, a \in \mathcal{A}} -Q_{iia} < \infty$) and the state space $\mathcal{S}$ is finite, then for any Markov strategy we obtain by the results discussed in Section 2.1 a Markov process with regular transition function $p^{\pi}(s, i, t, j)$, which satisfies for all $i, j \in \mathcal{S}$ and all $0 \leq s \leq t$ the Kolmogorov forward and backward equations

$$\begin{aligned}\frac{\partial p^{\pi}(s,i,t,j)}{\partial t} &= \sum_{k \in \mathcal{S}} p^{\pi}(s,i,t,k) Q^{\pi}_{kj}(t) \\ \frac{\partial p^{\pi}(s,i,t,j)}{\partial s} &= -\sum_{k \in \mathcal{S}} Q^{\pi}_{ik}(s) p^{\pi}(s,k,t,j).\end{aligned}$$

Given $0 < \beta < 1$ the controller aims to maximize the expected discounted reward, where the expected discounted reward of a strategy π given that the initial state is $i \in \mathcal{S}$ is defined as

$$V^{\pi}_i := \int_0^{\infty} e^{-\beta t} \sum_{j \in \mathcal{S}} r^{\pi}_j(t) p^{\pi}(0, i, t, j) dt, \tag{2.3}$$

where

$$r^{\pi}_i(t) := \int_{A(i)} r_{ia}\pi_{i\mathrm{d}a}(t).$$

The optimal discounted reward function is given by

$$V_i^* := \sup_\pi V_i^\pi$$

for every initial state $i \in \mathcal{S}$. A strategy π^* is said to be *discounted reward optimal* if $V_i^{\pi^*} = V_i^*$ for all $i \in \mathcal{S}$.

For the rest of this subsection we restrict our attention to finite state and action spaces. However, many of the following results also hold given suitable compactness and continuity assumptions. As a first step, we state the following result summarizing the main findings for continuous time Markov decision processes (see Guo and Hernández-Lerma (2009, Lemma 4.4, Remark 4.9, Theorem 4.10) and Kakumanu (1971, Theorem 3.3)):

Lemma 2.1. *The expected discounted reward of any discounted reward optimal strategy is given by the unique solution of the optimality equation*

$$\beta V_i^* = \max_{a \in A(i)} \left\{ r_{ia} + \sum_{j \in \mathcal{S}} Q_{ija} V_j^* \right\}.$$

Moreover, there is a deterministic discounted reward optimal stationary strategy.

In Kakumanu (1971) (Theorem 2.4) we furthermore find that the value given a certain stationary deterministic strategy d can be characterized as the unique solution of a certain linear equation:

Lemma 2.2. *Let d be a stationary deterministic strategy. The expected discounted reward of this strategy is the unique solution g of*

$$\beta g = r^d + Q^d g,$$

where r^d is the reward vector given the strategy d and Q^d the transition rate matrix given the strategy d, i.e. $r^d = (r_{id(i)})_{i \in \mathcal{S}}$ and $Q^d = (Q_{ijd(i)})_{i,j \in \mathcal{S}}$.

Remark 2.3. As some methods for discrete time models are better developed than those for continuous time models, the following tool known as *uniformization* is very helpful: Since the transition rates are bounded in our setting, we can transform our continuous time Markov decision process into an equivalent discrete time Markov decision process by setting

$$\alpha := \frac{||Q||}{\beta + ||Q||}, \quad \bar{r}_{ia} := \frac{r_{ia}}{\beta + ||Q||} \quad \text{and} \quad \bar{P}_{ija} := \frac{Q_{ija}}{||Q||} + \delta_{ij}$$

for all $i, j \in \mathcal{S}$ and $a \in A(i)$, where $||Q|| := \sup_{i \in \mathcal{S}} \sup_{a \in A(i)} -Q_{iia}$ (Guo and Hernández-Lerma, 2009, Remark 6.1). In Kakumanu (1977) it is shown that without adjusting the rewards (i.e. $\bar{r}_{ia} = r_{ia}$) the expected discounted reward of a certain stationary strategy in the continuous time model and the discrete time model are proportional. In our case with the adjusted rewards we obtain, relying on the same proof, that the expected discounted reward given a certain stationary strategy in the continuous time Markov decision process and the corresponding discrete time Markov decision process are the same. Thus, we obtain that a stationary strategy is optimal for the discrete time model if and only if it is optimal for the continuous time model. △

2.2.2. Discrete Time Markov Decision Processes

In Puterman (1994) the collection of objects

$$\{T, \mathcal{S}, (A(i))_{i\in\mathcal{S}}, P_{ija}, r_{ia}\}$$

is called a discrete time Markov decision process. It consists of a set of decision epochs, a state space, action sets, transition probabilities and reward functions. The set T is the set of decision epochs and will be $\mathbb{N}$ for our purpose, $\mathcal{S}$ is the state space, which we here assume to be finite, and $A(i)$ is the set of possible actions when the system is in state i. By $\mathcal{A} := \bigcup_{i\in\mathcal{S}} A(i)$ we denote the set of all possible actions and by $\mathcal{P}(A(i))$ we denote the collection of probability distributions on (Borel) subsets of $A(i)$. Agents are again allowed to randomize over actions, which corresponds to selecting a probability distribution in the probability simplex $\mathcal{P}(A(i))$.

The transition rates are functions P_{ija} ranging from $\mathcal{S} \times \mathcal{S} \times \mathcal{A}$ to $[0, 1]$ such that $\sum_{j\in\mathcal{S}} P_{ija} = 1$ for all $i \in \mathcal{S}$ and $a \in A(i)$. The reward functions r_{ia} range from $\mathcal{S} \times \mathcal{A}$ to $\mathbb{R}$ and we assume in the following that the rewards are uniformly bounded, that is we assume that we find an $M \in \mathbb{R}$ such that $|r_{ia}| \leq M$ for all $i \in \mathcal{S}$ and $a \in \mathcal{A}$.

Puterman (1994) allows in his investigations for more general strategies than Guo and Hernández-Lerma (2009), namely he allows for *history-dependent strategies* π. These are defined as measurable functions that take all states up to time n and all actions up to time $n-1$ into account when deciding for a probability distribution over $A(i)$ to be chosen at time n being in state i. He introduces *(Markovian) strategies* in a similar fashion as Guo and Hernández-Lerma (2009) as measurable functions $\pi : \mathcal{S} \times \mathbb{N} \to \mathcal{P}(\mathcal{A})$, which should satisfy $\sum_{a\in A(i)} \pi_{ia}(n) = 1$ for all $i \in \mathcal{S}, n \in \mathbb{N}$. Since we are only interested in Markovian strategies, we often write strategies instead on Markovian strategies. We define stationary strategies and deterministic strategies as in the continuous time case.

Informally, we can now describe the behaviour of the system as follows: As a result of choosing an action $a \in A(i)$ in state i at decision epoch n, the decision maker receives a reward r_{ia} and the system is in state j at time $n + 1$ with probability P_{ija}. A formal description of the system's behaviour can be obtained in a similar fashion as in Subsection 2.2.1, for details we refer to Section 2.1.6 in Puterman (1994).

The agent again maximizes the expected discounted reward with discount factor $0 < \alpha < 1$, which for a strategy π is given by

$$V_i^\pi = \mathbb{E}_i^\pi \left[\sum_{n=1}^{\infty} \alpha^{n-1} r_{X_n Y_n} \right],$$

where X_n is the system's state at time n and Y_n is the action chosen at time n. If π is a stationary strategy, then an alternative representation of the expected discounted reward is given by

$$V^\pi = \sum_{n=1}^{\infty} \alpha^{n-1} (P^\pi)^{n-1} r^\pi,$$

with

$$r_i^\pi = \sum_{a \in A(i)} \pi_{ia} r_{ia} \quad \text{and} \quad P_{ij}^\pi = \sum_{a \in A(i)} \pi_{ia} P_{ija}. \tag{2.4}$$

This second representation of the expected total discounted reward criterion gives rise to a powerful tool to compute the expected discounted reward for a stationary strategy (Puterman, 1994, Theorem 6.1.1), namely the policy evaluation equation:

Lemma 2.4. *For any stationary strategy π the expected total discounted reward function V^π is the unique solution in $\mathbb{R}^S$ of*

$$v = r^\pi + \alpha P^\pi v. \tag{2.5}$$

Furthermore, $(I - \alpha P^\pi)$ is invertible and we obtain

$$V^\pi = (I - \alpha P^\pi)^{-1} r^\pi.$$

The optimal discounted reward function, often called value function, is given by

$$V_i^* = \sup_\pi V_i^\pi,$$

where the supremum is taken over all history dependent strategies π. We say that a strategy π^* is discounted reward optimal whenever $V_i^{\pi^*} \geq V_i^\pi$ for all $i \in \mathcal{S}$ and all history-dependent strategies π.

The optimality equation

$$V_i^* = \sup_{a \in A(i)} \left\{ r_{ia} + \sum_{j \in \mathcal{S}} \alpha P_{ija} V_j^* \right\}$$

is the most important tool in theoretical investigations of optimal stationary strategies. One can show that there is at most one solution of this equation which is moreover the value function (Puterman (1994, Theorem 6.2.2)) and that a solution exists in case of bounded rewards and discrete state spaces (Puterman (1994, Theorem 6.2.5)). Furthermore, a history-dependent strategy is optimal if and only if V^π satisfies the optimality equation, which allows to easily verify whether a given strategy is optimal (Puterman, 1994, Theorem 6.2.6). Moreover, by Puterman (1994, Theorem 6.2.10) an optimal deterministic stationary strategy exists when $\mathcal{S}$ is discrete and $A(i)$ is finite for all $i \in \mathcal{S}$, ensuring existence of a deterministic stationary strategies in the cases we are interested in. Again, similar results can be obtained given suitable compactness and continuity assumptions for more general action spaces.

2.2.3. A Characterization of the Optimal Stationary Strategies of a Markov Decision Process

We prove in this section that in case of a finite state and action space (that is $\mathcal{S} = \{1, \dots, S\}$ and $\mathcal{A} = \{1, \dots, A\}$) the set of all optimal stationary strategies for a Markov decision process in discrete as well as continuous time is the convex hull of all optimal deterministic stationary strategies.

We start with the preliminary observation that every randomized stationary strategy can be written as a (not necessarily unique) convex combination of deterministic stationary strategies. Additionally, we provide one way to compute some set of coefficients in such a convex combination, which will prove to be helpful in characterizing the set of all stationary mean field equilibria in Chapter 5. We denote by Π^s the set of all stationary strategies and by D^s the set of all deterministic stationary strategies.

Lemma 2.5. *Every randomized stationary strategy $\pi \in \Pi^s$ can be represented as a convex combination of deterministic stationary strategies.*

The proof of this result is an application of Minkowski's Theorem (Theorem 5.10 in Brøndsted (1983)) which characterizes a compact convex set C in terms of its extreme points. A point x is an *extreme point* if $C \setminus \{x\}$ is again a convex set and the set of all extreme points is denoted by $\text{ext}(C)$.

Lemma 2.6 (Minkowski's Theorem)**.** *Let C be a compact and convex set in $\mathbb{R}^d$, then*

$$C = conv(ext(C)),$$

that is C is the convex hull of the extreme points of C.

Proof of Lemma 2.5. We start the proof by observing that deterministic strategies are indeed extreme points of Π^s: Let d be a deterministic strategy and let $\pi^1, \pi^2 \in \Pi^s \setminus \{d\}$ be arbitrary strategies and let $\lambda \in [0, 1]$. We show that $\lambda\pi^1 + (1 - \lambda)\pi^2$ again lies in

$\Pi^s \setminus \{d\}$. As Π^s is itself convex, it is clear that $\lambda\pi^1 + (1-\lambda)\pi^2 \in \Pi^s$. It remains to show that it is not d itself. This statement is trivial for $\lambda \in \{0,1\}$. For $\lambda \in (0,1)$ we note that for at least one state i it holds that $\pi^1_{i\cdot} \neq d_{i\cdot}$, which particularly means that for the action a that strategy d chooses, π^1 assigns a probability less than one. Now we have

$$\lambda\pi^1_{ia} + (1-\lambda)\pi^2_{ia} \leq \lambda\pi^1_{ia} + (1-\lambda) < \lambda + (1-\lambda) = 1.$$

Thus, the convex combination indeed does not equal d.

Next, we show that non-deterministic strategies are indeed no extreme points: Let π be any non-deterministic strategy, then there exists a state i for which there are two distinct actions a_1, a_2 such that $\pi_{ia_1} > 0$ and $\pi_{ia_2} > 0$. We choose δ such that $\min\{\pi_{ia_1}, \pi_{ia_2}\} > \delta > 0$ and $\min\{1-\pi_{ia_1}, 1-\pi_{ia_2}\} > \delta > 0$. Then $\pi^1 = \pi - \delta 1_{ia_1} + \delta 1_{ia_2}$ and $\pi^2 = \pi + \delta 1_{ia_1} - \delta 1_{ia_2}$ are again randomized strategies, but

$$\frac{1}{2}\pi^1 + \frac{1}{2}\pi^2 = \pi \notin \Pi^s \setminus \{\pi\}.$$

By Minkowski's Theorem (Lemma 2.6) we obtain that every randomized strategy can be written as a convex combination of deterministic strategies. □

We remark that there are strategies which can be represented by several convex combinations of deterministic strategies, for example the strategy π with $\pi_{1a_1} = \pi_{1a_2} = \pi_{2a_1} = \pi_{2a_2} = \frac{1}{2}$ can be written as convex combination of $d^1(1) = a_1, d^1(2) = a_1$ and $d^2(1) = a_2, d^2(2) = a_2$ with equal weights each as well as a convex combination of $d^3(1) = a_1, d^3(2) = a_2$ and $d^4(1) = a_2, d^4(2) = a_1$ again with equal weights.

For later use, we provide a way to obtain directly some set of coefficients for the convex combination:

Lemma 2.7. *Let $\pi \in \Pi^s$ be any stationary randomized strategy. Then*

$$\pi = \sum_{d \in D^s} \lambda_d d \quad \text{with} \quad \lambda_d = \pi_{1d(1)} \cdot \ldots \cdot \pi_{Sd(S)},$$

where $d(i)$ is the action chosen under d with certainty.

Proof. We start by enumerating all deterministic strategies by tuples from $\mathcal{A}^S$ with $d^{(a_1,\ldots,a_S)}$ being the strategy that chooses action a_i when being in state i, i.e. $d_i^{(a_1,\ldots,a_S)} = a_i$. We have to show that for all $a \in \mathcal{A}$ and $i \in \mathcal{S}$ we have

$$\pi_{ia} = \sum_{(a_1,\ldots,a_S) \in \mathcal{A}^S} \lambda_{d^{(a_1,\ldots,a_S)}} d_{ia}^{(a_1,\ldots,a_S)} :$$

As a first step, we note that

$$d_{ia}^{(a_1,\dots,a_S)} = \begin{cases} 1 & \text{if } a_i = a \\ 0 & \text{else} \end{cases}.$$

Thus, we have

$$\begin{aligned}
&\sum_{(a_1,\dots,a_S)\in\mathcal{A}^S} \lambda_{(a_1,\dots,a_S)} d_{ia}^{(a_1,\dots,a_S)} \\
&= \sum_{(a_1,\dots,a_{i-1},a_{i+1},\dots,a_S)\in\mathcal{A}^{S-1}} \lambda_{(a_1,\dots,a_{i-1},a,a_{i+1},\dots,a_S)} \\
&= \sum_{a_1\in\mathcal{A}} \cdots \sum_{a_{i-1}\in\mathcal{A}} \sum_{a_{i+1}\in\mathcal{A}} \cdots \sum_{a_S\in\mathcal{A}} \pi_{1a_1} \cdots \pi_{i-1,a_{i-1}} \pi_{ia} \pi_{i+1,a_{i+1}} \cdots \pi_{Sa_S} \\
&= \pi_{ia} \sum_{a_1\in\mathcal{A}} \pi_{1a_1} \cdot \ldots \left(\sum_{a_{i-1}\in\mathcal{A}} \pi_{i-1,a_{i-1}} \cdot \left(\sum_{a_{i+1}\in\mathcal{A}} \pi_{i+1,a_{i+1}} \cdot \ldots \cdot \left(\sum_{a_S\in\mathcal{A}} \pi_{Sa_S} \right)\right)\right) \\
&= \pi_{ia},
\end{aligned}$$

where the last line follows inductively from the previous line as $\sum_{a_j\in\mathcal{A}} \pi_{ja_j} = 1$ for all $j \in \mathcal{S}$. □

We are now equipped with all preparations necessary to prove the announced result. Since the uniformization procedure is applicable (see Remark 2.3) and, thus, the sets of optimal stationary strategies for both processes coincide, the main task is to prove the result for discrete time Markov decision process. It states that all optimal stationary strategies are a convex combination of the optimal deterministic stationary strategies, which in turn can be characterized as those strategies that choose the actions in state i that achieve the maximum in the optimality equation's i-line.

Theorem 2.8. *Let*

$$\{\mathbb{N}, \mathcal{S}, (A(i))_{i\in\mathcal{S}}, P_{ija}, r_{ia}\}$$

be a discrete time Markov decision process with the expected discounted reward criterion such that the state space $\mathcal{S}$ and the action space $\mathcal{A} = \bigcup_{i\in\mathcal{S}} A(i)$ are finite. Furthermore, let

$$\mathcal{D} = \{d : \mathcal{S} \to \mathcal{A} : d(i) \in O_i\}$$

with

$$O_i = argmax_{a\in A(i)} \left\{ r_{ia} + \sum_{j\in\mathcal{S}} \alpha P_{ija} V_i^* \right\}.$$

Then a stationary strategy is optimal if and only if it is a convex combination of strategies from $\mathcal{D}$.

Proof. We first note that by Puterman (1994, Corollary 6.2.8) any deterministic strategy in our set $\mathcal{D}$ is indeed optimal. Enumerate $\mathcal{D} = \{d_1, \ldots, d_n\}$ and let $\pi \in \Pi^s$ be a convex combination of strategies in $\mathcal{D}$, that is

$$\pi = \sum_{m=1}^{n} \lambda_m d_m \quad \text{with} \quad \lambda_m \geq 0 \quad \forall m \in \{1, \ldots, n\} \quad \text{and} \quad \sum_{m=1}^{n} \lambda_m = 1.$$

Since r^π and P^π are linear in π (see equation (2.4)), we can write the policy evaluation equation (2.5) as

$$V^\pi = r^\pi + \alpha P^\pi V^\pi = \sum_{m=1}^{n} \lambda_m (r^{d_m} + \alpha P^{d_m} V^\pi).$$

Since for all $m \in \{1, \ldots, n\}$ the strategy d_m is optimal, it follows that $V^\pi = V^*$ is the unique solution of the policy evaluation equation (2.5):

$$r^\pi + \alpha P^\pi V^* = \sum_{m=1}^{n} \lambda_m (r^{d_m} + \alpha P^{d_m} V^*) = \sum_{m=1}^{n} \lambda_m V^* = V^*.$$

By Puterman (1994, Theorem 6.2.5), which states that in our setting the unique solution of the optimality equation is V^*, and by Puterman (1994, Theorem 6.2.6), which states that a strategy is optimal if and only if the expected discounted reward given this strategy is a solution of the optimality equation, we obtain that the strategy π is optimal.

To show the converse implication we assume that π is not a convex combination of deterministic strategies from $\mathcal{D}$. Lemma 2.5 states that there is nonetheless a convex combination of (not necessarily optimal) deterministic strategies representing π. However, by assumption any convex combination of deterministic strategies representing the strategy π has a summand $d \notin \mathcal{D}$ with positive weight, which means that for the strategy d there is an action $i \in \mathcal{S}$ and an action $\tilde{a} \in A(i) \setminus O_i$ such that $d(i) = \tilde{a}$. This implies that also the strategy π chooses that action $\tilde{a}$ in state i with positive probability, that is $\pi_{i\tilde{a}} > 0$. By the policy evaluation equation (2.5) the i-th component of the strategy value satisfies

$$\begin{aligned} V_i^\pi &= r_i^\pi + \sum_{j \in \mathcal{S}} \alpha (P^\pi)_{ij} V_j^\pi \\ &= \sum_{a \in A(i)} \pi_{ia} \left(r_{ia} + \sum_{j \in \mathcal{S}} \alpha P_{ija} V_j^\pi \right) \\ &\leq \sum_{a \in A(i)} \pi_{ia} \left(r_{ia} + \sum_{j \in \mathcal{S}} \alpha P_{ija} V_j^* \right) \end{aligned}$$

$$< \sum_{a\in A(i)} \pi_{ia} \max_{a'\in A(i)} \left\{ r_{ia'} + \sum_{j\in\mathcal{S}} \alpha P_{ija'} V_j^* \right\}$$

$$= \max_{a'\in A(i)} \left\{ r_{ia'} + \sum_{j\in\mathcal{S}} \alpha P_{ija'} V_j^* \right\} = V_i^*,$$

where the second lines follows from $V^\pi \leq V^*$ and the third line follows from the fact that $\tilde{a} \notin O_i$ is chosen with positive probability $\pi_{i\tilde{a}}$. Since $V_i^\pi < V_i^*$ and since, by the finiteness of $\mathcal{S}$ and $\mathcal{A}$, an optimal strategy achieving value V^* exists, we obtain that π is not optimal. □

For continuous time Markov decision processes we obtain the following statement as a corollary:

Corollary 2.9. *Let*

$$\{\mathcal{S}, \mathcal{A}, (A(i))_{i\in\mathcal{S}}, Q_{ija}, r_{ia}\}$$

be a continuous time Markov decision process with the expected discounted reward criterion, where the state space $\mathcal{S}$ as well as the action space $\mathcal{A}$ are finite. Let $\mathcal{D} = \{d : \mathcal{S} \to \mathcal{A} : d(i) \in O_i\}$ with

$$O_i = \mathit{argmax}_{a\in A(i)} \left\{ r_{ia} + \sum_{j\in\mathcal{S}} Q_{ija} V_j^* \right\}.$$

Then a stationary strategy is optimal if and only if it is a convex combination of strategies from $\mathcal{D}$.

Proof. We note that the set of optimal strategies of the continuous time Markov decision process and the discrete time Markov decision process obtained through uniformization are the same (Remark 2.3). Thus, we have, by Theorem 2.8, that the statement holds with O_i defined as

$$O_i = \mathrm{argmax}_{a\in A(i)} \left\{ \bar{r}_{ia} + \sum_{j\in\mathcal{S}} \alpha \bar{P}_{ija} \tilde{V}_j^* \right\},$$

with $\tilde{V}^*$ being the value function of the discrete time Markov decision process and

$$\alpha = \frac{||Q||}{\beta + ||Q||}, \quad \bar{r}_{ia} = \frac{r_{ia}}{\beta + ||Q||}, \quad \bar{P}_{ija} = \frac{Q_{ija}}{||Q||} + \delta_{ij}.$$

Noting that V^* is the value function for the continuous time Markov decision process as well as the corresponding uniformized discrete time Markov decision process (see Remark 2.3), we obtain the following relation between the optimality equations:

$$
\begin{aligned}
V_i^* &= \max_{a \in A(i)} \left\{ \bar{r}_{ia} + \sum_{j \in \mathcal{S}} \alpha \bar{P}_{ija} V_j^* \right\} \\
&= \max_{a \in A(i)} \left\{ \frac{r_{ia}}{\beta + ||Q||} + \sum_{j \in \mathcal{S}} \frac{||Q||}{\beta + ||Q||} \left(\frac{Q_{ija}}{||Q||} + \delta_{ij} \right) V_j^* \right\} \\
&= \max_{a \in A(i)} \left\{ \frac{1}{\beta + ||Q||} \left(r_{ia} + \sum_{j \in \mathcal{S}} Q_{ija} V_j^* \right) + \frac{||Q||}{\beta + ||Q||} V_i^* \right\} \\
\Leftrightarrow \frac{\beta}{\beta + ||Q||} V_i^* &= \max_{a \in A(i)} \left\{ \frac{1}{\beta + ||Q||} \left(r_{ia} + \sum_{j \in \mathcal{S}} Q_{ija} V_j^* \right) \right\} \\
\Leftrightarrow \beta V_i^* &= \max_{a \in A(i)} \left\{ r_{ia} + \sum_{j \in \mathcal{S}} Q_{ija} V_j^* \right\}.
\end{aligned}
$$

This implies that we can characterize the sets O_i as those achieving the maximum in the right-hand side of the optimality equation (for the continuous time Markov decision process), which means

$$
O_i = \mathrm{argmax}_{a \in A(i)} \left\{ r_{ia} + \sum_{j \in \mathcal{S}} Q_{ija} V_j^* \right\}.
$$

□

Remark 2.10. We remark that in order to compute O_i it is necessary to determine the value function V^*. But also in case that this is not possible the theorem/corollary has an important implication: The set of all optimal strategies can always be characterized by a set $A_1 \times \ldots \times A_S$ with A_i being the optimal actions in state i, such that the set of all optimal strategies is a convex combination from $\{d \in D^s : d(i) \in A_i \text{ for all } i \in \mathcal{S}\}$. △

3. The Model

In this chapter we introduce the mean-field game model we are working with in this thesis. It is equivalent to the model introduced in Doncel et al. (2016a). However, we introduce it in a probabilistic way, which is crucial at several points in the thesis, in particular in Chapter 5. In Section 3.1 we formally introduce the model and illustrate the abstract definitions by formulating a simplified version of the corruption model of Kolokoltsov and Malafeyev (2017), which we will analyse later, in terms of our model. In Section 3.2 we thereafter discuss the model.

3.1. Formal Description of the Model

Let $\mathcal{S} = \{1, \dots, S\}$ $(S > 1)$ be the set of possible states of each player and let $\mathcal{A} = \{1, \dots, A\}$ be the set of possible actions. With $\mathcal{P}(\mathcal{S})$ we denote the probability simplex over $\mathcal{S}$ and with $\mathcal{P}(\mathcal{A})$ the probability simplex over $\mathcal{A}$. A *(mixed) strategy* is a measurable function $\pi : \mathcal{S} \times [0, \infty) \to \mathcal{P}(\mathcal{A})$, $(i, t) \mapsto (\pi_{ia}(t))_{a \in \mathcal{A}}$ with the interpretation that $\pi_{ia}(t)$ is the probability that at time t and in state i the player chooses action a. We say that a strategy $\pi = d : \mathcal{S} \times [0, \infty) \to \mathcal{P}(\mathcal{A})$ is *deterministic* if it satisfies for all $t \geq 0$ and for all $i \in \mathcal{S}$ that there is an $a \in \mathcal{A}$ such that $d_{ia}(t) = 1$ and $d_{ia'} = 0$ for all $a' \in \mathcal{A} \setminus \{a\}$. Throughout the presentation we often use the following equivalent representation, which is to represent a deterministic strategy as a function $d : \mathcal{S} \times [0, \infty) \to \mathcal{A}$, $(i, t) \mapsto d_i(t)$ with the interpretation that $d_i(t) = a$ states that at time t in state i action a is chosen. With Π we denote the set of all (mixed) strategies and with D the set of all deterministic strategies.

The individual dynamics of each player given a Lipschitz continuous flow of population distributions $m : [0, \infty) \to \mathcal{P}(\mathcal{S})$ and a strategy $\pi : \mathcal{S} \times [0, \infty) \to \mathcal{P}(\mathcal{A})$ are given as a Markov process $X^\pi(m)$ with given initial distribution $x_0 \in \mathcal{P}(\mathcal{S})$ and infinitesimal generator given by the $Q(t)$-matrix

$$(Q^\pi(m(t), t))_{ij} = \sum_{a \in \mathcal{A}} Q_{ija}(m(t)) \pi_{ia}(t),$$

where for all $a \in \mathcal{A}$ and $m \in \mathcal{P}(\mathcal{S})$ the matrices $(Q_{\cdot\cdot a}(m))_{a \in \mathcal{A}}$ are conservative generators, that is $Q_{ija}(m) \geq 0$ for all $i, j \in \mathcal{S}$ with $i \neq j$ and $\sum_{j \in \mathcal{S}} Q_{ija}(m) = 0$ for all $i \in \mathcal{S}$.

We remark that under the standing assumption A1 any flow of population distributions $m : [0,\infty) \to \mathcal{P}(\mathcal{S})$, which emerges from the population jointly using the strategy π, is Lipschitz continuous (see Lemma 4.2). Thus, the requirement that m is Lipschitz continuous does not impose an additional restriction on the considered model.

Given the initial condition $x_0 \in \mathcal{P}(\mathcal{S})$, the goal of each player is to maximize his expected discounted reward, which is given by

$$V_{x_0}(\pi^0, m) = \int_0^\infty \left(\sum_{i \in \mathcal{S}} \sum_{a \in \mathcal{A}} x_i^{\pi^0}(t) r_{ia}(m(t)) \pi_{ia}^0(t) \right) e^{-\beta t} \mathrm{d}t, \tag{3.1}$$

where $r : \mathcal{S} \times \mathcal{A} \times \mathcal{P}(\mathcal{S}) \to \mathbb{R}$ is a real-valued function, $\beta \in (0,1)$ is the discount rate and $x_i^{\pi^0}(t)$ is the probability that the individual player using strategy π^0 is in state i at time t, which is given as the marginal of the time-inhomogeneous Markov chain with generator matrix $Q^{\pi^0}(t, m(t))$ and initial condition x_0. That is, for a fixed flow of population distributions $m : [0,\infty) \mapsto \mathcal{P}(\mathcal{S})$ we face a Markov decision process with expected discounted reward criterion and time-inhomogeneous reward functions and transition rates.

We will work under the following continuity assumption, which will ensure that there is indeed a Markov process with inhomogeneous $Q(t)$-matrix $Q^\pi(m(t), t)$ in the sense of Section 2.1 (see Lemma A.4 for details):

Assumption A1. *For all $i, j \in \mathcal{S}$ and all $a \in \mathcal{A}$ the function $m \mapsto Q_{ija}(m)$ mapping from $\mathcal{P}(\mathcal{S})$ to $\mathbb{R}$ is Lipschitz-continuous in m . For all $i \in \mathcal{S}$ and all $a \in \mathcal{A}$ the function $m \mapsto r_{ia}(m)$ mapping from $\mathcal{P}(\mathcal{S})$ to $\mathbb{R}$ is continuous in m.*

With these preparations we define dynamic mean field equilibria:

Definition 3.1. Given an initial distribution m_0 a *mean field equilibrium* is a flow of population distributions $m : [0,\infty) \to \mathcal{P}(\mathcal{S})$ with $m(0) = m_0$ and a strategy $\pi : \mathcal{S} \times [0,\infty) \to \mathcal{P}(\mathcal{A})$ such that

- for all $t \geq 0$ the marginal distribution of the process $X^\pi(m)$ at time t is given by $m(t)$
- it holds that $V_{m_0}(\pi, m) \geq V_{m_0}(\pi', m)$ for all $\pi' \in \Pi$.

As in standard game theory, our concept of mean field equilibrium captures the intuitive idea that no player wants to deviate: Given that all players play according to strategy π, the population's distribution will be m. If an individual player evaluates whether he wants to deviate from playing π, he asks whether there is a strategy that yields a higher payoff given m. However, due to the second condition, this is not possible. Therefore, we indeed face an equilibrium in the classical economic sense.

Remark 3.2. Using the Kolmogorov forward equation (Guo and Hernández-Lerma, 2009, Proposition C.4) we see that the first condition implies the analytic condition used in Doncel et al. (2016a) to characterize mean field equilibria, which states that m is solution (in the sense of Caratheodory) of

$$\dot{m}_j(t) = \sum_{i \in \mathcal{S}} m_i(t) Q^{\pi}_{ij}(m(t), t) \quad \forall j \in \mathcal{S} \tag{3.2}$$

with initial condition $m(0) = m_0$. A proof of this equivalence as well as a short discussion of Caratheodory solutions can be found in the Appendix A. △

In order to define stationary mean field equilibria, we first introduce, analogous to Section 2.2, the notion of stationary strategies: A *stationary strategy* is a map $\pi : \mathcal{S} \times [0, \infty) \to \mathcal{P}(\mathcal{A})$ such that $\pi_{ia}(t) = \pi_{ia}$ for all $t \geq 0$. We denote by Π^s the set of all stationary strategies and by D^s the set of all deterministic stationary strategies.

Definition 3.3. A stationary mean field equilibrium is given by a stationary strategy $\pi \in \Pi^s$ and a vector $m \in \mathcal{P}(\mathcal{S})$ such that

- for all $t \geq 0$ the marginal distribution of the process $X^{\pi}(m)$ with initial condition $x_0 = m$ at time t is given by m
- for any initial distribution $x_0 \in \mathcal{P}(\mathcal{S})$ we have $V_{x_0}(\pi, m) \geq V_{x_0}(\pi', m)$ for all $\pi' \in \Pi$.

This notion is a sensible formalization of stationary equilibria: If $m(0) = m$ and the strategy is π, then the population's distribution will be m for all time points. An individual agent at a given time point can be in any state, however, if he evaluates whether he wants to deviate from playing π, the second condition ensures that this is not beneficial for him. Thus, he has no incentive to deviate from the equilibrium strategy π, which means that the population will indeed remain in the stationary equilibrium regime of playing π.

Remark 3.4. We remark that the matrix $Q^{\pi}_{ij}(m, t)$ does not depend on t in this context. Therefore, we write $Q^{\pi}_{ij}(m) := Q^{\pi}_{ij}(m, t)$. Using this, we obtain that the first condition is equivalent to

$$0 = \sum_{i \in \mathcal{S}} m_i Q^{\pi}_{ij}(m) \quad \forall j \in \mathcal{S}.$$

Furthermore, we point out, that the second condition seems stricter than it actually is: We require that there is a stationary strategy that achieves the highest possible value among all (also non-stationary) strategies for all initial conditions simultaneously. However, we will prove in Section 5.1 that any stationary strategy that is optimal for an associated Markov decision process satisfies this. △

Remark 3.5. In contrast to standard models, where the assumptions always imply that a unique optimal best response exists, the mean field equilibria we consider are not fully specified by the distribution because it might happen that several actions are simultaneously optimal and induce the same distribution. However, a dynamic mean field equilibrium is fully specified by describing the equilibrium strategy because we can show using standard techniques (Walter, 1998, Theorem 10.XX) that there is at most one Caratheodory solution to the differential equation (3.2) (see Lemma A.6 for details). For a stationary mean field equilibrium this again does not hold true as it might happen that given a strategy there are multiple stationary distributions. For this reason, we define mean field equilibria always as pairs of the equilibrium distribution and the equilibrium strategy. △

Remark 3.6. We are now in the position to explain why for non-trivial models (in the sense that there is not one action that maximizes the Hamiltonian for every population distribution) we always obtain population distributions at which several actions maximize the Hamiltonian: Since $Q(\cdot)$ and $r(\cdot)$ are continuous in m, also the Hamiltonian is continuous in m. Therefore, if we fix the costate variables, the sets of population distributions in which a particular action is a maximizer of the Hamiltonian are closed. Since the action space is finite and the set of all population distribution vectors is connected, we obtain that if there is more than one action that maximizes the Hamiltonian for some population distribution, then the set of population distributions where several actions simultaneously maximize the Hamiltonian is non-empty. This implies that for the case of finite action spaces the assumption that a unique maximizer of the Hamiltonian exists is violated in all interesting cases. Thus, new methods for the analysis of these models are necessary. △

Example. In Section 6.2 we analyse a simplified version of the corruption model in Kolokoltsov and Malafeyev (2017), which aims to analyse the effects of social pressure onto the spread of corruption in a society. In this model, a player can be in one of the three states honest (H), corrupt (C) and reserved (R). The corrupt players get the highest wage, the honest players a medium wage and the reserved players no wage. The players can choose, given that they are not reserved, whether they want to stay corrupt/honest or whether they want to change behaviour. Formally, the state space is given by $\mathcal{S} = \{C, H, R\}$ and the action space is given by $\mathcal{A} = \{change, stay\}$.

The dynamics of the individual player are influenced by two major sources: First, the individual decision to change states, in which case this happens with rate b, and second, the peer pressure, which works through two channels. On the one hand the more players are corrupt the higher is the rate that honest players become corrupt. On the other hand the more players are honest the higher is the risk of being convicted. Additionally, agents recover from state R with a constant rate r.

Formally, the transition rates of an individual player are as follows: In state C the player moves to state R (that is becomes convicted of being corrupt) with rate $q_{\text{soc}} m_H$, which increases in the share m_H of players in state H. The player moves from state C to state H with rate 0 if he chooses action *stay* and with rate b if he chooses action *change*. In state H the player becomes corrupt (that is moves to state C) with rate $q_{\text{inf}} m_C$ if he chooses the action *stay* and with rate $q_{\text{inf}} m_H + b$ if he chooses the action *change*. In state R the player moves to state H with rate r. Other transitions are not directly possible, that is the corresponding transition rates are zero. In total, this description yields to the following transition rate matrices

$$Q_{\cdot\cdot change} = \begin{pmatrix} -(b + q_{\text{soc}} m_H) & b & q_{\text{soc}} m_H \\ b + q_{\text{inf}} m_C & -(b + q_{\text{inf}} m_C) & 0 \\ 0 & r & -r \end{pmatrix}$$

$$Q_{\cdot\cdot stay} = \begin{pmatrix} -q_{\text{soc}} m_H & 0 & q_{\text{soc}} m_H \\ q_{\text{inf}} m_C & -q_{\text{inf}} m_C & 0 \\ 0 & r & -r \end{pmatrix}.$$

The individual player can then choose any measurable mapping $\pi : \mathcal{S} \times [0, \infty) \to \mathcal{P}(\mathcal{A})$ as strategy, for example for any $T \geq 0$ he could choose

$$(i,t) \mapsto \pi_i(t) = \begin{cases} (e^{-t}, 1 - e^{-t}) & \text{if } i = C, t \leq T \\ (0,1) & \text{else} \end{cases}.$$

Given this strategy the agent chooses the action *stay* whenever the agent is not corrupt or the time is larger than T. If the agent is in corrupt and the time is given by $t \leq T$, then the agent randomizes over the actions. More precisely, he chooses the action *change* with probability e^{-t} and the action *stay* with probability $1 - e^{-t}$. An example of a deterministic strategy would be

$$(i,t) \mapsto \pi_i(t) = \begin{cases} (1,0) & \text{if } i = C, t \leq T \\ (0,1) & \text{else} \end{cases}$$

and an example of a deterministic stationary strategy would be

$$(i,t) \mapsto \pi_i(t) = \begin{cases} (1,0) & \text{if } i = C \\ (0,1) & \text{else} \end{cases}.$$

The reward functions are given by

$$\begin{aligned} r_{C,stay}(m) &= r_{C,change}(m) = 10, \\ r_{H,stay}(m) &= r_{H,change}(m) = 5, \\ r_{R,stay}(m) &= r_{R,change}(m) = 0. \end{aligned}$$

We summarize the model description in a transition graph as in Figure 3.1.

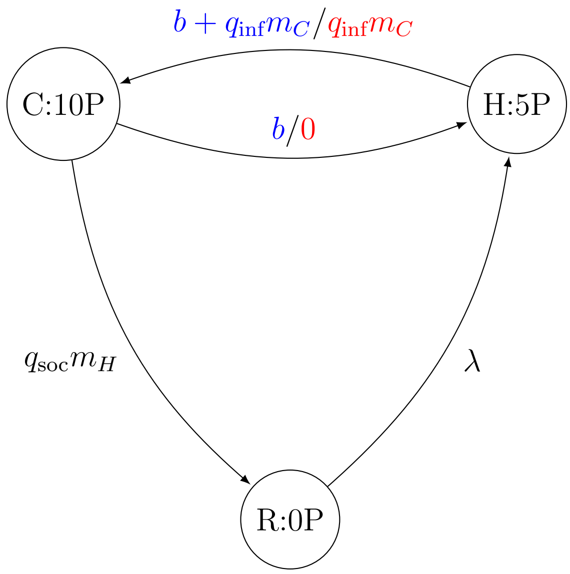

Figure 3.1.: Representation of the corruption model. The transition rates in black are action independent, the blue transition rates are those associated to choosing the action *change* and the red transition rates are those associated to choosing the action *stay*.

The individual control problem is now given as follows: Given a Lipschitz continuous function $m : [0, \infty) \to \mathcal{P}(\mathcal{S})$, a player searches for a measurable strategy π such that (3.1) is maximized, where his individual dynamics are given by the time-inhomogeneous Markov chain with initial distribution x_0 and generator

$$Q^\pi(m(t), t) = \begin{pmatrix} -(b\pi_{C\,change}(t) + q_{\text{soc}} m_H(t)) & b\pi_{C\,change}(t) & q_{\text{soc}} m_H(t) \\ \pi_{H\,change}(t) b + q_{\inf} m_C(t) & -(b\pi_{H\,change}(t) + q_{\inf} m_C(t)) & 0 \\ 0 & r & -r \end{pmatrix}.$$

$\triangle$

3.2. Discussion of the Model

This section discusses why we set certain assumptions in our model and in how far it is possible to modify them.

The central goal of the model introduced here is to twofold: We wish for a model that is closely related to the classical diffusion-based models, but at the same time we want to analyse non-trivial models with finite action spaces. We first note that in finite action space settings it is not easily possible to utilize differential calculus. The main reasons are that we do not face a unique maximizer of the Hamiltonian and that we now in general have to consider set-valued best-response maps that are not continuous. Thus, we restrict our attention to discrete state spaces. We cannot even assume that the state space is countable since the whole theory presented here crucially depends on the finite state space. In particular, the results on the computation of stationary equilibria in Chapter 5 require a finite state space.

We assume that the individual dynamics are given by a time-inhomogeneous Markov chain since we again wish that the individual dynamics are given by a Markov process. Moreover, we directly assume that the action space is finite. This assumption implies certain properties, like the existence of optimal strategies or the continuity of maxima over the finitely many deterministic stationary strategy. If we would drop the assumption that the action space is finite, it would still be possible to obtain some of the results, but we would need to set up additional assumptions to obtain the properties mentioned above and this would add technical difficulties. However, we can only characterize stationary equilibria in the case of finite action space models since, else, we would have to solve infinitely many (possibly completely different) fixed point problems. For these reasons we directly restrict to the case of finite action spaces here.

The infinite time horizon is mainly motivated by the fact that in economic applications the central interest are stationary solutions, in the sense that one considers equilibria that have a stationary equilibrium strategy and a stationary population distribution given this strategy, in which setting a finite time horizon would be artificial.

That we maximize the expected discounted reward is motivated by the fact that this is the central optimization criterion in most economic applications. However, we expect that it is possible to choose other optimization criteria. In particular for the expected average reward criterion it should be possible to obtain similar results since also for this optimality criterion the results in the literature indicate that a similar characterization of optimal strategies as presented in Subsection 2.2.3 is, under certain assumptions, possible (see Guo and Hernández-Lerma (2009, Section 5) and Puterman (1994, Section 8.4)).

The strategies, which we allow for, are rather simple. More precisely, we do not consider history-dependent strategies, but only Markovian strategies for several reasons. First, these strategies are sensible from an economic perspective (see Maskin and Tirole (2001)): They are simple, yet they capture rational decision making. Moreover, they are closely related to subgame perfection since these strategies yield that the past choices of the players influence the future play of the players only through the current state of the system. Second, this assumption is often made in the context of finite state mean field games also for practical reasons since then the individual dynamics given a certain strategy are given by a Markov process and not by a general marked point process, which yields that several technical details do not have to be considered (see Carmona and Delarue (2018a, Section 7.2)). Third, the restriction to a smaller strategy space always reduces the complexity of the necessary computations. This is particularly desirable in the context of mean field games where it is rarely possible to compute explicit solutions.

Another question that occurs with respect to the choice of strategies is the following: Why is the state variable on which the player bases his decision only his own state and not the state of the population? The answer to this question is given in Caines et al. (2017) for general mean field games and it is also valid in our setting: In equilibrium all players adopt the same strategy, thus knowing the strategy and the initial population distribution, the population distribution at all future times becomes predictable (see Lemma A.6). Since an equilibrium is characterized by the fact that no single player wants to deviate from the equilibrium strategy, a player has to find his optimal response given a certain strategy used by all other players. If the player could condition his actions on the (current) population distribution, deviations from the equilibrium strategy could directly be punished, which would yield to a problematic equilibrium notion since we would face non-credible threats as in the theory of sequential games.

Finally, let us remark that the stationary mixed strategy equilibria we obtain here do not allow for the classical criticism of mixed strategy equilibria for Nash equilibria. The notion of mixed strategies in classical games is criticised because in experiments it is not observed that player are indeed able to randomize as long as they are not experts in the game played (Walker and Wooders, 2008). This criticism is not applicable in our setting. Indeed, we obtain that a stationary strategy is optimal if and only if it is a convex combination of optimal deterministic stationary strategies (see Subsection 2.2.3). Thus, we can assume that in a stationary mean field equilibrium with a mixed equilibrium strategy the players are partitioned into different groups each playing a different optimal deterministic stationary strategy such that the aggregated behaviour is the mixed equilibrium strategy.

4. Existence Results

In this chapter we will establish existence of dynamic as well as stationary mean field equilibria in mixed strategies under Assumption A1 introduced in Chapter 3. In contrast to the finite state model presented in Gomes et al. (2013), we do not need any additional assumptions compared to the dynamic case to ensure existence of stationary equilibria.

The existence proofs will both rely on the standard idea to search for fixed points of the best-response correspondence (which is a set-valued map). This map usually maps to any strategy π those strategies that achieve the highest value given π. The problem in our setting is that we cannot ensure continuity/semi-continuity of this map. For this reason, Doncel et al. (2016a) sketched a proof of the dynamic existence result, where dynamic mean field equilibria were linked to fixed points of a best-response correspondence $\phi(\cdot)$ in flows of population distributions. More precisely, they considered a best-response map that maps to any flow of population distributions m those flows of population distributions induced by optimal strategies. Recently, a more detailed version of the proof has been published in Doncel et al. (2019).

In Section 4.1, we prove the dynamic existence result. We rely on the ideas sketched in Doncel et al. (2016a), but provide the complete proof including the topological and measure-theoretical details. In Section 4.2, we then develop based on the same ideas an existence proof for stationary equilibria. This proof heavily relies on the probabilistic formulation of our problem. In particular it relies on the investigations of the optimal stationary strategies started in Section 2.2.3 as well as classical results regarding stationary distribution of time-homogeneous Markov chains.

4.1. Existence of Dynamic Equilibria

This section is devoted to proving the following theorem:

Theorem 4.1. *Any mean field game that satisfies Assumption A1 has a mean field equilibrium.*

The proof of this theorem relies on the analytic characterization of mean field equilibria, which states that a mean field equilibrium is a pair (m, π) satisfying

- $V_{m_0}(\pi, m) \geq V_{m_0}(\pi^0, m)$ for all $\pi^0 \in \Pi$,
- m satisfies

$$\dot{m}_j(t) = \sum_{i\in\mathcal{S}} m_i(t) Q_{ij}^{\pi}(m(t), t) = \sum_{i\in\mathcal{S}} m_i(t) \sum_{a\in\mathcal{A}} Q_{ija}(m(t))\pi_{ia}(t), \quad \forall j \in \mathcal{S} \tag{4.1}$$

  for almost all $t \geq 0$ with initial condition $m(0) = m_0$.

More precisely, the proof consists of the following steps: First, we prove that the flow of population distributions of any mean field equilibrium has to lie in the set $\mathcal{M}$ consisting of all functions $m : [0,\infty) \to \mathcal{P}(\mathcal{S})$ that are Lipschitz continuous with Lipschitz constant L, which is the constant that uniformly bounds $Q_{ija}(m)$ for all $i, j \in \mathcal{S}$, $a \in \mathcal{A}$ and $m \in \mathcal{P}(\mathcal{S})$. With this observation, we then characterize the set of all mean field equilibria as the set of all fixed points of a best-response correspondence $\phi : \mathcal{M} \to \mathcal{M}$ that maps a flow of population distributions to the distribution(s) induced by optimal strategies. For this map $\phi(\cdot)$ we then want to apply a generalized version of Kakutani's fixed point theorem. Therefore, we verify that the necessary conditions are satisfied. More precisely, we show that the set $\mathcal{M}$ is a compact and convex subset of a locally convex space and that the function $\phi(\cdot)$ has non-empty, compact and convex values and is moreover upper semi-continuous. This then allows us to prove the desired existence result.

We start with the observation that any population distribution of a mean field equilibrium has to lie in the set $\mathcal{M}$:

Lemma 4.2. *All solutions of the dynamics equation* (4.1) *are Lipschitz continuous with Lipschitz constant L, where L is the constant that uniformly bounds $Q_{ija}(m)$ for all $i, j \in \mathcal{S}$, $a \in \mathcal{A}$ and $m \in \mathcal{P}(\mathcal{S})$.*

Proof. We have $\sum_{i\in\mathcal{S}} m_i(t) = 1$ for almost all $t \geq 0$, because $(Q_{ija}(m))_{i,j\in\mathcal{S}}$ is conservative for all $m \in \mathcal{P}(\mathcal{S})$ and all $a \in \mathcal{A}$ and therefore $\sum_{i\in\mathcal{S}} \dot{m}_i(t) = 0$. Using this we obtain that

$$\begin{aligned}|\dot{m}_j(t)| &= \left| \sum_{i\in\mathcal{S}} \sum_{a\in\mathcal{A}} m_i(t) Q_{ija}(m(t))\pi_{ia}(t) \right| \\ &\leq \sum_{i\in\mathcal{S}} \sum_{a\in\mathcal{A}} m_i(t) |Q_{ija}(t)| \pi_{ia}(t) \\ &\leq L \sum_{i\in\mathcal{S}} \sum_{a\in\mathcal{A}} m_i(t)\pi_{ia}(t) = L.\end{aligned}$$

The fundamental theorem of calculus for the Lebesgue integral (see Walter (1998, p.122)) yields that $m(\cdot)$ is a Caratheodory solution of (4.1) on $[s, t]$ if and only if

$$\sum_{i\in\mathcal{S}} \sum_{a\in\mathcal{A}} m_i(u) Q_{ija}(m(u))\pi_{ia}(u)$$

is integrable over $[s,t]$ (which is satisfied since the function is uniformly bounded by L) and

$$m(t) = m(s) + \int_s^t \sum_{i\in\mathcal{S}} \sum_{a\in\mathcal{A}} m_i(u) Q_{i\cdot a}(m(u)) \pi_{ia}(u) \mathrm{d}u,$$

where the integral is componentwise.

Thus, for $s \leq t$ we obtain

$$\begin{aligned} |(m(t) - m(s))_j| &= \left| \int_s^t \sum_{i\in\mathcal{S}} \sum_{a\in\mathcal{A}} m_i(u) Q_{ija}(m(u)) \pi_{ia}(u) \mathrm{d}u \right| \\ &\leq \int_s^t \left| \sum_{i\in\mathcal{S}} \sum_{a\in\mathcal{A}} m_i(u) Q_{ija}(m(u)) \pi_{ia}(u) \right| \mathrm{d}u \leq \int_s^t L \mathrm{d}u = L(t-s), \end{aligned}$$

which implies that

$$||m(t) - m(s)||_\infty = \max_{j\in\mathcal{S}} |m(t) - m(s)| \leq L(t-s).$$

□

With these preparations we define the best-response correspondence: We say that given population distributions $m, x \in \mathcal{M}$ there exists a feasible strategy π^0 for (m,x) if it solves the individual dynamics equation

$$\dot{x}_j(t) = \sum_{i\in\mathcal{S}} x_i(t) \sum_{a\in\mathcal{A}} Q_{ija}(m(t)) \pi^0_{ia}(t) \quad \text{for all} \quad j \in \mathcal{S}, \quad x(0) = m_0. \tag{4.2}$$

We then define

$$Y(m,x) = \sup_{\pi^0 \text{ feasible for } (m,x)} \int_0^\infty \left(\sum_{i\in\mathcal{S}} \sum_{a\in\mathcal{A}} x_i(t) \pi^0_{ia}(t) r_{ia}(m(t)) e^{-\beta t} \right) \mathrm{d}t, \tag{4.3}$$

which is the best possible payoff under (m,x) an individual can obtain. We set $Y(m,x) = -\infty$ if no feasible π^0 exists. We note that this definition is also valid if m does not satisfy the dynamics equation (4.1) for any π. The best-response correspondence $\phi : \mathcal{M} \to 2^{\mathcal{M}}$ is given by

$$m \mapsto \phi(m) := \mathrm{argmax}_{x\in\mathcal{M}} Y(m,x).$$

In order to show that this set-valued map indeed has a fixed point we apply Theorem 7.8.6 (p.169) of Granas and Dugundji (2003). For this we first need some definitions:

Let X and Y be subsets of topological vector space and let $F : X \to 2^Y$ be a *set-valued map (correspondence)*, where 2^Y is a short-hand notation for the power set of

Y. A set-valued map is called *upper semicontinuous* if for all $x \in X$ and all open sets $V \supseteq F(x)$ there exists an open set U such that $x \in U$ and for all $x' \in U$ it holds that $F(x') \subseteq V$. If Y is convex, then $F(\cdot)$ is called a *Kakutani map* provided that $F(\cdot)$ is upper semicontinuous with non-empty, compact and convex values.

Lemma 4.3. *Let C be a compact and convex subset of a locally convex space X, and let $F : C \to 2^C$ be a Kakutani-map. Then $F(\cdot)$ has a fixed point, that is there exists an $x \in C$ such that $x \in F(x)$.*

We first verify that $\mathcal{M}$ satisfies the conditions of the fixed point theorem. As a first step, we show that $\mathcal{M}$ is compact with respect to the topology of uniform convergence on compact sets. For this we use a version of the Arzela-Ascoli theorem presented in Kelley (1955) (Theorem 7.17, pp.233-234):

Lemma 4.4. *Let $\mathcal{C}(X, Y)$ be the family of all continuous functions from a regular locally compact topological space X to a Hausdorff uniform space Y, and let $\mathcal{C}(X, Y)$ be equipped with the topology of uniform convergence on compact sets. Then a subfamily $\mathcal{F}$ of $\mathcal{C}(X, Y)$ is compact if and only if*

(a) $\mathcal{F}$ is closed in $\mathcal{C}(X, Y)$,

(b) $\mathcal{F}[x] := \{f(x) : f \in \mathcal{F}\}$ has a compact closure for all $x \in X$, and

(c) the family $\mathcal{F}$ is equicontinuous.

In order to prove that the set $\mathcal{M}$ is compact, we note that the topology of uniform convergence on compact sets is induced by a norm, which implies that our space $\mathcal{M}$ is metrizable and a normed space. The proof of this statement is classical and can be found in Appendix A.3.

Lemma 4.5. *Consider the space $\mathcal{C}([0, \infty), \mathcal{P}(\mathcal{S}))$ equipped with the topology of uniform convergence on compact sets. The space is metrizable and the metric is induced by the norm*

$$||m|| = \sup_{t \geq 0, i \in \mathcal{S}} |m_i(t)| e^{-\beta t},$$

that is the corresponding metric is given by

$$d(m, m') = \sup_{i \in \mathcal{S}, t \geq 0} |m_i(t) - m'_i(t)| e^{-\beta t}.$$

With these preparations we prove that $\mathcal{M}$ is indeed a compact space:

Lemma 4.6. *The set $\mathcal{M}$ is compact in the topology of uniform convergence on compact sets.*

Proof. The set $\mathcal{M}$ is a subset of $\mathcal{C}([0,\infty),\mathcal{P}(\mathcal{S}))$ with $\mathcal{P}(\mathcal{S}) \subseteq \mathbb{R}^n$. We note that both spaces, $[0,\infty)$ and $\mathcal{P}(\mathcal{S})$, are equipped with the subspace topology induced by the standard topology on $\mathbb{R}$ and $\mathbb{R}^n$, respectively. Since both spaces are metric spaces, they are regular as well as Hausdorff (von Querenburg, 2001, p.84). Moreover, every metric space admits a uniform structure, whose topology equals the topology induced by the metric (von Querenburg, 2001, Example 11.6.a). Furthermore in von Querenburg (2001, Example 8.17) we find that $\mathbb{R}^n$ is locally compact and that the intersection of an open subset and a closed subset of a locally compact set is again locally compact. Thus, the set $[0,\infty)$, which is closed, is locally compact.

Therefore, in order to prove that the set $\mathcal{M}$ is compact, it remains to check whether the three conditions (a) to (c) from Lemma 4.4 are satisfied:

(a) We have to prove that $\mathcal{M}$ is closed in $\mathcal{C}([0,\infty),\mathcal{P}(\mathcal{S}))$. Since we consider a metric space, it is sufficient to show that for all $m \in \mathcal{C}([0,\infty),\mathcal{P}(\mathcal{S})) \setminus \mathcal{M}$ there exists an $\epsilon > 0$ such that for all $m' \in \mathcal{C}([0,\infty),\mathcal{P}(\mathcal{S}))$ with $d(m,m') < \epsilon$ it follows that $m' \in \mathcal{C}([0,\infty),\mathcal{P}(\mathcal{S})) \setminus \mathcal{M}$.

    Since $m \in \mathcal{C}([0,\infty),\mathcal{P}(\mathcal{S})) \setminus \mathcal{M}$, there exists a pair $s,t \geq 0$ such that $||m(s) - m(t)||_\infty > L|s-t|$. Choose

$$0 < \epsilon < \frac{1}{2}\left(||m(s) - m(t)||_\infty - L|s-t|\right) e^{\beta \min\{s,t\}}.$$

    This implies that for all $m' \in \mathcal{C}([0,\infty),\mathcal{P}(\mathcal{S}))$ with $d(m,m') < \epsilon$ we have

$$||m(s) - m'(s)||_\infty + ||m(t) - m'(t)||_\infty < ||m(s) - m(t)||_\infty - L|s-t|,$$

    thus

$$\begin{aligned} ||m'(s) - m'(t)||_\infty &= ||m'(s) - m(s) + m(s) - m(t) + m(t) - m'(t)||_\infty \\ &\geq \Big| ||m(s) - m(t)||_\infty - ||m'(s) - m(s) + m(t) - m'(t)||_\infty \Big|. \end{aligned}$$

    We note that

$$\begin{aligned} ||m'(s) - m(s) + m(t) - m'(t)||_\infty &\leq ||m'(s) - m(s)||_\infty + ||m(t) - m'(t)||_\infty \\ &< ||m(s) - m(t)||_\infty - L|s-t|, \end{aligned}$$

    which implies that

$$\begin{aligned} ||m'(s) - m'(t)||_\infty &\geq ||m(s) - m(t)||_\infty - ||m'(s) - m(s) + m(t) - m'(t)||_\infty \\ &> ||m(s) - m(t)||_\infty - (||m(s) - m(t)||_\infty - L|s-t|) \\ &= L|s-t|. \end{aligned}$$

    This in turn implies that $m'(\cdot)$ is also not Lipschitz continuous with Lipschitz constant L. Thus, we obtain $m' \in \mathcal{C}([0,\infty),\mathcal{P}(\mathcal{S})) \setminus \mathcal{M}$.

(b) Since $\mathcal{P}(\mathcal{S})$ is compact (it homeomorphic to the closed unit ball (Schubert, 1964, p.166)) and the closure of any set is closed, the set $\mathcal{M}[t]$ is, as a closed subset of a compact set, also compact.

(c) $\mathcal{M}$ is equicontinuous since a family of functions with a common Lipschitz constant is equicontinuous (von Querenburg, 2001, Example 14.18.a).

□

We start the consideration of $\phi(\cdot)$ by proving that $\phi(m)$ is non-empty for all $m \in \mathcal{M}$ and that for all elements in $\phi(m)$ a feasible strategy that achieves the maximum value exists. In a second step we prove convexity of $\phi(m)$. Afterwards we will prove the auxiliary statement that the graph of $\phi(\cdot)$ is closed in order to deduce that $\phi(m)$ is compact for all $m \in \mathcal{M}$ and that $\phi(\cdot)$ is upper semicontinuous.

Before we start with the investigations of the properties of $\phi(\cdot)$ itself, we show that the set of all feasible tuples (m, x, π) is compact in the space $\mathcal{M} \times \mathcal{M} \times \hat{L}^2$, where

$$\hat{L}^2 = L^2([0,\infty) \times \mathcal{S} \times \mathcal{A}, \hat{\lambda} \times \mu_S \times \mu_A)$$

with $\hat{\lambda}(\mathrm{d}t) = e^{-\frac{1}{2}\beta t}\mathrm{d}t$ and μ_S as well as μ_A being the counting measures on $\mathcal{S}$ and $\mathcal{A}$, respectively. The space $\mathcal{M}$ is again equipped with the topology of uniform convergence on compact sets and the space L^2 is equipped with the topology of weak convergence. By the Eberlein-Shmulyan Theorem (Werner, 2011, Theorem VIII.6.1), compactness with respect to the weak topology and weak sequential compactness, which means that every sequence in the set has a weakly converging subsequence with limit in the set, are equivalent on the Banach space L^2. We remind ourselves that a sequence $f^n \in L^2$ converges weakly to f if for all $g \in L^2$ we have

$$\lim_{n\to\infty} \int f^n g \,\mathrm{d}\left(\hat{\lambda} \times \mu_S \times \mu_A\right) = \int fg \,\mathrm{d}\left(\hat{\lambda} \times \mu_S \times \mu_A\right)$$

(Elstrodt, 2011, p.263-267). This allows us to infer something about the limit behaviour of the right-hand-side of the dynamics equation (4.2), in particular it allows us to obtain the desired compactness statement.

Lemma 4.7. *The set of all tuples (m, x, π) that are feasible is compact.*

Proof. We note that $\mathcal{M} \times \mathcal{M}$ is a metric space and that by Lemma VIII.6.2 in Werner (2011) the weak topology on any weakly compact subset is metrizable. Thus, by Lemma A.7 a set is compact in $\mathcal{M} \times \mathcal{M} \times \hat{L}^2$ if and only if it is sequentially compact. Therefore, it remains to prove that an arbitrary sequence (m^n, x^n, π^n) of feasible tuples has a converging subsequence, whose limit is again feasible.

Since by Lemma 4.6 the space $\mathcal{M}$ is a compact metric space, we find a converging subsequence $(m^{n_l^1})_{l\in\mathbb{N}}$ with limit $m \in \mathcal{M}$. Since the sequence $(x^{n_l^1})_{l\in\mathbb{N}}$ also lies in $\mathcal{M}$, we find a converging subsequence $(x^{n_l^2})_{l\in\mathbb{N}}$ with limit $x \in \mathcal{M}$.

Also the sequence $(\pi^{n_l^2})_{l\in\mathbb{N}} \in (\hat{L}^2)^{\mathbb{N}}$ has a weakly converging subsequence: Indeed, we note that the sequence $(e^{-\frac{1}{2}\beta t}\pi^{n_l^2})_{l\in\mathbb{N}}$ lies in L^2 and is bounded by $e^{-\frac{1}{2}\beta t}\cdot 1$. Thus, by Theorem III.3.7 in Werner (2011), we find a weakly converging subsequence $(e^{-\frac{1}{2}\beta t}\pi^{n_l^3})_{l\in\mathbb{N}}$ with limit $e^{-\frac{1}{2}\beta t}\pi$. It remains to check whether the limit π is again a strategy, that is a Borel measurable function that satisfies $\sum_{a\in\mathcal{A}}\pi_{ia}(t) = 1$. For this we note that the set of all rescaled strategies is convex. It is furthermore closed because it is the preimage of the one-point-set containing the function $(i,t) \to e^{-\frac{1}{2}\beta t}\cdot 1$ under the map $f \mapsto ((i,t) \mapsto \sum_{a\in\mathcal{A}} f_{ia}(t))$. Thus, by Theorem III.3.8 in Werner (2011) the limit is again a strategy of the desired form.

As a last step we verify that the tuple (m, x, π) is feasible. For this we integrate (4.2) and obtain for all $l \in \mathbb{N}$

$$x_j^{n_l^3}(t) = \int_0^t \sum_{i\in\mathcal{S}}\sum_{a\in\mathcal{A}} x_i^{n_l^3}(s) Q_{ija}(m^{n_l^3}(s))\pi_{ia}^{n_l^3}(s)ds.$$

Moreover, we have

$$\begin{aligned}
&\left|\int_0^t \sum_{i\in\mathcal{S}}\sum_{a\in\mathcal{A}} x_i(s)Q_{ija}(m(s))\pi_{ia}(s)ds - x_j(t)\right| \\
&\leq \left|\int_0^t \sum_{i\in\mathcal{S}}\sum_{a\in\mathcal{A}} x_i(s)Q_{ija}(m(s))\pi_{ia}(s)ds - \int_0^t \sum_{i\in\mathcal{S}}\sum_{a\in\mathcal{A}} x_i(s)Q_{ija}(m(s))\pi_{ia}^{n_l^3}(s)ds\right| \\
&\quad + \left|\int_0^t \sum_{i\in\mathcal{S}}\sum_{a\in\mathcal{A}} x_i(s)Q_{ija}(m(s))\pi_{ia}^{n_l^3}(s)ds - \int_0^t \sum_{i\in\mathcal{S}}\sum_{a\in\mathcal{A}} x_i^{n_l^3}(s)Q_{ija}(m^{n_l^3}(s))\pi_{ia}^{n_l^3}(s)ds\right| \\
&\quad + \left|\int_0^t \sum_{i\in\mathcal{S}}\sum_{a\in\mathcal{A}} x_i^{n_l^3}(s)Q_{ija}(m^{n_l^3}(s))\pi_{ia}^{n_l^3}(s)ds - x_j(t)\right| \\
&= \left|\int_0^t \sum_{i\in\mathcal{S}}\sum_{a\in\mathcal{A}} x_i(s)Q_{ija}(m(s))\pi_{ia}(s)ds - \int_0^t \sum_{i\in\mathcal{S}}\sum_{a\in\mathcal{A}} x_i(s)Q_{ija}(m(s))\pi_{ia}^{n_l^3}(s)ds\right| \\
&\quad + \int_0^t \sum_{i\in\mathcal{S}}\sum_{a\in\mathcal{A}} \left|x_i(s)Q_{ija}(m(s)) - x_i^{n_l^3}(s)Q_{ija}(m^{n_l^3}(s))\right| \pi_{ia}^{n_l^3}(s)ds \\
&\quad + \left|x_j^{n_l^3}(t) - x_j(t)\right|.
\end{aligned}$$

By construction of the sequence $(m^{n_l^3}, x^{n_l^3}, \pi^{n_l^3})_{l\in\mathbb{N}}$ every summand converges to zero: The first summand converges to zero since $\pi^{n_l^3}$ converges weakly to π in $\hat{L}^2$. The second summand converges to zero because, first, $x^{n_l^3}(\cdot)$ converges uniformly on $[0,t]$ to $x(\cdot)$

and $m^{n_l^3}(\cdot)$ converges uniformly on $[0,t]$ to $m(\cdot)$ and, second, $Q_{ija}(\cdot)$ is continuous. The third summand converges to zero again because $x^{n_l^3}(\cdot)$ converges uniformly on $[0,t]$ to $x(\cdot)$. Thus, we obtain that

$$x_j(t) = \int_0^t \sum_{i\in\mathcal{S}} \sum_{a\in\mathcal{A}} x_i(s) Q_{ija}(m(s)) \pi_{ia}(s) ds,$$

which proves the desired claim. □

Lemma 4.8. *For all $m \in \mathcal{M}$ the set $\phi(m)$ is non-empty. Furthermore, for all $x \in \phi(m)$ there exists a strategy $\pi^0 \in \Pi$ such that π^0 is feasible for (m,x) and this strategy achieves the maximum value in $Y(m,x)$.*

Proof. We first of all note that given $m \in \mathcal{M}$ there exists an $x \in \mathcal{M}$ together with a strategy $\pi \in \Pi$ that solve equation (4.2) and for which the cost integral is finite: Choosing $\pi_{ia}(t) = \delta_{a'}(a)$ for an arbitrary $a' \in \mathcal{A}$ ensures that $(B(t))_{ij} = \sum_{a\in\mathcal{A}} Q_{ija}(m(t))\pi_{ia}(t)$ is a matrix of Lipschitz continuous functions. Thus, for any initial condition $x(0) = x_0 \in \mathcal{P}(\mathcal{S})$, Theorem 10.VI in Walter (1998) guarantees that the system $\dot{x}(t) = B(t)x(t)$ admits a solution x, which is equivalent to the statement that (m, x, π) is a feasible tuple. Furthermore, we note that

$$\int_0^\infty \left(\sum_{i\in\mathcal{S}} \sum_{a\in\mathcal{A}} x_i(t) \pi_{ia}^0(t) r_{ia}(m(t)) e^{-\beta t} \right) \mathrm{d}t$$

is finite for all strategies $\pi^0 \in \Pi$ since $r_{ia}(m)$ is uniformly bounded for all $i \in \mathcal{S}$, $a \in \mathcal{A}, m \in \mathcal{P}(\mathcal{S})$ because $r_{ia}(\cdot)$ is a continuous function on the compact space $\mathcal{P}(\mathcal{S})$ and $\mathcal{S}$ as well as $\mathcal{A}$ are finite. This yields that the integrand is dominated by $Ce^{-\beta t}$ for some constant $C \in (0, \infty)$, which yields that the integral itself is finite. This in total yields that $Y(m,x)$ is finite.

By compactness of the set of all feasible tuples (m, x, π) it follows that the supremum in the definition of $Y(m,x)$ is indeed a maximum and that we find an π^0 such that $V_{m_0}(\pi^0, m) = \sup_{x\in\mathcal{M}} Y(m,x)$. □

Lemma 4.9. *For all $m \in \mathcal{M}$ the set $\phi(m)$ is convex.*

Proof. A straight-forward computation yields the following equivalent description of the problem: Let $x \in \mathcal{M}$ be arbitrary and replace the quantity $x_i(t)\pi_{ia}^0(t)$ by $z_{ia}(t)$ in (4.3), that is we can rewrite $Y(m,x)$ as

$$Y(m,x) = \max_z \int_0^\infty \left(\sum_{i\in\mathcal{S}} \sum_{a\in\mathcal{A}} z_{ia}(t) r_{ia}(t)(m(t)) e^{-\beta t} \right) \mathrm{d}t,$$

where z satisfies for all $t \geq 0$

$$\begin{cases} \sum_{a\in\mathcal{A}} z_{ja}(t) = x_j(t) & \forall j \in \mathcal{S} \\ z_{ja}(t) \geq 0 & \forall j \in \mathcal{S}, a \in \mathcal{A} \, . \\ \dot{x}_j(t) = \sum_{i\in\mathcal{S}} \sum_{a\in\mathcal{A}} z_{ia}(t) Q_{ija}(m(t)) & \forall j \in \mathcal{S} \end{cases}$$

Let $x_1, x_2 \in \phi(m)$ and let $x_3 = \lambda x_1 + (1-\lambda)x_2$ for some $\lambda \in (0,1)$. By the previous observation, we find z_1 and z_2 that minimize our modified problem for x_1 and x_2, respectively, while satisfying the constraints. Since the constraints as well as the objective function are linear in z and x, it is immediate that also $z_3 = \lambda z_1 + (1-\lambda)z_2$ satisfies the same constraints and minimizes the objective functional. Thus, we obtain $x_3 \in \phi(m)$. □

Lemma 4.10. *The graph of $m \mapsto \phi(m)$ is closed, that is for any sequences $(m^n)_{n\in\mathbb{N}} \in \mathcal{M}^{\mathbb{N}}$ and $x^n \in \phi(m^n)$, $n \in \mathbb{N}$, such that $\lim_{n\to\infty} m^n = m$ and $\lim_{n\to\infty} x^n = x$, we have $x \in \phi(m)$.*

Proof. Define $\tilde{Y} : \mathcal{M} \times \mathcal{M} \times \hat{L}^2 \to \mathcal{M} \times \mathbb{R}$ via

$$(m, x, \pi) \mapsto \left(m, \int_0^\infty \sum_{i\in\mathcal{S}} \sum_{a\in\mathcal{A}} x_i(t)\pi_{ia}(t) r_{ia}(m(t)) e^{-\beta t} \mathrm{d}t \right).$$

We note that the functionals $f \mapsto \int f g d\mu$ with $g \in L^2$ are exactly those functionals that are continuous with respect to the weak topology (Corollary VIII.3.4 in Werner (2011) and Lemma VII.3.1 in Elstrodt (2011)). This and the observation that

$$(m, x) \mapsto \left(t \mapsto x_i(t) r_{ia}(m(t) e^{-\frac{1}{2}\beta t} \right)$$

is a continuous function that maps the pair (m, x) to a function in L^2 yields that $\tilde{Y}(\cdot,\cdot,\cdot)$ is continuous. Furthermore, by Lemma 4.7 and Lemma A.8, the function

$$m \mapsto V(m) := \max_{(m,x,\pi)\text{feasible}} \int_0^\infty \sum_{i\in\mathcal{S}} \sum_{a\in\mathcal{A}} x_i(t)\pi_{ia}(t) r_{ia}(m(t)) e^{-\beta t} \mathrm{d}t$$

is continuous. Thus, we obtain that $\Gamma_V := \{(m, v) \in \mathcal{M} \times \mathbb{R} : V(m) = v\}$ is closed. This implies that $\tilde{Y}^{-1}(\Gamma_V)$ is closed. It furthermore holds that

$$\begin{aligned} \Gamma_\phi &= \{(m, x) \in \mathcal{M} \times \mathcal{M} : x \in \phi(m)\} \\ &= \{(m, x) \in \mathcal{M} \times \mathcal{M} : \exists \pi \in \hat{L}^2 : (m, x, \pi) \in G\}, \end{aligned}$$

with

$$G = \tilde{Y}^{-1}(\Gamma_V) \cap \{(m, x, \pi) \in \mathcal{M} \times \mathcal{M} \times \hat{L}^2 : (m, x, \pi) \text{ is feasible}\}.$$

Using this we obtain that Γ_ϕ is closed: Indeed, let $(m^n, x^n)_{n\in\mathbb{N}} \in (\Gamma_\phi)^{\mathbb{N}}$ be a converging sequence with limit (m, x). Then, by Lemma 4.8, we find strategies $\pi^n \in \Pi$ ($n \in \mathbb{N}$) such that $(m^n, x^n, \pi^n) \in G$. As in the proof of Lemma 4.7, Lemma VIII.6.2 in Werner (2011), the compactness of G and Lemma A.7 imply that the sequence $(m^n, x^n, \pi^n)_{n\in\mathbb{N}}$ has a converging subsequence whose limit is (m, x, π) for some $\pi \in \Pi$ and which moreover lies in G. Thus, also (m, x) lies in Γ_ϕ. □

Lemma 4.11. *For all $m \in \mathcal{M}$ the set $\phi(m)$ is compact.*

Proof. Since $\mathcal{M}$ is a compact space, it is sufficient to prove that $\phi(m)$ is closed in order to show that it is compact. Since $\mathcal{M}$ is furthermore a metric space, it is sufficient to show that every limit of a sequence in $\phi(m)$ again lies in $\phi(m)$. This follows from Lemma 4.10 if we consider the sequence $m^n = m$ for all $n \in \mathbb{N}$. □

In order to prove upper semicontinuity we use Theorem 231.2 of Heuser (2008):

Lemma 4.12. *Let X, Y be metric spaces. The compact-valued set-valued map $F : X \to 2^Y$ is upper semicontinuous if and only if for every convergent sequence $(x^n)_{n\in\mathbb{N}}$ with limit x and each sequence $(y^n)_{n\in\mathbb{N}}$ satisfying $y^n \in F(x^n)$ for all $n \in \mathbb{N}$ there exists a convergent subsequence of $(y^n)_{n\in\mathbb{N}}$ whose limit lies in $F(x)$.*

Lemma 4.13. *The set-valued map $m \mapsto \phi(m)$ is upper semicontinuous.*

Proof. We show that the sequence characterization for upper semicontinuity presented in Lemma 4.12 is satisfied. For this we first note that by Lemma 4.11 our set-valued map is compact-valued. Let $(m^n)_{n\in\mathbb{N}}$ be a sequence with limit m and choose $(x^n)_{n\in\mathbb{N}}$ such that $x^n \in \phi(m^n)$ for all $n \in \mathbb{N}$. Since $\mathcal{M}$ is compact and a metric space, the sequence $(x^n)_{n\in\mathbb{N}}$ has a convergent subsequence $(x_{n_k})_{k\in\mathbb{N}}$ with limit x. Now we apply Lemma 4.10 and obtain that $x \in \phi(m)$. Thus, by Lemma 4.12 we obtain that $\phi(\cdot)$ is upper semicontinuous. □

Proof of Theorem 4.1. In order to prove the existence of a fixed point we have to check that we can indeed apply the fixed point theorem presented in Lemma 4.3:

For this we first have to check that our space $\mathcal{M}$ is indeed compact, convex and a subset of a locally convex space. The compactness has been proved in Lemma 4.6. Moreover, we note that $\mathcal{M}$ is convex since it is the set of all Lipschitz continuous functions with Lipschitz constant L and convex combinations of two such functions are again Lipschitz continuous with Lipschitz constant L. Since $\mathcal{M} \subseteq \mathcal{C}([0,\infty), \mathcal{P}(S))$ and $\mathcal{C}([0,\infty), \mathcal{P}(S))$ is locally convex given the topology of uniform convergence on compact sets (Werner (2011, pp.398-399)), also the third condition regarding $\mathcal{M}$ is satisfied.

As a second step we have to check that $\phi(\cdot)$ is a Kakutani-map: For this we note that by Lemma 4.13 our map $\phi(\cdot)$ is upper semicontinuous, by Lemma 4.8 the values of $\phi(\cdot)$ are non-empty, by Lemma 4.9 the values are convex and by Lemma 4.11 the values are compact. Thus, we indeed satisfy the conditions of Lemma 4.3 and find a fixed point of the map $\phi(\cdot)$. This fixed point is indeed a mean field equilibrium: By definition of $\phi(\cdot)$ and $Y(\cdot,\cdot)$ as well as Lemma 4.8 we have that (m, π) satisfies the dynamics equation (4.1) (since (m, m, π) is feasible) and that π is a maximizer of $V_{m_0}(\cdot, m)$. □

4.2. Existence of Stationary Equilibria

In Section 5.1 we prove (independent of this section) two results regarding the individual control problem that are crucial for the following proof. Namely, we show that

(a) For any $m \in \mathcal{P}(\mathcal{S})$ there is a deterministic stationary strategy that maximizes $V_{x_0}(\cdot, m)$ for all $x_0 \in \mathcal{P}(\mathcal{S})$ among all (also time-dependent) strategies.

(b) A stationary strategy is optimal for $m \in \mathcal{P}(\mathcal{S})$ if and only if it is a convex combination from $\mathcal{D}(m)$, which is the set of all deterministic stationary strategies that are optimal.

Using this, we will prove that whenever the Assumption A1 holds there exists a stationary mean field equilibrium. For this we will adapt the idea used in the dynamic existence proof, namely to prove the existence of a fixed point of an associated best response map in the dynamics. The map we consider here maps to each point $m \in \mathcal{P}(\mathcal{S})$ all stationary points of $Q^\pi(m)$ given that $\pi \in \Pi^s$ is an optimal strategy for m. In contrast to the proof of the existence of dynamic equilibria presented in Doncel et al. (2016a) we do not only rely on standard calculus arguments, instead the proof crucially relies on our probabilistic representation of the problem. In particular, the insights regarding the structure of optimal strategies discussed in Section 5.1 rely on this representation.

We define the best response map $\phi : \mathcal{P}(\mathcal{S}) \to 2^{\mathcal{P}(\mathcal{S})}$ by setting

$$\phi(m) := \{x \in \mathcal{P}(\mathcal{S}) | \exists \pi \in \text{conv}(\mathcal{D}(m)) : 0 = x^T Q^\pi(m)\}.$$

We will show that this map has a fixed point and that each fixed point of this map induces a stationary mean field equilibrium. More precisely, we prove:

Theorem 4.14. *Any continuous time mean field game with discounted costs that satisfies Assumption A1 has a stationary mean field equilibrium.*

As in the case of dynamic mean field equilibria the main tool is Kakutani's fixed point theorem, which can be found in Heuser (2008, Theorem 232.1):

Lemma 4.15 (Kakutani's fixed point theorem)**.** *Let C be a non-empty, convex and compact subset of a normed space X and let the set-valued map $F : C \to 2^C$ be with non-empty, convex values and have a closed graph (that is for any sequence $(x^n)_{n\in\mathbb{N}} \in C^{\mathbb{N}}$ and $y^n \in F(x^n)$, $n \in \mathbb{N}$, such that $\lim_{n\to\infty} x^n = x$ and $\lim_{n\to\infty} y^n = y$, we have $y \in F(x)$). Then $F(\cdot)$ has at least one fixed point.*

Proof of Theorem 4.14. Since $\mathcal{P}(\mathcal{S})$ is convex, compact and a subset of the normed space $\mathbb{R}^n$, we satisfy all assumptions regarding the domain and range of the set-valued map. Therefore, it remains to check whether the map $\phi(\cdot)$ satisfies the desired properties.

We first note that $\phi(m)$ is *non-empty* for all $m \in \mathcal{P}(\mathcal{S})$. By fact (a) the set $\mathcal{D}(m)$ is non-empty. Since any continuous time Markov chain with finite state space has at least one stationary distribution, there exists an $x \in \mathcal{P}(\mathcal{S})$ such that $0 = x^T Q^{\pi}(m)$, which yields that $x \in \phi(m)$.

Furthermore, for each $m \in \mathcal{P}(\mathcal{S})$ the set $\phi(m)$ is *convex*: Let $x^1, x^2 \in \phi(m)$ be two distinct points. Then, by definition of $\phi(\cdot)$, we find two strategies $\pi^1, \pi^2 \in \text{conv}(\mathcal{D}(m))$ such that

$$0 = \sum_{i\in\mathcal{S}} \sum_{a\in\mathcal{A}} x_i^1 Q_{ija}(m) \pi_{ia}^1 \quad \text{and} \quad 0 = \sum_{i\in\mathcal{S}} \sum_{a\in\mathcal{A}} x_i^2 Q_{ija}(m) \pi_{ia}^2.$$

Define $z_{ia}^1 := x_i^1 \pi_{ia}^1$ and $z_{ia}^2 = x_i^2 \pi_{ia}^2$, which satisfy

$$0 = \sum_{i\in\mathcal{S}} \sum_{a\in\mathcal{A}} Q_{ija}(m) z_{ia}^1 \quad \text{and} \quad 0 = \sum_{i\in\mathcal{S}} \sum_{a\in\mathcal{A}} Q_{ija}(m) z_{ia}^2.$$

Now let $\alpha \in [0,1]$ be arbitrary and define $x^3 = \alpha x^1 + (1-\alpha)x^2$. Then $z^3 = \alpha z^1 + (1-\alpha) z^2$ satisfies

$$0 = \sum_{i\in\mathcal{S}} \sum_{a\in\mathcal{A}} Q_{ija}(m) z_{ia}^3,$$

which means that

$$\pi_{ia}^3 := \begin{cases} 0 & \text{if } x_i^3 = 0 \\ z_{ia}^3 / x_i^3 & \text{if } x_i^3 > 0 \end{cases}$$

satisfies $0 = (x^3)^T Q^{\pi^3}(m)$. It remains to verify that $\pi^3 \in \text{conv}(\mathcal{D}(m))$. For this we note that $\pi_{ia}^3 > 0$ if and only if $z_{ia}^3 > 0$, which in turn is equivalent to the requirement that $z_{ia}^1 > 0$ or $z_{ia}^2 > 0$. This can only happen if $\pi_{ia}^1 > 0$ or $\pi_{ia}^2 > 0$. Thus, since $\pi^1, \pi^2 \in \text{conv}(\mathcal{D}(m))$, also $\pi^3 \in \text{conv}(\mathcal{D}(m))$.

We now verify that ϕ has a *closed graph*, that is that for any sequence $(m^n)_{n\in\mathbb{N}} \in \mathcal{P}(\mathcal{S})^{\mathbb{N}}$ and $x^n \in \phi(m^n)$ for all $n \in \mathbb{N}$ with $\lim_{n\to\infty} m^n = m$ and $\lim_{n\to\infty} x^n = x$ we indeed have $x \in \phi(m)$: Let $(m^n, x^n)_{n\in\mathbb{N}}$ be a converging sequence satisfying $x^n \in \phi(m^n)$ for all $n \in \mathbb{N}$ and denote its limit by (m, x). By definition of $\phi(\cdot)$, we find a sequence $\pi^n \in \text{conv}(\mathcal{D}(m^n))$ such that $0 = x^n Q^{\pi^n}(m^n)$. By compactness of Π^s, we find a converging subsequence $(\pi^{n_k})_{k\in\mathbb{N}}$ with limit $\pi \in \Pi^s$.

For any $k \in \mathbb{N}$ let $A_1^k \times \ldots \times A_S^k \subseteq \mathcal{A}^S$ be such that $\pi_{ia}^{n_k} > 0$ for all $i \in \mathcal{S}, a \in A_i^k$ and $\pi_{ia} = 0$ for all $i \in \mathcal{S}, a \notin A_i^k$. Since $\mathcal{A}^S$ is finite, we find a set $A_1 \times \ldots \times A_S$ that occurs infinitely often in the sequence $(A_1^k \times \ldots \times A_S^k)_{k\in\mathbb{N}}$. From this we obtain that $\pi_{ia} = 0$ for all $i \in \mathcal{S}$ and $a \notin A_i$. Moreover, since $\pi^n \in \text{conv}(\mathcal{D}(m^n))$ we obtain that for all $l \in \mathbb{N}$ such that $A_1 \times \ldots \times A_S = A_1^l \times \ldots \times A_S^l$ we have $V^d(m^{n_l}) = V^*(m^{n_l})$ for all deterministic strategies satisfying $d(i) \in A_i$ for all $i \in \mathcal{S}$. By continuity of $V^d(\cdot)$ and $V^*(\cdot)$ (see Lemma 5.3), we obtain $V^d(m) = V^*(m)$. Thus, $\pi \in \text{conv}(\mathcal{D}(m))$.

Furthermore, by continuity of $Q(\cdot)$, we obtain that

$$0 = \sum_{i \in \mathcal{A}} \sum_{a \in \mathcal{A}} x_i^{n_k} Q_{ija}(m^{n_k}) \pi_{ia}^{n_k} \leftarrow \sum_{i \in \mathcal{A}} \sum_{a \in \mathcal{A}} x_i Q_{ija}(m) \pi_{ia},$$

which shows that $x \in \phi(m)$.

Now Kakutani's fixed point theorem (Lemma 4.15) yields a fixed point $m \in \phi(m)$, which induces a stationary mean field equilibrium. Indeed, for any fixed point m we find a strategy $\pi \in \mathrm{conv}(\mathcal{D}(m))$ such that $0 = m^T Q^\pi(m)$. Since, by the facts (a) and (b), we moreover have that $V_{x_0}(\pi, m) \geq V_{x_0}(\pi', m)$ for all $x_0 \in \mathcal{P}(\mathcal{S})$ and all $\pi' \in \Pi$, the pair (m, π) constitutes a stationary mean field equilibrium. □

5. A Toolbox for Finding Stationary Mean Field Equilibria

This chapter presents several tools for finding stationary mean field equilibria. The task of finding mean field equilibria consists, as in the classical theory for diffusion-based models, of an individual control problem (Given a fixed population distribution which strategies are optimal?) and a fixed point problem (Is m a stationary point of $Q^{\pi}(\cdot)$ with π being optimal for m?). In our model formulation the central problem is that a priori there are infinitely many different strategies that might be optimal and, even worse, that each yields to a fixed point problem that is in general not related to the others. For this reason in the applications considered in the literature (for example Kolokoltsov and Bensoussan (2016) and Kolokoltsov and Malafeyev (2017)) the authors restricted themselves to equilibria in deterministic strategies.

In this chapter we will present several tools to reduce the complexity of the two problems, which in turn allows also to compute equilibria in mixed strategies: First, we show that a stationary strategy is optimal if and only if it is a convex combination of optimal deterministic stationary strategies. These optimal deterministic stationary strategies can moreover be determined as those maximizing the right-hand side the optimality equation of the associated Markov decision process. Second, we derive a simple, yet powerful reformulation of the equations describing the fixed point, the so-called cut criterion, which is based on the observation that in many models agents cannot influence the transitions between two subsets of states.

Furthermore, we will derive under the assumption that $Q^d(m)$ is irreducible for all deterministic strategies $d \in D^s$ and all population distributions $m \in \mathcal{P}(\mathcal{S})$ an explicit characterization of the stationary point given $Q^{\pi}(m)$. Using this we obtain the following programme, consisting of three steps, to compute all stationary mean field equilibria: First we have to find the optimal deterministic stationary strategies for each population distribution $m \in \mathcal{P}(\mathcal{S})$, then we have to find all fixed points of a suitably defined set-valued map and, finally, we have to solve for each fixed point a linear equation. Additionally, we provide, under the irreducibility assumption, a criterion to verify when there is a unique stationary point given $Q^{\pi}(\cdot)$.

In Section 5.1 we investigate the individual control problem and in Section 5.2 we derive the cut criterion. In Section 5.3 we derive that under an irreducibility assumption there is a set-valued map whose fixed points yield all stationary mean field equilibria. In Section 5.4 we derive a sufficient criterion for the uniqueness of the stationary point given $Q^{\pi}(\cdot)$. Section 5.5 concludes the chapter with a short discussion why we cannot expect sensible uniqueness criteria in our setting.

5.1. The Individual Optimal Control Problem

In this section, we show that the individual's control problem is equivalent to a continuous time Markov decision process with the expected discounted reward criterion. Using the equivalence, we then characterize the set of all optimal strategies as the convex hull of all optimal deterministic strategies, which are given as those maximizing the right-hand side of the optimality equation of the associated Markov decision process. This insights, which were announced in Section 4.2, conclude the proof that a stationary mean field equilibrium exists whenever Assumption A1 holds. We will conclude the section with a formal argument that in all non-trivial games there will be points for which more than one deterministic stationary strategy, and thus infinitely many stationary (mixed) strategies, are optimal. This fact means that in all non-trivial games we have to consider a priori infinitely many fixed point problems.

The following lemma states that the individual's optimization problem is equivalent to a continuous time Markov decision process with expected discounted reward criterion.

Lemma 5.1. *Let $m \in \mathcal{P}(\mathcal{S})$ be a population distribution. A Markovian randomized strategy $\pi^0 \in \Pi$ is optimal in our model given m for all initial conditions $x^0 \in \mathcal{P}(\mathcal{S})$, i.e. achieves the maximum value of $V_{x_0}(\cdot, m)$, if and only if it is discounted reward optimal for the continuous time Markov decision process with rates $Q_{ija}(m)$ and rewards $r_{ia}(m)$.*

Proof. We remind ourselves that Assumption A1 ensures that $r_{ia}(\cdot)$ is uniformly bounded as it is a continuous function on a compact space. Therefore, the value function is finite for every population distribution function, every individual strategy and every initial distribution. Thus, we can rewrite the value function using the representation $x_i^{\pi^0}(t) = \sum_{k \in \mathcal{S}} x_k^0 \cdot p^{\pi^0}(0, k, t, i)$ as well as the definition (2.3) of the value function of a continuous time Markov decision process to obtain:

$$\begin{aligned} V_{x^0}(\pi^0, m) &= \int_0^\infty \sum_{i \in \mathcal{S}} \sum_{a \in \mathcal{A}} x_i^{\pi^0}(t) r_{ia}(m) \pi_{ia}^0(t) e^{-\beta t} \mathrm{d}t \\ &= \int_0^\infty \sum_{i \in \mathcal{S}} \sum_{a \in \mathcal{A}} \sum_{k \in \mathcal{S}} x_k^0 p^{\pi^0}(0, k, t, i) r_{ia}(m) \pi_{ia}^0(t) e^{-\beta t} \mathrm{d}t \end{aligned}$$

$$
\begin{aligned}
&= \sum_{k\in\mathcal{S}} x_k^0 \int_0^\infty e^{-\beta t} \sum_{i\in\mathcal{S}} \sum_{a\in\mathcal{A}} r_{ia}(m)\pi_{ia}^0(t) p^{\pi^0}(0,k,t,i)\mathrm{d}t \\
&= \sum_{k\in\mathcal{S}} x_k^0 V_k^{\pi^0}(m).
\end{aligned}
$$

Let $\pi^0 \in \Pi$ be a strategy that maximizes $V_{x^0}(\cdot, m)$ for all $x^0 \in \mathcal{P}(\mathcal{S})$ and assume that it is not discounted reward optimal for the MDP. Then, since there is at least one discounted reward optimal strategy, there exists a strategy π for which $V_i^\pi(m) \geq V_i^{\pi^0}(m)$ for all $i \in \mathcal{S}$, where the inequality is strict for at least one i. Thus, for any initial distribution x^0 such that $x_k^0 > 0$ for all $k \in \mathcal{S}$ we would arrive at

$$
V_{x^0}(\pi, m) = \sum_{k\in\mathcal{S}} x_k^0 V_k^\pi(m) > \sum_{k\in\mathcal{S}} x_k^0 V_k^{\pi^0}(m) = V_{x^0}(\pi^0, m),
$$

which is contradiction to the fact that π^0 maximizes $V_{x^0}(\cdot, m)$.

Similarly if we are given a strategy π^0 that is discounted reward optimal for the MDP and assume that it does not maximize $V_{x^0}(\cdot, m)$ for some $x^0 \in \mathcal{P}(\mathcal{S})$, then there exists a strategy π, which achieves a larger value, which means that

$$
V_{x^0}(\pi, m) = \sum_{k\in\mathcal{S}} x_k^0 V_k^\pi(m) > \sum_{k\in\mathcal{S}} x_k^0 V_k^{\pi^0}(m) = V_{x^0}(\pi^0, m).
$$

However, at the same time, as π^0 is discounted reward optimal, we have $V_k^\pi(m) \leq V_k^{\pi_0}(m)$ for all $k \in \mathcal{S}$, which is again a contradiction. □

This result together with Corollary 2.9 immediately allows us to characterize the set of all stationary strategies as the convex hull of all deterministic optimal strategies:

Theorem 5.2. *Let $m \in \mathcal{P}(\mathcal{S})$ be a population distribution. Furthermore, define*

$$
O_i(m) = argmax_{a\in\mathcal{A}} \left\{ r_{ia}(m) + \sum_{j\in\mathcal{S}} Q_{ija}(m) V_j^*(m) \right\}
$$

with $V_j^*(m)$ being the solution of the optimality equation for the Markov decision process with transition rates $Q_{ija}(m)$ and rewards $r_{ia}(m)$. Then the set of all optimal deterministic stationary strategies is given by

$$
\mathcal{D}(m) := \{d : \mathcal{S} \to \mathcal{A} | d(i) \in O_i(m)\}
$$

and a randomized stationary strategy is optimal if and only if it is a convex combination of strategies from $\mathcal{D}(m)$. In particular, there is an optimal deterministic stationary strategy.

This theorem yields that we can partition $\mathcal{P}(\mathcal{S})$ into finitely many sets $\text{Opt}(A_1 \times \ldots \times A_S)$, which consist of all those $m \in \mathcal{P}(\mathcal{S})$ for which the set of all optimal deterministic strategies is $\{d \in D^s : d(i) \in A_i \text{ for all } i \in \mathcal{S}\}$ and then consider the fixed point problem for each of these sets separately.

The sets $\text{Opt}(A_1 \times \ldots \times A_S)$ are computed as follows: First determine the unique value function $V(m)$ by solving the optimality equation

$$\beta V_i(m) = \max_{a \in \mathcal{A}} \left\{ r_{ia}(m) + \sum_{j \in \mathcal{S}} Q_{ija}(m) V_j(m) \right\}, \quad i \in \mathcal{S} \tag{5.1}$$

for each m. Then determine the optimal actions, which are those actions that achieve the maximum on the right hand side of the equation.

To address the question why games without a unique optimal strategy for each $m \in \mathcal{P}(\mathcal{S})$ yield to infinitely many fixed point problems and for future qualitative considerations the following lemma is of principal significance. However, the result is in general not sensible for computing the optimality sets since the computational effort is often much higher than in the approach using the optimality equation.

Lemma 5.3. *The value function of the associated Markov decision problem $V : \mathcal{P}(\mathcal{S}) \to \mathbb{R}^S$ can be represented as follows*

$$V(m) = \max_{d \in D^s} V^d(m)$$

with $V^d(m)$ being the discounted reward given the deterministic stationary strategy d. Furthermore, for each $d \in D^s$ the matrix $(\beta I - Q^d(m))$ is invertible and

$$V^d(m) = (\beta I - Q^d(m))^{-1} r^d(m),$$

with $Q^d(m) = (Q_{ijd(i)}(m))_{i,j \in \mathcal{S}}$ and $r^d(m) = (r_{id(i)})_{i \in \mathcal{S}}$. Moreover, the value function $V(\cdot)$ is continuous.

Proof. The first part of the statement directly follows from the existence of a deterministic stationary optimal strategy for Markov decision processes (Lemma 2.1).

To compute the discounted reward given a deterministic strategy d one can rely on Lemma 2.2, which states that the discounted reward is the unique solution of

$$\beta g = r^d(m) + Q^d(m) g \quad \text{i.e.} \quad (\beta I - Q^d(m)) g = r^d(m).$$

Thus, in order to prove the statement we have to show that $\beta I - Q^d$ is indeed invertible. For this we again rely on uniformization (Remark 2.3): Namely, applying $Q^d(m) = ||Q(m)||(\bar{P}^d(m) - I)$ we obtain

$$\begin{aligned}\beta I - Q^d(m) &= \beta I - ||Q(m)||(\bar{P}^d(m) - I)\\ &= (\beta + ||Q(m)||)I - ||Q(m)||\bar{P}^d(m)\\ &= (\beta + ||Q(m)||)\left(I - \frac{||Q(m)||}{\beta + ||Q(m)||}\bar{P}^d(m)\right).\end{aligned}$$

Noting that the spectral radius of $\frac{||Q(m)||}{\beta+||Q(m)||}\bar{P}^d(m)$ is less than 1 (Asmussen, 2003, Proposition I.6.2), Corollary C.4 in Puterman (1994) yields that the matrix is invertible. This proves the representation for $V^d(m)$.

It remains to show that V is indeed continuous: For this note that $Q^d(\cdot)$ and $r^d(\cdot)$ are continuous functions in m for all $d \in D^s$. Since the matrix inversion is continuous, it follows that $V^d(\cdot)$ is continuous. Since the maximum over finitely many continuous functions is continuous, we obtain the desired result. □

Remark 5.4. This lemma allows to show that a game that is not trivial in the sense that there are two different population distributions such that different strategies are optimal for each of them, has a closed, non-empty set of points where infinitely many strategies are optimal. Thus, we indeed have to consider infinitely many (potentially different) fixed point problems in order to compute all stationary mean field equilibria.

Indeed, the set of all $m \in \mathcal{P}(\mathcal{S})$ for which $d \in D^s$ is one (but possibly not the only) optimal stationary strategy are closed: Since $V^d(\cdot)$ and $V(\cdot)$ are continuous, also the map $\sum_{i\in\mathcal{S}}\left(V_i^d(m) - V_i(m)\right)$ is continuous. The pre-image of $[0, \infty)$ yields all points for which d, possibly together with some other strategies, is optimal. By continuity, this set is closed. Now let $d^1, d^2 \in D^s$ be two distinct strategies such that there are two points $m^1, m^2 \in \mathcal{P}(\mathcal{S})$ such that d^1 is optimal for m^1 and d^2 is optimal for m^2. Since $\mathcal{P}(\mathcal{S})$ is connected, we obtain that there is a closed non-empty set for which several strategies are optimal simultaneously. △

5.2. A Cut Criterion to Find Fixed Points of the Dynamics

To solve the fixed point problem one has to solve the nonlinear equation $m^T Q^\pi(m) = 0$ for all strategies π that are optimal for some population distribution and then one has to check whether π is indeed optimal for these solutions. In many settings (for example in the corruption model introduced in Section 3.1 and analysed in Section 6.2) this task is not simple. However, often a cut criterion similar to the one used for Markov chains is useful, although it is just a reformulation of the balance equation $m^T Q^\pi(m) = 0$ (see Kelly (1979, Lemma 1.4) and Hübner (2009, Theorem 7.3) for a description of the criterion for standard continuous time Markov chains). The criterion states that if we partition the state space of the Markov chain into two sets, then the probability flow from one set to the other has to be equal the probability flow from this other set to the first. The particular use of the criterion is that in most models that have been considered so far there has always been a set of states for which the dynamics to and from this set is independent of the chosen strategy. This means that any mean field equilibrium irrespective of the chosen strategy has to satisfy certain equations coming from the cut criterion, which could be obtained from the standard balance equations $m^T Q^\pi(m) = 0$ only by sensible rearrangements. In Chapter 6 we will show that the criterion indeed simplifies the search for fixed points in the examples.

Theorem 5.5. *Let π be a stationary strategy and let $\mathcal{T} \subseteq \mathcal{S}$. Then any stationary population distribution satisfies*

$$\sum_{j \in \mathcal{T}} \sum_{i \in \mathcal{S} \setminus \mathcal{T}} m_i Q^\pi_{ij}(m) = \sum_{j \in \mathcal{T}} \sum_{i \in \mathcal{S} \setminus \mathcal{T}} m_j Q^\pi_{ji}(m).$$

Proof. The stationarity condition reads for all $j \in \mathcal{S}$

$$0 = \sum_{i \in \mathcal{S}} m_i Q^\pi_{ij}(m).$$

Furthermore, since $Q^\pi(m)$ is conservative, we have for all $j \in \mathcal{S}$

$$m_j \sum_{i \in \mathcal{S}} Q^\pi_{ji}(m) = 0.$$

This yields for all $j \in \mathcal{S}$

$$\sum_{i \in \mathcal{S}} m_i Q^\pi_{ij}(m) = m_j \sum_{i \in \mathcal{S}} Q^\pi_{ji}(m).$$

Summing this identity over all $j \in \mathcal{T}$ yields

$$\sum_{j \in \mathcal{T}} \sum_{i \in \mathcal{S}} m_i Q^\pi_{ij}(m) = \sum_{j \in \mathcal{T}} \sum_{i \in \mathcal{S}} m_j Q^\pi_{ji}(m).$$

Subtracting the identity

$$\sum_{j \in \mathcal{T}} \sum_{i \in \mathcal{T}} m_i Q^\pi_{ij}(m) = \sum_{j \in \mathcal{T}} \sum_{i \in \mathcal{T}} m_j Q^\pi_{ji}(m)$$

yields the desired result. □

5.3. A Characterization of $\phi(\cdot)$ Purely Relying on Deterministic Strategies

This section is devoted to proving the main theorem of the chapter. This theorem states that under an irreducibility assumption the set $\phi(m)$, which has been introduced in Section 4.2, has an explicit characterization. More precisely, we obtain that $\phi(m)$ is the set of all convex combinations of $x^d(m)$, which denotes the unique stationary point given $Q^d(m)$, for all those deterministic stationary strategies d that are optimal for m. This result yields to the following programme to obtain all mean field equilibria:

1. Determine for all $m \in \mathcal{P}(\mathcal{S})$ the set of optimal deterministic stationary strategies $\mathcal{D}(m)$.

2. Determine all fixed points of $m \mapsto \phi(m) := \{x^d(m) : d \in \mathcal{D}(m)\}$.

3. Solve for each fixed point m of $\phi(\cdot)$ the linear equation $0 = Q^\pi(m)m$ for solutions $\pi \in \Pi^s$ such that π is a convex combination of $\mathcal{D}(m)$. Any solution of this system yields a stationary mean field equilibrium (m, π).

For the rest of the section we assume irreducibility of $Q^\pi(m)$ for all strategies $\pi \in \Pi^s$ as this yields to a unique stationary point given $Q^\pi(m)$ (see Asmussen (2003, Section II.4)) and, moreover, to an explicit representation of it. We note that it is sufficient to verify irreducibility for all deterministic strategies $d \in D^s$ since any stationary strategy π is a convex combination of deterministic strategies and thus $Q^\pi(m) = \sum_{d \in D^s} \lambda_d Q^d(m)$ is also irreducible. With this observation we formulate the main theorem, which is proven in the rest of the section:

Theorem 5.6. *Let $m \in \mathcal{P}(\mathcal{S})$ such that $Q^d(m)$ is irreducible for all $d \in D^s$. Furthermore, let $\{d^1, \dots, d^n\}$ be the set of all optimal strategies for m. Then*

$$\phi(m) = conv(x^{d^1}(m), \dots, x^{d^n}(m)),$$

with $x^{d^k}(m)$ being the unique solution x of $0 = \sum_{i \in \mathcal{S}} \sum_{a \in \mathcal{A}} x_i Q_{ija}(m) d^k_{ia}$.

The proof of this theorem relies on the idea to characterize the stationary distribution $x^\pi(m)$ of the CTMC with irreducible generator $Q^\pi(m)$ by a closed form expression and to show thereafter that $x^\pi(m)$ is a convex combination of $(x^d(m))_{d \in D^s}$. In order to follow this programme, we have to investigate several properties of the generator matrix Q. We start with the following lemma regarding the structural properties of an irreducible, conservative generator $Q \in \mathbb{R}^{S \times S}$, more precisely, regarding the minor Q'_{SS}, which arises from Q by deleting the last row and column:

Lemma 5.7. *Let $Q \in \mathbb{R}^{S \times S}$ be an irreducible, conservative generator matrix. Then all eigenvalues of the minor Q'_{SS} have negative real part. Consequently Q'_{SS} has full rank and*

$$sign(\det(Q'_{SS})) = (-1)^{S+1}.$$

For the proof we rely on the results presented in Berman and Plemmons (1979, Chapter 6) on M-matrices, which are matrices of the form $A = sI - B$ with B being a non-negative matrix (that is $B_{ij} \geq 0$ for all $i, j \in \mathcal{S}$) and s being greater or equal the spectral radius of B. Given that a matrix A is from the set

$$Z^{S \times S} = \{A \in \mathbb{R}^{S \times S} : A_{ij} \leq 0 \quad \text{for all } i \neq j\}$$

the authors provide 50 criteria that are equivalent to A being a non-singular M-matrix (Berman and Plemmons, 1979, Theorem 6.2.3). For our purpose the following two are relevant:

I_{28} There exists a vector $y \geq 0$ with $y_i > 0$ for at least one $i \in \mathcal{S}$ such that $Ay > 0$, where the inequality signs hold pointwise.

G_{20} A is positive stable, that is, the real part of each eigenvalue of A is positive.

The technical proof of Lemma 5.7, which can be found in Appendix A.4, basically consists of the following arguments: We first note that $-Q'_{SS}$ is indeed a matrix fitting in the framework of Berman and Plemmons (1979), then we propose a method to find a vector y such that I_{28} is satisfied, from this we conclude that G_{20} is satisfied and thereafter deduce the desired statement.

The previous lemma has a simple consequence, which follows from the observation that the columns of the generator are linearly dependent (the last column is the negative sum of the other columns):

Corollary 5.8. *Let $Q \in \mathbb{R}^{S \times S}$ be an irreducible, conservative generator matrix. Then the rank of Q is $S - 1$.*

Using these results we can explicitly characterize the stationary distribution given a stationary strategy π and a population distribution m:

Lemma 5.9. *Let $\pi \in \Pi^s$ be a stationary strategy and let $m \in \mathcal{P}(\mathcal{S})$ such that $Q^\pi(m)$ is irreducible. Let $\tilde{Q}^\pi(m)$ be the matrix $(Q^\pi(m))^T$ where the last row is replaced by $(1, \dots, 1)$. Then $\tilde{Q}^\pi(m)$ is invertible and we have that the unique stationary distribution $x^\pi(m)$ is given by*

$$x^\pi(m) = (\tilde{Q}^\pi(m))^{-1} \cdot (0, \dots, 0, 1)^T. \tag{5.2}$$

Proof. The unique stationary distribution of our process is uniquely determined by

$$0 = \sum_{i \in \mathcal{S}} x_i^\pi(m) Q_{ij}^\pi(m) \quad \text{for all } j \in \mathcal{S} \quad \text{and} \quad \sum_{i \in \mathcal{S}} x_i^\pi(m) = 1 \tag{5.3}$$

(Asmussen, 2003, Theorem II.4.2). Since the last equation of the system $0 = x^\pi(m)^T Q^\pi(m)$ is the negative sum of all other equations, we can replace the last column of Q by the one vector and obtain that the system (5.3) is equivalent to

$$\begin{pmatrix} 0 \\ \vdots \\ 0 \\ 1 \end{pmatrix} = \begin{pmatrix} Q_{11}^\pi(m) & \dots & Q_{1,S-1}^\pi(m) & 1 \\ \vdots & & \vdots & \vdots \\ Q_{S,1}^\pi(m) & & Q_{S,S-1}^\pi(m) & 1 \end{pmatrix}^T x^\pi(m),$$

which by definition is $(0, \dots, 0, 1)^T = \tilde{Q}^\pi(m) x^\pi(m)$.

We now show that the rank of the matrix $\tilde{Q}^\pi(m)$ is S because in this case we can invert the matrix: As in Resnick (1992, p.137-139) we rely on the existence of the stationary distribution given $Q^\pi(m)$. In order to show that $\tilde{Q}^\pi(m)$ has full rank, we show that $y^T \tilde{Q}^\pi(m) = 0$ implies that $y = 0$. For the stationary distribution $x^\pi(m)$ given $Q^\pi(m)$ we have

$$0 = \left(y^T \tilde{Q}^\pi(m) \right) x^\pi(m) = y^T \left(\tilde{Q}^\pi(m) x^\pi(m) \right) = y^T (0, \dots, 1)^T = y_S.$$

Thus,

$$0 = y^T \tilde{Q}^\pi(m) = y^T \begin{pmatrix} Q_{11}^\pi(m) & \dots & Q_{S1}^\pi(m) \\ \vdots & \dots & \vdots \\ Q_{1,S-1}^\pi(m) & \dots & Q_{S,S-1}^\pi(m) \\ 1 & \dots & 1 \end{pmatrix}$$

implies that

$$0 = (y_1, \dots, y_{S-1}) (Q^\pi(m))'_{SS}.$$

Since by Lemma 5.7 the matrix $(Q^\pi(m))'_{SS}$ has full rank, we obtain that $y_1 = \dots = y_{S-1} = 0$, which proves that $\tilde{Q}^\pi(m)$ has full rank. □

Remark 5.10. We conclude from the previous lemma that given any stationary strategy $\pi \in \Pi^s$ the function $m \mapsto x^\pi(m)$ is continuous: Since by Cramer's rule the matrix inversion can be rewritten as a quotient of polynomials with non-vanishing denominator, the matrix inversion is a continous function. Thus, the function $x^\pi(m)$ is, as a concatenation of continuous functions, also continuous. △

From the representation in Lemma 5.9 we now deduce a formula which is simpler in terms of computational effort: It simply follows from Cramer's rule, which allows to compute the inverse of a matrix in terms of the determinants of the $(S-1) \times (S-1)$-minors. More precisely the inverse of a matrix $A = (a_{ij})_{i,j \in \mathcal{S}}$ is given by

$$(A^{-1})_{ij} = \frac{1}{\det(A)}(-1)^{i+j} \det(A'_{ji}),$$

with A'_{ji} being the matrix A where we deleted the j-th row and the i-th column. In our case we obtain:

Theorem 5.11. *Let $\pi \in \Pi^s$ be a stationary strategy and let $m \in \mathcal{P}(\mathcal{S})$ be a population distribution such that $Q^\pi(m)$ is irreducible. Let $\tilde{Q}^\pi(m)$ be the matrix defined in Lemma 5.9. Then*

$$x^\pi(m)_i = \frac{1}{\det(\tilde{Q}^\pi(m))}(-1)^{S+i} \det(Q^\pi(m))'_{iS},$$

with

$$\det(\tilde{Q}^\pi(m)) = \sum_{k=1}^{S}(-1)^{S+k} \det((Q^\pi(m))'_{kS}.$$

Proof. The statement follows from Cramer's rule together with the Laplace expansion of $\det(\tilde{Q}^\pi(m))$ along the last line

$$\det(\tilde{Q}^\pi(m)) = \sum_{k=1}^{S}(-1)^{S+k} \cdot 1 \cdot \det((\tilde{Q}^\pi(m))'_{Sk})$$

as well as the observation that

$$\det((\tilde{Q}^\pi(m))'_{Sk}) = \det(((Q^\pi(m))^T)'_{Sk}) = \det((Q^\pi(m))'_{kS}),$$

as $\tilde{Q}^\pi(m)$ differs from $Q^\pi(m)^T$ only in the S-th row. □

In order to establish the desired result on further characterizing the convex set $\phi(m)$ one final preparation has to be made, which is to show that the determinant of $\tilde{Q}^\pi(m)$ has uniform sign over all $\pi \in \Pi^s$: Write $d^{(a_1,\dots,a_S)}$ for the deterministic strategy satisfying $d(i) = a_i$ for all $i \in \mathcal{S}$. Then it holds that

$$\tilde{Q}^\pi(m) = \sum_{(a_1,\dots,a_S) \in \mathcal{A}^S} \left(\prod_{i=1}^{S} \pi_{ia_i} \right) \tilde{Q}^{d^{(a_1,\dots,a_S)}}(m).$$

Now a simple application of the intermediate value theorem yields the desired result.

Lemma 5.12. *Let $m \in \mathcal{P}(\mathcal{S})$ be a population distribution such that $Q^d(m)$ is irreducible for all $d \in D^s$. Then $\det(\tilde{Q}^\pi(m))$ has uniform sign over Π^s.*

Proof. We note that the map $\pi \mapsto \tilde{Q}^\pi(m)$, which ranges from Π^s to $\mathbb{R}^{S\times S}$, is continuous. Since the determinant is also a continuous function, we obtain that $\pi \mapsto \det(\tilde{Q}^\pi(m))$ is a continuous function. By Lemma 5.9 we have that $\det(\tilde{Q}^\pi(m)) \neq 0$ for all $\pi \in \Pi^s$. If there would be a strategy π_1 and a strategy π_2 such that $\det(\tilde{Q}^{\pi_1}(m)) < 0$ and $\det(\tilde{Q}^{\pi_2}(m)) > 0$, then, by the intermediate value theorem, there would be a $\pi \in \Pi^s$ such that $\det(\tilde{Q}^\pi(m)) = 0$, which would be a contradiction. □

With all these preparations we prove the characterization result stated in the beginning of the section:

Proof of Theorem 5.6. For readability we suppress the dependence of Q and x on m.

Let $\pi = \sum_{(a_1,\dots,a_n)\in\mathcal{A}^S} \lambda_{(a_1,\dots,a_n)} d^{(a_1,\dots,a_n)}$. We note that by Theorem 5.2 the factors $\lambda_{(a_1,\dots,a_S)}$ are zero for all non-optimal strategies. With $\hat{Q}$ being the matrix Q^T without the last row, we write $\tilde{Q}^\pi$ as follows:

$$\tilde{Q}^\pi = \begin{pmatrix} \sum_{a_1\in\mathcal{A}} \pi_{1a_1}\hat{Q}_{1,\cdot,a_1} & \dots & \sum_{a_S\in\mathcal{A}} \pi_{Sa_S}\hat{Q}_{S,\cdot,a_S} \\ 1 & \dots & 1 \end{pmatrix}.$$

As the determinant is linear in columns and we have $\sum_{a_i\in\mathcal{A}} \pi_{ia_i} = 1$ for all $i \in \mathcal{S}$ we obtain

$$\begin{aligned}
\det(\tilde{Q}^\pi) &= \sum_{a_1\in\mathcal{A}} \pi_{1a_1} \det \begin{pmatrix} \hat{Q}_{1,\cdot,a_1} & \sum_{a_2\in\mathcal{A}} \pi_{2a_2}\hat{Q}_{2,\cdot,a_2} & \dots & \sum_{a_S\in\mathcal{A}} \pi_{Sa_S}\hat{Q}_{S,\cdot,a_S} \\ 1 & 1 & \dots & 1 \end{pmatrix} \\
&= \dots \\
&= \sum_{(a_1,\dots,a_S)\in\mathcal{A}^S} \pi_{1a_1}\dots\pi_{Sa_S} \det \begin{pmatrix} \hat{Q}_{1,\cdot,a_1} & \dots & \hat{Q}_{S,\cdot,a_S} \\ 1 & \dots & 1 \end{pmatrix} \\
&= \sum_{(a_1,\dots,a_S)\in\mathcal{A}^S} \pi_{1a_1}\dots\pi_{Sa_S} \det(\tilde{Q}^{d^{(a_1,\dots,a_S)}}) \qquad (5.4)
\end{aligned}$$

Similarly, we obtain for any $a_i \in \mathcal{A}$ that

$$\det(\tilde{Q}^\pi)'_{Si} = \sum_{(a_1,\dots,a_{i-1},a_{i+1},\dots,a_S)\in\mathcal{A}^{S-1}} \pi_{1a_1}\cdot\ldots\cdot\pi_{i-1,a_{i-1}}\pi_{i+1a_{i+1}}\cdot\ldots\cdot\pi_{Sa_S} \det(\tilde{Q}^{d^{(a_1,\dots,a_S)}})'_{Si}.$$

This implies

$$\begin{aligned}
&(x^\pi)_i \\
&= \frac{1}{\det\tilde{Q}^\pi}(-1)^{S+i}\det((\tilde{Q}^\pi)'_{Si}) = \frac{(-1)^{S+i}}{\det\tilde{Q}^\pi}\sum_{a_i\in\mathcal{A}} \pi_{ia_i}\det((\tilde{Q}^\pi)'_{Si}) \\
&= \frac{(-1)^{S+i}}{\det\tilde{Q}^\pi}\sum_{(a_1,\dots,a_S)\in\mathcal{A}^S} \pi_{ia}\cdot\pi_{1a_1}\cdot\ldots\cdot\pi_{i-1,a_{i-1}}\pi_{i+1a_{i+1}}\cdot\ldots\cdot\pi_{Sa_S}\det(\tilde{Q}^{d^{(a_1,\dots,a_S)}})'_{Si}
\end{aligned}$$

$$= \frac{1}{\det \tilde{Q}^\pi} \sum_{(a_1,\dots,a_S)\in\mathcal{A}^S} \pi_{1a_1} \cdot \ldots \cdot \pi_{Sa_S} (-1)^{S+i} \det(\tilde{Q}^{d^{(a_1,\dots,a_S)}})'_{Si}$$

$$= \sum_{(a_1,\dots,a_S)\in\mathcal{A}^S} \frac{\pi_{1a_1} \cdot \ldots \cdot \pi_{Sa_S} \det(\tilde{Q}^{d^{a_1,\dots,a_S}})}{\det(\tilde{Q}^\pi)} \cdot x^{d^{(a_1,\dots,a_S)}}.$$

Thus, x^π is a linear combination of $x^{d^{(a_1,\dots,a_S)}}$ for any stationary strategy π and furthermore the coefficients are given by

$$\lambda_{(a_1,\dots,a_S)} = \frac{\pi_{1a_1} \cdot \ldots \cdot \pi_{Sa_S} \det(\tilde{Q}^{d^{(a_1,\dots,a_S)}})}{\det(\tilde{Q}^\pi)}.$$

From (5.4) we obtain that

$$\sum_{(a_1,\dots,a_S)\in\mathcal{A}^S} \lambda_{(a_1,\dots,a_S)} = 1.$$

Furthermore, from Lemma 5.12 we note that the signs of the determinants $\tilde{Q}^\pi$ and $\tilde{Q}^{\pi'}$ are the same for any two strategies π, π'. Thus, we obtain, as $\pi_{ia} \geq 0$ for all $i \in \mathcal{S}, a \in \mathcal{A}$, that $\lambda_{(a_1,\dots,a_S)} \geq 0$ for all $(a_1, \dots, a_S) \in \mathcal{A}^S$. Therefore, every point in $\phi(m)$ is a convex combination of $x^{d_1}, \dots, x^{d_n}$ where $\mathcal{D}(m) = \{d_1, \dots, d_n\}$. Similarly, every convex combination of $x^{d_1}, \dots, x^{d_n}$ is a stationary point given a strategy $\pi \in \text{conv}(\mathcal{D}(m))$, which proves the desired result. □

This result yields that the following programme is sufficient to compute all mean field equilibria: Determine for every $m \in \mathcal{P}(\mathcal{S})$ all optimal deterministic stationary strategies, compute all fixed points of the set-valued map $\phi(m) = \text{conv}(x^{d^1}(m), \dots, x^{d^k}(m))$ with $x^{d^i}(m)$ being the feasible point for d^i with $d^1, \dots, d^k$ being the optimal strategies for m and thereafter solve for every fixed point m the linear equation $0 = \tilde{Q}^\pi(m)m$ for strategies $\pi \in \text{conv}(\mathcal{D}(m))$.

Formally, we can describe this as follows: By $FP(f)$ we denote the set of all fixed points of the set-valued map f. Moreover, let, as introduced on page 54, the set $\text{Opt}(A_1 \times \ldots \times A_S)$ contain all points $m \in \mathcal{P}(\mathcal{S})$ for which $\mathcal{D}(m) = \{d \in D^s : d(i) \in A_i\}$.

Theorem 5.13. *Assume that there is a set $\mathcal{T} \subseteq \mathcal{P}(\mathcal{S})$ such that for all $m \in \mathcal{T}$ and all $\pi \in \Pi^s$ the matrix $Q^\pi(m)$ is irreducible. Then the set of all distributions lying in $\mathcal{T}$ induced by some stationary mean field equilibrium is given by*

$$\bigcup_{A_1\times\ldots\times A_S\subseteq\mathcal{A}^S} FP(conv\{x^{(a_1,\dots,a_S)}(\cdot) : (a_1, \dots, a_s) \in A_1 \times \ldots \times A_S\}) \cap Opt(A_1 \times \ldots \times A_S) \cap \mathcal{T}.$$

Proof. We first note that (m, π) is a stationary mean field equilibrium if and only if m is a fixed point of $\phi(\cdot)$ and π is a strategy such that $\pi \in \text{conv}(\mathcal{D}(m))$ and $0 = m^T Q^\pi(m)$. Theorem 5.6 then yields the desired result. □

Remark 5.14. In order to find the fixed points of the set-valued maps

$$\text{conv}\{x^{(a_1,\dots,a_S)}(m) : (a_1,\dots,a_s) \in A_1 \times \ldots \times A_S\}$$

that lie in the set $\text{Opt}(A_1 \times \ldots \times A_S)$ it is often helpful to work with the explicit characterization of $x^\pi(m)$ derived in Theorem 5.11 given by

$$x^\pi(m)_i = \frac{1}{\sum_{i=1}^S (-1)^{S+i} \det((Q^\pi(m))'_{iS})} (-1)^{S+i} \det(Q^\pi(m)'_{iS}).$$

If the optimality sets are characterized by a specific relation among the m_i (which is often the case) then if $x(\cdot)_i$ does not preserve this relation for all m in an optimality set, we cannot face a fixed point in this optimality set. This type of arguement often reduces the set of candidates immensely (see for example Section 6.3). △

In case of constant dynamics (that is $Q_{ija}(m) = Q_{ija}$ for all $m \in \mathcal{P}(\mathcal{S})$) the fixed point problem is trivial since the maps $x^d(\cdot)$ are constant with value $(\tilde{Q}^d)^{-1} \cdot (0,\dots,0,1)^T$. Thus, we can characterize the set of all mean field equilibria as explicitly by only computing the optimality sets and the stationary points given Q^d as follows:

Corollary 5.15. *Let the dynamics be constant, that is $Q_{ija}(m) = Q_{ija}$ for all $m \in \mathcal{P}(\mathcal{S})$, and let the generators Q^d given any deterministic stationary strategy $d \in D^s$ be irreducible. Then the set of all distributions induced by some stationary mean field equilibrium is given by*

$$\bigcup_{A_1\times\ldots\times A_s \subseteq \mathcal{A}^S} \Big(Opt(A_1 \times \ldots \times A_S) \cap \\ conv\{(\tilde{Q}^{(a_1,\dots,a_S)})^{-1} \cdot (0,\dots,0,1)^T : (a_1,\dots,a_s) \in A_1 \times \ldots \times A_S\}\Big).$$

5.4. A Sufficient Criterion for the Uniqueness of the Fixed Point of $x^\pi(\cdot)$

In this section we propose, relying on Brouwer degree theory, a sufficient condition for $x^\pi(\cdot)$ having a unique fixed point.

Theorem 5.16. *Let $\pi \in \Pi^s$ be a stationary strategy. Assume that $Q^\pi(m)$ is irreducible for all $m \in \mathcal{P}(\mathcal{S})$, that $f^\pi(m) := x^\pi(m) - m$ is continuously differentiable and that the matrix*

$$M^\pi(m) := \begin{pmatrix} \frac{\partial f_1^\pi(m)}{\partial m_1} & \cdots & \frac{\partial f_1^\pi(m)}{\partial m_{S-1}} \\ \vdots & \ddots & \vdots \\ \frac{\partial f_{S-1}^\pi(m)}{\partial m_1} & \cdots & \frac{\partial f_{S-1}^\pi(m)}{\partial m_{S-1}} \end{pmatrix} - \begin{pmatrix} \frac{\partial f_1^\pi(m)}{\partial m_S} & \cdots & \frac{\partial f_1^\pi(m)}{\partial m_S} \\ \vdots & \ddots & \vdots \\ \frac{\partial f_{S-1}^\pi(m)}{\partial m_S} & \cdots & \frac{\partial f_{S-1}^\pi(m)}{\partial m_S} \end{pmatrix}$$

is non-singular for all $m \in \mathcal{P}(\mathcal{S})$. Then the mapping $x^\pi(\cdot)$ has a unique fixed point.

The proof idea is to show that $m \mapsto x^\pi(m) - m$ and the shifted antipodal map $m \mapsto c - m$ are homotopic and then use standard facts regarding the Brouwer degree to show uniqueness of the fixed point of $x^\pi(\cdot)$. This programme is inspired by Chenault (1986), where sufficient criteria for the uniqueness of Nash equilibria are discussed. Before we start with the proof we review the necessary results on the Brouwer degree, for which Deimling (1985, Chapter 1) is a reference:

In Section 1.1 and 1.2 in Deimling (1985) it is shown that there is a unique mapping $\deg(\cdot)$ that maps to every tuple (f, Ω, y) consisting of an open bounded set $\Omega \subseteq \mathbb{R}^n$, a continuous function $f : \overline{\Omega} \to \mathbb{R}^n$ and a value $y \in \mathbb{R}^n \setminus f(\partial\Omega)$ an integer such that the following properties are satisfied:

(d1) $\deg(id, \Omega, y) = 1$ for $y \in \Omega$

(d2) $\deg(f, \Omega, y) = \deg(f, \Omega_1, y) + \deg(f, \Omega_2, y)$ whenever Ω_1, Ω_2 are disjoint open subsets of Ω such that $y \notin f(\bar{\Omega} \setminus (\Omega_1 \cup \Omega_2))$

(d3) $\deg(h(t, \cdot), \Omega, y(t))$ is independent of $t \in [0, 1]$ whenever $h : [0, 1] \times \bar{\Omega} \to \mathbb{R}^n$ is continuous, $y : [0, 1] \to \mathbb{R}^n$ is continuous and $y(t) \notin h(t, \partial\Omega)$ for all $t \in [0, 1]$.

If Ω is an open, bounded set, f is continuous on $\bar{\Omega}$ as well as continuously differentiable on Ω, and $y \notin f(\partial\Omega)$ is a non-critical value, that is the determinant of the Jacobian for all $x \in f^{-1}(\{y\})$ is non-vanishing, then we can compute the degree by

$$\deg(f, \Omega, y) = \sum_{x \in f^{-1}(\{y\})} \operatorname{sgn} \det \left(\frac{\partial}{\partial x} f(x) \right), \tag{5.5}$$

with the standard convention that $\sum_\emptyset = 0$ (see Section 1.2 in Deimling (1985)).

In order to prove the theorem we cannot directly apply the results presented so far to our setting since $\mathcal{P}(\mathcal{S})$ has an empty interior if we consider it as a subset of $\mathbb{R}^S$. Since $m_S = 1 - \sum_{i=1}^{S-1} m_i$ we have that $\mathcal{P}(\mathcal{S})$ is homeomorphic to

$$\Omega' = \left\{ m \in \mathbb{R}^{S-1} : m_i \geq 0 \quad \forall i \in \{1, \dots, S-1\} \quad \text{and} \quad \sum_{i=1}^{S-1} m_i \leq 1 \right\},$$

which satisfies $\Omega' = \overline{\Omega}$ for

$$\Omega = \left\{ m \in \mathbb{R}^{S-1} : m_i > 0 \quad \forall i \in \{1, \dots, S-1\} \quad \text{and} \quad \sum_{i=1}^{S-1} m_i < 1 \right\},$$

which is the setting necessary to apply the previously presented results. With these preparations we prove the theorem:

Proof of Theorem 5.16. We first note that fixed points of $x^\pi(\cdot)$ cannot lie on the boundary of $\mathcal{P}(\mathcal{S})$: Indeed, the boundary consists of all those points for which at least one component is zero, but by Asmussen (2003, Theorem II.4.2) all points $x^\pi(m)$ have only non-zero components. Thus, we can restrict our attention to ensuring uniqueness of a fixed point in $\text{int}(\mathcal{P}(\mathcal{S}))$.

We note that $m \in \mathcal{P}(\mathcal{S})$ is a fixed point of $x^\pi(m)$ if and only if it is a zero of the map $f^\pi(m) = x^\pi(m) - m$, which is, by Remark 5.10, continuous on $\mathcal{P}(\mathcal{S})$. Let us define

$$\phi : \bar{\Omega} \to \mathcal{P}(\mathcal{S}), \quad (m_1, \ldots, m_{S-1}) \mapsto \left(m_1, \ldots, m_{S-1}, 1 - \sum_{i=1}^{S-1} m_i\right)$$

and

$$\psi : \mathbb{R}^S \to \mathbb{R}^{S-1}, \quad (m_1, \ldots, m_{S-1}, m_S) \mapsto (m_1, \ldots, m_{S-1}).$$

Let us furthermore define $\bar{f}^\pi : \bar{\Omega} \to \mathbb{R}^{S-1}$ by $m \mapsto \psi(f^\pi(\phi(m)))$.

First, we verify that 0 is a non-critical value. By the chain rule we obtain

$$\begin{aligned}
\frac{\partial \bar{f}^\pi}{\partial m}(m) &= \frac{\partial \psi}{\partial m}(f^\pi(\phi(m))) \cdot \frac{\partial f^\pi}{\partial m}(\phi(m)) \cdot \frac{\partial \phi}{\partial m}(m) \\
&= \underbrace{\begin{pmatrix} 1 & 0 & \ldots & 0 & 0 \\ 0 & 1 & \ddots & \vdots & \vdots \\ \vdots & \ddots & \ddots & 0 & 0 \\ 0 & \ldots & 0 & 1 & 0 \end{pmatrix}}_{\in \mathbb{R}^{(S-1)\times S}} \cdot \begin{pmatrix} \frac{\partial f_1^\pi(m)}{\partial m_1} & \ldots & \frac{\partial f_1^\pi(m)}{\partial m_S} \\ \vdots & \ddots & \vdots \\ \frac{\partial f_S^\pi(m)}{\partial m_1} & \ldots & \frac{\partial f_S^\pi(m)}{\partial m_S} \end{pmatrix} \cdot \begin{pmatrix} 1 & 0 & \ldots & 0 \\ 0 & 1 & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \ldots & 0 & 1 \\ -1 & -1 & \ldots & -1 \end{pmatrix} \\
&= \begin{pmatrix} \frac{\partial f_1^\pi(m)}{\partial m_1} & \ldots & \frac{\partial f_1^\pi(m)}{\partial m_{S-1}} \\ \vdots & \vdots & \vdots \\ \frac{\partial f_{S-1}^\pi(m)}{\partial m_1} & \ldots & \frac{\partial f_{S-1}^\pi(m)}{\partial m_{S-1}} \end{pmatrix} - \begin{pmatrix} \frac{\partial f_1^\pi(m)}{\partial m_S} & \ldots & \frac{\partial f_1^\pi(m)}{\partial m_S} \\ \vdots & \ddots & \vdots \\ \frac{\partial f_{S-1}^\pi(m)}{\partial m_S} & \ldots & \frac{\partial f_{S-1}^\pi(m)}{\partial m_S} \end{pmatrix} = M^\pi(m),
\end{aligned}$$

which by assumption is non-singular for all $m \in \mathcal{P}(\mathcal{S})$. This in particular implies that $M^\pi(m)$ is non-singular for all $m \in (\bar{f}^\pi)^{-1}(0)$. Furthermore, it yields that

$$\det\left(\frac{\partial \bar{f}^\pi(m)}{\partial m}\right) \neq 0$$

for all $m \in \Omega$. Indeed, since $\phi(\cdot), \psi(\cdot), \det(\cdot)$ and $x^\pi(\cdot)$ are continuous (see Remark 5.10), we note that $\det\left(\frac{\partial \bar{f}^\pi(\cdot)}{\partial m}\right) : \Omega \to \mathbb{R}$ is continuous. Thus, the intermediate value theorem yields that $\det\left(\frac{\partial \bar{f}^\pi(m)}{\partial m}\right)$ has uniform sign over Ω.

We apply $(d3)$ in order to show that the degree at the regular value 0 of the shifted antipodal map $\frac{S-1}{S}(1,\ldots,1)^T - m$ and our map $\bar{f}^\pi$ are the same. For this let

$$\bar{h}(t,m) = t \cdot \bar{f}^\pi(m) + (1-t) \cdot \left(\frac{(S-1)}{S}(1,\ldots,1)^T - m\right)$$
$$= t \cdot \psi(x^\pi(\phi(m))) + (1-t) \cdot \frac{S-1}{S}(1,\ldots,1)^T - m,$$

which is continuous as a product of continuous functions. Let $y(t) \equiv 0$, which is also a continuous function. Thus, it remains to show that $y(t) \notin \bar{h}(t,\partial\Omega)$: Assume that $m \in \partial\Omega$. Then either there is some $i \in \{1,\ldots,S-1\}$ such that $m_i = 0$ or we have $\sum_{i=1}^{S-1} m_i = 1$. By Asmussen (2003, Theorem II.4.2) the vector $x^\pi(m)$ satisfies $x_j^\pi(m) > 0$ for all $j \in \{1,\ldots,S\}$. In the first case, we have that $\bar{h}_i(t,m) > 0$ since $x_i^\pi(m) > 0$ and $\frac{S-1}{S} > 0$. In the second case, we have that by the previous observation the sum of all components of $\psi(x^\pi(\phi(m)))$ is less than one and also the sum of all components of $\frac{S-1}{S}(1,\ldots,S)^T$ is less than one, thus the sum of all components of $t\dot{\psi}(x^\pi(\phi(m))) + (1-t) \cdot \frac{S-1}{S}(1,\ldots,S)^T$ is less than one. However, we assumed that the sum of all entries of $-m$ is exactly -1, so in total we obtain that the sum of all entries of $\bar{h}(t,m)$ has to be less than zero, which especially implies that $\bar{h}(t,m) \neq 0$.

Thus, (d3) yields that

$$\deg\left(\frac{S-1}{S}(1,\ldots,1)^T - m, \Omega, 0\right) = \deg(\bar{f}^\pi, \Omega, 0).$$

From equation (5.5) we then concludes that $\deg(\frac{S-1}{S}(1,\ldots,1)^T - m, \Omega, 0) = (-1)^{S-1}$ and thus we have

$$(-1)^{S-1} = \deg(\bar{f}^\pi) = \sum_{m \in (\bar{f}^\pi)^{-1}(\{0\})} \operatorname{sgn}\det\left(\frac{\partial}{\partial m}\bar{f}^\pi(m)\right).$$

Since the sign of the determinant is uniform over $\Omega \supseteq (\bar{f}^\pi)^{-1}(0)$, we obtain that $(\bar{f}^\pi)^{-1}(\{0\})$ consists of one element. Thus, we have a unique fixed point. □

In the case of constant dynamics, i.e. $Q_{ija}^\pi(m) = Q_{ija}^\pi$, the assumption regarding the matrix $M^\pi(m)$ can be verified directly, indeed, the matrix equals $-I$. For other cases the representation presented in Theorem 5.11 gives rise to a rather direct representation of the partial derivatives of $x^\pi(\cdot)$. More precisely, by deriving the determinants of the minors $(Q^\pi(m))'_{iS}$ as well as their partial derivatives we can compute the partial derivatives of $x^\pi(\cdot)$:

Lemma 5.17. *Let $\pi \in \Pi^s$ be a stationary strategy and assume that $Q^\pi(m)$ is irreducible for all $m \in \mathcal{P}(\mathcal{S})$. Assume furthermore, that $Q^\pi_{ija}(m)$ is continuously differentiable for all $m \in \mathcal{P}(\mathcal{S})$ and all $i, j \in \mathcal{S}$. Then*

$$\frac{\partial x_j^\pi(m)}{\partial m_i} = \frac{(-1)^{S+j}}{\det(\tilde{Q}^\pi(m))^2} \left(\sum_{k=1}^{S} (-1)^{S+k} \left(\left(\frac{\partial}{\partial m_i} \det(Q^\pi(m))'_{jS} \right) \det(Q^\pi(m))'_{kS} - \det(Q^\pi(m))'_{jS} \left(\frac{\partial}{\partial m_i} \det(Q^\pi(m))'_{kS} \right) \right) \right),$$

with

$$\det(\tilde{Q}^\pi(m)) = \sum_{k=1}^{S} (-1)^{S+k} \det((Q^\pi(m))'_{kS}).$$

Proof. Applying the quotient rule on the explicit characterization of $x^\pi(m)$ derived in Theorem 5.11, we obtain

$$\begin{aligned}
\frac{\partial x_j(m)}{\partial m_i} &= \frac{(-1)^{S+j}}{(\det(\tilde{Q}^\pi(m)))^2} \left[\left(\frac{\partial}{\partial m_i} \det(Q^\pi(m))'_{jS} \right) \det(\tilde{Q}^\pi(m)) \right. \\
&\quad \left. - \det(Q^\pi(m))'_{jS} \left(\frac{\partial}{\partial m_i} \det(\tilde{Q}^\pi(m)) \right) \right] \qquad (5.6) \\
&= \frac{(-1)^{S+j}}{(\det(\tilde{Q}^\pi(m)))^2} \left[\left(\frac{\partial}{\partial m_i} \det(Q^\pi(m))'_{jS} \right) \left(\sum_{k=1}^{S} (-1)^{S+k} \det(Q^\pi(m))'_{kS} \right) \right. \\
&\quad \left. - \det(Q^\pi(m))'_{jS} \left(\frac{\partial}{\partial m_i} \left(\sum_{k=1}^{S} (-1)^{S+k} \det(Q^\pi(m))'_{kS} \right) \right) \right] \\
&= \frac{(-1)^{S+j}}{(\det(\tilde{Q}^\pi(m)))^2} \left[\sum_{k=1}^{S} (-1)^{S+k} \left(\left(\frac{\partial}{\partial m_i} \det(Q^\pi(m))'_{jS} \right) \det(Q^\pi(m))'_{kS} \right. \right. \\
&\quad \left. \left. - \det(Q^\pi(m))'_{jS} \left(\frac{\partial}{\partial m_i} \det(Q^\pi(m))'_{kS} \right) \right) \right]. \qquad (5.7)
\end{aligned}$$

□

Remark 5.18. A helpful tool to decide whether a matrix is non-singular is presented in Horn and Johnson (2013, Theorem 6.1.10 and 6.1.11). It states that a matrix A is non-singular whenever it is either strictly diagonal dominant, which means that for all $i = 1, \ldots, S$ we have

$$|a_{ii}| > \sum_{j \neq i} |a_{ij}|,$$

or whenever it satisfies that all diagonal entries are non-zero, for all $i = 1, \ldots, S$ we have

$$|a_{ii}| \geq \sum_{j \neq i} |a_{ij}|$$

and the inequality is strict for at least $S - 1$ values $i \in \{1, \ldots, S\}$. △

5.5. Uniqueness of Equilibria

In classical mean field game theory for diffusion-based models, uniqueness criteria have only been introduced for two small sets of mean field games and it is conjectured in Guéant et al. (2011) that these are the only general uniqueness regimes. More precisely, one either needs a small time horizon or the so called Lasry-Lion monotonicity conditions. This condition requires that the cost function is monotonic in a certain sense and furthermore that the dynamics of an individual do not depend on the population distribution (see Guéant et al. (2011), Carmona and Delarue (2018a, Section 3.4)). Both of these approaches have also been successfully applied to finite state models with continuous action space: In Cecchin and Fischer (2018) as well as Gomes et al. (2013) uniqueness for small time horizons is shown and Gomes et al. (2013) provide a suitable monotonicity condition that ensures uniqueness of dynamic equilibria. We remark that in order to prove uniqueness of stationary equilibria Gomes et al. (2013) additionally assumes contractivity, which requires a certain relation between the Hamiltonian and the value function.

In our setting the methods used previously are not applicable. In the case of stationary solutions, it is not possible to consider a small time horizon, so these approaches cannot be applied. Also the approach using monotonicity conditions cannot be replicated in our setting as these approaches rely on the existence of a unique optimal control for the individual player.

In order to ensure that there is a unique mean field equilibrium in our setting we would have to ensure that there is only one point such that it lies in $\text{Opt}(A_1 \times \ldots \times A_S)$ and that it is simultaneously a fixed point of exactly one strategy that randomizes over the actions from $A_1 \times \ldots \times A_S$. Since stationary points given deterministic and randomized strategies do not relate in general dynamics models, also in our case it would only be possible to come up with uniqueness criteria for constant dynamics. However, also in the case of constant dynamics we will not obtain any satisfactory criteria. The main reason is that in order to describe whether the stationary points given certain dynamics lie in a specific optimality set or not we have to compute both optimality sets and stationary points, which means that we have (almost) computed the equilibria themselves.

6. Examples

This chapter presents three examples fitting into our model and illustrates the application of the techniques presented before. The first two examples have been considered in similar version (and with different purposes) in the literature before (see Gomes et al. (2014b) and Kolokoltsov and Malafeyev (2017)). The first example covers consumer choice in the mobile phone sector, where the utility increases in the share of individuals with the same provider, and has been used as an illustration for numerical methods for dynamic equilibria. The second example models the spread of corruption with peer pressure and for this model all deterministic stationary equilibria have been computed. The third model has not been considered so far and describes consumer choice with congestion effects, namely, the more individuals share the same provider the higher is the risk of losing the service for some time.

As we have seen in Chapter 5 we first have to solve the optimal control problem of the individual agent. For this we will follow in all three examples the same approach: First, we will investigate the optimality equation to eliminate never optimal strategies. Then we will derive the expected discounted reward of the remaining strategies. Thereafter, we will verify that they indeed solve the optimality equation.

The second step, namely, the analysis of the fixed point problem, is solved differently in all three examples. The consumer choice model of Section 6.1 is a model with constant dynamics, which means that Corollary 5.15 directly yields the distributions induced by the existing stationary mean field equilibria. The corruption model of Section 6.2 has non-irreducible dynamics, thus the fixed point problem is solved relying on the cut criterion presented in Theorem 5.5 as well as sensible rearrangements of the balance equations. The consumer choice model with congestion effects in Section 6.3 has irreducible dynamics, which means that we can apply Theorem 5.13. However, due to the non-constant dynamics we cannot directly compute the fixed points of $x^{\pi}(\cdot)$. Instead, we restrict the set of candidates by using the cut criterion and prove thereafter using Theorem 5.16 that a unique fixed point exists. Using this we then derive which points can be stationary points given a mixed strategy. Moreover, we show using the explicit representation of $x^{d}(\cdot)$ that there is at most one equilibrium with a deterministic equilibrium strategy. Putting all the observations together, we then obtain a full characterization of all stationary mean field equilibria in the third example.

6.1. An Example with Constant Dynamics: Consumer Choice

The first example we present is a model with constant dynamics, which means that the dynamics do not depend on the current population distribution. The model is similar to the model introduced by Gomes et al. (2014b) as a toy example on which the authors demonstrated numerical methods for a specific class of finite state mean field games, namely those yielding systems of hyperbolic partial differential equations. However, the action space in their model as well as the possible strategies of each player differ systematically from the model considered here.

We model consumer choice in the mobile phone sector. The utility of using a certain provider is increasing in the share of consumers using it, whereas the costs are constant. We assume that the utility coming from the other customers sharing the same provider $i \in \{1,2\}$ is given by the isoelastic utility function $\ln(m_i) + s_i$, with $s_i \geq 0$. Since $\ln(m_i)$ is always negative for our choice of m_i, one cannot interpret s_i as costs directly, but one rather has to think of s_i consisting of two components: the costs themselves and some base utility from service provision. The players choose in our model whether to stick to their provider i or whether to switch to the other provider, in this case the player additionally faces a time-unit switching cost $c \geq 0$. For technical reasons it is important that the player always faces, independent of the chosen strategy, a small risk of going to the other provider (we can think of this risk as the possibility of another family member changing the provider although one does not want to).

The choice options substantially differ from the model in Gomes et al. (2014b), where the players could continuously control the rates at which they switch to the other state and were facing costs corresponding to the square of the rate. From an applied point of view it is questionable, in particular when agents are not experts in the game at hand, that players indeed understand what it means to control the transition rates of a Markov chain. Indeed, economic experiments show that most people cannot understand the true effect of random devices even in simple settings (Walker and Wooders, 2008). Furthermore, even if the players understand these control options, it is unreasonable that such a precise control can be put into practise.

The formal description of the model is given as follows: The state space is $\mathcal{S} = \{1,2\}$, the action space is given by $\mathcal{A} = \{change, stay\}$. For technical reasons (mainly because $\ln(0)$ is not defined) we define for $\delta > 0$ the function $f_\delta : \mathbb{R} \to \mathbb{R}$ given by

$$y \mapsto \begin{cases} \frac{1}{4\delta} y^2 + \frac{1}{2} y + \frac{1}{4}\delta & \text{if } y \leq \delta \\ y & \text{if } y \geq \delta \end{cases}$$

and we note that $f_\delta(\cdot)$ is continuously differentiable and strictly increasing on $(-\delta, \infty)$.

Using $f_\delta(\cdot)$ we can properly define the transition rates and reward functions by

$$Q_{\cdot\cdot change} = \begin{pmatrix} -b & b \\ b & -b \end{pmatrix} \qquad Q_{\cdot\cdot stay} = \begin{pmatrix} -\epsilon & \epsilon \\ \epsilon & -\epsilon \end{pmatrix}$$
$$r_{\cdot change} = \begin{pmatrix} \ln(f_\delta(m_1)) + s_1 - c \\ \ln(f_\delta(m_2)) + s_2 - c \end{pmatrix} \qquad r_{\cdot stay} = \begin{pmatrix} \ln(f_\delta(m_1)) + s_1 \\ \ln(f_\delta(m_2)) + s_2 \end{pmatrix},$$

where $0 < \epsilon < b$, $s_1, s_2, c > 0$ and $\delta > 0$ is small. We remark that the model, by definition of $f_\delta(\cdot)$ exhibits a utility function that is almost isoelastic, only for very small values of m_1 and m_2, i.e. $m_1 < \delta$ or $m_2 < \delta$ this is not the case.

In order to analyse the model we first consider the optimality equation (5.1), which in our setting reads

$$\begin{aligned} \beta V_1(m) = \max\{&\ln(f_\delta(m_1)) + s_1 - c - bV_1(m) + bV_2(m), \\ &\ln(f_\delta(m_1)) + s_1 - \epsilon V_1(m) + \epsilon V_2(m)\} \\ \beta V_2(m) = \max\{&\ln(f_\delta(m_2)) + s_2 - c + bV_1(m) - bV_2(m), \\ &\ln(f_\delta(m_2)) + s_2 + \epsilon V_1(m) - \epsilon V_2(m)\}. \end{aligned}$$

It yields that it is optimal to change in state 1 if and only if $V_1(m) - V_2(m) \leq -\frac{c}{b-\epsilon}$ and that it is optimal to change in state 2 if and only if $V_1(m) - V_2(m) \geq \frac{c}{b-\epsilon}$. Thus, choosing the action *change* in both states simultaneously is never optimal. Therefore, we focus on the three potentially optimal strategies $\{change\} \times \{stay\}$, $\{stay\} \times \{stay\}$ and $\{stay\} \times \{change\}$. By Lemma 5.3 the expected discounted reward given these strategies is

$$\begin{aligned} &V^{\{change\}\times\{stay\}}(m) \\ &= \left(\beta I - Q^{\{change\}\times\{stay\}}(m)\right)^{-1} r^{\{change\}\times\{stay\}}(m) \\ &= \begin{pmatrix} \beta + b & -b \\ -\epsilon & \beta + \epsilon \end{pmatrix}^{-1} \begin{pmatrix} \ln(f_\delta(m_1)) + s_1 - c \\ \ln(f_\delta(m_2)) + s_2 \end{pmatrix} \\ &= \frac{1}{\beta(\beta + b + \epsilon)} \begin{pmatrix} \beta + \epsilon & b \\ \epsilon & \beta + b \end{pmatrix} \begin{pmatrix} \ln(f_\delta(m_1)) + s_1 - c \\ \ln(f_\delta(m_2)) + s_2 \end{pmatrix} \\ &= \frac{1}{\beta(\beta + b + \epsilon)} \begin{pmatrix} (\ln(f_\delta(m_1)) + s_1) \cdot (\beta + \epsilon) - c \cdot (\beta + \epsilon) + (\ln(f_\delta(m_2)) + s_2) \cdot b \\ (\ln(f_\delta(m_1)) + s_1) \cdot \epsilon - c \cdot \epsilon + (\ln(f_\delta(m_2)) + s_2) \cdot (\beta + b) \end{pmatrix} \\ &V^{\{stay\}\times\{stay\}}(m) \\ &= \begin{pmatrix} \beta + \epsilon & -\epsilon \\ -\epsilon & \beta + \epsilon \end{pmatrix}^{-1} \begin{pmatrix} \ln(f_\delta(m_1)) + s_1 \\ \ln(f_\delta(m_2)) + s_2 \end{pmatrix} \\ &= \frac{1}{\beta^2 + 2\beta\epsilon} \begin{pmatrix} (\ln(f_\delta(m_1)) + s_1) \cdot (\beta + \epsilon) + (\ln(f_\delta(m_2)) + s_2) \cdot \epsilon \\ (\ln(f_\delta(m_1)) + s_1) \cdot \epsilon + (\ln(f_\delta(m_2)) + s_2) \cdot (\beta + \epsilon) \end{pmatrix} \\ &V^{\{stay\}\times\{change\}}(m) \\ &= \begin{pmatrix} \beta + \epsilon & -\epsilon \\ -b & \beta + b \end{pmatrix}^{-1} \begin{pmatrix} \ln(f_\delta(m_1)) + s_1 \\ \ln(f_\delta(m_2)) + s_2 - c \end{pmatrix} \end{aligned}$$

$$= \frac{1}{\beta(\beta+b+\epsilon)} \begin{pmatrix} (\ln(f_\delta(m_1)) + s_1) \cdot (\beta+b) - c \cdot \epsilon + (\ln(f_\delta(m_2)) + s_2) \cdot \epsilon \\ (\ln(f_\delta(m_1)) + s_1) \cdot b - c \cdot (\beta+\epsilon) + (\ln(f_\delta(m_2)) + s_2) \cdot (\beta+\epsilon) \end{pmatrix}.$$

In order to understand which strategy is optimal, we compute the differences $V_1^d(m) - V_2^d(m)$:

$$\begin{aligned}
&V_1^{\{change\}\times\{stay\}}(m) - V_2^{\{change\}\times\{stay\}}(m) \\
&= \frac{1}{\beta+b+\epsilon}\left(-c + \ln(f_\delta(m_1)) - \ln(f_\delta(m_2)) + s_1 - s_2\right) \\
&V_1^{\{stay\}\times\{stay\}}(m) - V_2^{\{stay\}\times\{stay\}}(m) \\
&= \frac{1}{\beta+2\epsilon}\left(\ln(f_\delta(m_1)) - \ln(f_\delta(m_2)) + s_1 - s_2\right) \\
&V_1^{\{stay\}\times\{change\}}(m) - V_2^{\{stay\}\times\{change\}}(m) \\
&= \frac{1}{\beta+b+\epsilon}\left(c + \ln(f_\delta(m_1)) - \ln(f_\delta(m_2)) + s_1 - s_2\right).
\end{aligned}$$

Let us write

$$k_1 = \frac{\exp\left(\frac{-c(\beta+2\epsilon)}{b-\epsilon} - s_1 + s_2\right)}{1 + \exp\left(\frac{-c(\beta+2\epsilon)}{b-\epsilon} - s_1 + s_2\right)} \quad \text{and } k_2 = \frac{\exp\left(\frac{c(\beta+2\epsilon)}{b-\epsilon} - s_1 + s_2\right)}{1 + \exp\left(\frac{c(\beta+2\epsilon)}{b-\epsilon} - s_1 + s_2\right)}$$

and note that $k_1 < k_2$. Now choose $\delta > 0$ to be small enough (i.e. $\delta \ll \min\{k_1, 1-k_2\}$). Then we obtain using $m_1 = 1 - m_2$ and that $f_\delta(\cdot)$ is strictly increasing

$$\begin{aligned}
&V_1^{\{change\}\times\{stay\}}(m) - V_2^{\{change\}\times\{stay\}}(m) \leq -\frac{c}{b-\epsilon} \\
&\Leftrightarrow -c + \ln(f_\delta(m_1)) - \ln(f_\delta(m_2)) + s_1 - s_2 \leq \frac{-c(\beta+b+\epsilon)}{b-\epsilon} \\
&\Leftrightarrow \ln\left(\frac{f_\delta(m_1)}{f_\delta(1-m_1)}\right) \leq \frac{-c(\beta+2\epsilon)}{b-\epsilon} - s_1 + s_2 \\
&\Leftrightarrow \ln\left(\frac{m_1}{1-m_1}\right) \leq \frac{-c(\beta+2\epsilon)}{b-\epsilon} - s_1 + s_2 \\
&\Leftrightarrow m_1 \leq \frac{\exp\left(\frac{-c(\beta+2\epsilon)}{b-\epsilon} - s_1 + s_2\right)}{1 + \exp\left(\frac{-c(\beta+2\epsilon)}{b-\epsilon} - s_1 + s_2\right)},
\end{aligned} \tag{6.1}$$

which yields that $V^{\{change\}\times\{stay\}}(\cdot)$ is the unique solution of (5.1) if and only if (6.1) holds.

Analogously, we obtain

$$\frac{-c}{b-\epsilon} \leq V_1^{\{stay\}\times\{stay\}}(m) - V_2^{\{stay\}\times\{stay\}}(m) \leq \frac{c}{b-\epsilon}$$

$$\Leftrightarrow \frac{\exp\left(\frac{-c(\beta+2\epsilon)}{b-\epsilon} - s_1 + s_2\right)}{1+\exp\left(\frac{-c(\beta+2\epsilon)}{b-\epsilon} - s_1 + s_2\right)} \leq m_1 \leq \frac{\exp\left(\frac{c(\beta+2\epsilon)}{b-\epsilon} - s_1 + s_2\right)}{1+\exp\left(\frac{c(\beta+2\epsilon)}{b-\epsilon} - s_1 + s_2\right)}$$

as well as

$$\frac{c}{b-\epsilon} \leq V_1^{\{stay\}\times\{change\}}(m) - V_2^{\{stay\}\times\{change\}}(m)$$
$$\Leftrightarrow \frac{\exp\left(\frac{c(\beta+2\epsilon)}{b-\epsilon} - s_1 + s_2\right)}{1+\exp\left(\frac{c(\beta+2\epsilon)}{b-\epsilon} - s_1 + s_2\right)} \leq m_1.$$

By the previous investigations we have the following non-empty optimality sets

$$\begin{aligned}
\mathrm{Opt}(\{change\}\times\{stay\}) &= \{m \in \mathcal{P}(\mathcal{S}) : m_1 < k_1\} \\
\mathrm{Opt}(\{change, stay\}\times\{stay\}) &= \{(k_1, 1-k_1)\} \\
\mathrm{Opt}(\{stay\}\times\{stay\}) &= \{m \in \mathcal{P}(\mathcal{S}) : m_1 \in (k_1, k_2)\} \\
\mathrm{Opt}(\{change\}\times\{change, stay\}) &= \{(k_2, 1-k_2)\} \\
\mathrm{Opt}(\{stay\}\times\{change\}) &= \{m \in \mathcal{P}(\mathcal{S}) : m_1 > k_2\}
\end{aligned}$$

As a next step we compute the fixed points given the three strategies $\{change\}\times\{stay\}$, $\{stay\}\times\{stay\}$ and $\{stay\}\times\{change\}$. We remark that we do not need to consider the fixed point given $\{change\}\times\{change\}$ because it is never optimal to randomize in both states. Since we face standard continuous time Markov chains, this task is simple and we obtain for each strategy a unique stationary point:

$$x^{\{change\}\times\{stay\}} = \begin{pmatrix} \frac{\epsilon}{b+\epsilon} \\ \frac{b}{b+\epsilon} \end{pmatrix}, \quad x^{\{stay\}\times\{stay\}} = \begin{pmatrix} \frac{1}{2} \\ \frac{1}{2} \end{pmatrix}, \quad x^{\{stay\}\times\{change\}} = \begin{pmatrix} \frac{b}{b+\epsilon} \\ \frac{\epsilon}{b+\epsilon} \end{pmatrix}.$$

Finally, we obtain the exact number and position of equilibria by coefficient comparison. The strategies that are used in mixed strategy equilibria are then obtained by solving for each distribution m of a stationary mean field equilibrium the system $m^T Q^\pi(m) = 0$ for strategies $\pi \in \mathrm{conv}(\mathcal{D}(m))$.

(i) If $k_1 < \frac{\epsilon}{b+\epsilon}$ and $k_2 < \frac{1}{2}$, then $x^{\{stay\}\times\{change\}}$ together with the deterministic strategy $\{stay\}\times\{change\}$ is the unique stationary mean field equilibrium.

(ii) If $k_1 < \frac{\epsilon}{b+\epsilon}$ and $\frac{1}{2} \leq k_2 \leq \frac{b}{b+\epsilon}$, then $x^{\{stay\}\times\{stay\}}$ together with the deterministic strategy $\{stay\}\times\{stay\}$ and $x^{\{stay\}\times\{change\}}$ together with the deterministic strategy $\{stay\}\times\{change\}$ are the deterministic equilibria. Furthermore, there is one mixed strategy equilibrium with $m_1 = k_2$ that randomizes over *stay* and *change* in the second state. If $k_2 \in \{\frac{1}{2}, \frac{b}{b+\epsilon}\}$, then the mixed strategy equilibrium coincides with the pure strategy equilibrium with the same population distribution.

(iii) If $k_1 < \frac{\epsilon}{b+\epsilon}$ and $\frac{b}{b+\epsilon} < k_2$, then the only equilibrium is given by $x^{\{stay\}\times\{stay\}}$ together with the deterministic strategy $\{stay\} \times \{stay\}$.

(iv) If $\frac{\epsilon}{b+\epsilon} \leq k_1 \leq \frac{1}{2}$ and $k_2 < \frac{1}{2}$, then $x^{\{stay\}\times\{change\}}$ together with the deterministic strategy $\{stay\} \times \{change\}$ and $x^{\{change\}\times\{stay\}}$ together with the deterministic strategy $\{change\} \times \{stay\}$ are the deterministic equilibria. Furthermore, there is a mixed strategy equilibrium at the point with $m_1 = k_1$ that randomizes over *stay* and *change* in the first state. If $k_1 = \frac{\epsilon}{b+\epsilon}$, then the mixed strategy equilibrium coincides with the pure strategy equilibrium with the same population distribution.

(v) If $\frac{\epsilon}{b+\epsilon} \leq k_1 \leq \frac{1}{2}$ and $\frac{1}{2} \leq k_2 \leq \frac{b}{b+\epsilon}$, then $x^{\{stay\}\times\{change\}}$ together with the deterministic strategy $\{stay\} \times \{change\}$, $x^{\{stay\}\times\{stay\}}$ together with the deterministic strategy $\{stay\} \times \{stay\}$ and $x^{\{change\}\times\{stay\}}$ together with the deterministic strategy $\{change\} \times \{stay\}$ are the deterministic equilibria. Furthermore, there is one mixed strategy equilibrium with $m_1 = k_1$ that randomizes over *stay* and *change* in the first state and one mixed strategy equilibrium with $m_1 = k_2$ that randomizes over *stay* and *change* in the second state. If $k_1 \in \{\frac{\epsilon}{b+\epsilon}, \frac{1}{2}\}$ or $k_2 \in \{\frac{1}{2}, \frac{b}{b+\epsilon}\}$, then the mixed strategy equilibrium coincides with the pure strategy equilibrium with the same population distribution.

(vi) If $\frac{\epsilon}{b+\epsilon} \leq k_1 \leq \frac{1}{2}$ and $\frac{b}{b+\epsilon} < k_2$, then $x^{\{stay\}\times\{stay\}}$ together with the deterministic strategy $\{stay\} \times \{stay\}$ and $x^{\{change\}\times\{stay\}}$ together with the deterministic strategy $\{change\} \times \{stay\}$ are the deterministic equilibria. Furthermore, there is one mixed strategy equilibrium with $m_1 = k_1$ that randomizes over *stay* and *change* in the first state. If $k_1 \in \{\frac{\epsilon}{b+\epsilon}, \frac{1}{2}\}$, then the mixed strategy equilibrium coincides with the pure strategy equilibrium with the same population distribution.

(vii) If $\frac{1}{2} < k_1$ and $\frac{1}{2} < k_2 \leq \frac{b}{b+\epsilon}$, then $x^{\{stay\}\times\{change\}}$ together with the deterministic strategy $\{stay\} \times \{change\}$ and $x^{\{change\}\times\{stay\}}$ together with the deterministic strategy $\{change\} \times \{stay\}$ are the deterministic equilibria. Furthermore, there is a mixed strategy equilibrium at the point with $m_1 = k_2$ that randomizes over *stay* and *change* in the second state. If $k_2 = \frac{b}{b+\epsilon}$, then the mixed strategy equilibrium coincides with the pure strategy equilibrium with the same population distribution.

(viii) If $\frac{1}{2} < k_1$ and $\frac{b}{b+\epsilon} < k_2$, then then $x^{\{change\}\times\{stay\}}$ together with the deterministic strategy $\{change\} \times \{stay\}$ is the unique stationary mean field equilibrium.

6.2. A First Example with Non-Constant Dynamics: A Simplified Corruption Model

We now consider a simplified version of the corruption model presented in Kolokoltsov and Malafeyev (2017), which has been introduced in Section 3.1 and which goal is to capture the effect of social pressure on the spread of corruption in a society. For convenience we shortly review the model: A player can be in one of the three states honest (H), corrupt (C) and reserved (R) and he can choose, given that he is not reserved, whether he wants to stay corrupt/honest or whether he wants to switch behaviour. The social pressure acts in two ways: First, the more players are corrupt the higher is the pressure (one cannot escape) to also become corrupt. Second, the more players are honest the higher is the rate to become convicted to be corrupt.

The formal characterization is given by $\mathcal{S} = \{C, H, R\}$ and $\mathcal{A} = \{\textit{change}, \textit{stay}\}$ together with

$$Q_{\cdot\cdot \textit{change}} = \begin{pmatrix} -(b + q_{\text{soc}} m_H) & b & q_{\text{soc}} m_H \\ b + q_{\text{inf}} m_C & -(b + q_{\text{inf}} m_C) & 0 \\ 0 & \lambda & -\lambda \end{pmatrix}$$

$$Q_{\cdot\cdot \textit{stay}} = \begin{pmatrix} -q_{\text{soc}} m_H & 0 & q_{\text{soc}} m_H \\ q_{\text{inf}} m_C & -q_{\text{inf}} m_C & 0 \\ 0 & \lambda & -\lambda \end{pmatrix}$$

and $r_{\cdot \textit{change}} = r_{\cdot \textit{stay}} = (10, 5, 0)^T$, where all parameters b, q_{inf}, q_{soc} and λ are strictly positive.

We remark that there are two significant simplifications compared to the model of Kolokoltsov and Malafeyev (2017): First, there is no principal agent that convicts players. Second, in their model the wages are not fixed, but arbitrary with values $w_C > w_H > w_R$.

We start with computing the value function for given $m \in \mathcal{P}(\mathcal{S})$ as the unique solution of

$$\begin{aligned} \beta V_C(m) &= \max\{10 - (b + q_{\text{soc}} m_H) V_C(m) + b V_H(m) + q_{\text{soc}} m_H V_R(m), \\ &\qquad 10 - q_{\text{soc}} m_H V_C(m) + q_{\text{soc}} m_H V_R(m)\} \\ \beta V_H(m) &= \max\{5 + (b + q_{\text{inf}} m_C) V_C(m) - (b + q_{\text{inf}} m_C) V_H(m), \\ &\qquad 5 + q_{\text{inf}} m_C V_C(m) - q_{\text{inf}} m_C V_H(m)\} \\ \beta V_R(m) &= \lambda V_H(m) - \lambda V_R(m). \end{aligned}$$

We obtain that the agent should choose the action *change* in state C if $V_C(m) \leq V_H(m)$ and the action *stay* in state C if $V_C(m) \geq V_H(m)$. Analogously, we obtain that the agent should chose the action *change* in state H if $V_C(m) \geq V_H(m)$ and the action *stay* in state H if $V_C(m) \leq V_H(m)$. A straightforward calculation as in Section 6.1 yields that the optimality sets are given by

$$\begin{aligned}
\mathrm{Opt}(\{\mathit{stay}\} \times \{\mathit{change}\}) &= \left\{ m \in \mathcal{P}(\mathcal{S}) : m_H < \frac{\lambda + \beta}{q_{\mathrm{soc}}} \right\} \\
\mathrm{Opt}(\{\mathit{change}\} \times \{\mathit{stay}\}) &= \left\{ m \in \mathcal{P}(\mathcal{S}) : m_H > \frac{\lambda + \beta}{q_{\mathrm{soc}}} \right\} \\
\mathrm{Opt}(\{\mathit{stay}, \mathit{change}\} \times \{\mathit{stay}, \mathit{change}\}) &= \left\{ m \in \mathcal{P}(\mathcal{S}) : m_H = \frac{\lambda + \beta}{q_{\mathrm{soc}}} \right\}.
\end{aligned}$$

Thus, depending on the choice of parameters the quantity $\frac{\lambda+\beta}{q_{\mathrm{soc}}}$ is greater than one, exactly one or less than one, which means that there are one, two or three non-empty-optimality sets. In the first case only $\mathrm{Opt}(\{\mathit{stay}\} \times \{\mathit{change}\})$ is non-empty, in the second case $\mathrm{Opt}(\{\mathit{stay}\} \times \{\mathit{change}\})$ and $\mathrm{Opt}(\{\mathit{stay}, \mathit{change}\} \times \{\mathit{stay}, \mathit{change}\})$ are non-empty and in the third case all optimality sets are non-empty.

As a second step we use the cut criterion from Theorem 5.5 for the set $\mathcal{T} = \{R\}$, which yields the equation $q_{\mathrm{soc}} m_H m_C = \lambda m_R$ for all stationary strategies. Together with the equation $1 = m_R + m_H + m_C$ we obtain

$$\begin{aligned}
q_{\mathrm{soc}} m_H m_C = \lambda(1 - m_H - m_C), \text{ i.e. } & m_H(q_{\mathrm{soc}} m_C + \lambda) = \lambda(1 - m_C), \\
\text{i.e. } & m_H = \frac{\lambda(1 - m_C)}{q_{\mathrm{soc}} m_C + \lambda}
\end{aligned}$$

or equivalently $m_C = \lambda(1 - m_H)/(q_{\mathrm{soc}} m_H + \lambda)$. This representation implies that $m_C \in [0, 1]$ if and only if $m_H \in [0, 1]$. Furthermore, we obtain that the sum of m_C and m_H is always less than one. Thus, any mean field equilibrium is uniquely characterized by describing m_C. More precisely, any stationary mean field equilibrium has a distribution of the form

$$\left(m_C, \frac{\lambda - \lambda m_C}{q_{\mathrm{soc}} m_C + \lambda}, \frac{q_{\mathrm{soc}} - q_{\mathrm{soc}} m_C^2}{q_{\mathrm{soc}} m_C + \lambda} \right). \tag{6.2}$$

It remains to consider the fixed point problems given the possible optimal strategies: We start by noting that stationary points of the dynamics given the strategy $\{\mathit{stay}\} \times \{\mathit{change}\}$ exists whenever (6.2) and

$$-q_{\mathrm{soc}} m_C m_H + q_{\mathrm{inf}} m_C m_H + b m_H = 0, \quad \text{i.e.} \quad m_H(m_C(q_{\mathrm{soc}} - q_{\mathrm{inf}}) - b) = 0$$

is satisfied, which is true if $m_H = 0$ or $m_C = b/(q_{\mathrm{soc}} - q_{\mathrm{inf}})$. Thus, whenever these points lie in $\mathrm{Opt}(\{\mathit{stay}\} \times \{\mathit{change}\})$ or $\mathrm{Opt}(\{\text{stay, change}\} \times \{\text{stay, change}\})$ we have a deterministic mean field equilibrium given the strategy $\{\mathit{stay}\} \times \{\mathit{change}\}$.

Similarly, stationary points of the dynamics given the strategy $\{\mathit{change}\} \times \{\mathit{stay}\}$ have to satisfy (6.2) and

$$-b m_C - q_{\mathrm{soc}} m_H m_C + q_{\mathrm{inf}} m_C m_H = 0, \quad \text{i.e.} \quad m_C(m_H(q_{\mathrm{inf}} - q_{\mathrm{soc}}) - b) = 0,$$

which is true if either $m_C = 0$ or $m_H = b/(q_{\text{inf}} - q_{\text{soc}})$. Thus whenever these points lie in $\text{Opt}(\{\textit{change}\} \times \{\textit{stay}\})$ or $\text{Opt}(\{\text{stay, change}\} \times \{\text{stay, change}\})$ we have a deterministic mean field equilibrium given the strategy $\{\textit{change}\} \times \{\textit{stay}\}$.

When searching for stationary points given the dynamics of mixed strategies, which might be equilibria, we can restrict to those that lie inside the set $\text{Opt}(\{\text{stay, change}\} \times \{\text{stay, change}\})$. In this set all equilibria have to satisfy $m_H = \frac{\lambda+\beta}{q_{\text{soc}}}$, that is

$$\frac{\lambda+\beta}{q_{\text{soc}}} = \frac{\lambda(1-m_C)}{q_{\text{soc}} m_C + \lambda} \Leftrightarrow m_C = \frac{\lambda(q_{\text{soc}} - \lambda - \beta)}{(2\lambda+\beta) q_{\text{soc}}}.$$

It remains to check whether there is a strategy such that the point

$$\left(\frac{\lambda(q_{\text{soc}} - \lambda - \beta)}{(2\lambda+\beta) q_{\text{soc}}}, \frac{\lambda+\beta}{q_{\text{soc}}}, \frac{(\lambda+\beta)(q_{\text{soc}} - \lambda - \beta)}{(2\lambda+\beta) q_{\text{soc}}} \right) \tag{6.3}$$

is indeed a fixed point for the individual dynamics equation given strategy π, which means that we have to find constants $\pi_{1,\textit{change}}$ and $\pi_{2,\textit{change}}$ that satisfy for this point

$$(-\pi_{1,\textit{change}} b - q_{\text{soc}} m_H) m_C + \pi_{2,\textit{change}} b m_H + q_{\text{inf}} m_C m_H = 0.$$

As in the previous example, we would need to perform a case analysis to obtain the exact set of mean field equilibria for all possible equilibrium constellations. Additionally, we would need to solve the balance equations $m^T Q^\pi(m) = 0$ for π for the point m given by (6.3), which is the only candidate for a randomized equilibrium. Both tasks are simple, but tedious, and we omit them here.

6.3. A Second Example with Non-Constant Dynamics: Consumer Choice Again

The second example with non-constant dynamics we consider consists of two "good" states, where a positive reward is earned, and one "bad" state, where no reward is earned. The agents in the "good" state face congestion effects, namely there is a risk, increasing in the share of individuals in that state, to go to the "bad" state. The control options are to switch between the two good states. One can interpret this model as a stylized version to model the choice between two mobile phone providers, where the utility to be customer of one of these providers differs and one has the additional risk of a breakdown in connection. More precisely, the more people share the same provider, the higher is the risk of an eventual failure of the service associated with a loss of utility. For simplicity, we assume that agents in the "bad" state have no choice option, but recover into each of the two states with equal probability.

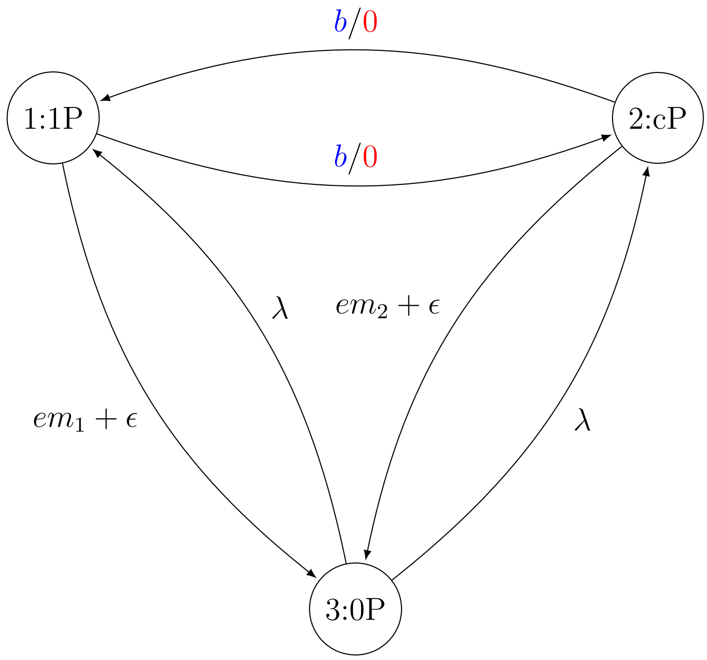

Figure 6.1.: Representation of the consumer choice model discussed in Section 6.3. The transition rates in black are action independent, the blue transition rates are those associated to choosing the action *change* and the red transition rates are those associated to choosing the action *stay*.

The formal characterization is given by $\mathcal{S} = \{1, 2, 3\}$ and $\mathcal{A} = \{change, stay\}$ together with

$$Q_{\cdot\cdot change} = \begin{pmatrix} -(b + em_1 + \epsilon) & b & em_1 + \epsilon \\ b & -(b + em_2 + \epsilon) & em_2 + \epsilon \\ \lambda & \lambda & -2\lambda \end{pmatrix}$$

$$Q_{\cdot\cdot stay} = \begin{pmatrix} -(em_1 + \epsilon) & 0 & em_1 + \epsilon \\ 0 & -(em_2 + \epsilon) & em_2 + \epsilon \\ \lambda & \lambda & -2\lambda \end{pmatrix}$$

and $r_{\cdot stay} = r_{\cdot change} = (1, c, 0)$, where all constants are strictly positive and additionally $c \leq 1$. A visualization of the model is given in Figure 6.3.

We again start by analysing the optimality equation

$$\begin{aligned} \beta V_1(m) &= \max\{1 - (b + em_1 + \epsilon)V_1(m) + bV_2(m) + (em_1 + \epsilon)V_3(m), \\ &\qquad 1 - (em_1 + \epsilon)V_1(m) + (em_1 + \epsilon)V_3(m)\} \\ \beta V_2(m) &= \max\{c + bV_1(m) - (b + em_2 + \epsilon)V_2(m) + (em_2 + \epsilon)V_3(m), \\ &\qquad c - (em_2 + \epsilon)V_2(m) + (em_2 + \epsilon)V_3(m)\} \\ \beta V_3(m) &= \lambda V_1(m) + \lambda V_2(m) - 2\lambda V_3(m). \end{aligned}$$

We directly see that the optimal action in state 1 is *change* if $V_1(m) \leq V_2(m)$ and *stay* if $V_1(m) \geq V_2(m)$. Similarly, the optimal action in state 2 is *change* if $V_1(m) \geq V_2(m)$ and *stay* if $V_1(m) \leq V_2(m)$.

This implies that at least one of the strategies $\{change\} \times \{stay\}$ and $\{stay\} \times \{change\}$ is optimal. As in the previous examples, we compute using Lemma 5.3 the expected discounted reward given these strategies and investigate thereafter for which parameters these rewards are indeed solutions of the optimality equation. This yields to the following non-empty optimality sets:

$$\begin{aligned}\mathrm{Opt}(\{stay\} \times \{change\}) &= \left\{ m \in \mathcal{P}(\mathcal{S}) : \frac{1}{c} m_2 + \frac{1-c}{ce}(\beta + \epsilon + 2\lambda) > m_1 \right\} \\ \mathrm{Opt}(\{change\} \times \{stay\}) &= \left\{ m \in \mathcal{P}(\mathcal{S}) : \frac{1}{c} m_2 + \frac{1-c}{ce}(\beta + \epsilon + 2\lambda) < m_1 \right\} \\ \mathrm{Opt}(\{stay, change\} \times \{stay, change\}) &= \left\{ m \in \mathcal{P}(\mathcal{S}) : \frac{1}{c} m_2 + \frac{1-c}{ce}(\beta + \epsilon + 2\lambda) = m_1 \right\}.\end{aligned}$$

We remark that in the symmetric case $c = 1$ we indeed obtain that

$$\begin{aligned}\mathrm{Opt}(\{stay\} \times \{change\}) &= \{m \in \mathcal{P}(\mathcal{S}) : m_2 > m_1\} \\ \mathrm{Opt}(\{change\} \times \{stay\}) &= \{m \in \mathcal{P}(\mathcal{S}) : m_2 < m_1\} \\ \mathrm{Opt}(\{stay, change\} \times \{stay, change\}) &= \{m \in \mathcal{P}(\mathcal{S}) : m_2 = m_1\}.\end{aligned}$$

To obtain the stationary points of the dynamics for the strategies of interest we start the analysis of the fixed point problem by applying the cut criterion for $\mathcal{T} = \{1, 2\}$, which yields

$$\lambda m_3 + \lambda m_3 = (em_1 + \epsilon) \cdot m_1 + (em_2 + \epsilon) \cdot m_2. \tag{6.4}$$

Combing this equation with $1 = m_1 + m_2 + m_3$ we obtain

$$\lambda - (2\lambda + \epsilon)m_1 - em_1^2 = em_2^2 + (2\lambda + \epsilon)m_2 - \lambda. \tag{6.5}$$

This means that the set of all remaining candidates lies on the curve implicitly defined by the previous equation with the additional conditions that $m_1, m_2 \in [0, 1]$ and $m_1 + m_2 \leq 1$.

An alternative equivalent description of all these points is that we consider $m_1, m_2 \in [0, 1]$ that satisfy $f(m_1) = -f(m_2)$ and $m_1 + m_2 \leq 1$ where

$$f(y) = ey^2 + (2\lambda + \epsilon)y - \lambda. \tag{6.6}$$

The parabola given by $f(\cdot)$ is upward-sloping with vertex $(-\frac{2\lambda+\epsilon}{2e}, -\frac{4\lambda e+(2\lambda+\epsilon)^2}{4e})$ and roots

$$y_1 = -\frac{2\lambda + \epsilon}{2e} + \sqrt{\frac{(2\lambda + \epsilon)^2 + 4\lambda e}{4e^2}} \quad \text{and} \quad y_2 = -\frac{2\lambda + \epsilon}{2e} - \sqrt{\frac{(2\lambda + \epsilon)^2 + 4\lambda e}{4e^2}}. \tag{6.7}$$

Additionally we note that

- f is increasing on $[0, \infty)$ (because the vertex has x-coordinate less than 0)
- $y_2 < 0$ and $y_1 \in (0, \frac{1}{2})$ as $f(0) = -\lambda < 0$ and $f(\frac{1}{2}) = \frac{1}{4}e + \frac{1}{2}\epsilon > 0$.

Therefore, any point (m_1, m_2) that is a stationary point of an equilibrium satisfies $\min\{m_1, m_2\} \in [0, y_1]$. Furthermore, since f is strictly increasing on $[0, 1]$, we have $f(\max\{m_1, m_2\}) \leq \lambda$, which implies that $\max\{m_1, m_2\} \leq \bar{y}$ for $\bar{y} > 0$ such that $f(\bar{y}) = \lambda$. We note that $\bar{y} < 1$ since $f(1) = e + 2\lambda + \epsilon - \lambda > \lambda$. Thus, any value $m_1 \in [0, \bar{y}]$ yields a candidate for a stationary point of the dynamics via $f(m_1) = -f(m_2)$ and $1 = m_1 + m_2 + m_3$. Moreover, this way yields all points which could be stationary points of the dynamics given any strategy π.

As a next step, we show that there is a unique stationary point given the deterministic strategies $\{change\} \times \{stay\}$ and $\{stay\} \times \{change\}$. With this result, we then characterize all points that are stationary points given randomized strategies as well as the strategies for which this happens. Thereafter we combine all results and characterize the set of all stationary equilibria.

From now on we will work under the following assumptions:

- $\lambda > e$: The rate with which agents are pressed out of a good state is smaller than the rate of discovery.
- $b > e$: The rate with which the individual decision is implemented is larger than the rate of contagion effects.

We remark that both assumptions are reasonable and that these conditions are not necessary but yield easier computations, which is sufficient for the purpose of this chapter of illustrating the use of the tools presented in Chapter 5.

Lemma 6.1. *Given the deterministic strategies $\{change\} \times \{stay\}$ and $\{stay\} \times \{change\}$ we find exactly one stationary point of the dynamics given this strategy.*

Proof. We present the argument for the strategy $\{change\} \times \{stay\}$. This suffices to prove the claim because the dynamics given the strategy $\{change\} \times \{stay\}$ are the dynamics given the strategy $\{stay\} \times \{change\}$ if we interchange the states 1 and 2.

Since $\epsilon > 0$, the matrix $Q^{\{change\} \times \{stay\}}(m)$ is irreducible for all $m \in \mathcal{P}(\mathcal{S})$. Thus, by Lemma 5.9 we obtain an explicit characterization of the unique stationary distribution given $Q^{\{change\} \times \{stay\}}(m)$, which is, as before, denoted by $x^{\{change\} \times \{stay\}}(m)$. Moreover, by Remark 5.10, the function $x^{\{change\} \times \{stay\}} : \mathcal{P}(\mathcal{S}) \to \mathcal{P}(\mathcal{S})$ is continuous. Therefore, the Brouwer fixed point theorem yields that there is at least one fixed point of the mapping $x^{\{change\}\times\{stay\}}(\cdot)$. It remains to prove that the fixed point is unique. For this we will, for the sake of readability, write for the rest of the proof $x(m)$ instead of $x^{\{change\}\times\{stay\}}(m)$ as well as $Q(m)$ instead of $Q^{\{change\}\times\{stay\}}(m)$.

We will apply Theorem 5.16. Since $Q(m)$ and hence $x(m)$ are continuously differentiable, we have to verify that the determinant of the matrix

$$M^{\{change\}\times\{stay\}}(m) = \begin{pmatrix} \frac{\partial x_1(m)}{\partial m_1} - 1 & \frac{\partial x_1(m)}{\partial m_2} \\ \frac{\partial x_2(m)}{\partial m_1} & \frac{\partial x_2(m)}{\partial m_2} - 1 \end{pmatrix} - \begin{pmatrix} \frac{\partial x_1(m)}{\partial m_3} & \frac{\partial x_1(m)}{\partial m_3} \\ \frac{\partial x_2(m)}{\partial m_3} & \frac{\partial x_2(m)}{\partial m_3} \end{pmatrix}$$

is non-zero for all $m \in \mathcal{P}(\mathcal{S})$.

In order to compute the Jacobian we will rely on Lemma 5.17, that is we compute

$$\begin{aligned} \frac{\partial x_j(m)}{\partial m_i} = \frac{(-1)^{S+j}}{\det(\tilde{Q}(m))^2} \Bigg(\sum_{k=1}^{S} (-1)^{S+k} \Bigg(\left(\frac{\partial}{\partial m_i} \det(Q(m))'_{jS} \right) \det(Q(m))'_{kS} \\ - \det(Q(m))'_{jS} \left(\frac{\partial}{\partial m_i} \det(Q(m))'_{kS} \right) \Bigg) \Bigg). \end{aligned}$$

For this we first make the following preliminary observations

$$\begin{aligned} \det(Q(m))'_{1S} &= \lambda e m_2 + \lambda \epsilon \\ \det(Q(m))'_{2S} &= -2b\lambda - \lambda e m_1 - \lambda \epsilon \\ \det(Q(m))'_{3S} &= b e m_2 + b\epsilon + e^2 m_1 m_2 + e m_1 \epsilon + e m_2 \epsilon + \epsilon^2 \\ \det(\tilde{Q}(m)) &= \sum_{k=1}^{S} (-1)^{S+k} \det(Q(m))'_{kS} \\ &= e^2 m_1 m_2 + \lambda e m_1 + e m_1 \epsilon + \lambda e m_2 + b e m_2 + e m_2 \epsilon + 2\lambda\epsilon + 2\lambda b + b\epsilon + \epsilon^2, \end{aligned}$$

which yield

$$\begin{aligned} \frac{\partial}{\partial m_1} \det(Q(m))'_{1S} &= 0 & \frac{\partial}{\partial m_2} \det(Q(m))'_{1S} &= \lambda e \\ \frac{\partial}{\partial m_1} \det(Q(m))'_{2S} &= -\lambda e & \frac{\partial}{\partial m_2} \det(Q(m))'_{2S} &= 0 \\ \frac{\partial}{\partial m_1} \det(Q(m))'_{3S} &= e^2 m_2 + e\epsilon & \frac{\partial}{\partial m_2} \det(Q(m))'_{3S} &= be + e\epsilon + e^2 m_1 \end{aligned}$$

and
$$\frac{\partial}{\partial m_3} \det(Q(m))'_{1S} = \frac{\partial}{\partial m_3} \det(Q(m))'_{2S} = \frac{\partial}{\partial m_3} \det(Q(m))'_{3S} = 0,$$

as well as

$$\frac{\partial}{\partial m_1}\det(\tilde{Q}(m)) = e^2 m_2 + \lambda e + e\epsilon$$
$$\frac{\partial}{\partial m_2}\det(\tilde{Q}(m)) = e^2 m_1 + \lambda e + be + e\epsilon.$$

With these preparations we compute the partial derivatives of $x(m)$:

$$\begin{aligned}
\frac{\partial x_1(m)}{\partial m_1} &= \frac{(-1)^{3+1}}{(\det(\tilde{Q}(m)))^2}\big((-1)^{3+1}\cdot 0 + (-1)^{3+2}(0-(\lambda e m_2+\lambda\epsilon)\cdot(-\lambda e))\\
&\quad +(-1)^{3+3}(0-(\lambda e m_2+\lambda\epsilon)\cdot(e\epsilon+e^2m_2))\big)\\
&= \frac{-1}{(\det(\tilde{Q}(m))^2}\big((\lambda e m_2+\lambda\epsilon)(\lambda e+e\epsilon+e^2m_2)\big)\\
\frac{\partial x_1(m)}{\partial m_2} &= \frac{(-1)^{3+1}}{(\det(\tilde{Q}(m)))^2}\big((-1)^{3+1}\cdot 0 + (-1)^{3+2}(\lambda e\cdot(-2b\lambda-\lambda e m_1-\lambda\epsilon)-0)\\
&\quad +(-1)^{3+3}(\lambda e\cdot(bem_2+b\epsilon+e^2m_1m_2+em_1\epsilon+em_2\epsilon+\epsilon^2)\\
&\quad -(\lambda e m_2+\lambda\epsilon)\cdot(be+e\epsilon+e^2m_1))\big)\\
&= \frac{1}{\det(\tilde{Q}(m))^2}\big(2\lambda^2eb+\lambda^2e^2m_1+\lambda^2e\epsilon\big)\\
\frac{\partial x_2(m)}{\partial m_1} &= \frac{(-1)^{3+2}}{(\det(\tilde{Q}(m))^2)}\big((-1)^{3+1}(-\lambda e\cdot(\lambda e m_2+\lambda\epsilon)-0)+(-1)^{3+2}\cdot 0\\
&\quad +(-1)^{3+3}(-\lambda e\cdot(bem_2+b\epsilon+e^2m_1m_2+em_1\epsilon+em_2\epsilon+\epsilon^2)\\
&\quad -(-2b\lambda-\lambda e m_1-\lambda\epsilon)\cdot(e\epsilon+e^2m_2))\big)\\
&= \frac{1}{\det(\tilde{Q}(m))^2}\big((\lambda-b)(\lambda e^2m_2+\lambda e\epsilon)\big)\\
\frac{\partial x_2(m)}{\partial m_2} &= \frac{(-1)^{3+2}}{(\det(\tilde{Q}(m)))^2}\big((-1)^{3+1}(0-(-2b\lambda-\lambda e m_1-\lambda\epsilon)\lambda e)+(-1)^{3+2}\cdot 0\\
&\quad +(-1)^{3+3}(0-(-2b\lambda-\lambda e m_1-\lambda\epsilon)(be+e\epsilon+e^2m_1)\big)\\
&= \frac{-1}{(\det\tilde{Q}(m))^2}\big((2\lambda b+\lambda e m_1+\lambda\epsilon)(\lambda e+be+e\epsilon+e^2m_1\big)
\end{aligned}$$

and

$$\frac{\partial x_1(m)}{\partial m_3} = \frac{\partial x_2(m)}{\partial m_3} = \frac{\partial x_3(m)}{\partial m_3} = 0.$$

In order to apply Theorem 5.16 it is sufficient to prove that the matrix

$$\begin{pmatrix} \frac{\partial x_1(m)}{\partial m_1}-1 & \frac{\partial x_1(m)}{\partial m_2}\\ \frac{\partial x_2(m)}{\partial m_1} & \frac{\partial x_2(m)}{\partial m_2}-1\end{pmatrix}$$

is non-singular for all $m \in \mathcal{P}(\mathcal{S})$. To verify this, we prove that our matrix is strictly diagonal dominant, which means that we prove that

$$\left|\frac{\partial x_1(m)}{\partial m_1} - 1\right| > \left|\frac{\partial x_1(m)}{\partial m_2}\right| \quad \text{and} \quad \left|\frac{\partial x_2(m)}{\partial m_2} - 1\right| > \left|\frac{\partial x_2(m)}{\partial m_1}\right|,$$
$$\text{i.e.} \quad \left|\frac{\partial x_1(m)}{\partial m_1} - 1\right| - \left|\frac{\partial x_1(m)}{\partial m_2}\right| > 0 \quad \text{and} \quad \left|\frac{\partial x_2(m)}{\partial m_2} - 1\right| - \left|\frac{\partial x_2(m)}{\partial m_1}\right| > 0$$

holds. Indeed, we have

$$\begin{aligned}&\left(\left|\frac{\partial x_1(m)}{\partial m_1} - 1\right| - \left|\frac{\partial x_1(m)}{\partial m_2}\right|\right) \cdot \det(\tilde{Q}(m))^2 \\ &= \underbrace{\Big((\lambda e m_2 + \lambda\epsilon)(\lambda e + e\epsilon + e^2 m_2) + \det(\tilde{Q}(m))^2 - (2\lambda^2 eb + \lambda^2 e^2 m_1 + \lambda^2 e\epsilon)\Big)}_{=:g_1(m_1,m_2)}.\end{aligned}$$

And analogously, since

$$\begin{aligned} g_1(0,0) &= \lambda\epsilon(\lambda e + e\epsilon) + (2\lambda\epsilon + 2\lambda b + b\epsilon + \epsilon^2)^2 - 2\lambda^2 eb - \lambda^2 e\epsilon \\ &= \lambda e\epsilon^2 + (2\lambda\epsilon + 2\lambda b + b\epsilon + \epsilon^2)^2 - 2\lambda e^2 b \\ &> 2(\lambda b)^2 - 2\lambda e^2 b \geq 0 \end{aligned}$$

and

$$\begin{aligned} \frac{\partial g_1(m_1,m_2)}{\partial m_1} &= \left(\frac{\partial}{\partial m_1}\det\left(\tilde{Q}(m)\right)\right) \cdot 2 \cdot \det\left(\tilde{Q}(m)\right) - \lambda^2 e^2 \\ &= (e^2 m_2 + \lambda e + e\epsilon) \cdot 2 \cdot \det\left(\tilde{Q}(m)\right) - \lambda^2 e^2 \\ &> \lambda e \cdot 2 \cdot 2\lambda b - \lambda^2 e^2 \geq 0 \\ \frac{\partial g_1(m_1,m_2)}{\partial m_2} &= \lambda e(\lambda e + e\epsilon + e^2 m_2) + e^2(\lambda e m_2 + \lambda\epsilon) \\ &\quad + \left(\frac{\partial}{\partial m_2}\det\left(\tilde{Q}(m)\right)\right) \cdot 2 \cdot \det\left(\tilde{Q}(m)\right) \\ &> 0 \end{aligned}$$

for all $m_1, m_2 \geq 0$, we obtain that $g_1(m_1, m_2) > 0$ for all $m_1, m_2 \geq 0$.

Similarly, we have

$$\begin{aligned}&\left(\left|\frac{\partial x_2(m)}{\partial m_2} - 1\right| - \left|\frac{\partial x_1(m)}{m_2}\right|\right) \cdot \det(\tilde{Q}(m))^2 \\ &= \underbrace{\Big((2\lambda b + \lambda e m_1 + \lambda\epsilon)(\lambda e + be + e\epsilon + e^2 m_1) + \det(\tilde{Q}(m))^2 - |\lambda - b|(\lambda e^2 m_2 + \lambda e\epsilon)\Big)}_{=:g_2(m_1,m_2)}.\end{aligned}$$

Since

$$
\begin{aligned}
g_2(0,0) &= (2\lambda b + \lambda\epsilon)(\lambda e + be + e\epsilon) + (2\lambda\epsilon + 2\lambda b + b\epsilon + \epsilon^2)^2 - |\lambda - b|\lambda e\epsilon \\
&> 2\lambda be\epsilon + \lambda^2 e\epsilon - |\lambda - b|\lambda e\epsilon \geq 0
\end{aligned}
$$

and

$$
\begin{aligned}
\frac{\partial g_2(m_1,m_2)}{\partial m_1} &= \lambda e(\lambda e + be + e\epsilon + e^2 m_1) + e^2(2\lambda b + \lambda e m_1 + \lambda\epsilon) \\
&\quad + \left(\frac{\partial}{\partial m_1}\det\left(\tilde{Q}(m)\right)\right)\cdot 2\det\left(\tilde{Q}(m)\right) > 0 \\
\frac{\partial g_2(m_1,m_2)}{\partial m_2} &= \left(\frac{\partial}{\partial m_2}\det\left(\tilde{Q}(m)\right)\right)\cdot 2\cdot\det\left(\tilde{Q}(m)\right) - |\lambda - b|\lambda e^2 \\
&= (e^2 m_1 + \lambda e + be + e\epsilon)\cdot 2\cdot\det\left(\tilde{Q}(m)\right) - |\lambda - b|\lambda e^2 \\
&> (\lambda e + be)\cdot 2\cdot 2\lambda b - |\lambda - b|\lambda e^2 \\
&= 4\lambda^2 be + 4b^2\lambda e - |\lambda - b|\lambda e^2 \geq 0
\end{aligned}
$$

for all $m_1, m_2 \geq 0$, we obtain that $g_2(m_1,m_2) > 0$ for all $m_1, m_2 \geq 0$. This shows that the matrix $M(m)$ is indeed strictly diagonal dominant, which by Remark 5.18 yields that the matrix $M(m)$ is non-singular. Thus, Theorem 5.16 yields that there is a unique stationary point given $Q^{\{change\}\times\{stay\}}(\cdot)$. □

We now characterize all stationary points given randomized strategies. For this we will write, for the sake of readability, *cs* instead of $\{change\}\times\{stay\}$ and *sc* instead of $\{stay\}\times\{change\}$. By the previously discussed result, we have that there is a unique stationary point m^{sc} given the strategy *sc* and a unique stationary point m^{cs} given the strategy *cs*. These satisfy that $(m_1^{cs}, m_2^{cs}, m_3^{cs}) = (m_2^{sc}, m_1^{sc}, m_3^{sc})$ because $m^T Q^{sc}(m) = 0$ is the same equation as $m^T Q^{cs}(m) = 0$ with the roles of m_1 and m_2 interchanged. Furthermore, we have $m_1^{sc} > m_2^{sc} = m_1^{cs}$ since

$$
\begin{aligned}
bm_2^{sc} + (em_2^{sc} + \epsilon)m_2^{sc} = \lambda m_3^{sc} \Leftrightarrow & (em_2^{sc} + \epsilon)m_2^{sc} = \lambda m_3^{sc} - bm_2^{sc} \\
& (em_1^{sc} + \epsilon)m_1^{sc} = \lambda m_3^{sc} + bm_2^{sc}.
\end{aligned}
$$

With these preparations we can fully characterize the stationary points given randomized strategies:

Lemma 6.2. *Let* (m_1, m_2, m_3) *be a solution of* (6.4) *and* $1 = m_1 + m_2 + m_3$.

- *If* $m_1 \in (m_1^{cs}, m_1^{sc})$*, then* (m_1, m_2, m_3) *is a stationary point given* $Q^\pi(\cdot)$ *if and only if*

$$
\pi_{1,change} = g(m_1, m_2) + \pi_{2,change}\frac{m_2}{m_1} \tag{6.8}
$$

$$\pi_{2,change} \in \left[\max\left\{0, -\frac{g(m_1,m_2)m_1}{m_2}\right\}, \min\left\{1, \frac{(1-g(m_1,m_2))m_1}{m_2}\right\}\right] \tag{6.9}$$

with

$$g(m_1,m_2) = \frac{1}{2b}\left(\frac{(em_2+\epsilon)m_2}{m_1} - (em_1+\epsilon)\right).$$

In particular, there are always infinitely many such strategies.

- *If* $m_1 = m_1^{sc}$ *or* $m_1 = m_1^{cs}$*, respectively, then* (m_1, m_2, m_3) *is a stationary point given* $Q^\pi(\cdot)$ *if and only if* $\pi = \{stay\} \times \{change\}$ *or* $\pi = \{change\} \times \{stay\}$*, respectively.*

- *If* $m_1 \notin [m_1^{cs}, m_1^{sc}]$*, then there is no strategy* $\pi \in \Pi^s$ *such that* (m_1, m_2, m_3) *is a stationary point given* $Q^\pi(\cdot)$*.*

Proof. Assume that m satisfies (6.4) and $1 = m_1 + m_2 + m_3$. Then any strategy that satisfies the balance equation $m^T Q^\pi(m) = 0$ satisfies

$$\begin{cases}(em_1+\epsilon)m_1 + \pi_{1,change}bm_1 = \lambda m_3 + \pi_{2,change}bm_2 \\ (em_2+\epsilon)m_2 + \pi_{2,change}bm_2 = \lambda m_3 + \pi_{1,change}bm_1\end{cases},$$

which by the cut criterion (6.4) is equivalent to

$$\begin{cases}\pi_{1,change}bm_1 = \frac{1}{2}((em_2+\epsilon)m_2 - (em_1+\epsilon)m_1) + \pi_{2,change}bm_2 \\ \pi_{2,change}bm_2 = \frac{1}{2}((em_1+\epsilon)m_1 - (em_2+\epsilon)m_2) + \pi_{1,change}bm_1\end{cases}. \tag{6.10}$$

Thus, $\pi_{1,change}$ and $\pi_{2,change}$ must satisfy

$$\pi_{1,change} = \underbrace{\frac{1}{2b}\left(\frac{(em_2+\epsilon)m_2}{m_1} - (em_1+\epsilon)\right)}_{=g(m_1,m_2)} + \pi_{2,change}\frac{m_2}{m_1}. \tag{6.11}$$

Let us first collect some basic properties of $g(\cdot,\cdot)$: It holds that $g(m_1^{sc}, m_2^{sc}) = -\frac{m_2^{sc}}{m_1^{sc}} < 0$ since

$$0 = g(m_1^{sc}, m_2^{sc}) + \frac{m_2^{sc}}{m_1^{sc}}$$

and that $g(m_1^{cs}, m_2^{cs}) = 1$ since

$$1 = g(m_1^{cs}, m_2^{cs}) + 0 \cdot \frac{m_2^{cs}}{m_1^{cs}}.$$

Furthermore, we directly see that $g : (0,\infty) \times (0,\infty) \to \mathbb{R}$ is strictly increasing in m_2 and strictly decreasing in m_1. This and $g(m_1, m_1) = 0$ for all $m_1 \in (0,1]$ further implies that $g(m_1,m_2) < 0$ yields that $m_2 < m_1$ and that $g(m_1,m_2) > 0$ yields that $m_2 > m_1$.

Let us now assume that $m_1 \in (m_1^{cs}, m_1^{sc})$. We show that $\pi_{1,change}$ and $\pi_{2,change}$ defined by (6.8) and (6.9) indeed lie in $[0,1]$: If $g(m_1,m_2) \leq 0$ we have that $m_1 \leq m_2$ and thus

$$\pi_{2,change} \in \left[\frac{-g(m_1,m_2)m_1}{m_2}, 1\right] \subseteq [0,1].$$

Moreover,

$$\pi_{1,change} \leq g(m_1,m_2) + \frac{m_2}{m_1} \leq 1$$

since we have that $m_1 < m_1^{sc}$ and by (6.4) also $m_2 > m_2^{sc}$, which yields that

$$-\frac{m_2}{m_1} < -\frac{m_2^{sc}}{m_1^{sc}} = g(m_1^{sc}, m_2^{sc}) < g(m_1,m_2).$$

And additionally,

$$\pi_{1,change} \geq g(m_1,m_2) + \frac{-g(m_1,m_2)m_1}{m_2} \cdot \frac{m_2}{m_1} = 0,$$

which yields that for all strategies π given by (6.8) and (6.9) the point (m_1,m_2,m_3) is a stationary point given $Q^\pi(\cdot)$. Since

$$\frac{-g(m_1,m_2)m_1}{m_2} < \frac{m_2^{sc}}{m_1^{sc}}\frac{m_1}{m_2} < 1$$

by previous observations and since $m_1 < m_1^{sc}$ and $m_2 > m_2^{sc}$, we see that there are infinitely many possible choices for $\pi_{2,change}$, which shows that there are infinitely many strategies π such that (m_1,m_2,m_3) is a stationary point given $Q^\pi(\cdot)$.

Similarly, if $g(m_1,m_2) \geq 0$ we see that

$$\pi_{2,change} \in \left[0, \frac{(1-g(m_1,m_2))m_1}{m_2}\right] \subseteq [0,1]$$

with $(1-g(m_1,m_2))m_1/m_2 > 0$ since $g(m_1,m_2) < 1$ and $m_2 > m_1$ by previous observations. Thus,

$$\pi_{1,change} \leq g(m_1,m_2) + \frac{(1-g(m_1,m_2))m_1}{m_2} \cdot \frac{m_2}{m_1} = 1$$

and

$$\pi_{1,change} \geq g(m_1,m_2) \geq 0.$$

Furthermore, for $m_1 \in [m_1^{cs}, m_1^{sc}]$ and for any strategy π that is not defined by (6.8) and (6.9) the point m is no stationary point given $Q^\pi(\cdot)$: If $g(m_1,m_2) \leq 0$, then

$$\pi_{2,change} < -g(m_1,m_2)\frac{m_1}{m_2} \overset{(6.11)}{\Rightarrow} \pi_{1,change} < 0$$

and if $g(m_1, m_2) \geq 0$, then

$$\pi_{2,change} > \frac{(1 - g(m_1, m_2))m_1}{m_2} \overset{(6.11)}{\Rightarrow} \pi_{1,change} > 1.$$

If $m_1 > m_1^{sc}$, then by (6.4) it holds that $m_2 < m_2^{sc}$, which yields

$$\pi_{1,change} \overset{(6.11)}{=} g(m_1, m_2) + \pi_{2,change}\frac{m_2}{m_1} < g(m_1^{sc}, m_2^{sc}) + \pi_{2,change}\frac{m_2^{sc}}{m_1^{sc}} \leq 0,$$

which implies that we cannot find any strategy, which yields that m is a stationary point given $Q^\pi(\cdot)$. Similarly, if $m_1 < m_1^{cs}$, then $m_2 > m_2^{cs}$, which yields that

$$\pi_{1,change} \overset{(6.11)}{=} g(m_1, m_2) + \pi_{2,change}\frac{m_2}{m_1} > g(m_1^{cs}, m_2^{cs}) + \pi_{2,change}\frac{m_2^{cs}}{m_1^{cs}} \geq 1,$$

which implies that we cannot find any strategy, which yields that m is a stationary point given $Q^\pi(\cdot)$. □

As a final preparation we prove the following lemma:

Lemma 6.3. *There is no stationary mean field equilibrium in Opt($\{change\} \times \{stay\}$).*

Proof. All points in the optimality set Opt($\{change\} \times \{stay\}$) satisfy

$$\frac{1}{c}m_2 + \frac{1-c}{ce}(\beta + \epsilon + 2\lambda) < m_1,$$

which especially means that $m_2 < m_1$. However, by Theorem 5.11 the stationary point $x^{\{change\}\times\{stay\}}(m)$ of $Q^{\{change\}\times\{stay\}}(m)$ is the normalized vector

$$\begin{pmatrix} \lambda e m_2 + \lambda\epsilon \\ \lambda e m_1 + 2b\lambda + \lambda\epsilon \\ ebm_2 + e^2 m_1 m_2 + b\epsilon + em_1\epsilon + em_2\epsilon + \epsilon^2 \end{pmatrix},$$

which especially satisfies that $m_1 > m_2$ implies $x_1^{\{change\}\times\{stay\}}(m) < x_2^{\{change\}\times\{stay\}}(m)$. Thus, there is no fixed point of $x^{\{change\}\times\{stay\}}(\cdot)$ in Opt($\{change\} \times \{stay\}$). □

With this lemma we completely characterize the set of all mean field equilibria: For this it is central to investigate at which points the linear function $m_2 \mapsto m_1 := \frac{1}{c}m_2 + \frac{1-c}{c}\frac{\beta+\epsilon+2\lambda}{e}$ and the implicit curve defined by

$$\{(m_1, m_2) \in [0, \bar{y}]^2 : f(m_1) = -f(m_2)\}, \tag{6.12}$$

where $f : \mathbb{R} \to \mathbb{R}$ is defined by (6.6), intersect. Since $f(\cdot)$ is increasing on $[0, 1]$, the curve is decreasing in m_2 on $[0, 1]$. Since at the same time the linear function is increasing in m_2, the curve and the linear function intersect at most once in $[0, 1]$.

Furthermore, we obtain that the point of intersection $z = (z_1, z_2)$ satisfies $z_2 \leq y_1$, with y_1 being defined as the positive root of $f(\cdot)$ (see equation (6.7)). Indeed, the function $c \mapsto \frac{1}{c}m_2 + \frac{1-c}{c}\frac{\beta+\epsilon+2\lambda}{e}$ is decreasing in c on $(0, 1]$ because

$$\frac{\partial}{\partial c}\left(\frac{1}{c}m_2 + \frac{1-c}{c}\frac{\beta+\epsilon+2\lambda}{e}\right) = -\frac{1}{c^2}\left(m_2 + \frac{\beta+\epsilon+2\lambda}{e}\right) < 0.$$

Therefore, the linear functions $m_2 \mapsto \frac{1}{c}m_2 + \frac{1-c}{c}\frac{\beta+\epsilon+2\lambda}{e}$, $c \in (0, 1)$, dominate $m_2 \mapsto m_2$. Thus, $z_2 \leq y_1$ and $z_1 \geq y_1 > m_1^{cs}$.

If $z_1 < m_1^{sc}$, then by Lemma 6.2 there are infinitely many randomized equilibria and m^{sc} is not an equilibrium. Indeed, $z_2 > m_2^{sc}$ by (6.5) and this yields

$$m_1^{sc} > z_1 = \frac{1}{c}z_2 + \frac{1-c}{c}\frac{\beta+\epsilon+2\lambda}{e} > \frac{1}{c}m_2^{sc} + \frac{1-c}{c}\frac{\beta+\epsilon+2\lambda}{e},$$

which shows that m^{sc} does not lie in $\mathrm{Opt}(\{stay\} \times \{change\})$. If $z_1 \geq m_1^{sc}$, then $m^{sc} \in \mathrm{Opt}(\{stay\} \times \{change\})$ is by Lemma 6.2 the only mean field equilibrium.

Let us illustrate this method for four different parameter sets that cover all cases that are structurally different. For this we plot the necessary functions in the Figures 6.2 - 6.5 as follows: The implicit curve defined by (6.12) is plotted in red. The implicit curve describing the third equation that a stationary point given the strategy $\{change\}\times\{stay\}$ has to satisfy besides $1 = m_1 + m_2 + m_3$ and the equation derived from the cut criterion is given by (6.11) with $\pi_{1,change} = 1$ and $\pi_{2,change} = 0$ and is plotted in green. Similarly, the implicit curve given by the third equation that a stationary point given the strategy $\{stay\}\times\{change\}$ has to satisfy besides $1 = m_1+m_2+m_3$ and the equation derived from the cut criterion is given by (6.11) with $\pi_{1,change} = 0$ and $\pi_{2,change} = 1$ and is plotted in blue. The black line collects all those points for which both strategies $\{change\}\times\{stay\}$ and $\{stay\} \times \{change\}$ are optimal, which is given by $m_1 = \frac{1}{c}m_2 + \frac{(1-c)}{ce}(\beta + \epsilon + 2\lambda)$. From the previous discussion, we note that we are in the first place interested whether the m_1-value of the intersection of the black line and the red curve lie before or after the m_1-value of the intersection of the red and the blue curve because this yields the number and position of the equilibria.

Consider the following numerical example: Let $\lambda = 0.5$, $e = 0.4$, $b = 0.7$, $\beta = 0.8$ and $\epsilon = 0.01$. We will vary c in order to obtain all the cases that can occur. First, let $c = 0.80$ and consider the plot in Figure 6.2. The black line does not intersect $[0, 1]^2$, which yields that there is no mean field equilibrium with a mixed equilibrium strategy, instead we obtain that $(m^{sc}, \{stay\} \times \{change\})$ is the unique mean field equilibrium.

Second, let $c = 0.84$ and consider the plot in Figure 6.3. Here we obtain that the black line intersects $[0,1]^2$, but does not intersect the red curve, which again yields that there is no mean field equilibrium with a mixed equilibrium strategy, instead we obtain that $(m^{sc}, \{stay\} \times \{change\})$ is the unique mean field equilibrium.

Third, let $c = 0.88$, the plot in Figure 6.4 reveals that there is an intersection of the red and the black line in $[0,1]^2$ and that the point of intersection satisfies $z_1 < m_1^{\text{sc}}$, where the latter coordinate is given by the intersection of the blue curve and the black line. Thus, we again obtain that there is a unique stationary mean field equilibrium with equilibrium distribution m^{cs} and equilibrium strategy $\{stay\} \times \{change\}$.

Finally, let $c = 0.92$, we see in the plot in Figure 6.5 that there is an intersection of the red and the black line in $[0,1]^2$ and that the point of intersection satisfies $z_1 > m_1^{\text{sc}}$, where the latter coordinate is again given by the intersection of the blue curve and the black line. Thus, we obtain that there are infinitely many stationary equilibria with equilibrium distribution (z_1, z_2), where the equilibrium strategies are given by (6.8) and (6.9).

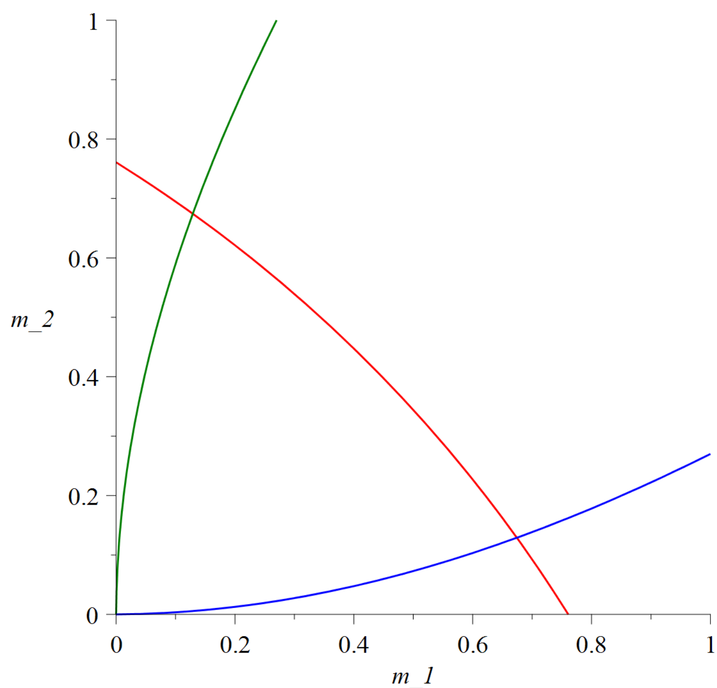

Figure 6.2.: The plot for the case $c = 0.80$.

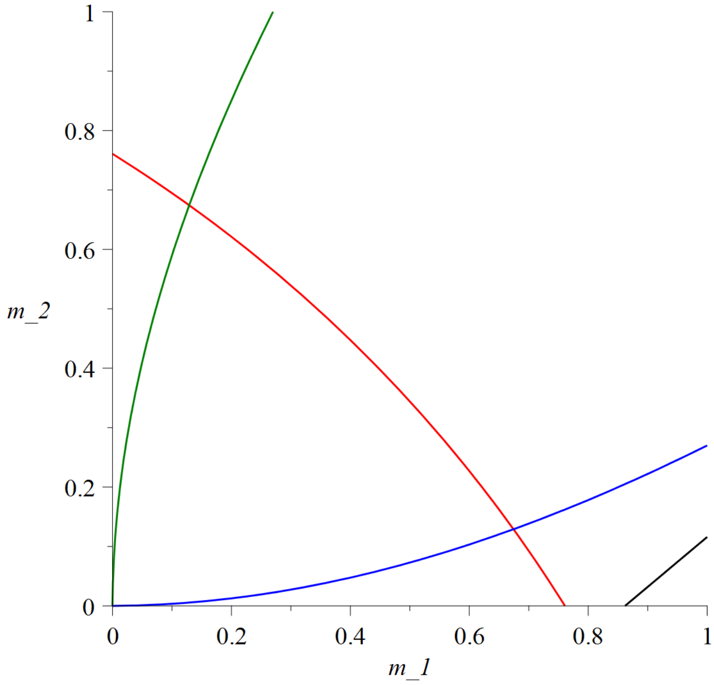

Figure 6.3.: The plot for the case $c = 0.84$.

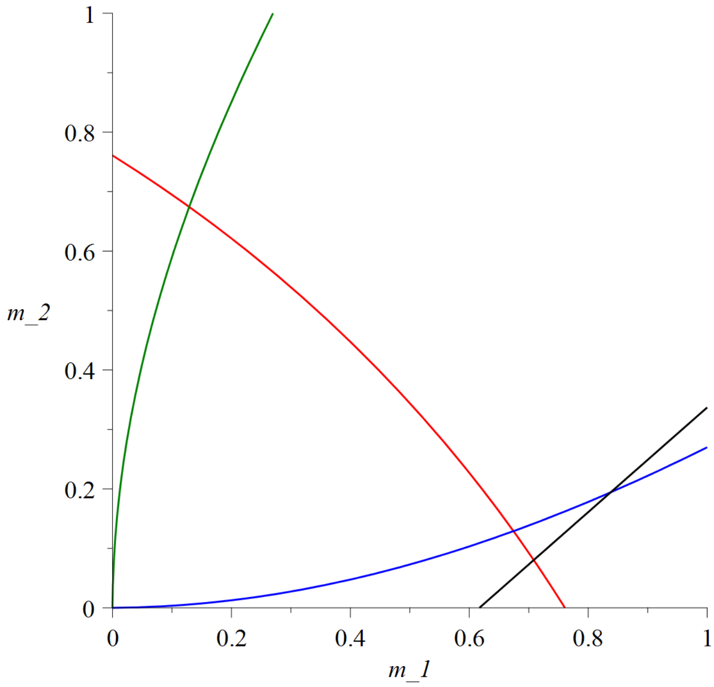

Figure 6.4.: The plot for the case $c = 0.88$.

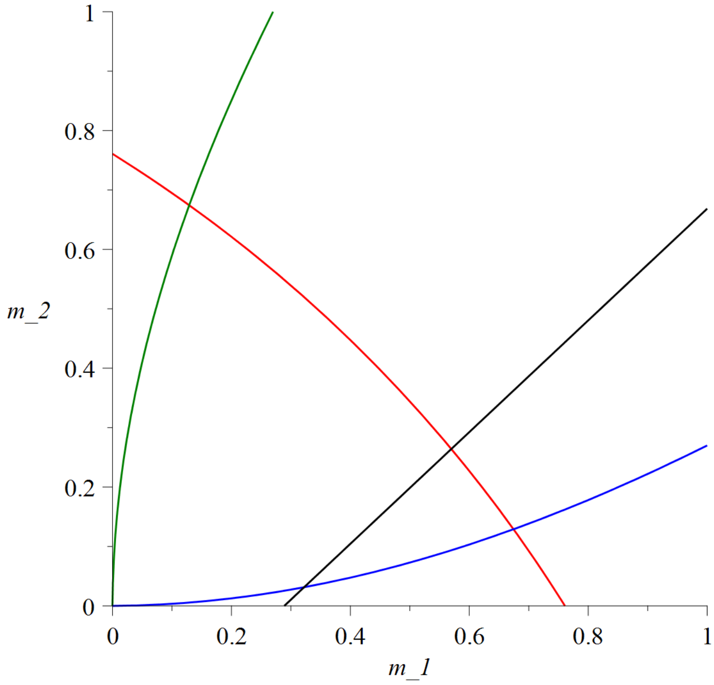

Figure 6.5.: The plot for the case $c = 0.92$.

7. Stability with Respect to Model Perturbations

In this chapter we study whether stationary mean field equilibria are stable against slight perturbations of the model. This is a classical question in the context of normal form games (van Damme (1991)) and has recently been discussed also for Markov perfect equilibria of dynamic stochastic games (Doraszelski and Escobar (2010)). However, up to the knowledge of the author, such investigations have not been performed for mean field games. The closest stability discussions for diffusion-based models are presented in Guéant (2009b) where perturbations of the initial and terminal conditions are considered. For other stability considerations we can only cite Adlakha and Johari (2013) and Light and Weintraub (2018) where comparative statics results for oblivious equilibria for a parameter for which the transition probabilities (and payoffs) are monotonic have been considered.

We will first define the concept of essential equilibria for mean field games in a way similar to the definitions for Nash equilibria of normal form games (Wen-Tsün and Jia-He, 1962) and for Markov perfect equilibria of dynamic stochastic games (Doraszelski and Escobar, 2010). Furthermore, we introduce the notion of strongly stable equilibria in a similar fashion as Kojima et al. (1985). As we have seen in Chapter 5, finding stationary mean field equilibria consists of two steps, namely the optimal control problem and the fixed point problem. Thus, we will also first analyse the two problems individually and, thereafter, draw the conclusions regarding the stability of stationary mean field equilibria.

The main results are the following: We will see that there is a close relation of the stability of stationary points of the dynamics and essential fixed points. Since there are only partial characterization results for essential fixed points, it is clear that we do not obtain a simple full characterization of all those dynamics with only essential stationary points. However, we still obtain several characterization results similar to those for essential fixed points. With respect to the stability of equilibria, we obtain that equilibria are not stable whenever we randomize over more than two deterministic strategies or whenever the deterministic equilibrium strategy is not the unique optimal strategy for the equilibrium distribution. If the deterministic equilibrium strategy is the unique optimal strategy and if the distribution is an essential stationary point of the dynamics, then the equilibrium is essential. For equilibria that have an equilibrium strategy that is a convex combination of two deterministic strategies we show that they are essential if we face constant irreducible dynamics and furthermore in any neighbourhood of the equilibrium distribution we find two points such that each of the two deterministic strategies is the unique optimal strategies for one of them point. Strong stability is only possible if we restrict our attention to the class of games with constant dynamics, which means that we allow only for perturbed games with constant dynamics. Furthermore, we will show that the set of all games of which all stationary equilibria are essential is a residual subset of the set of all games. This means that essentiality is a (topological) generic property.

In Section 7.1 we will set up the notation and introduce the relevant notions. Thereafter, in Section 7.2 we will first analyse the stability of the stationary points of the dynamics and in Section 7.3 we will discuss how optimality sets change under small perturbations. In Section 7.4 we will then put the results together and derive the statements regarding essential and strongly stable equilibria. Finally, in Section 7.5 we discuss to what extend essentiality is a generic property.

7.1. Definitions and Notation

The notion of essential equilibria has been introduced in Wen-Tsün and Jia-He (1962). It originally addresses the question whether an equilibrium is stable to slight payoff perturbations by calling an equilibrium ϕ of a game G essential if any game that is near G (with respect to some metric) has an equilibrium near ϕ (with respect to some metric). As mentioned in Doraszelski and Escobar (2010) this question is becoming more and more important as nowadays many models are estimated using econometric techniques and, thus, the estimated game characteristics will be erroneous. Therefore, it is central to understand whether equilibria "vanish" under small perturbations.

In van Damme (1991) the case of normal form games is discussed. For given fixed action sets the game is then identified by a vector in $\mathbb{R}^{n \cdot m}$ collecting the payoffs of all n player given each of the m action profiles. Similarly, equilibria can be identified as vectors in $\prod_{i=1}^{n} \mathcal{P}(A_i)$ with A_i being the action sets of all players. To make the notion of closeness precise we equip $\mathbb{R}^{n \cdot m}$ and $\mathbb{R}^m$, respectively, with the standard topology. Also in the context of Markov perfect equilibria in stochastic dynamic games the games and the equilibria are identified by vectors in $\mathbb{R}^n$ (Doraszelski and Escobar, 2010). The only difference is that the vectors describing the game now collect rewards, discount factors (which might differ among individuals in their model) and transition rates. In both settings, an equilibrium ϕ of a game G is essential if for any $\epsilon > 0$ there exists a $\delta > 0$ such that all games G' in the δ-neighbourhood $N_\delta(G)$ of G have an equilibrium in the ϵ-neighbourhood $N_\epsilon(\phi)$ of ϕ.

The concept of essential equilibria has been refined by Kojima et al. (1985) and has also been formulated for Markov perfect equilibria in stochastic dynamic games in Doraszelski and Escobar (2010). Instead of just requiring that an equilibrium does not vanish under a small perturbation it is now required that the equilibrium "changes continuously and uniquely for any small perturbation to payoff functions of players" (Kojima et al., 1985, p.650). This equilibrium concept is motivated from the desire not only to know that an equilibrium will be sustained given small perturbations, but also to know whether bifurcation happens or not. Moreover, for strongly stable equilibria we can understand the effects of small parameter changes on the equilibria, which allows for comparative statics analysis. Formally an equilibrium ϕ of a normal form game G is strongly stable if there exist neighbourhoods U of G and V of ϕ such that

- for all $G' \in U$ exactly one equilibrium of G' lies in the set V,
- the mapping $s : U \to V$ that maps to each game G' the unique equilibrium in V is continuous.

In our setting given fixed sets of states and actions, the game is described not only by the rewards, but also by the transition rate matrices. Thus, there are two sources for perturbations: Both transition rate matrices and reward functions can be inaccurate. From the economic motivation that essential equilibria should be stable against slight errors in model set-up/estimation, we allow simultaneous perturbations of both, the transition rate matrices and the rewards. A difficulty compared to the case of normal form games/Markov perfect equilibria, is that instead of reward vectors, the transition rate matrices and the rewards are functions. Formally, we define essential and strongly stable equilibria as follows:

Fix $\mathcal{S}$, $\mathcal{A}$ and β for the rest of the chapter. We remark that we could also allow for perturbations in β and the results could also be formulated in this setting, but this would increase the complexity of the considered problem without yielding any additional insights.

Let $\mathcal{G}$ be the set of all games described in Chapter 3 and denote a particular game as a pair (Q, r). These games satisfy, as described in Chapter 3, that $Q(\cdot) = ((Q_{ija}(\cdot))_{i,j\in\mathcal{S}})_{a\in\mathcal{A}}$ is a Lipschitz continuous family of functions that satisfy for all $a \in \mathcal{A}$ and $m \in \mathcal{P}(\mathcal{S})$, that the matrix $Q_{\cdot\cdot a}(m)$ is a generator and $r(\cdot)$ is a function in $\mathcal{C}(\mathcal{P}(\mathcal{S}), \mathbb{R}^{S\times A})$. We equip $\mathcal{G}$ with the following metric

$$d((Q,r),(Q',r')) = \sup_{i,j\in\mathcal{S},a\in\mathcal{A},m\in\mathcal{P}(\mathcal{S})} |Q_{ija}(m) - Q'_{ija}(m)| + \sup_{i\in\mathcal{S},a\in\mathcal{A},m\in\mathcal{P}(\mathcal{S})} |r_{ia}(m) - r'_{ia}(m)|,$$

although other choices might be possible (and do not necessarily yield to equivalent results). Our choice simplifies computations at some steps due to its close relation to the maximum norm. Moreover, it is sensible in the light of quantifying errors in model prediction. We emphasize that we denote by d the distance of two games as well as a deterministic strategy. However, it will always be clear from the context what d represents at that point.

Any stationary equilibrium (m, π) lies in the candidate set $\mathcal{P}(\mathcal{S}) \times \mathcal{P}(\mathcal{A})^S$, which is a subset of $\mathbb{R}^{S(1+A)}$ and which we equip with the maximum norm. Furthermore, we define the map $\text{SMFE} : \mathcal{G} \to 2^{\mathcal{P}(\mathcal{S})\times\mathcal{P}(\mathcal{A})^S}$ as the map that maps every game (Q, r) to the set of all stationary mean field equilibria.

In order to define and analyse essential and strongly stable equilibria, we introduce the following general notation for an arbitrary metric space (X, d): For $A \subseteq X$, $\epsilon > 0$ we define the set

$$N_\epsilon(A) = \{x \in X | \exists y \in A : d(x, y) < \epsilon\},$$

consisting of all points $x \in X$ that have a distance less than ϵ from some point in A. Moreover, for $A, B \subseteq X$ we define the Hausdorff metric

$$H(A, B) := \inf\{\epsilon > 0 : A \subseteq N_\epsilon(B) \quad \text{and} \quad B \subseteq N_\epsilon(A)\}$$

on the space 2^X. With these preparations we are able to define the central notions of this chapter:

Definition 7.1. Let (m, π) be a stationary mean field equilibrium of the game (Q, r). We say that (m, π) is *essential* if for every $\epsilon > 0$ there exists a $\delta > 0$ such that for all games $(Q', r') \in N_\delta(Q, r)$ we have $(m, \pi) \in N_\epsilon(\text{SMFE}(Q', r'))$, which means that there is a stationary mean field equilibrium of (Q', r') that lies in $N_\epsilon(m, \pi)$.

Definition 7.2. We say that a stationary mean field equilibrium (m, π) of a game (Q, r) is *strongly stable* if there is a neighbourhood U of (Q, r) and a neighbourhood V of (m, π) such that

- $|\text{SMFE}(Q', r') \cap V| = 1$ for all $(Q', r') \in U$,
- the mapping $s : U \to V$ defined by $\{s(Q', r')\} = \text{SMFE}(Q', r') \cap V$ is continuous.

We remark that, as in the theory of normal form games (Corollary 2.4.5 in van Damme (1991)), a strongly stable equilibrium is automatically essential.

In economic applications there might be some knowledge regarding the considered models, for example one might know that one faces constant dynamics or that the recurrence classes are fixed. In this case we sometimes consider the question whether equilibria of games (Q, r) in some set $\tilde{\mathcal{G}} \subseteq \mathcal{G}$ are essential (strongly stable) with respect to changes $(Q', r') \in N_\delta(Q, r) \cap \tilde{\mathcal{G}}$. In this case we say that an equilibrium (m, π) of (Q, r) is essential (strongly stable) in $\tilde{\mathcal{G}}$. In particular, we will consider the set of games where $Q(\cdot)$ is constant in m and, moreover, irreducible for all $d \in D^s$, which we will denote by $\mathcal{G}^{ci}$.

7.2. Essential Stationary Points of the Dynamics

In this section we analyse what happens to stationary points given a transition rate matrix function $Q(\cdot)$, which are exactly the solutions of $0 = m^T Q(m)$, under small perturbations. In order to simplify notation, we introduce the space TR of all standard transition rate matrix functions, which are all those functions $Q : \mathcal{P}(\mathcal{S}) \to \mathbb{R}^{S \times S}$ such that $Q(m)$ is a generator matrix for all $m \in \mathcal{P}(\mathcal{S})$ and $Q_{ij}(\cdot)$ is Lipschitz continuous for all $i, j \in \mathcal{S}$. Moreover, we equip the space TR with the metric

$$\hat{d}(Q, Q') = \sup_{i,j \in \mathcal{S}, m \in \mathcal{P}(\mathcal{S})} |Q_{ij}(m) - Q'_{ij}(m)|.$$

We can characterize stationary points given $Q(\cdot)$ equivalently as the fixed points of the map $x(\cdot) : \mathcal{P}(\mathcal{S}) \to 2^{\mathcal{P}(\mathcal{S})}$ that maps to every m the set of all stationary points of the dynamics given $Q(m)$, which are the solutions of $0 = x^T Q(m)$. By Iosifescu (1980, Theorem 5.3) this map has non-empty and convex values, and by continuity of $Q_{ij}(\cdot)$ it has a closed graph. Thus, by Kakutani's fixed point theorem (Lemma 4.15) such a map has a fixed point, which is in turn a stationary point of the dynamics given $Q(\cdot)$.

Let m be a stationary point given $Q(\cdot)$. We say that m is an *essential stationary point of the dynamics* $Q(\cdot)$ if for all $\epsilon > 0$ there is a $\delta > 0$ such that for all $Q' \in N_\delta(Q)$ there is a stationary point given $Q'(\cdot)$ in $N_\epsilon(m)$. We say that m is a *strongly stable stationary point of the dynamics* $Q(\cdot)$ if the map STAT : TR $\to 2^{\mathcal{P}(\mathcal{S})}$, which maps $Q(\cdot) \in$ TR to the set of stationary points given $Q(\cdot)$, satisfies for a neighbourhood U of $Q(\cdot)$ and a neighbourhood V of m that

- $|\text{STAT}(Q') \cap V| = 1$ for all $Q' \in U$,
- the function $s : U \to V$ given by $\{s(Q')\} = STAT(Q') \cap V$ is continuous.

In general, we cannot hope for a full characterization of all those maps of which all or some of the stationary points of the dynamics are essential. The reason is that finding stationary points of the dynamics is closely linked to finding essential fixed points and that characterizing essential fixed points is hard. Indeed, Del Prete et al. (1999) characterize essential fixed points and maps on the space $[0, 1]$. However, similar characterizations for more complex spaces like $[0, 1]^n$, $n > 1$ are, up to the knowledge of the author, still open. Thus, also this section only collects partial results. More precisely, we will obtain the following:

We will first assume that $Q(m)$ is irreducible for all $m \in \mathcal{P}(\mathcal{S})$ and show that whenever the map $x(\cdot)$ has an essential fixed point m in the sense of Fort (1950), then this point is also an essential point of the dynamics given $Q(\cdot)$. We use this to apply the known characterizations of essential fixed points to obtain criteria for essential stationary points given irreducible dynamics. We obtain that the unique stationary point given a transition rate matrix function that satisfies the irreducibility assumption and has only one stationary point is essential. Moreover, we obtain that at least one stationary point given a transition rate matrix function that satisfies the irreducibility assumption and has a totally disconnected set of stationary points is essential.

As a second step we will show for general dynamics that the set of all transition rate matrices for which all stationary points are essential is residual in TR. This shows that there are "many" transition rate matrix functions with only essential stationary points. Moreover, we obtain that any transition rate matrix function can be approximated arbitrarily close by a transition rate matrix function with only essential stationary points. Furthermore, we obtain as corollary of the proof that the stationary point given any (also reducible) transition rate matrix function with exactly one stationary point is essential.

The third part of the section is devoted to analysing constant transition rates: We first show that any stationary point of the dynamics given a constant transition rate matrix with more than one recurrence class is not essential. However, if the perturbations yield the same recurrence classes, then the stationary points will again be essential. Finally, we prove that if the transition rate matrix is irreducible then the stationary point of the dynamics is even strongly stable in the class of all constant transition rate matrices. Moreover, we provide an example showing that this does not hold for the class of all transition rate matrix functions.

7.2.1. Relation to Essential Fixed Points

In this subsection we assume that $Q(m)$ is irreducible for all $m \in \mathcal{P}(\mathcal{S})$ and prove that essential fixed points of $x(\cdot)$ are indeed essential stationary points of the dynamics $Q(\cdot)$. Afterwards we establish easy to verify criteria for essential stationary points.

The notion of essential fixed points has been introduced in Fort (1950): Let (X, d) be a compact metric space such that every continuous mapping from X into X has a fixed point. Equip the set $\mathcal{C}(X, X)$ of all continuous functions from X into X with the metric

$$\rho(f, g) = \sup_{x \in X} d(f(x), g(x)).$$

We say that x is an *essential fixed point* of f if for each neighbourhood U of x there is an $\delta > 0$ such that any $g \in \mathcal{C}(X, X)$ satisfying $\rho(f, g) < \delta$ has a fixed point in U.

With these definitions we formulate the desired result:

Theorem 7.3. *Let $Q(\cdot) \in TR$. Assume that $Q(m)$ is irreducible for all $m \in \mathcal{P}(\mathcal{S})$ and assume that m is an essential fixed point of the map $x(\cdot)$, then m is an essential stationary point of the dynamics $Q(\cdot)$.*

We want to prove that given any $\epsilon > 0$ there is a $\delta > 0$ such that for all $Q' \in N_\delta(Q)$ there is a stationary point given $Q'(\cdot)$ in $N_\epsilon(m)$. Since m itself is an essential fixed point of $x(\cdot)$, there is some $\tilde{\delta} > 0$ such that whenever $\rho(x, x') < \tilde{\delta}$ a fixed point of $x'(\cdot)$ lies in $N_\epsilon(m)$. Thus, we aim to find a $\delta > 0$ such that $Q' \in N_\delta(Q)$ implies $\rho(x, x') < \tilde{\delta}$. The proof of this statement relies on the fact that whenever $Q(m)$ is irreducible for all $m \in \mathcal{P}(\mathcal{S})$ then we can compute the stationary distribution given $Q(m)$ as the unique solution of

$$\tilde{Q}(m)x(m) = (0, \dots, 0, 1)^T$$

with $\tilde{Q}(m)$ being the matrix $Q(m)^T$, where the last row is replaced by $(1, \dots, 1)$ (see Lemma 5.9). This and the following theorem from numerical linear algebra found in Wendland (2018, Theorem 2.43) then yields the desired result:

Lemma 7.4. *Denote by $||\cdot||$ the maximum norm on $\mathbb{R}^n$ as well as the associated matrix norm, which is the row sum norm. Let $A \in \mathbb{R}^{n\times n}$ be an invertible matrix, $\Delta A \in \mathbb{R}^{n\times n}$ be a perturbation matrix that satisfies $||A^{-1}||||\Delta A|| < 1$ and let $b, \Delta b \in \mathbb{R}^n$ be vectors with $b \neq 0$. Then the estimate*

$$\frac{||\Delta x||}{||x||} \leq \kappa(A) \left(1 - \kappa(A)\frac{||\Delta A||}{||A||}\right)^{-1} \left(\frac{||\Delta b||}{||b||} + \frac{||\Delta A||}{||A||}\right)$$

holds, where $\kappa(A) := ||A|| \cdot ||A^{-1}||$ is the conditioning number of the matrix A.

Proof of Theorem 7.3 . As explained before we have to find for any $\tilde{\delta} > 0$ an $\delta > 0$ such that $Q' \in N_\delta(Q)$ implies $\rho(x, x') < \tilde{\delta}$.

We start with some preliminary observations:

- $S\hat{d}(Q, Q') \geq ||Q(m) - Q'(m)||$ for all $m \in \mathcal{P}(\mathcal{S})$ since

$$\begin{aligned}\hat{d}(Q, Q') &= \sup_{m\in\mathcal{P}(\mathcal{S})} \max_{i,j\in\mathcal{S}} |Q_{ij}(m) - Q'_{ij}(m)| \geq \sup_{m\in\mathcal{P}(\mathcal{S}), i\in\mathcal{S}} \frac{1}{S}\sum_{j=1}^{S} |Q_{ij}(m) - Q'_{ij}(m)| \\ &= \sup_{m\in\mathcal{P}(\mathcal{S})} \frac{1}{S}||Q(m) - Q'(m)||.\end{aligned}$$

- $\inf_{m\in\mathcal{P}(\mathcal{S})} ||\tilde{Q}(m)|| =: L_1 > 0$ as $\tilde{Q}(m)$ is an invertible matrix for all $m \in \mathcal{P}(\mathcal{S})$ (see Lemma 5.9) and the set $\mathcal{P}(\mathcal{S})$ is compact.

- $||x(m)|| \leq 1$ for all $m \in \mathcal{P}(\mathcal{S})$ as $x(m) \in \mathcal{P}(\mathcal{S})$.

- $\sup_{m\in\mathcal{P}(\mathcal{S})} \kappa(\tilde{Q}(m)) =: L_2 < \infty$ since

$$\sup_{m\in\mathcal{P}(\mathcal{S})} \kappa(\tilde{Q}(m)) \leq \sup_{m\in\mathcal{P}(\mathcal{S})} ||\tilde{Q}(m)|| \cdot \sup_{m\in\mathcal{P}(\mathcal{S})} ||(\tilde{Q}(m))^{-1}|| =: L_2,$$

where the finiteness of L_2 follows from the fact that $\mathcal{P}(\mathcal{S})$ is compact, $||\cdot||$ as well as $(\cdot)^{-1}$ are continuous functions and $\tilde{Q}(m)$ is invertible for all $m \in \mathcal{P}(\mathcal{S})$ (see Lemma 5.9).

Let

$$0 < \delta < \min\left\{\frac{1}{2S||(\tilde{Q}(m))^{-1}||}, \frac{\tilde{\delta}L_1}{2SL_2}\right\} \tag{7.1}$$

and choose $Q' \in N_\delta(Q)$, then

$$\begin{aligned}
&\left(1 - \kappa(\tilde{Q}(m))\frac{||\tilde{Q}(m) - (\tilde{Q}')(m)||}{||\tilde{Q}(m)||}\right)^{-1} \\
&= \left(1 - ||(\tilde{Q}(m))^{-1}|| \cdot ||\tilde{Q}(m) - \tilde{Q}'(m)||\right)^{-1} \\
&< \left(1 - ||(\tilde{Q}(m))^{-1}|| \cdot \frac{1}{2||(\tilde{Q}(m))^{-1}||}\right)^{-1} \\
&= \left(1 - \frac{1}{2}\right)^{-1} = 2.
\end{aligned}$$

Furthermore, since $x(m)$ is the unique solution of

$$\tilde{Q}(m)x(m) = (0, \dots, 0, 1)^T,$$

and $\delta < \frac{1}{2S||(\tilde{Q}(m))^{-1}||}$ holds, we can compute the difference between $x(\cdot)$ and $x'(\cdot)$ by using Lemma 7.4:

$$\begin{aligned}
||x(m) - x'(m)|| &\leq ||x(m)||\kappa(\tilde{Q}(m))\left(1 - \kappa(\tilde{Q}(m))\frac{||\tilde{Q}(m) - \tilde{Q}'(m)||}{||\tilde{Q}(m)||}\right)^{-1} \\
&\quad \cdot \frac{||\tilde{Q}(m) - \tilde{Q}'(m)||}{||\tilde{Q}(m)||} \\
&\leq 1 \cdot L_2 \cdot 2 \cdot \frac{S\hat{d}(Q, Q')}{L_1} = \frac{2SL_2}{L_1}\hat{d}(Q, Q') \leq \tilde{\delta}.
\end{aligned}$$

□

Fort (1950) provides some results to characterize essential fixed points, which by the previous result carry over to our setting: Namely, Theorem 2 in Fort (1950) states that for any map with a single fixed point this point is essential. Theorem 3 in Fort (1950) states that for any space X that is additionally an n-dimensional manifold either with or without a boundary and any function f whose set of fixed points is totally disconnected (that is the connected components are one-point sets) has at least one essential fixed point. In our case, this together with the fact that if m is a fixed point of $x(\cdot)$ then m is a stationary point given $Q(\cdot)$ yields the following corollaries:

Corollary 7.5. *Let $Q(\cdot) \in TR$. Assume that $Q(m)$ is irreducible for any $m \in \mathcal{P}(\mathcal{S})$ and assume further that there is only one stationary point given the dynamics $Q(\cdot)$. Then this point is an essential stationary point.*

Corollary 7.6. *Let $Q(\cdot) \in TR$. Assume that $Q(m)$ is irreducible for any $m \in \mathcal{P}(\mathcal{S})$ and assume further that the set of stationary points of dynamics is totally disconnected. Then at least one stationary point of the dynamics is an essential stationary point.*

Example. In the two examples with irreducible dynamics presented in Chapter 6 we directly obtain that several stationary points are essential: In the consumer choice model of Section 6.1 we obtain that the stationary point given any transition rate Q^π is essential since given constant dynamics with an irreducible transition rate matrix we always have a unique fixed point. In the consumer choice model of Section 6.3 we obtain that the unique stationary points given the strategies $\{stay\} \times \{change\}$ and $\{change\} \times \{stay\}$ are essential. Furthermore, we can show that given any other strategy π any equilibrium candidates has to be a zero of a fourth order polynomial, which means that the set of all stationary points is totally disconnected. Thus, there is an essential stationary point given $Q^\pi(\cdot)$. △

Remark 7.7. The notion of essential fixed points has also been introduced for set-valued maps (see Jia-He (1962)), where the the metric ρ is replaced with

$$\rho(f, g) = \sup_{x \in X} H(f(x), g(x)).$$

We remark that it is no longer possible to consider the set of all upper semi-continuous set-valued maps with non-empty, compact values on a compact space X because not all maps in this set have a fixed point. Thus, one restricts to the set $\tilde{\mathcal{C}}$ of all maps with at least one fixed point and defines essential fixed points as those points x for which any perturbation in $\tilde{\mathcal{C}}$ has a fixed point close to m.

However, we cannot directly prove that any essential fixed point of $x(\cdot)$ is also an essential stationary point of the dynamics since it is no longer true that for any $\tilde{\delta} > 0$ there is a $\delta > 0$ such that if $\hat{d}(Q, Q') < \delta$ then $\rho(x, x') < \tilde{\delta}$: Assume that Q is a generator with two recurrence classes, and let i be an element of the first recurrence class and j be an element of the second recurrence class. Then $Q' = Q + A$ with A being the matrix with the following non-zero entries $A_{ii} = A_{jj} = -\delta$ and $A_{ij} = A_{ji} = \delta$ is a slight perturbation of Q and has a unique stationary point. As Q itself has infinitely many stationary points that form a convex set, the maximum distance of stationary points given Q and Q' will be uniformly (in δ) bounded from below. Thus, for small enough $\tilde{\delta} > 0$ the statement will not be true.

Nonetheless, a result similar to Theorem 7.3 would not be helpful in any case. Indeed, the only characterization result for essential fixed points of set-valued maps is that a unique fixed point is an essential fixed point. However, this is a corollary of the main theorem of the next subsection. △

7.2.2. Having Essential Stationary Points of the Dynamics is a Generic Property

Although finding essential stationary points of the dynamics is difficult and we cannot characterize many of them, we will prove that having only essential stationary points of the dynamics is a generic property in the topological sense. More precisely, we prove that the set of all $Q(\cdot) \in TR$ having only essential stationary points is a residual set, which is a countable intersection of sets with a dense interior. Moreover, we obtain that the set is dense, which means that for any transition rate matrix function we find arbitrarily close transition rate matrix functions with only essential stationary points. The proof moreover yields as a corollary that every unique stationary point is essential.

Theorem 7.8. *The set of all transition rate matrices $Q(\cdot)$ for which all stationary points of the dynamics given $Q(\cdot)$ are essential is a residual set in the set of all transition rate matrix functions TR. In particular, this set lies dense in the set of all transition rate matrix functions TR.*

The statement is similar to Theorem 1 in Fort (1950), which states that the set of functions with only essential fixed points is dense in $\mathcal{C}(X, X)$, and to Theorem 1 in Jia-He (1962), which states that the set of set-valued maps with only essential stationary points is dense in $\tilde{\mathcal{C}}$. However, we cannot apply these result to prove our statement directly since we cannot establish a sensible relation between $x(\cdot)$ and $Q(\cdot)$. We will rather rely on similar ideas as used in Fort (1950). More precisely, our proof will work as follows: Let STAT : TR $\to 2^{\mathcal{P}(\mathcal{S})}$ be the map that maps $Q(\cdot)$ to the stationary points of the dynamics (which in turn are the fixed points of $x(\cdot)$). Then, we show that STAT$(\cdot)$ is upper semi-continuous (in a similar fashion to Wehausen (1945)) and that each stationary point of $Q(\cdot)$ is essential if and only if $Q(\cdot)$ is a point of continuity (with respect to the metric $H(\cdot, \cdot)$) of STAT$(\cdot)$. Thereafter, we apply a theorem found in Fort (1949), which states the set of all points of continuity indeed forms a G_δ-residual set, which is a countable intersection of dense open sets. By Baire's Theorem (Kelley, 1955, p.200-201) this set is dense.

Proof of Theorem 7.8. We will follow the previously sketched programme by first showing that the function STAT$(\cdot)$ is upper semi-continuous: Let $Q(\cdot) \in$ TR and $\epsilon > 0$. We want to find a $\delta > 0$ such that for all $Q'(\cdot) \in$ TR that satisfy $\hat{d}(Q, Q') < \delta$ we have STAT$(Q') \subseteq N_\epsilon($STAT$(Q))$.

Assume by way of contradiction that such a $\delta > 0$ does not exist. In this case we find sequences $(Q^n)_{n\in\mathbb{N}}$ and $(m^n)_{n\in\mathbb{N}}$ such that for all $n \in \mathbb{N}$

- $\hat{d}(Q^n, Q) < \frac{1}{n}$

- $m^n \in \mathrm{STAT}(Q^n)$ because there is a fixed point of $x^n(\cdot)$ by Kakutani's fixed point theorem (Lemma 4.15) (the map $x^n(\cdot)$ is a set-valued map with non-empty, compact and convex values and a closed graph),

- $m^n \notin N_\epsilon(\mathrm{STAT}(Q))$.

Since $(m^n)_{n\in\mathbb{N}}$ is a sequence in the compact space $\mathcal{P}(\mathcal{S})$, it has a converging subsequence $(m^{n_k})_{k\in\mathbb{N}}$. Denote its limit by m. As $Q^n(\cdot)$ converges uniformly to $Q(\cdot)$ we obtain that

$$\sum_{i\in\mathcal{S}} Q_{ij}^{n_k}(m^{n_k})m_i^{n_k} = 0 \quad \forall j \in \mathcal{S}, k \in \mathbb{N} \Rightarrow \sum_{i\in\mathcal{S}} Q_{ij}(m)m_i = 0 \quad \forall j \in \mathcal{S}, \tag{7.2}$$

which yields that $m \in \mathrm{STAT}(Q)$. At the same time $m \notin N_\epsilon(\mathrm{STAT}(Q))$ because it is the limit of a subsequence of (m^{n_k}) that, by construction, completely lies in the closed set $\mathcal{P}(\mathcal{S}) \setminus N_\epsilon(\mathrm{STAT}(Q))$, a contraction. Thus, the set-valued map $\mathrm{STAT}(\cdot)$ is upper semi-continuous.

As a second step we show that each stationary point given $Q(\cdot) \in \mathrm{TR}$ is essential if and only if $Q(\cdot)$ is a point of continuity of $\mathrm{STAT}(\cdot)$ with respect to the metric $H(\cdot,\cdot)$: Assume that each stationary point given $Q(\cdot)$ is essential. Let $\epsilon > 0$. Since all fixed points are essential, we find for each $m \in \mathrm{STAT}(Q)$ a neighbourhood V_m of $Q(\cdot)$ such that all $Q'(\cdot) \in V_m$ have a stationary point in the $\frac{\epsilon}{2}$ neighbourhood of m.

We note that $\mathrm{STAT}(Q)$ is compact: Indeed, let $(m^n)_{n\in\mathbb{N}}$ be a sequence in $\mathrm{STAT}(Q) \subseteq \mathcal{P}(\mathcal{S})$. Since $\mathcal{P}(\mathcal{S})$ is compact, we find a converging subsequence $(m^{n_k})_{k\in\mathbb{N}}$ with limit $m \in \mathcal{P}(\mathcal{S})$. By continuity of $Q(\cdot)$ we obtain that

$$0 = \sum_{i\in\mathcal{S}} Q_{ij}(m^{n_k})m_i^{n_k} \to \sum_{i\in\mathcal{S}} Q_{ij}(m)m_i \quad \text{for all} \quad j \in \mathcal{S},$$

which yields that $m \in \mathrm{STAT}(Q)$. Thus, by compactness, there exists a finite set of points $\{m_1, \ldots, m_n\} \subseteq \mathrm{STAT}(Q)$ such that each point in $\mathrm{STAT}(Q)$ lies within the $\frac{\epsilon}{2}$ neighbourhood of one of the m_k.

Let $V = \bigcap_{k=1}^n V_{m_k}$. Then $Q'(\cdot) \in V$ implies that $\mathrm{STAT}(Q) \subseteq N_\epsilon(\mathrm{STAT}(Q'))$. Indeed, because $Q'(\cdot)$ has a stationary point at most $\frac{\epsilon}{2}$ away from each of $\{m_1, \ldots, m_n\}$ and these points are at most $\frac{\epsilon}{2}$ away from any other point in $\mathrm{STAT}(Q)$, we obtain by the triangle inequality that each stationary point of the dynamics given $Q(\cdot)$ lies at most ϵ away from some stationary point of the dynamics given $Q'(\cdot)$. Together with the first part of the proof this implies that $\mathrm{STAT}(\cdot)$ is continuous at $Q(\cdot)$.

Assume that $\mathrm{STAT}(\cdot)$ is continuous at $Q(\cdot)$. Let $m \in \mathrm{STAT}(Q)$ and let $\epsilon > 0$. We show that there is a $\delta > 0$ such that for all $Q' \in N_\delta(Q)$ there is a stationary point of $Q'(\cdot)$ in $N_\epsilon(m)$. Since $\mathrm{STAT}(\cdot)$ is continuous, we can choose $\delta > 0$ such that $\hat{d}(Q, Q') < \delta$ implies that $H(\mathrm{STAT}(Q), \mathrm{STAT}(Q')) < \epsilon$. This implies for all $Q' \in \mathrm{TR}$ satisfying $\hat{d}(Q, Q') < \delta$

that $\text{STAT}(Q) \subseteq N_\epsilon(\text{STAT}(Q'))$. Thus, there is a stationary point of the dynamics given $Q'(\cdot)$ that lies in $N_\epsilon(m)$, which means that m is essential.

The theorem found in Fort (1950, Section 7) states that the points of continuity of any upper semi-continuous function $F(\cdot)$ ranging from a topological space to the power set of a separable metric space equipped with metric $H(\cdot, \cdot)$ is a G_δ-residual set in the topological space. This theorem yields that the set of all $Q(\cdot)$ for which all stationary points of the dynamics are essential is a G_δ-residual set. Since TR is a closed subset of a complete metric space (see Munkres (2014, §43)), we obtain by Baire's Theorem (Kelley, 1955, p.200) that this set is dense in TR. □

The continuity properties of the map $\text{STAT}(\cdot)$ allow us, as in Fort (1950) and Jia-He (1962), to generalize Corollary 7.5 for arbitrary (also reducible) transition rate matrix functions $Q(\cdot) \in \text{TR}$:

Lemma 7.9. *Let $Q(\cdot) \in TR$ be such that there is exactly one stationary point given the dynamics $Q(\cdot)$. Then this point is essential.*

Proof. Let $\epsilon > 0$. Since $\text{STAT}(\cdot)$ is upper semi-continuous we find a $\delta > 0$ such that for any $Q'(\cdot) \in N_\delta(Q)$ we have $\text{STAT}(Q') \subseteq N_\epsilon(\text{STAT}(Q))$. Since $\text{STAT}(Q)$ is a singleton and by Theorem 4.14 the set $\text{STAT}(Q')$ is non-empty, we obtain that $\text{STAT}(Q) \subseteq N_\epsilon(\text{STAT}(Q'))$. Thus, $\text{STAT}(\cdot)$ is continuous at $Q(\cdot)$, which in turn implies that the unique stationary point given $Q(\cdot)$ is essential. □

7.2.3. Essential Stationary Points in the Case of Constant Dynamics

If we consider constant transition rate matrices the investigation of essential and strongly stable stationary points in this class of games becomes simpler by using results for classical Markov chains. More precisely, we show that the stationary points given any non-irreducible generator matrix are not essential even in the class of all games with constant dynamics. If we only allow for perturbations that again yield constant dynamics with the same recurrence classes, then we establish that all stationary points are essential. Such restrictions might be sensible if we know the context of our model. The stationary points given irreducible generator matrices are, as we have seen, essential. Now, we furthermore prove that if we restrict the perturbations to those that yield another constant dynamics generator matrix then these stationary points are even strongly stable.

We start with the discussion of constant dynamics given by non-irreducible generator matrices: Let Q be a transition rate matrix with r recurrence classes $R_1, \dots, R_r$. Then by Iosifescu (1980, Theorem 5.3) any stationary distribution is a convex combination of the unique stationary distributions given each recurrence class. Let us denote the stationary distribution given recurrence class R_i by x^{R_i}. We remark that $x_l^{R_i}$ is positive if and only if $l \in R_i$. With these preparations we obtain:

Lemma 7.10. *Let Q be any non-irreducible generator matrix. Then for any stationary point x there is an $\epsilon > 0$ such that for any $\delta > 0$ there is a generator matrix Q' such that $d(Q, Q') < \delta$ and for any stationary point x' of Q' we have $\hat{d}(x, x') > \epsilon$.*

Proof. Since Q is non-irreducible there are at least two recurrence classes R_1 and R_2. Assume without loss of generality that there is some $k \in R_1$ such that $x_k > 0$. Let $i \in R_1$ and $j \in R_2$ and consider $Q' = Q + A$ with $A \in \mathbb{R}^{S \times S}$ having the following non-zero entries $A_{ij} = \frac{1}{2}\delta$ and $A_{ii} = -\frac{1}{2}\delta$. Then it holds that $\hat{d}(Q, Q') = \frac{1}{2}\delta$. Furthermore, any stationary distribution given Q' assigns zero mass to any state from R_1 since these states are transient given Q'. Thus, the distance of any stationary point given Q' and x is at least $\max_{k \in R_1} x_k$, thus for $\epsilon = \frac{1}{2} \max_{k \in R_1} x_k$ we find a perturbation $Q' \in N_\delta(Q)$ such that there is no stationary point given Q' that lies in $N_\epsilon(x)$. □

If we know that our model has constant dynamics with certain fixed recurrence classes, that is we are only uncertain about the size of the non-zero transition rates, then we obtain a positive result:

Lemma 7.11. *Let Q be a non-irreducible generator matrix and let x be a stationary point given Q. Then for any $\epsilon > 0$ we find a $\delta > 0$ such that any generator matrix Q', which has the same recurrence classes as Q and furthermore satisfies $\hat{d}(Q, Q') < \delta$ has a stationary point in $N_\epsilon(x)$.*

Proof. Let $R_1, \ldots, R_r$ be the recurrence classes of Q. Since the recurrence classes of Q' are still $R_1, \ldots, R_r$, we obtain that any stationary point given Q' is a convex combination of the unique stationary points $(x')^{R_i}$ for the distinct recurrence classes $R_1, \ldots, R_r$. These stationary points are uniquely determined through the restriction of Q' to R_i, which we will denote by $(Q')^{R_i}$. Since by Corollary 7.5 the points x^{R_i} $(i = 1, \ldots, r)$ are essential, we obtain that given $\epsilon > 0$ there is a $\delta_i > 0$ such that if $(Q')^{R_i} \in N_{\delta_i}(Q^{R_i})$ then $(x')^{R_i} \in N_\epsilon(x^{R_i})$. We set $\delta = \min_{i=1,\ldots,r} \delta_i$ and let $Q' \in N_\delta(Q)$. Thus, we have that $(Q')^{R_i} \in N_\delta(Q^{R_i}) \subseteq N_{\delta_i}(Q^{R_i})$.

We remind ourselves that the stationary point x can be written as $x = \sum_{i=1}^r \lambda_i x^{R_i}$ with $\lambda_i \geq 0$ for all $i = 1, \ldots, r$ and $\sum_{i=1}^r \lambda_i = 1$. Thus, we obtain for $x' = \sum_{i=1}^r \lambda_i (x')^{R_i}$ that

$$d(x, x') = \left\| \sum_{i=1}^r \lambda_i x^{R_i} - \sum_{i=1}^r \lambda_i (x')^{R_i} \right\|_\infty \leq \max_{k \in \mathcal{S}} \sum_{i=1}^r \left| \lambda_i (x_k^{R_i} - (x')_k^{R_i} \right|$$

$$= \sum_{i=1}^{r} \lambda_i \max_{k \in \mathcal{S}} \left| x_k^{R_i} - (x')_k^{R_i} \right| < \sum_{i=1}^{r} \lambda_i \epsilon = \epsilon,$$

which yields that $x' \in N_\epsilon(x)$. □

For irreducible constant dynamics we directly infer from Corollary 7.5 that the stationary point of the dynamics is essential even in the larger class TR consisting of all standard nonlinear generators. Moreover, we show that whenever we restrict our attention to perturbations which again yield constant dynamics this point is indeed strongly stable:

Lemma 7.12. *Let TR^{ci} be the set of all transition rate matrices that are irreducible. Then the map $STAT\colon TR^{ci} \to \mathcal{P}(\mathcal{S})$ which maps every irreducible transition rate matrix Q to the unique stationary point given Q is continuous.*

The proof relies on arguments presented in Heidergott et al. (2010a,b) to show that if transition rate matrices depend continuously on a single real parameter then the stationary distributions are also a continuous function of this parameter. The main tool for the proof is an update formula, which allows to compute the stationary distribution given an irreducible and positive recurrent generator matrix Q' from the irreducible and positive recurrent generator matrix Q as well as its stationary distribution x and the deviation matrix D. We remark that in the case of finite state spaces, which we consider here, any irreducible generator matrix is automatically positive recurrent. More precisely, we have

$$X' = X + X'(Q' - Q)D \tag{7.3}$$

with X and X' being the ergodic matrices given Q and Q', respectively, which are the matrices whose rows equal the unique stationary distribution. The matrix D is the deviation matrix defined by

$$D_{ij} = \int_0^\infty (p_{ij}(t) - x_j)\mathrm{d}t, \quad i, j \in \mathcal{S}$$

and it exists for finite state spaces (see Section 4 of Heidergott et al. (2010a)).

Proof. It suffices to prove that for any sequence of irreducible generator matrices $(Q^n)_{n\in\mathbb{N}}$ converging to Q we also have that the stationary distributions x^n given Q^n converge to the stationary distribution x given Q. Without loss of generality assume that

$$|Q_{ij}^n - Q_{ij}| < \frac{1}{n},$$

else we chose a suitable subsequence.

The update formula (7.3) can be rewritten as

$$X^n - X = X^n(Q^n - Q)D = X(Q^n - Q)D + (X^n - X)(Q^n - Q)D.$$

We show that the right-hand side converges to zero, which implies that the stationary distributions also converge. More precisely, we show that the norm of the summands converges to zero: We note that

$$|((Q^n - Q)D)_{ij}| \leq \sum_{k \in \mathcal{S}} |Q^n_{ik} - Q_{ik}| \cdot |D_{kj}| = \frac{1}{n} \sum_{k \in \mathcal{S}} |D_{kj}| \to 0,$$

which implies that $|(X(Q^n - Q)D)_{ij}| \to 0$. This yields that $||X(Q^n - Q)D|| \to 0$, which in turn implies that $X(Q^n - Q)D \to 0$.

Similarly, we obtain that

$$\begin{aligned} |((X^n - X)(Q^n - Q)D)_{ij}| &\leq |(X^n(Q^n - Q)D)_{ij}| + |(X(Q^n - Q)D)_{ij}| \\ &\leq \sum_{l,k \in \mathcal{S}} X^n_{il} \cdot |Q^n_{lk} - Q_{lk}| \cdot |D_{kj}| + \sum_{l,k \in \mathcal{S}} X_{il} \cdot |Q^n_{lk} - Q^n_{lk}| \cdot |D_{kj}| \\ &= \frac{1}{n} S x^n_i \sum_{k \in \mathcal{S}} |D_{kj}| + \frac{1}{n} S x_i \sum_{k \in \mathcal{S}} |D_{kj}| \leq \frac{2}{n} S \sum_{k \in \mathcal{S}} |D_{kj}| \to 0, \end{aligned}$$

which also implies that $(X^n - X)(Q^n - Q)D$ converges to the zero matrix, which yields the desired result. □

Remark 7.13. In general it does not hold that stationary points given constant generators are strongly stable in the class TR. This in particular means that strong stability is indeed a stricter notion than essentiality. Indeed, the statement is intuitive if we remind ourselves of the degrees of freedom we have for choosing the dynamics – we can basically bend the dynamics as we like. To illustrate this consider

$$Q = \begin{pmatrix} -1 & 1 \\ 1 & -1 \end{pmatrix},$$

which yields that the unique stationary point given Q is $(\frac{1}{2}, \frac{1}{2})$. When we perturb the dynamics close to this point by $a(m_1)$ as follows

$$Q(m) = \begin{pmatrix} -(1 + a(m_1))^2 & (1 + a(m_1))^2 \\ 1 & -1 \end{pmatrix},$$

the stationary point $x(m)$ given m is given by

$$\left(\frac{1}{1 + (1 + a(m_1))^2}, \frac{(1 + a(m_1))^2}{1 + (1 + a(m_1)^2)} \right).$$

If the stationary point given Q would be strongly stable then for some neighbourhood $N_\delta(Q)$ of Q and some neighbourhood $N_\epsilon(m)$ of m any function $a(m_1)$ such that $Q'(\cdot)$ lies in $N_\delta(Q)$ should yield to a unique stationary point in $N_\epsilon(m)$. However, it always possible to construct perturbations that violate this:

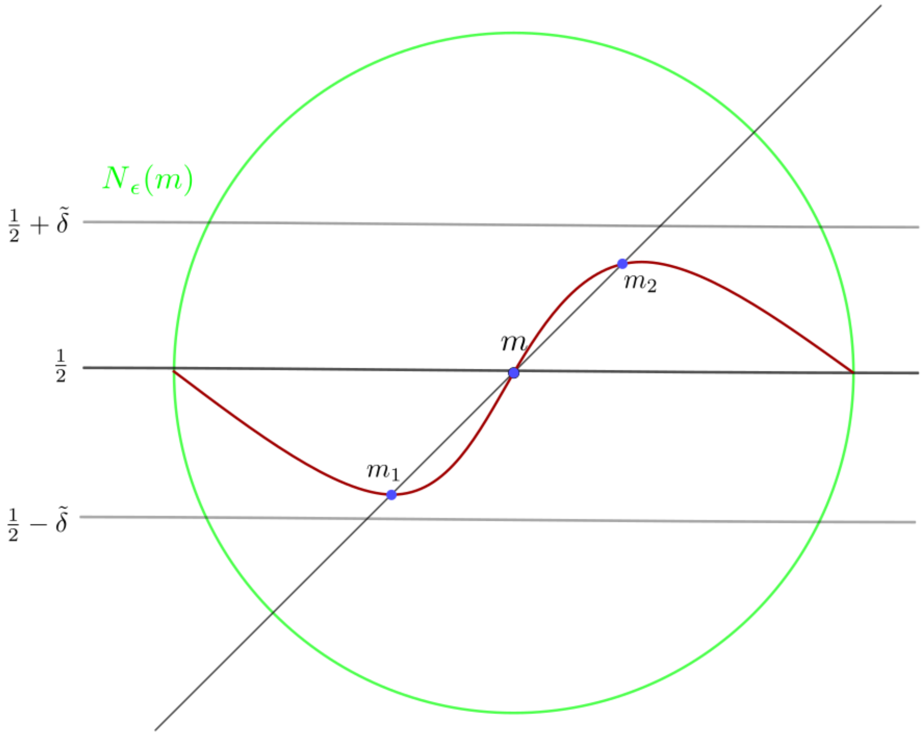

Figure 7.1.: A possible choice of $x_1(m_1)$ that yields more than one stationary point of the dynamics: The green circle depicts the ϵ-neighbourhood of m, the grey vertical lines depict the corridor in which we can choose $x(m)$. The red line is one possible choice for $x(m)$ and any intersection of this line with the identity map is a stationary point of the dynamics. Whenever there is more than one intersection in $N_\epsilon(m)$ strong stability is violated.

Let $\delta, \epsilon > 0$ and choose $\tilde{\delta} \in (0, 0.5)$ such that $\delta > \frac{2\tilde{\delta}}{\frac{1}{2}-\tilde{\delta}}$. Then choose a function $x_1(m_1)$ as in Figure 7.1. This function describes the unique stationary point given the transition rate matrix function $Q(m_1)$ of the previously defined type with

$$a(m_1) = \sqrt{\frac{1}{x_1(m_1)} - 1} - 1.$$

However, the perturbed matrix $Q'(\cdot)$ is by choice of $\tilde{\delta}$ at most $\delta > 0$ away from Q and at the same time there are three stationary points given $Q'(\cdot)$ in $N_\epsilon(m)$. This shows that $\left(\frac{1}{2}, \frac{1}{2}\right)$ is not a strongly stable stationary point of Q. $\triangle$

7.3. Changes of Optimality Sets given Slight Perturbations

In this section, we analyse in how far the optimality of strategies changes under small perturbations. We are mainly interested in two questions: First, under which conditions can we ensure that a strategy d that is optimal for m is also optimal for m in a slightly perturbed game (Q', r')? Second, under which conditions can we infer from the fact that a strategy d is not optimal for m in the game (Q, r) that d is also not optimal for the slightly perturbed game (Q', r')? We will provide general conditions such that the desired properties hold and demonstrate why these conditions are indeed necessary for such a result to hold.

The main tool for our investigations is the following lemma regarding the distance of the expected discounted reward given the original game and the slightly perturbed game. The proof relies on the fact that the expected discounted reward given a deterministic stationary strategy $d \in D^s$ and a population distribution $m \in \mathcal{P}(\mathcal{S})$ is the unique solution $V^d(m)$ of

$$(\beta I - Q^d(m))V^d(m) = r^d(m),$$

which allows us to use Lemma 7.4.

Lemma 7.14. *Let $(Q, r) \in \mathcal{G}$ be a game and let $\gamma > 0$. Then there exists a $\delta > 0$ such that for all $(Q', r') \in N_\delta(Q, r)$ and all deterministic stationary strategies $d \in D^s$ the distance of the expected discounted reward given d at any point $m \in \mathcal{P}(\mathcal{S})$ in the game (Q, r) and in the perturbed game (Q', r') is at most γ.*

Proof. We assume without loss of generality that all reward vectors are positive: Assume that this is not the case and consider new reward vectors $\bar{r}_{ia}(m) = r_{ia}(m) + c$, with

$$c = \left(- \inf_{m \in \mathcal{P}(\mathcal{S}), i \in \mathcal{S}, a \in \mathcal{A}} r_{ia}(m) + 1 \right).$$

Then the expected discounted reward for any deterministic stationary strategy $d \in D^s$ given $(\bar{r}_{ia}(\cdot))_{i \in \mathcal{S}, a \in \mathcal{A}}$ is $V^d(\cdot) + \bar{V}$ with $\bar{V} = \frac{1}{\beta} c(1, \ldots, 1)^T$. Indeed, since $Q^d(m)$ has row sum zero and $\bar{V}$ is a constant vector, we have

$$\begin{aligned}(\beta I - Q^d(m))(V^d(m) + \bar{V}) &= r^d(m) + (\beta I - Q^d(m))\bar{V} \\ &= r^d(m) + c \cdot (1, \ldots, 1)^T = \bar{r}^d(m),\end{aligned}$$

which yields by Lemma 2.2 that $V^d(\cdot) + \bar{V}$ is the expected discounted reward given strategy d. Thus, the optimal strategies in the original game and the modified game with rewards $\bar{r}_{ia}(\cdot)$ coincide. Therefore, also the games (Q, r) and $(Q, \bar{r})$ have exactly the same mean field equilibria.

Throughout this proof we denote by $||\cdot||$ the maximum norm on $\mathbb{R}^S$ as well as the associated matrix norm, which is the row sum norm.

Similar to the proof of Theorem 7.3 we have

$$||Q^d(m) - (Q')^d(m)|| \leq S \cdot d((Q,r),(Q',r'))$$

for all $m \in \mathcal{P}(\mathcal{S})$ since

$$\begin{aligned} d((Q,r),(Q',r')) &\geq \sup_{m \in \mathcal{P}(\mathcal{S})} \max_{i,j \in \mathcal{S}, a \in \mathcal{A}} |Q_{ija}(m) - Q'_{ija}(m)| \\ &\geq \sup_{m \in \mathcal{P}(\mathcal{S}), d \in D^s} \max_{i \in \mathcal{S}} \frac{1}{S} \sum_{j=1}^{S} |Q^d_{ij}(m) - (Q')^d_{ij}(m)| \\ &= \sup_{m \in \mathcal{P}(\mathcal{S})} \max_{d \in D^s} \frac{1}{S} ||Q^d(m) - (Q')^d(m)||. \end{aligned}$$

Furthermore, we make the following preliminary observations:

- By Lemma 5.3 the matrix $\beta I - Q^d(m)$ is invertible for any $m \in \mathcal{P}(\mathcal{S})$, $d \in D^s$, $Q^d(\cdot) \in \mathrm{TR}$ and $\beta \in (0,1)$. Thus, it holds that

  $$||\beta I - Q^d(m)|| > 0 \quad \text{and} \quad ||(\beta I - Q^d(m))^{-1}|| > 0$$

  for all $d \in D^s$, $Q^d(\cdot) \in \mathrm{TR}$, $m \in \mathcal{P}(\mathcal{S})$ and $\beta \in (0,1)$. Since the norm is a continuous mapping and $\mathcal{P}(\mathcal{S})$ is compact, we note that there are uniform (in $m \in \mathcal{P}(\mathcal{S})$) upper and lower bounds for both of these terms and that the lower bounds are strictly positive. Since the set of deterministic strategies is finite, we note that there are also uniform upper/lower bounds in $m \in \mathcal{P}(\mathcal{S})$ and $d \in D^s$ simultaneously, where again the lower bounds are strictly positive. For future use let us write

  $$L_1 := \inf_{m \in \mathcal{P}(\mathcal{S}), d \in D^s} ||\beta I - Q^d(m)|| > 0.$$

- Using that there are uniform upper bounds (in m and d) for $||\beta I - Q^d(m)||$ and $||(\beta I - Q^d(m))^{-1}||$ we obtain for the conditioning number:

  $$\begin{aligned} &\sup_{m \in \mathcal{P}(\mathcal{S}), d \in D^s} \kappa(\beta I - Q^d(m)) \\ &= \sup_{m \in \mathcal{P}(\mathcal{S}), d \in D^s} ||\beta I - Q^d(m)|| \cdot ||(\beta I - Q^d(m))^{-1}|| \\ &\leq \left(\sup_{m \in \mathcal{P}(\mathcal{S}), d \in D^s} ||\beta I - Q^d(m)|| \right) \cdot \left(\sup_{m \in \mathcal{P}(\mathcal{S}), d \in D^s} ||(\beta I - Q^d(m))^{-1}|| \right) \\ &=: L_2 < \infty. \end{aligned}$$

- Since $r^d(m) > 0$ for all $m \in \mathcal{P}(\mathcal{S})$ and $d \in D^s$ there is a strictly positive uniform lower bound of $||r^d(m)||$, which we will call L_3.

- Moreover, since for all $d \in D^s$ the function $V^d(\cdot)$ is continuous with compact domain (see Lemma 5.3) and $|| \cdot ||$ is also a continuous mapping, we obtain from (3.1) that $||V^d(m)||$ has a finite uniform upper and lower bounds in $m \in \mathcal{P}(\mathcal{S})$ and $d \in D^s$. Since $r^d(m) > 0$ for all $m \in \mathcal{P}(\mathcal{S})$ we obtain by (3.1) that $V^d(m) > 0$ for all $m \in \mathcal{P}(\mathcal{S})$ and $d \in D^s$, which in total implies that

$$L_4 := \sup_{m \in \mathcal{P}(\mathcal{S}), d \in D^s} ||V^d(m)|| > 0.$$

Let

$$0 < \delta < \min_{d \in D^s} \left\{ \frac{1}{2S||(\beta I - Q^d(m))^{-1}||}, \frac{\gamma L_1 L_3}{2 L_2 L_4 (S L_3 + L_1)} \right\}$$

and let $(Q', r') \in N_\delta(Q, r)$. Then we obtain the following estimate

$$\begin{aligned}
&\left(1 - \kappa(\beta I - Q^d(m)) \frac{||Q^d(m) - (Q')^d(m)||}{||\beta I - Q^d(m)||}\right)^{-1} \\
&= \left(1 - ||\beta I - Q^d(m))^{-1}|| \cdot ||Q^d(m) - (Q')^d(m)||\right)^{-1} \\
&\leq \left(1 - ||(\beta I - Q^d(m))^{-1}|| \cdot \frac{1}{2||(\beta I - Q^d(m))^{-1}||}\right)^{-1} \\
&= \left(1 - \frac{1}{2}\right)^{-1} = 2.
\end{aligned}$$

Using this and Lemma 7.4, we obtain for any $d \in D^s$ and any $m \in \mathcal{P}(\mathcal{S})$ that the distance of the expected discounted reward given d for the game (Q, r) (denoted by $V^d(m)$) and the expected discounted reward given d for the game (Q', r') (denoted by $(V')^d(m)$) satisfies:

$$\begin{aligned}
&||(V')^d(m) - V^d(m)|| \\
&\leq ||V^d(m)|| \kappa(\beta I - Q^d(m)) \left(1 - \kappa(\beta I - Q^d(m)) \frac{||(Q')^d(m) - Q^d(m)||}{||\beta I - Q^d(m)||}\right)^{-1} \\
&\quad \cdot \left(\frac{||(Q')^d(m) - Q^d(m)||}{||\beta I - Q^d(m)||} + \frac{||(r')^d(m) - r^d(m)||}{||r^d(m)||} \right) \\
&\leq L_4 \cdot L_2 \cdot 2 \cdot \left(\frac{||(Q')^d(m) - Q^d(m)||}{L_1} + \frac{||(r')^d(m) - r^d(m)||}{L_3} \right) \\
&\leq L_4 \cdot L_2 \cdot 2 \cdot \left(\frac{S d((Q,r),(Q',r'))}{L_1} + \frac{d((Q,r),(Q',r'))}{L_3} \right) \\
&= d((Q,r),(Q',r')) \cdot \frac{2 L_2 L_4 (S L_3 + L_1)}{L_1 \cdot L_3} \leq \gamma.
\end{aligned}$$

□

Now we are able to answer the questions raised in the beginning of the section. We start with a lemma stating under which conditions we can ensure that a strategy d that is optimal for m given game (Q, r) is also optimal for m in any slightly perturbed game. The condition that d is the unique optimal strategy for the point m is necessary for the general statement to hold, as we will illustrate in an example after the proof. Under this condition, we can even prove that d is the unique optimal strategy in any slightly perturbed game for a neighbourhood of m.

Lemma 7.15. *Let $(Q, r) \in \mathcal{G}$ be a game. Assume that $d \in D^s$ is the unique optimal deterministic strategy for m, that is $\mathcal{D}(m) = \{d\}$. Then there is an $\epsilon > 0$ such that for all $m' \in \overline{N_\epsilon(m)}$ we have $\mathcal{D}(m') = \{d\}$. Furthermore, for any such ϵ there is a $\delta > 0$ such that for all $(Q', r') \in N_\delta(Q, r)$ and all $m' \in N_\epsilon(m)$ we have $\mathcal{D}'(m') = \{d\}$.*

Proof. Since $\mathcal{D}(m) = \{d\}$ we have that $V^d(m) > V^{\hat{d}}(m)$ pointwise for all $\hat{d} \in D^s \setminus \{d\}$. By continuity of $V^d(\cdot)$ and $V^{\hat{d}}(\cdot)$ we furthermore find an $\epsilon > 0$ such that

$$V^d(m') > V^{\hat{d}}(m')$$

pointwise for all $m' \in \overline{N_\epsilon(m)}$ and all $\hat{d} \in D^s \setminus \{d\}$. Since $\overline{N_\epsilon(m)}$ is compact, we have that

$$\inf_{m' \in \overline{N_\epsilon(m)}} ||V^d(m') - V^{\hat{d}}(m')|| = \gamma > 0.$$

By Lemma 7.14 we choose $\delta > 0$ such that for all games $(Q', r') \in N_\delta(Q, r)$, all strategies $\hat{d} \in D^s$ and all $m' \in \mathcal{P}(\mathcal{S})$ we have

$$||V^{\hat{d}}(m') - (V')^{\hat{d}}(m')|| < \frac{\gamma}{3}.$$

With this it holds pointwise for $\hat{d} \in D^s \setminus \{d\}$ and $m' \in N_\epsilon(m)$ that

$$(V')^d(m') - (V')^{\hat{d}}(m') > V^d(m') - \frac{\gamma}{3}1 - V^{\hat{d}}(m') - \frac{\gamma}{3}1 = V^d(m') - V^{\hat{d}}(m') - \frac{2\gamma}{3} > 0,$$

which yields that d is the only optimal strategy in the perturbed game for m'. □

Remark 7.16. The assumption that d is the unique optimal strategy is necessary because otherwise we cannot ensure that d is optimal under a small perturbation: Let us consider two strategies d^1 and d^2 such that d^1 is optimal for all points and d^2 is optimal at exactly one point m. For instance choose a game with $Q^{d^1}(\cdot) = Q^{d^2}(\cdot)$ and $r^{d^1}(\cdot)$ to be constant and $r^{d^2}(\cdot)$ such that its global maximum is attained at m and equals $r^{d^1}(\cdot)$. Then if we would consider $r^{d^2}(\cdot) - \delta$ instead of $r^{d^2}(\cdot)$, this would be a small perturbation such that d^2 is not optimal at m and, moreover, d^2 is not optimal for any point close to m. △

For the case that d is not optimal given m, we obtain a similar result. However, we do not need to require any extra conditions. It suffices that d is not optimal at m to ensure that it is not optimal in a neighbourhood of m for any slightly perturbed game.

Lemma 7.17. *Let $(Q, r) \in \mathcal{G}$ be a game. Assume that $d \in D^s$ is not optimal for m, that is $d \notin \mathcal{D}(m)$. Then there is an $\epsilon > 0$ such that for all $m' \in \overline{N_\epsilon(m)}$ we have $d \notin \mathcal{D}(m')$. Furthermore, for any such ϵ there is a $\delta > 0$ such that for all $(Q', r') \in N_\delta(Q, r)$ and all $m' \in N_\epsilon(m)$ we have $d \notin \mathcal{D}'(m')$.*

Proof. Since $d \notin \mathcal{D}(m)$ we have that $V^d(m) < V^*(m)$ pointwise. By continuity of $V^d(\cdot)$ and $V^*(\cdot)$ we furthermore find an $\epsilon > 0$ such that

$$V^d(m') < V^*(m')$$

pointwise for all $m' \in \overline{N_\epsilon(m)}$. Since $\overline{N_\epsilon(m)}$ is compact, we have that

$$\inf_{m' \in \overline{N_\epsilon(m)}} ||V^d(m') - V^*(m')|| = \gamma > 0.$$

By Lemma 7.14 we choose $\delta > 0$ such that for all games $(Q', r') \in N_\delta(Q, r)$, all $\hat{d} \in D^s$ and all $m' \in \mathcal{P}(\mathcal{S})$ we have

$$||V^{\hat{d}}(m') - (V')^{\hat{d}}(m')|| < \frac{\gamma}{3}.$$

With this it holds pointwise for all games $(Q', r') \in N_\delta(Q, r)$, all strategies $\hat{d} \in D^s$ and all $m' \in N_\epsilon(m)$ that

$$(V')^*(m') - (V')^d(m') > V^*(m') - \frac{\gamma}{3}1 - V^d(m') - \frac{\gamma}{3}1 = V^*(m') - V^d(m') - \frac{2\gamma}{3} > 0,$$

which yields that d is not optimal in the perturbed game for m'. □

7.4. Which Equilibria are Essential? Which Equilibria are Strongly Stable?

This section discusses in detail for which classes of mean field games we can expect essentiality and strong stability and provides examples and intuitions, why this is not possible to obtain this for other classes. The main results are as follows: Deterministic stationary mean field equilibria such that the equilibrium strategy is the unique optimal strategy are essential if the equilibrium point is an essential stationary point of the dynamics. They are moreover strongly stable in the class of all games with constant dynamics if the generator is irreducible. Randomized mean field equilibria, where the randomization happens only for two actions, and where the dynamics are given by a constant irreducible generator are essential in the class of all games with constant dynamics. For all other cases we provide counterexamples and intuitions that in general it cannot be obtained that an equilibrium is essential or strongly stable.

7.4.1. Equilibria with a Deterministic Equilibrium Strategy

We start with the positive result for equilibria with deterministic equilibrium strategy and general dynamics:

Theorem 7.18. *Let $(Q, r) \in \mathcal{G}$ be a game and let (m, d) be an equilibrium of (Q, r) such that $d \in D^s$ is the unique optimal deterministic strategy given m (that is $\mathcal{D}(m) = \{d\}$) and m is an essential stationary point given $Q(\cdot)$. Then (m, d) is essential.*

Proof. Since d is the unique optimal strategy given m there is an $\bar{\epsilon}$ such that for all $m' \in \overline{N_{\bar{\epsilon}}(m)}$ we have $\mathcal{D}(m') = \{d\}$. Let $\epsilon \in (0, \bar{\epsilon})$. By Lemma 7.15, we find a $\delta_1 > 0$ such that for any game $(Q', r') \in N_{\delta_1}(m)$ and all $m' \in N_\epsilon(m)$ we have $\mathcal{D}(m') = \{d\}$.

Since m is an essential stationary point of the dynamics, there is a $\delta_2 > 0$ such that for any $(Q', r') \in N_{\delta_2}(Q, r)$ there is a stationary point of the dynamics given $(Q')^d(\cdot)$ in $N_\epsilon(m)$.

Choosing $\delta = \min\{\delta_1, \delta_2\}$ yields that the strategy d is optimal in $N_\epsilon(m)$ and furthermore in this set lies a stationary point given $(Q')^d(\cdot)$, which is, by construction, a stationary equilibrium given (Q', r'). □

Remark 7.19. It is important to assume that m is an essential stationary point of the dynamics given $Q(\cdot)$ since otherwise we can construct counterexamples (a counterexample in the case of constant reducible dynamics can be found in Subsection 7.2.3). However, also the condition that d is the unique optimal strategy is necessary. Indeed, we cannot ensure without this assumption that the strategy d remains optimal (see also Section 7.3): Consider the consumer choice model presented in Section 6.1 with parameters such that $k_1 < \epsilon/(b + \epsilon)$ and $k_2 = b/(b + \epsilon)$. The point $m = (b/(b + \epsilon), \epsilon/(b + \epsilon))$ together with the strategy $\{stay\} \times \{change\}$ is an equilibrium, but the strategy $\{stay\} \times \{stay\}$ is also optimal at the point m. If we increase s_2 slightly by $\delta > 0$, then we obtain that $k_2 > \frac{b}{b+\epsilon}$, which yields that only the point $(1/2, 1/2)$ together with the strategy $\{stay\} \times \{stay\}$ is an mean field equilibrium. However, it is not "close" to the original equilibrium $(b/(b + \epsilon), \epsilon/(b + \epsilon))$. △

The previous result together with Lemma 7.11 yields the following corollary characterizing essential equilibria in the class of all games with $Q^d(\cdot)$ being constant and with the same set of recurrence classes:

Corollary 7.20. *Let $(Q, r) \in \mathcal{G}$. Let (m, d) be a stationary mean field equilibrium such that d is the unique optimal deterministic strategy given m (that is $\mathcal{D}(m) = \{d\}$) and assume that $Q^d(\cdot)$ is constant over $\mathcal{P}(\mathcal{S})$. Then for any $\epsilon > 0$ there is an $\delta > 0$ such that any game $(Q', r') \in N_\delta(Q, r)$ such that $(Q')^d$ is also constant and the recurrence classes of Q and Q' coincide has an equilibrium with distribution in $N_\epsilon(m)$ and equilibrium strategy d.*

From Theorem 7.18 it furthermore follows using Lemma 7.12 that if we restrict our attention to constant irreducible dynamics we obtain that equilibria are even strongly stable in the set of all games with constant irreducible dynamics:

Corollary 7.21. *Let* $(Q, r) \in \mathcal{G}^{ci}$*. Let* (m, d) *be a stationary mean field equilibrium such that* d *is the unique optimal deterministic strategy given* m *(that is* $\mathcal{D}(m) = \{d\}$*), then* (m, d) *is strongly stable in* $\mathcal{G}^{ci}$*.*

7.4.2. Equilibria with a Randomized Equilibrium Strategy

In general, randomized equilibria will not be essential: Consider an equilibrium (m, π) where randomization happens over at least three state-action pairs, that is there are three distinct pairs (i, a) such that $\pi_{i,a} \in (0, 1)$. We note that this covers both the case that in one state randomization happens over three actions and the case that for two states randomization happens over at least two actions. In this situation, it is not clear whether all actions that have positive weight in the equilibrium strategy are again simultaneously optimal for a point close to m under a small perturbation. If this does not happen the equilibrium is not essential since for $\epsilon < \min_{(i,a) \text{ s.t. } \pi_{ia} \in (0,1)} \pi_{ia}$ the distance between the equilibrium strategies of the original game and the perturbed game will be too large.

This is not only a theoretical possibility but happens under perturbations that are reasonable from a practical viewpoint: Let us consider a slightly modified version of the consumer choice model in Section 6.1, namely, assume that $c_1 = c_2 = 0$ and $s_1 = s_2$. Then only the strategies $\{\mathit{stay}\} \times \{\mathit{change}\}$ and $\{\mathit{change}\} \times \{\mathit{stay}\}$ will be strategies that are unique optimal strategies for some m and that the first strategy is optimal whenever $m_1 \geq m_2$ and the second strategy is optimal whenever $m_1 \leq m_2$. Thus, $(1/2, 1/2)$ is an equilibrium point for all strategies that satisfy $\pi_{1,\mathit{change}} = \pi_{2,\mathit{change}}$. If we consider the slightly perturbed game $c_1 = c_2 = \delta$, the we will obtain that $(1/2, 1/2)$ is an equilibrium point only given the pure strategy $\{\mathit{stay}\} \times \{\mathit{stay}\}$, which is not "close" to most randomized equilibrium strategies of the original game. We remark that although in our example the equilibrium distributions are close, this does not have to hold in general.

If there are only two state-action pairs (i, a) satisfying $\pi_{ia} \in (0, 1)$ it cannot happen that the strategies will not be simultaneously optimal for some point close to m under suitable assumptions due to the intermediate value theorem. However, there are more obstacles: It is no longer sufficient to require that m is an essential stationary point of the dynamics $Q^\pi(\cdot)$. Indeed, it might happen that the equilibrium strategy π changes, in which case we would search for stationary points given a different randomized strategy $\tilde{\pi}$. In the case of general dynamics we do not know how the stationary points under the perturbation and the new strategy relate to those given $Q^\pi(\cdot)$. Thus, we cannot formulate an essentiality result for general dynamics.

However, this problem does not occur in the constant dynamics setting, where the relation of the fixed points of $Q^\pi(\cdot)$ are clear. Thus, we formulate the following positive result in this setting:

Theorem 7.22. *Let $(Q, r) \in \mathcal{G}^{ci}$ be a game with constant irreducible dynamics and assume that (m, π) is an equilibrium. Moreover, assume that $\mathcal{D}(m) = \{d_1, d_2\}$ such that for m_1 being the stationary point given Q^{d_1} and m_2 being the stationary point given Q^{d_2} there is an $\lambda \in (0, 1)$ such that $m = \lambda m_1 + (1 - \lambda)m_2$. Further assume that $\beta_1, \beta_2 > 0$ are such that*

- *the strategy d_1 and not the strategy d_2 is optimal for $m_1(\lambda \mp y) + m_2(1 - \lambda \pm y)$ for all $y \in (0, \beta_1)$*
- *the strategy d_2 and not the strategy d_1 is optimal for $m_1(\lambda \pm y) + m_2(1 - \lambda \mp y)$ for all $y \in (0, \beta_2)$.*

Then the equilibrium (m, π) is essential in $\mathcal{G}^{ci}$.

The condition that requires that d_1 is optimal on one side and that d_2 is optimal on the other side of our equilibrium value is necessary. Without this condition the construction in Remark 7.16 would yield a counterexample.

The following proof combines the ideas presented in the previous two sections. Since it involves several steps, Figure 7.2 illustrates the situation of the proof.

Proof. Without loss of generality we assume that

- the strategy d_1 and not the strategy d_2 is optimal for $m_1(\lambda - y) + m_2(1 - \lambda + y)$ for all $y \in (0, \beta_1)$
- the strategy d_2 and not the strategy d_1 is optimal for $m_1(\lambda + y) + m_2(1 - \lambda - y)$ for all $y \in (0, \beta_2)$.

It suffices to prove the statement for all $\epsilon \in (0, \bar{\epsilon})$ with $\bar{\epsilon} > 0$ such that

- for all $m' \in \overline{N_{\bar{\epsilon}}(m)}$ we have $\mathcal{D}(m') \subseteq \{d_1, d_2\}$ (this is possible since at m only the strategies d_1 and d_2 as well as convex combinations thereof are optimal).
- $\bar{\epsilon} \le \min\{\beta_1, \beta_2\}$.

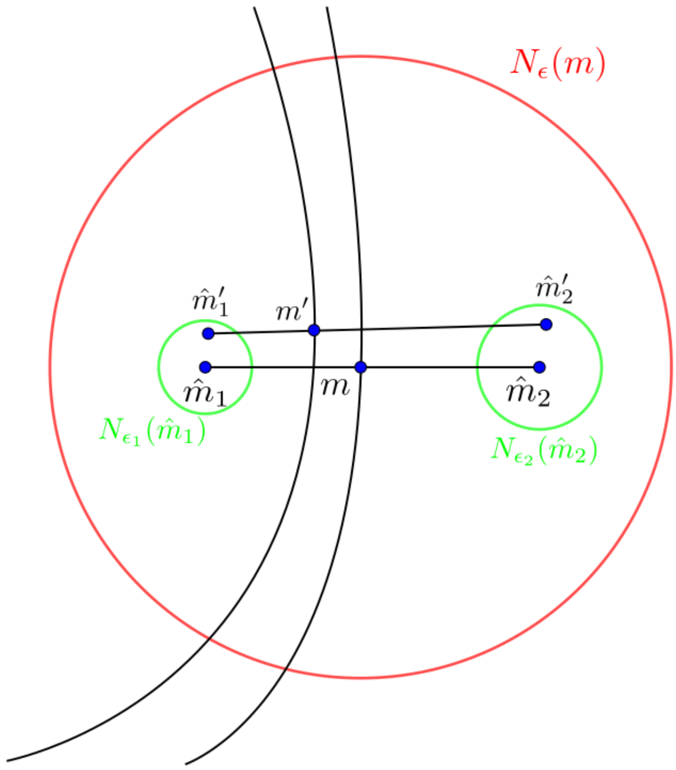

Figure 7.2.: The situation in the proof of Theorem 7.22. The black curved line through m represents the points where both d_1 and d_2 are optimal in the game (Q, r). The black curved line through $\tilde{m}$ represents the points where both d_1 and d_2 are optimal in the game (Q', r'). The black straight lines are the lines on which all relevant stationary points given Q^π and $(Q')^\pi$, respectively, with $\pi \in \text{conv}(d_1, d_2)$ lie.

Let $\epsilon \in (0, \bar{\epsilon})$ be given and choose

$$\hat{m}_1 \in \{m_1(\lambda - y) + m_2(1 - \lambda + y) : y \in (0, \epsilon)\}$$
$$\hat{m}_2 \in \{m_1(\lambda + y) + m_2(1 - \lambda - y) : y \in (0, \epsilon)\}$$

Since d_1 is the unique optimal strategy for $\hat{m}_1$ and d_2 is the unique optimal strategy for $\hat{m}_2$, there are $\epsilon_1, \epsilon_2 > 0$ such that in $\overline{N_{\epsilon_1}(\hat{m}_1)}$ only the strategy d_1 is optimal and in $\overline{N_{\epsilon_2}(\hat{m}_2)}$ only the strategy d_2 is optimal. By Lemma 7.17 there is a $\delta_1 > 0$ such that for all perturbed games $(Q', r') \in N_{\delta_1}(Q, r)$ all strategies $d \in D^s \setminus \{d_1, d_2\}$ are not optimal in $N_\epsilon(m)$. Furthermore, by Lemma 7.15 there are a $\delta_{2,i} > 0$ $(i = 1, 2)$ such that for all perturbed games $(Q', r') \in N_{\delta_{2,i}}(Q, r)$ the strategy d_i is the unique optimal strategy in $N_{\epsilon_i}(\hat{m}_i)$.

Since $\hat{m}_i$ is the unique stationary point given the dynamics of the constant irreducible generator matrix $Q^{\hat{\pi}_i}$ for some strategy $\hat{\pi}_i \in \text{conv}\{d_1, d_2\}$ it is an essential stationary point of the dynamics. Thus, for some $\delta_{3,i} > 0$ we have that the stationary point $\hat{m}_i'$ given strategy $\hat{\pi}_i$ in any perturbed game $(Q', r') \in N_{\delta_{3,i}}(Q, r) \cap \mathcal{G}^{ci}$ lies in $N_{\epsilon_i}(\hat{m}_i)$.

Thus for $\delta = \min\{\delta_1, \delta_{2,1}, \delta_{2,2}, \delta_{3,1}, \delta_{3,2}\}$ we obtain that for all perturbed games $(Q', r') \in N_\delta(Q, r) \cap \mathcal{G}^{\text{ci}}$

- all strategies $d \in D^s \setminus \{d_1, d_2\}$ are not optimal in $N_\epsilon(m)$
- in $N_{\epsilon_1}(\hat{m}_1)$ only the strategy d_1 is optimal
- in $N_{\epsilon_2}(\hat{m}_2)$ only the strategy d_2 is optimal
- $\hat{m}_1' \in N_{\epsilon_1}(\hat{m}_1)$ and $\hat{m}_2' \in N_{\epsilon_2}(\hat{m}_2)$.

Since
$$f(\lambda) = (V')^{d_1}(\lambda \hat{m}_1' + (1-\lambda)\hat{m}_2') - (V')^{d_2}(\lambda \hat{m}_1' + (1-\lambda)\hat{m}_2')$$
is continuous in $\lambda \in [0, 1]$ and satisfies that $f(0) < 0$ and $f(1) > 0$, there is a $\lambda' \in (0, 1)$ for which both strategies are simultaneously optimal. Denote by m' some point at which $(V')^{d_1}(m')$ and $(V')^{d_2}(m')$ coincide. Since $N_\epsilon(m)$ is convex, we have that $m' \in N_\epsilon(m)$. Since we are in the constant dynamics setting this point is the stationary point given a randomized strategy π' that lies between $\hat{\pi}_1$ and $\hat{\pi}_2$. By construction, the distance of π' to π itself is at most ϵ, which proves the claim. □

Remark 7.23. In the case of randomized mean field equilibria, strong stability cannot be obtained since we can again bend the boundary between the two optimality sets as we like, in particular, we could obtain a situation as in Remark 7.13.

7.5. Is Essentiality a Generic Property?

We have seen that there are games that have equilibria, which are not essential. However, for most equilibrium concepts this happens only in "exceptional cases" and one can often prove that for a large (”generic”) set of games all equilibria are essential (see van Damme (1991); Doraszelski and Escobar (2010)). The question, which properties are generic for certain classes of games, is a classical one: The first investigations were carried out in settings, where games are described by vectors in $\mathbb{R}^n$, in which case a generic property is a property that holds for a set with full Lebesgue measure. This genericity notion captures well the intuitive idea that a randomly chosen game exhibits the property. Genericity considerations of that type have been discussed for normal form games in Harsanyi (1973) and for Markov perfect equilibria in Doraszelski and Escobar (2010) as well as Haller and Lagunoff (2000).

For more complex spaces like function spaces there is no direct analogue of Lebesgue almost every that has similar properties, which include that a measure zero set has no interior and that translates of measure zero sets have zero measure. A partially satisfactory notion is the topological genericity notion, where a set is generic if it is open and dense or, in case of complete spaces, if it is residual. A genericity result of that type for games with convex compact polyhedra as strategy spaces can be found in Dubey (1986).

However, the topological genericity notion is less sharp than the measure theoretic notion since in $\mathbb{R}^n$ there are open dense sets with arbitrarily small measure (Mas-Colell, 1985). Hunt et al. (1992) then proposed a measure theoretic genericity notion for complete metric spaces, called prevalence, which is sharper and exhibits the desired properties. See also Ott and Yorke (2005) for more details and examples. Anderson and Zame (2001) extend this notion to convex subsets of vector spaces, which is the case of interest in economics, and provide several economic genericity statements. These include the statement that in static games with strategy spaces being compact polyhedra the set of games with finitely many pure strategy equilibria is prevalent.

In our case it is not clear a priori if inessential equilibria are "exceptional": Inessential equilibria (m, d) with a deterministic equilibrium strategy either have to satisfy that d is not the unique optimal strategy given m or that m is not an essential stationary point of the dynamics $Q^d(\cdot)$. These properties are indeed somewhat uncommon (see Theorem 7.8). For inessential equilibria (m, π) with a randomized strategy the case is not clear since we do not understand why inessentiality occurs in detail. In particular, we cannot relate for general dynamics stationary points of $Q^{d^1}(\cdot)$ and $Q^{d^2}(\cdot)$ with those of $Q^{\pi}(\cdot)$ with π being a convex combination of d^1 and d^2. Fortunately, we can nonetheless prove topological genericity relying on the proof idea of Fort (1950), which has also been used in Subsection 7.2.2:

Theorem 7.24. *The set of all games for which all equilibrium points are essential is residual in the set of all games $\mathcal{G}$. Moreover, this set is also dense in $\mathcal{G}$, which means that we approximate any game arbitrarily close by a game with only essential equilibria.*

Proof. The proof follows the same line of argument as the proof of Theorem 7.8: Let us consider the map SMFE : $\mathcal{G} \to 2^{\mathcal{P}(\mathcal{S}) \times \mathcal{P}(\mathcal{A})^S}$ that maps any game (Q, r) to the set of all mean field equilibria. We want to utilize the results of Section 7 in Fort (1949), which states that the points of continuity of any upper semi-continuous function $F(\cdot)$ ranging from a topological space to the power set of a separable metric space equipped with metric $H(\cdot, \cdot)$ is a G_δ-residual set in the topological space. For this we have to prove that SMFE$(\cdot)$ is upper-semicontinuous and that (Q, r) is a point of continuity if and only if all equilibria are essential. Since the proof of the characterization of the points of continuity applies mutatis mutandis, we only prove here that the map SMFE$(\cdot)$ is upper semi-continuous:

Let $(Q,r) \in \mathcal{G}$ be a game and let $\epsilon > 0$. We want to find a $\delta > 0$ such that any game (Q', r') satisfying $d((Q,r),(Q',r')) < \delta$ satisfies $\text{SMFE}(Q',r') \subseteq N_\epsilon(\text{SMFE}(Q,r))$.

Assume by way of contradiction that such a $\delta > 0$ does not exist. Then we find sequences $(Q^n, r^n)_{n\in\mathbb{N}}$ and $(m^n, \pi^n)_{n\in\mathbb{N}}$ such that

- $d((Q^n, r^n), (Q, r)) < \frac{1}{n}$
- $(m^n, \pi^n) \in \text{SMFE}(Q^n, r^n)$ for all $n \in \mathbb{N}$ since any game has a stationary mean field equilibrium by Theorem 4.14,
- $(m^n, \pi^n) \notin N_\epsilon(\text{SMFE}(Q, r))$.

Since $\mathcal{P}(\mathcal{S}) \times \mathcal{P}(\mathcal{A})^S$ is compact we find a converging subsequence $(m^{n_k^1}, \pi^{n_k^1})_{k\in\mathbb{N}}$ of $(m^n, \pi^n)_{n\in\mathbb{N}}$. Let (m, π) be its limit. Let $A_1^k \times \ldots \times A_S^k$ be such that $\pi_{ia}^k > 0$ for all $i \in \mathcal{S}, a \in A_i^k$ and $\pi_{ia} = 0$ for all $i \in \mathcal{S}, a \notin A_i^k$. Since $\mathcal{A}$ is finite we find a set $A_1 \times \ldots \times A_S$ that occurs infinitely often. Moreover, the subsequence $(m^{n_l^2}, \pi^{n_l^2})_{l\in\mathbb{N}}$ that runs through all indices n_k^1 such that $A_1^{n_k^1} \times \ldots \times A_S^{n_k^1} = A_1 \times \ldots \times A_S$ converges towards (m, π). As a next step we note that for all $d \in D^s$ such that $d(i) \in A_i$ for all $i \in \mathcal{S}$ it holds that $V^d(m^{n_l}) = V^*(m^{n_l})$. By continuity of $V^d(\cdot)$ and $V^*(\cdot)$, this implies that $V^d(m) = V^*(m)$. Since $\pi_{ia}^{n_l} = 0$ for all $i \in \mathcal{S}, a \notin A_i$ and $l \in \mathbb{N}$, we have $\pi_{ia} = 0$. Thus, the previous observations yield together with Theorem 5.2 that π is optimal for m.

Furthermore, by uniform convergence, we have that

$$\sum_{i\in\mathcal{S}} \sum_{a\in\mathcal{A}} Q_{ija}(m) m_i \pi_{ia} \leftarrow \sum_{i\in\mathcal{S}} \sum_{a\in\mathcal{A}} Q_{ij}^{n_l}(m^{n_l}) m_i^{n_l} \pi_{ia}^{n_l} = 0.$$

Thus, m is a stationary point given $Q^\pi(\cdot)$, which in total yields that $(m, \pi) \in \text{SMFE}(Q, r)$. However, $m^n \notin N_\epsilon(\text{SMFE}(Q, r))$ implies $m \notin N_\epsilon(\text{SMFE}(Q, r))$, a contradiction. □

It would be desirable by the discussion in the beginning of the section to obtain also a measure-theoretic genericity result. However, the previous proof crucially relies on the close link of essentiality and the topological notion of continuity. Thus, in order to obtain a measure-theoretic genericity result it is necessary to come up with a completely new proof idea. However, up to the knowledge of the author, not even for maps with only essential fixed points any measure-theoretic genericity results are known. Since in the case of arbitrary dynamics fixed point problems are a part of the solution, prevalence for this problem should be considered first before considering prevalence of the set of games with only essential equilibria.

For constant dynamics, we have that if the dynamics are given by irreducible generators the equilibria are essential whenever the deterministic equilibrium strategy is unique or whenever we consider a truly randomizing equilibrium strategy that satisfies the condition of Theorem 7.22. Since the set of all irreducible generator matrices lies dense in the set of all generator matrices it is only necessary to obtain the desired properties of equilibria for irreducible transition rate matrices. One approach, relying on the previous results would be the following: Ensure for a prevalent set that

1. the unique stationary point x^d given the strategy $d \in D^s$ satisfies that if d is optimal for x^d, then d is the unique optimal stationary strategy at x^d;

2. for any point x^π that is a stationary point given a stationary strategy π that randomizes over three or more deterministic stationary strategies not all of these strategies are optimal at x^π;

3. for any point x^π that is a stationary point given a stationary strategy π such that there are two deterministic stationary strategies d^1 and d^2 satisfying $\pi = \lambda d_1 + (1-\lambda)d^2$ and such that d^1 and d^2 are both optimal at x^π (which implies by the first point that $\lambda \in (0,1)$) we have that there are $\beta_1, \beta_2 > 0$ such that for all points $(\lambda + y)d_1 + (1-\beta-y)d^2$ with $y \in (0, \beta_1)$ one of the strategies is the unique optimal strategy and for all points $(\lambda - y)d_1 + (1-\beta+y)d^2$ with $y \in (0,\beta_2)$ the other strategy is optimal.

Since by Ott and Yorke (2005, Section 3) a countable intersection of prevalent sets is again prevalent we basically have to show that the following three sets are prevalent:

- First, fix two distinct strategies $d^1, d^2 \in D^s$ and consider the set of reward functions $(r_{ia}(\cdot))_{i\in\mathcal{S},a\in\mathcal{A}}$ such that for any point $\tilde{x}^\pi = \pi x^{d^1} + (1-\pi)x^{d^2}$ satisfying $V^{d^1}(\tilde{x}^\pi) = V^{d^2}(\tilde{x}^\pi)$ we have that closely before $\tilde{x}^\pi$ the functions satisfy $V^{d^1} < V^{d^2}$ and closely after $\tilde{x}^\pi$ the functions satisfy $V^{d^2} < V^{d^1}$ (or with the roles interchanged). It should be possible to prove that these sets are prevalent using the same techniques that have been used to establish Proposition 3 in Hunt et al. (1992).

- As a next step we would have to show that the set of reward functions $(r_{ia}(\cdot))_{i\in\mathcal{S},a\in\mathcal{A}}$ such that inside $\mathrm{conv}(x^{d^1}, x^{d^2}, x^{d^3})$ never all three strategies d^1, d^2 and d^3 are optimal simultaneously, is prevalent. For this we would have to construct a probe ourselves.

- As a final step we would have to show that the set of reward functions $(r_{ia}(\cdot))_{i\in\mathcal{S},a\in\mathcal{A}}$ such that at x^d the function $V^d(\cdot)$ does not equal any other function $V^{\hat{d}}(\cdot)$ with $\hat{d} \in D^s \setminus \{d\}$ is prevalent. For this we would again have to find a probe ourselves.

8. A Myopic Adjustment Process: Definition, Existence and Local Convergence to Stationary Equilibria

In the next two chapters, we will discuss a different decision mechanism in the context of our mean field game model that is closely related to the literature on learning in classical game theory. The main motivation lies in the fact, that it is not clear that agents are able to compute dynamic equilibria. More precisely, in the case of an infinite time horizon the individual control problem given a fixed flow of population distribution is a Markov decision process with non-stationary transition rates and rewards and there are no general results on how to compute optimal strategies. Although in the case of a finite time horizon dynamic equilibria can be characterized through a set of forward-backward systems of ordinary differential equations (see Belak et al. (2019)), it is again not clear whether agents can compute equilibria since the forward-backward systems are non-standard and it is often only possible to compute solutions of these systems numerically. Therefore, instead of assuming that the agents are fully rational and able to compute the equilibria, we introduce "myopic" decision making of the agents. We assume that the agents believe that the population distribution stays constant and, thus, choose their optimal best response given the current distribution. If the population distribution changes the agents re-evaluate their decision.

We will show that for any initial condition such a process exists and discuss local and global convergence of this process: In this chapter, we will show that we have local convergence towards equilibria (m, d) with a deterministic equilibrium strategy d such that d is the unique optimal strategy for the equilibrium distribution for two regimes. For the first regime, where the dynamics are constant and irreducible, we can explicitly describe a neighbourhood in which convergence holds. For the second regime, where the Jacobian $\frac{\partial}{\partial m}(Q^d(m))^T m$ has a certain structure, we can only characterize the neighbourhood implicitly. In Chapter 9, we will then prove under a set of assumptions that we even obtain (almost) global convergence of the myopic adjustment process. As in the classical learning theory for normal form games, the convergence results reinforce the predictive power of stationary mean field equilibria because they yield that stationary equilibria are indeed emerging through agents' behaviour although the current distribution is not the equilibrium distribution.

In this chapter we will first review in Section 8.1 the related literature on learning in standard game theory and learning in mean field game theory. Thereafter we will define and discuss in Section 8.2 our myopic adjustment process. In Section 8.3 we prove existence of a trajectory of this process for any initial condition. Section 8.4 contains the results on local convergence towards stationary mean field equilibria with a unique deterministic equilibrium strategy. Section 8.5 concludes with the application of these results to the examples of Chapter 6.

8.1. Literature Review

The theory of learning in games originates in the early beginnings of game theory in the work of Cournot (Nachbar, 2009). Nowadays it is a rather large branch of modern game theory, for an overview consider the survey by Nachbar (2009) or the monograph by Fudenberg and Levine (1998). It is mainly concerned with the question why the behaviour of agents in games constitutes an (Nash) equilibrium, although in many situations we cannot assume this behaviour arises from "introspection and calculation" for several conceptual and empirical reasons (Fudenberg and Levine, 2016, p.153): First, in complex games we cannot assume that agents are indeed able to compute the equilibria. Second, if multiple equilibria arise, it is not clear how agents coordinate on one of them. Third, in laboratory experiments of repeated static games it is usually observed that in early rounds of the game, the choices of the agents do not constitute Nash equilibria. However, if the agents receive feedback then the choices of the agents will converge in later rounds towards a Nash equilibrium.

The theory of learning yields an alternative explanation that in the long run agents play equilibrium strategies by providing processes in which agents are partially rational and which converge towards an equilibrium. Most of the literature in classical learning theory is focussed on infinitely repeated static games with perfect monitoring and large populations where players are matched randomly in each round (Nachbar, 2009). We remark that in this setting agents cannot alter their opponents choice by using a certain strategy. Thus, their only goal is to maximize their payoff. In this context the following three processes have received the most attention: In *fictitious play* the players choose their best response given the historical frequencies of actions. In *partial best-response dynamics* at each period a fixed (randomly chosen) share of the population readjust their action towards the best response given the previous period's frequencies of actions. In the *replicator dynamics* the share of players using a strategy grows proportional to the strategy's current payoff, yielding that strategies with a high current payoff are used more often, whereas strategies with a low current payoff are used less often. However, also more sophisticated learning rules that detect simple cycles, which clever agents would exploit, and stochastic versions of the processes, which overcome the problem that a small change in beliefs yields to a large discontinuous change in strategies, have been considered. See the monograph of Fudenberg and Levine (1998) for details and results.

In the context of mean field games two different types of learning rules have been considered. Cardaliaguet and Hadikhanloo (2017) introduce fictitious play for classical diffusion-based mean field games with time horizon $[0, T]$, which are repeated as follows: All players start with a common smooth initial belief of the distribution of the agents over the time horizon $[0, T]$ and play the corresponding optimal control to that belief. In all subsequent rounds, the players play the same game and choose the optimal strategy given the average of the observations of the previous rounds. They prove that in potential games fictitious play converges towards stationary mean field equilibria. Moreover, in Hadikhanloo (2017) convergence of fictitious play given the classical Lasry-Lion monotonicity assumption is proven and in Briani and Cardaliaguet (2018) it is shown that for potential games any stable equilibrium is a local attractor for fictitious play. Additionally, in Hadikhanloo (2018) fictitious play for mean field games with finite discrete time horizon and finite state space as in Gomes et al. (2010) is introduced and convergence for monotone games is proved.

Besides fictitious play in repeated mean field games, Mouzouni (2018) considers a completely different approach in the classical diffusion-based mean field games model, which is connected to our approach. At each time the agents assume that their environment is not changing and play their unique best response given this environment. Whenever the environment changes, they do not anticipate this, but just readjust their decision. It is then shown that there always exists a unique solution of the forward-forward system describing the process, that the mean field system is indeed the limit of the corresponding N-player system and that given the monotonicity condition and a quadratic Hamiltonian the process converges locally towards stationary mean field equilibria. We remark, that the underlying mean field model differs substantially and thus also the emerging mathematical theory is completely different.

Finally, we mention Adlakha and Johari (2013), where in a discrete time many player game with complements also a myopic adjustment process has been defined by stating that in each time step the individual agent computes his best response given the current population distribution. However, the proof exploits the monotonicity of the game and therefore again the emerging mathematical theory differs substantially.

8.2. Definition and Discussion

In this section we present the precise definition of the myopic adjustment process as well as its economic justification.

As we consider an infinite time horizon, we cannot assume that the game is played repeatedly, but instead we assume, as in Mouzouni (2018), that the agents adjust their choices at any time reacting to the current situation of the game. We assume that the agent only observes the population distributions $m(t)$ and not the actions of the other players. Moreover, we remind ourselves that the influence of the player onto the game characteristics and in consequence on the play of the others is negligible in mean field games. Thus, it is reasonable that we assume that an agent does not try to influence the other players' actions by his choices, but that he only optimizes his own payoff.

Furthermore, we assume that each agent chooses stationary strategies because most of the time the system's state will change immediately after the decision and thus only the instantaneously chosen action is relevant for the dynamics of the population distribution. Moreover, we assume that the players choose Markovian strategies that only depend on the current population distribution. The reasons for this are the same as in the justification of the notion of Markov perfect equilibria in Maskin and Tirole (2001): First, stationary Markovian strategies are the simplest form of strategies that capture rational decision making. Second, stationary Markovian strategies are closely related to subgame perfection, namely with this choice of strategies we achieve that any two games that have equivalent preferences and action spaces are played in the same way. Third, the choice of stationary Markovian strategies reduces the large number of possible best responses and thus increases predictive power.

As a next step, we have to answer the following questions: What is a good prediction of the population distribution of the agents on $[0, \infty)$? What is a corresponding best-response? Surprisingly, without a deep investigations of the particular example it is m_0 and some action from $\mathcal{D}(m_0)$: For this we first note that solving a Markov decision process with infinite time horizon and non-stationary rewards and transition rates is a very difficult problem, which cannot explicitly be handled. This means that agents with a "sophisticated" belief of the other players' distribution cannot profit from that belief. Second and related to the first reason, it is not clear how partially rational agents should come up with a sophisticated belief if they cannot compute best responses for arbitrary flows of population distributions.

This discussion leads to the following formal definition of the myopic adjustment process: For each $m \in \mathcal{P}(\mathcal{S})$, which is the current population distribution, the individual agent computes the set of optimal stationary Markovian strategies. By Theorem 5.2 the set of all optimal stationary Markovian strategies is the set of all convex combinations of the set $\mathcal{D}(m)$ containing all optimal deterministic stationary Markovian strategies.

Thus, the instantaneous change of the population distribution is a convex combination of $\left(\sum_{i\in\mathcal{S}} m_i Q^d_{ij}(m)\right)_{d\in\mathcal{D}(m)}$, which means that

$$\dot{m} \in F(m) := \operatorname{conv}\left\{ \left(\sum_{i\in\mathcal{S}} \sum_{a\in\mathcal{A}} m_i Q_{ija}(m) d_{ia} \right)_{j\in\mathcal{S}} : d \in \mathcal{D}(m) \right\}. \tag{8.1}$$

Since not all patterns of changes are reasonable, we require them to be mildly consistent. More precisely, we require the agents to choose their instantaneous best responses such that the myopic adjustment process is the solution of the set-valued differential equation in the sense of Deimling (1992) (see the next section for definitions)

$$\dot{m}(t) \in F(m(t)) \quad \text{for almost all } t \geq 0, \quad m(0) = m_0. \tag{8.2}$$

This restriction, which is mostly technically motivated, can be justified by stating that the agent should follow a path of instantaneous best responses that is continuous whenever possible.

Remark 8.1. By definition, a point is a stationary point of (8.1) and (8.2) if and only if it is the distribution of a stationary mean field equilibrium. Thus, if the process converges, then the limit is a stationary mean field equilibrium, which gives rise to a heuristic solution method. We note that it is a priori not clear whether this process converges. We are now interested in conditions ensuring convergence, mainly not to use them as solution methods, but to increase the explanatory power of the equilibrium concept.

8.3. Existence

The first natural question that occurs is whether such an adjustment process exists for every game and every initial distribution. For this we rely on the theory of differential inclusions as presented in Deimling (1992):

Given a closed set $D \subseteq \mathbb{R}^S$, a set-valued map $F : D \to 2^{\mathbb{R}^S} \setminus \{\emptyset\}$ and $m_0 \in D$, we search for an absolutely continuous function $m : [0, \infty) \to D$ such that

$$\dot{m}(t) \in F(m(t)) \text{ for almost all } t \in [0, \infty), \quad m(0) = m_0. \tag{8.3}$$

The main tool to discuss existence of solutions that stay in a closed set D is the Bouligand tangent cone

$$T_D(m) = \left\{ y \in \mathbb{R}^S : \liminf_{h\downarrow 0} \frac{d(m + hy, D)}{h} = 0 \right\},$$

where $\liminf_{h\downarrow 0}$ is the right-sided limes inferior and $d(x, A) := \inf_{y\in A} d(x, y)$ for $d(\cdot, \cdot)$ being induced by some norm on $\mathbb{R}^S$.

In general it holds that $0 \in T_D(m)$ for all $m \in D$ and that if $m \in \text{int}(D)$ then $T_D(m) = \mathbb{R}^S$ (Deimling, 1992, Proposition 4.1). We are particularly interested in the tangent cone of $\mathcal{P}(\mathcal{S})$, for which we obtain the following characterization:

Lemma 8.2. *Let $I(m) := \{i \in \mathcal{S} : m_i = 0\}$. Then for all $m \in \mathcal{P}(\mathcal{S})$ we have $y \in T_{\mathcal{P}(\mathcal{S})}(m)$ if and only if $y_i \geq 0$ for all $i \in I(m)$ and $\sum_{i \in \mathcal{S}} y_i = 0$. In particular, the set $T_{\mathcal{P}(\mathcal{S})}(m)$ is convex.*

Proof. By Deimling (1992, Proposition 4.1) we obtain that for any convex set D it holds that

$$T_D(m) = \overline{\{h(y - m) : h \geq 0, y \in D\}}$$

and that $T_D(m)$ is convex. By Aubin and Cellina (1984, Proposition 5.1.7) we then obtain the desired characterization. □

The existence result suitable for our setting can be found in Deimling (1992, Lemma 5.1):

Lemma 8.3. *Let $D \subseteq \mathbb{R}^S$ be a closed set and let $F : D \to 2^{\mathbb{R}^S} \setminus \{\emptyset\}$ satisfy the following properties*

(i) $F(\cdot)$ is upper semi-continuous,

(ii) for all $m \in D$ the set $F(m)$ is closed and convex,

(iii) there is a constant $c > 0$ such that for all $m \in D$ it holds that

$$||F(m)|| := \sup\{||y|| : y \in F(m)\} \leq c(1 + ||m||),$$

(iv) for all $m \in D$ we have $F(m) \cap T_D(m) \neq \emptyset$.

Then for every $m_0 \in D$ the system (8.3) *has a solution $m : [0, \infty) \to D$.*

With these preparations we prove that there exists a population distribution function emerging from the myopic adjustment process described in Section 8.2:

Theorem 8.4. *The differential inclusion defined by* (8.1) *and* (8.2) *admits a solution $m : [0, \infty) \to \mathcal{P}(\mathcal{S})$.*

Proof. We show that the conditions of Lemma 8.3 are satisfied because this directly proves the desired claim. We note that $\mathcal{P}(\mathcal{S})$ is closed.

We first show that F is upper semi-continuous, which means that condition (i) is satisfied: Let $m \in \mathcal{P}(\mathcal{S})$ and let $N \subseteq \mathbb{R}^S$ be an open set such that $F(m) \subseteq N$. Since $2^{\mathbb{R}^S}$ is a metric space given $H(\cdot,\cdot)$ we find an $\epsilon > 0$ such that $N_\epsilon(F(m)) \subseteq N$. We now want to find a $\delta > 0$ such that for all $\tilde{m} \in N_\delta(m)$ we have $F(\tilde{m}) \in N_\epsilon(F(m))$.

The set $\mathcal{D}(m)$ consists of all strategies $d \in D^s$ that satisfy $V^d(m) = V^*(m)$. Thus, for any $d' \in D^s \setminus \mathcal{D}(m)$ we find a constant $c_{d'} > 0$ such that $V^{d'}(m) < V^*(m) - c_{d'}$. In particular, since D^s is finite, the constant $c := \min_{d' \in D^s \setminus \mathcal{D}(m)} c_{d'}$ satisfies $V^{d'}(m) < V^*(m) - c$ for all $d' \in D^s \setminus \mathcal{D}(m)$. Since for any $d \in D^s$ the function $V^d : \mathcal{P}(\mathcal{S}) \to \mathbb{R}^S$ is continuous, we find a $\delta_d > 0$ such that for all $\tilde{m} \in N_{\delta_d}(m)$ we have $|V^d(\tilde{m}) - V^d(m)| < \frac{c}{3}$. In particular for $\delta_1 = \min_{d \in D^s} \delta_d$ we have for all $\tilde{m} \in N_{\delta_1}(m)$ and all strategies $d' \in D^s \setminus \mathcal{D}(m)$ that

$$\begin{aligned} V^*(\tilde{m}) - V^{d'}(\tilde{m}) &\geq (V^*(\tilde{m}) - V^*(m)) + (V^*(m) - V^{d'}(m)) + V^{d'}(m) - V^{d'}(\tilde{m})) \\ &> -\frac{c}{3} + c - \frac{c}{3} = \frac{c}{3} > 0 \end{aligned}$$

holds pointwise. This implies that $d' \notin D(\tilde{m})$, which proves that $\mathcal{D}(\tilde{m}) \subseteq \mathcal{D}(m)$.

For each $d \in \mathcal{D}(m)$ we furthermore note that

$$F^d : \mathcal{P}(\mathcal{S}) \to \mathbb{R}^S, m \mapsto \left(\sum_{i \in \mathcal{S}} \sum_{a \in \mathcal{A}} m_i Q_{ija}(m) d_{ia} \right)_{j \in \mathcal{S}}$$

is continuous. Thus, there is a $\delta_{2,d} > 0$ such that all $\tilde{m} \in N_{\delta_{2,d}}(m)$ satisfy $|F^d(\tilde{m}) - F^d(m)| < \epsilon$. Therefore, the choice $\delta = \min\{\delta_1, \min_{d \in \mathcal{D}(m)} \delta_{2,d}\}$ has the desired properties.

Since for all $m \in \mathcal{P}(\mathcal{S})$ the set $F(m)$ is a convex polytope, the set $F(m)$ is closed and convex for all $m \in \mathcal{P}(\mathcal{S})$ (Brøndsted, 1983, Corollary 2.9). Thus, condition (ii) is satisfied.

Next, we check that condition (iii) is satisfied, that is we check that $||F(m)||$ is indeed bounded by $c(1 + ||m||)$ for some $c > 0$: Since $Q_{ija}(m)$ is uniformly bounded by L for all $i, j \in \mathcal{S}$, $a \in \mathcal{A}$ and $m \in \mathcal{P}(\mathcal{S})$ we have that

$$||F^d(m)||_1 = \sum_{j \in \mathcal{S}} \left| \sum_{i \in \mathcal{S}} \sum_{a \in \mathcal{A}} m_i Q_{ija}(m) d_{ia} \right| \leq \sum_{j \in \mathcal{S}} \sum_{i \in \mathcal{S}} \sum_{a \in \mathcal{A}} m_i L d_{ia} = SL,$$

which yields that $c = SL$ is a suitable choice.

Finally, we show that the tangency condition (iv) is satisfied, that is for any $m \in \mathcal{P}(\mathcal{S})$ we have $F(m) \cap T_{\mathcal{P}(\mathcal{S})}(m) \neq \emptyset$. For this we utilize the characterization of $T_{\mathcal{P}(\mathcal{S})}(m)$ from Lemma 8.2: We first note that for all points $m \in \text{int}(\mathcal{P}(\mathcal{S}))$ we have $T_{\mathcal{P}(\mathcal{S})} = \mathbb{R}^n$ as well as $F(m) \neq \emptyset$ because $\mathcal{D}(m) \neq \emptyset$ by Theorem 5.2. Thus, it suffices to consider the boundary points of $\mathcal{P}(\mathcal{S})$. Let $m \in \mathcal{P}(\mathcal{S})$ be a boundary point. Then there is at least one $j \in \mathcal{S}$ such that $m_j = 0$. Thus, for any strategy $d \in D^s$ the vector $\left(\sum_{i \in \mathcal{S}} \sum_{a \in \mathcal{A}} m_i Q_{ija}(m) d_{ia}\right)_{j \in \mathcal{S}}$ will have non-negative entries at each $j \in \mathcal{S}$ for which $m_j = 0$ because the only non-positive column entry of $Q_{\cdot ja}(m) d_{\cdot a}$ in row j will get weight $m_j = 0$. Furthermore, since $Q(m)$ is conservative we obtain that

$$\sum_{j \in \mathcal{S}} \sum_{i \in \mathcal{S}} \sum_{a \in \mathcal{A}} m_i Q_{ija}(m) d_{ia} = \sum_{i \in \mathcal{S}} \sum_{a \in \mathcal{A}} m_i \underbrace{\left(\sum_{j \in \mathcal{S}} Q_{ija}(m) \right)}_{=0} d_{ia} = 0.$$

Thus, we have shown that all those points of which we build the convex combinations indeed lie in $T_{\mathcal{P}(\mathcal{S})}(m)$. Since $T_{\mathcal{P}(\mathcal{S})}(m)$ is convex itself, this proves the claim. □

8.4. Local Convergence towards Deterministic Stationary Equilibria

This section discusses the question, whether the myopic adjustment process converges locally to a stationary mean field equilibrium, which means that whenever we start the process "close" to a certain stationary equilibrium it will converge towards this equilibrium. This question is of interest because it justifies the consideration of stationary mean field equilibria: Indeed, such a local convergence result yields that stationary mean field equilibria are not only a good prediction of behaviour whenever the population distribution is exactly the stationary distribution in the equilibrium, but also for a whole neighbourhood of that equilibrium.

Local convergence of the myopic adjustment process will in general not happen if the equilibrium lies on the boundary of the optimality set. Indeed, in this case there are points arbitrary close to the equilibrium at which a different strategy is optimal and for which we cannot tell anything about the behaviour of the trajectories given this strategy – it might happen that we are pushed away from the equilibrium. The same argument also yields that we cannot expect convergence towards randomized equilibria in general because also in this case several strategies are optimal in an arbitrarily small neighbourhood of the equilibrium. Again several trajectories which might push the processes away from the equilibrium are possible.

Thus, it remains to verify whether we have local convergence towards a stationary deterministic mean field equilibrium $(\bar{m}, d)$ where d is the unique optimal strategy at $\bar{m}$. Since d is the unique optimal strategy given $\bar{m}$, there is an $\epsilon > 0$ such that for all $\tilde{m} \in N_\epsilon(\bar{m})$ we have $\mathcal{D}(\tilde{m}) = \{d\}$. This yields that for all $\tilde{m} \in N_\epsilon(\bar{m})$ we have $F(\tilde{m}) = \{(Q^d(\tilde{m}))^T \tilde{m}\}$. Thus, it suffices to investigate whether there is a $\delta > 0$ such that for all $m_0 \in \mathcal{P}(\mathcal{S})$ satisfying $|m_0 - \bar{m}| < \delta$ we have that the solution of $\dot{m}(t) = (Q^d(m(t)))^T m(t)$ lies in $N_\epsilon(\bar{m})$ for all $t \geq 0$ and converges towards $\bar{m}$.

This question is closely linked to the notion of asymptotically stable solutions of autonomous ordinary differential equations $\dot{x} = f(x)$ (see §29 in Walter (1998) for definitions and classical results). However, our notion is indeed weaker since we do not allow arbitrary starting points $m_0 \in \mathbb{R}^S$, but only starting points that lie in $\mathcal{P}(\mathcal{S})$. Moreover, our ordinary differential equations are not of a type that allow to apply classical theorems from the theory of ordinary equations. Indeed, our differential equation only admits solutions that stay in the set $\mathcal{P}(\mathcal{S})$, which in particular implies that the Jacobian always has a zero eigenvalue. For this reason we have to derive our own methods to prove that there is a $\delta > 0$ such that for any starting value $m_0 \in \mathcal{P}(\mathcal{S}) \cap N_\delta(\bar{m})$ the solution of $\dot{m}(t) = (Q^d(m(t)))^T m(t)$ stays in $N_\epsilon(\bar{m})$ for all $t \geq 0$ and converges towards $\bar{m}$.

In case of constant dynamics, we prove a first result relying on the explicit characterization of solutions of homogeneous differential equations with constant coefficients.

Theorem 8.5. *Assume that $\bar{m}$ is a stationary deterministic mean field equilibrium with equilibrium strategy d such that d is the unique optimal strategy at $\bar{m}$ (that is $\mathcal{D}(\bar{m}) = \{d\}$). Furthermore, assume that Q^d is a constant irreducible generator matrix. Then there is a $\delta > 0$ such that any solution of the myopic adjustment process $\dot{m}(t) \in F(m(t))$ with initial condition $m_0 \in N_\delta(\bar{m}) \cap \mathcal{P}(\mathcal{S})$ converges exponentially fast to $\bar{m}$, which means that we find constants $C_1, C_2 > 0$ such that*

$$||m(t) - \bar{m}|| \leq C_1 e^{-C_2 t} \quad \text{for all } t \geq 0.$$

More precisely, let $\lambda_1, \ldots, \lambda_n$ be the eigenvalues of $(Q^d)^T$ with $\lambda_1 = 0$ being a simple eigenvalue with eigenvector $\bar{m}$ and for all $i \in \{2, \ldots, n\}$ let $(v_i^0, \ldots, v_i^{m(\lambda_i)-1})$ be a basis of the generalized eigenspace $Eig(\lambda_i) := ker((Q^d)^T - \lambda_i I)^{m(\lambda_i)}$, where $m(\lambda_i)$ denotes the algebraic multiplicity of the eigenvalue λ_i. Define

$$C_i^k := \sum_{l=0}^{m(\lambda_i)-1} e^{-l} \frac{l^l}{l!(-Re(\lambda_i))^l} ||((Q^d)^T - \lambda_i I)^l v_i^k||$$

and let $\epsilon > 0$ be such that for all $\tilde{m} \in N_\epsilon(\bar{m})$ it holds that $\mathcal{D}(\tilde{m}) = \{d\}$, then we have that

$$\delta = \frac{\epsilon/2 \min_{i,k} ||v_i^k||}{\max_{i,k} C_i^k}$$

is a suitable choice.

We note that the ordinary differential equation $\dot{m} = (Q^d)^T m, m(0) = m_0$ is an autonomous homogeneous differential equation with constant coefficients, for which the solution is given by $t \mapsto \exp((Q^d)^T \cdot t) \cdot m_0$. By Logemann and Ryan (2014, Theorem 2.11) we furthermore have that if λ is an eigenvector of $(Q^d)^T$ with algebraic multiplicity $m(\lambda)$ and m_0 lies in generalized eigenspace $\text{Eig}(\lambda)$, then the solution is given by

$$x(t) = \exp((Q^d)^T \cdot t) m_0 = e^{\lambda t} \sum_{l=0}^{m(\lambda)-1} \frac{t^l}{l!} ((Q^d)^T - \lambda I)^l m_0.$$

With these preparations we prove the theorem:

Proof of Theorem 8.5. Let $\epsilon > 0$ be such that for all $\tilde{m} \in N_\epsilon(m)$ we have $\mathcal{D}(\tilde{m}) = \{d\}$.

Since the generator matrix $(Q^d)^T$ is irreducible it has zero as simple eigenvalue with corresponding eigenvalue $\bar{m}$ and all other eigenvalues have strictly negative real parts (Asmussen, 2003, Corollary II.4.9). Let $\lambda_1, \lambda_2, \ldots, \lambda_n$ be the distinct eigenvalues of $(Q^d)^T$ with $\lambda_1 = 0$ and $\mathrm{Re}(\lambda_i) < 0$ for all $i \geq 2$. For each $i \geq 2$ let $v_i^0, \ldots, v_i^{m(\lambda_i)-1}$ be a basis of $\mathrm{Eig}(\lambda_i)$. Then $(\bar{m}, v_2^0, \ldots, v_2^{m(\lambda_2)-1}, \ldots, v_n^0, \ldots, v_n^{m(\lambda_n)-1})$ is a basis of $\mathbb{R}^S$. Thus, for each initial condition $m_0 \in \mathbb{R}^S$ we find a unique representation

$$m_0 = \alpha_1 \bar{m} + \sum_{i=2}^{n} \sum_{k=0}^{m(\lambda_i)-1} \alpha_i^k v_i^k.$$

By the previously stated explicit representation of the solution of the ordinary differential equation, we obtain that

$$\begin{aligned} m(t) = e^{(Q^d)^T \cdot t} m_0 &= \alpha_1 e^{(Q^d)^T \cdot t} \bar{m} + \sum_{i=2}^{n} \sum_{k=0}^{m(\lambda_i)-1} \alpha_i^k e^{(Q^d)^T \cdot t} v_i^k \\ &= \alpha_1 \bar{m} + \sum_{i=2}^{n} e^{\lambda_i t} \sum_{k=0}^{m(\lambda_i)-1} \alpha_i^k \sum_{l=0}^{m(\lambda_i)-1} \frac{t^l}{l!} ((Q^d)^T - \lambda_i I)^l v_i^k, \end{aligned}$$

where the last line follows from the fact that $\bar{m}$ is the eigenvector for the eigenvalue 0. Since the continuous time Markov chain with generator Q^d is ergodic, we have that $m(t) \to \bar{m}$. Thus, it holds that $\alpha_1 = 1$.

Using this we obtain that

$$\begin{aligned} ||m(t) - \bar{m}|| &= \left\| \sum_{i=2}^{n} e^{\lambda_i t} \sum_{k=0}^{m(\lambda_i)-1} \alpha_i^k \sum_{l=0}^{m(\lambda_i)-1} \frac{t^l}{l!} ((Q^d)^T - \lambda_i I)^l v_i^k \right\| \\ &\leq \sum_{i=2}^{n} \sum_{k=0}^{m(\lambda_i)-1} |\alpha_i^k| \left(\sum_{l=0}^{m(\lambda_i)-1} e^{\mathrm{Re}(\lambda_i) t} \frac{t^l}{l!} || \left((Q^d)^T - \lambda_i I \right)^l v_i^k || \right). \end{aligned} \tag{8.4}$$

Since for fixed $l \in \mathbb{N}$ the function $t \mapsto e^{\mathrm{Re}(\lambda_i) t} \frac{t^l}{l!}$ has a unique global maximum point in $[0, \infty)$ at $t = -l/\mathrm{Re}(\lambda_i)$ and the maximal value is given by $e^{-l} \frac{l^l}{l!(-\mathrm{Re}(\lambda_i))^l}$, we obtain that

$$||m(t) - \bar{m}|| \leq \sum_{i=2}^{n} \sum_{k=0}^{m(\lambda_i)-1} |\alpha_i^k| C_i^k$$

with

$$C_i^k = \sum_{l=0}^{m(\lambda_i)-1} e^{-l} \frac{l^l}{l!(-\mathrm{Re}(\lambda_i))^l} ||((Q^d)^T - \lambda_i I)^l v_i^k||,$$

which does not depend on t anymore. Now whenever the constants α_i^k are such that

$$\sum_{i=2}^{n} \sum_{k=0}^{m(\lambda_i)-1} |\alpha_i^k| C_i^k < \epsilon,$$

then $m(t)$ will stay in $N_\epsilon(\bar{m})$ for all $t \geq 0$. This condition is especially satisfied if

$$||m_0 - \bar{m}|| = \sum_{i=2}^{n} \sum_{k=0}^{m(\lambda_i)-1} |\alpha_i^k| ||v_i^k|| < \frac{\epsilon \min_{i,k} ||v_i^k||}{\max_{i,k} C_i^k} =: \delta.$$

The exponential convergence follows from (8.4). □

In the case of more general dynamics it is necessary to require that $Q^d(\cdot)$ is defined on an open set $O \supseteq \mathcal{P}(\mathcal{S})$ in order to utilize differential calculus. Moreover, the necessary eigenvalue structure does not directly follow from the irreducibility of $Q(m)$, but we have to require it.

Theorem 8.6. *Let $(\bar{m}, d)$ be a stationary deterministic mean field equilibrium with equilibrium strategy d such that d is the unique optimal strategy at $\bar{m}$ (that is $\mathcal{D}(\bar{m}) = \{d\}$). Let $O \supseteq \mathcal{P}(\mathcal{S})$ be an open set such that $Q^d : O \to \mathbb{R}^{S \times S}$ is componentwise Lipschitz continuous, the matrix $Q^d(m)$ is a transition rate matrix for all $m \in \mathcal{P}(\mathcal{S})$ and the function $f^d : O \to \mathbb{R}^S, m \mapsto (Q^d(m))^T m$ is continuously differentiable in m. Assume further that the Jacobian $\frac{\partial}{\partial m} f^d(m)$ has a zero eigenvalue with eigenvector $\bar{m}$ and all other eigenvalues have strictly negative real part. Then there is an $\delta > 0$ such that any solution of the myopic adjustment process with $m_0 \in N_\delta(\bar{m}) \cap \mathcal{P}(\mathcal{S})$ converges exponentially fast to $\bar{m}$, which means that we find constants $C_1, C_2 > 0$ such that*

$$||m(t) - \bar{m}|| \leq C_1 e^{-C_2 t} \quad \text{for all } t \geq 0.$$

Remark 8.7. This result covers the previous result for case of constant dynamics. However, the proof of this result is non-constructive. In particular, we cannot explicitly compute δ, which has been possible in the constant dynamics setting. △

As in the discussion of nonlinear sinks in Hirsch and Smale (1974, Section 9.1) we define a suitable inner product $\langle \cdot, \cdot \rangle_B$ on $\mathbb{R}^S$ and bound $\langle f^d(m), m \rangle_B$ in order to bound $\frac{\partial}{\partial t} ||m(t)||_B$ by $-c||m(t)||_B$, which then allows to prove exponential convergence. In order to follow this programme we first remind ourselves that we can describe an inner product $\langle \cdot, \cdot \rangle_B$ on $\mathbb{R}^S$ by specifying a basis $B = (b_1, \ldots, b_n)$ as well as $\langle b_i, b_j \rangle_B$ for all $1 \leq i \leq j \leq S$. The following lemma found in Hirsch and Smale (1974, Chapter 7) is a main tool to prove the theorem:

Lemma 8.8. *Let $A \in \mathbb{R}^{S\times S}$ and let the Jordan normal form of A consist of blocks $A_1, \dots, A_J$ where λ_i is the eigenvalue of block A_i. Then there is a basis*

$$B = (b_1^0, \dots, b_1^{m(\lambda_1)-1}, \dots, b_J^0, \dots, b_J^{m(\lambda_J)-1})$$

and an inner product $\langle \cdot, \cdot \rangle_B$ such that

(i) for all $i \in \{1, \dots, J\}$ the family $(b_i^0, \dots, b_i^{m(\lambda_i)-1})$ is a basis of the generalized eigenspace Eig(λ_i),

(ii) $\langle b_i^l, b_j^k \rangle_B = 0$ whenever $i \neq j$ or $l \neq k$ as well as $\langle b_i^l, b_i^l \rangle_B = 1$ for all $i \in \{1, \dots, J\}$, $l \in \{0, \dots, m(\lambda_i) - 1\}$, and

(iii) for all $i \in \{1, \dots, J\}$ and $x \in$ span$(b_i^0, \dots, b_i^{m(\lambda_i)-1})$ and all $\beta >$ Re(λ_i) we have

$$\langle Ax, x \rangle_B \leq \beta ||x||_B^2.$$

Proof of Theorem 8.6. Denote by $A := \frac{\partial}{\partial m} f^d(\bar{m})$ the Jacobian matrix of $f^d(\cdot)$ at the point $\bar{m}$. Let $\lambda_1, \lambda_2, \dots, \lambda_J$ be the eigenvalues of A and let $\lambda_1 = 0$; by assumption $\mathrm{Re}(\lambda_i) < 0$ for all $i \in \{2, \dots, J\}$. Thus, there are constants $b, c > 0$ such that $\mathrm{Re}(\lambda_i) < -b < -c$ for all $i \in \{2, \dots, J\}$.

The Jordan normal form of A has a block A_1 for the eigenvalue 0 with multiplicity 1 and eigenvector $\bar{m}$ and blocks A_i $(i = 2, \dots, J)$ with eigenvalue λ_i satisfying $\mathrm{Re}(\lambda_i) < -b$. Then by Lemma 8.8 we find a basis $B = (b_1^0, b_2^0, \dots, b_2^{m(\lambda_2)-1}, \dots, b_J^0, \dots, b_J^{m(\lambda_J)-1})$ with $b_1^0 = \bar{m}$ such that

- $(b_i^0, \dots, b_i^{m(\lambda_i)-1})$ is a basis of the generalized eigenspace $\mathrm{Eig}(\lambda_i)$ for all $i \in \{2, \dots, J\}$,
- $\langle b_i^l, b_j^k \rangle_B = 0$ whenever $i \neq j$ or $l \neq k$ as well as $\langle b_i^l, b_i^l \rangle_B = 1$ for all $i \in \{1, \dots, J\}$, $l \in \{0, \dots, m(\lambda_i) - 1\}$, and
- for all vectors $x = 0 \cdot \bar{m} + \sum_{i=2}^{J} \sum_{l=0}^{m(\lambda_i)-1} \alpha_i^l b_i^l$ we have

  $$\langle Ax, x \rangle_B \leq -b ||x||_B^2,$$

  since

$$\begin{aligned}
\langle Ax, x \rangle_B &= \left\langle A \cdot \left(\sum_{i=2}^{J} \sum_{l=0}^{m(\lambda_i)-1} \alpha_i^l b_i^l \right), \sum_{i=2}^{J} \sum_{l=0}^{m(\lambda_i)-1} \alpha_i^l b_i^l \right\rangle_B \\
&= \sum_{i=2}^{J} \left\langle A \left(\sum_{l=0}^{m(\lambda_i)-1} \alpha_i^l b_i^l \right), \sum_{l=0}^{m(\lambda_i)-1} \alpha_i^l b_i^l \right\rangle_B
\end{aligned}$$

$$
\begin{aligned}
&+\sum_{i=2}^{J}\sum_{j=2,j\neq i}^{J}\left\langle \sum_{l=0}^{m(\lambda_i)-1}\alpha_i^l A b_i^l, \sum_{k=0}^{m(\lambda_j)-1}\alpha_j^k b_j^k \right\rangle_B \\
&\leq -b\sum_{i=2}^{J}\left\|\left(\sum_{l=0}^{m(\lambda_i)-1}\alpha_i^l b_i^l\right)\right\|_B^2 + 0 \\
&= -b\sum_{i=2}^{J}\left\langle \sum_{l=0}^{m(\lambda_i)-1}\alpha_i^l b_i^l, \sum_{l=0}^{m(\lambda_i)-1}\alpha_i^l b_i^l \right\rangle_B \\
&= -b\left\langle \sum_{i=2}^{J}\sum_{l=0}^{m(\lambda_i)-1}\alpha_i^l b_i^l, \sum_{i=2}^{J}\sum_{l=0}^{m(\lambda_i)-1}\alpha_i^l b_i^l \right\rangle_B \\
&= -\left\|\sum_{i=2}^{J}\sum_{l=0}^{m(\lambda_i)-1}\alpha_i^l b_i^l\right\|_B^2 .
\end{aligned}
$$

We remark that the first inequality follows from property (iii) of the basis B as well as the fact that $Ab_i^j \in \mathrm{Eig}(\lambda_i)$ together with property (ii) of the basis. The penultimate equality follows again from property (ii) of the basis.

Let $x = \alpha_1^0\bar{m}+\sum_{i=2}^{J}\sum_{l=0}^{m(\lambda_i)-1}\alpha_i^l b_i^l \in (\mathcal{P}(\mathcal{S})-\bar{m}) := \{x \in \mathbb{R}^S : \exists m \in \mathcal{P}(\mathcal{S}) : x = m-\bar{m}\}$, then we have

$$
\begin{aligned}
\langle Ax, x\rangle_B &= \left\langle A\left(\alpha_1^0\bar{m}+\sum_{i=2}^{J}\sum_{l=0}^{m(\lambda_i)-1}\alpha_i^l b_i^l\right), \alpha_1^0\bar{m}+\sum_{i=2}^{J}\sum_{l=0}^{m(\lambda_i)-1}\alpha_i^l b_i^l\right\rangle_B \\
&= (\alpha_1^0)^2\langle A\bar{m},\bar{m}\rangle_B + \left\langle \alpha_1^0 A\bar{m}, \sum_{i=2}^{J}\sum_{l=0}^{m(\lambda_i)-1}\alpha_i^l b_i^l\right\rangle_B + \left\langle \sum_{i=2}^{J}\sum_{l=0}^{m(\lambda_i)-1}\alpha_i^l A b_i^l, \alpha_0^1\bar{m}\right\rangle_B \\
&\quad + \left\langle A\cdot\sum_{i=2}^{J}\sum_{l=0}^{m(\lambda_i)-1}\alpha_i^l b_i^l, \sum_{i=2}^{J}\sum_{l=0}^{m(\lambda_i)-1}\alpha_i^l b_i^l\right\rangle_B \\
&= 0+0+0+\left\langle A\cdot\sum_{i=2}^{J}\sum_{l=0}^{m(\lambda_i)-1}\alpha_i^l b_i^l, \sum_{i=2}^{J}\sum_{l=0}^{m(\lambda_i)-1}\alpha_i^l b_i^l\right\rangle_B \\
&\leq -b\cdot\left\|\sum_{i=2}^{J}\sum_{l=0}^{m(\lambda_i)-1}\alpha_i^l b_i^l\right\|_B^2 \\
&= -b\cdot\underbrace{\frac{\left\|\sum_{i=2}^{J}\sum_{l=0}^{m(\lambda_i)-1}\alpha_i^l b_i^l\right\|_B^2}{\left\|\alpha_1^0\bar{m}+\sum_{i=2}^{J}\sum_{l=0}^{m(\lambda_i)-1}\alpha_i^l b_i^l\right\|_B^2}}_{=:D(\alpha)>0}\cdot\left\|\underbrace{\alpha_1^0\bar{m}+\sum_{i=2}^{J}\sum_{l=0}^{m(\lambda_i)-1}\alpha_i^l b_i^l}_{=x}\right\|_B^2 .
\end{aligned}
$$

We remark that the third equality follows from $A\bar{m} = 0$ and since $Ab_i^l \in \mathrm{Eig}(\lambda_i)$, which by property (iii) yields that $\left\langle A \cdot b_i^l, \bar{m} \right\rangle_B = 0$. Furthermore, we note that $D(\alpha) > 0$ holds because $\bar{m} \notin \mathcal{P}(\mathcal{S}) - \bar{m}$.

In total, we obtain that for any $x = \alpha_1^0 \bar{m} + \sum_{i=2}^J \sum_{l=0}^{m(\lambda_i)-1} \alpha_i^l b_i^l \in \mathcal{P}(\mathcal{S}) - \bar{m}$ it holds that

$$\langle Ax, x \rangle_B \leq -b \cdot D(\alpha) \cdot ||x||_B^2.$$

Since $\mathcal{P}(\mathcal{S}) - \bar{m}$ is compact and

$$\alpha \mapsto \alpha_1^0 \bar{m} + \sum_{i=2}^{J} \sum_{l=0}^{m(\lambda_i)-1} \alpha_i^l b_i^l$$

is a homeomorphism, also

$$P := \Big\{ \alpha = (\alpha_1^0, \alpha_2^0, \ldots, \alpha_2^{m(\lambda_2)-1}, \ldots, \alpha_J^0, \ldots, \alpha_J^{m(\lambda_J)-1}) \in \mathbb{R}^S : \\ \alpha_1^0 \bar{m} + \sum_{i=2}^{J} \sum_{l=0}^{m(\lambda_i)-1} \alpha_i^l b_i^l \in \mathcal{P}(\mathcal{S}) - \bar{m} \Big\}$$

is compact. This implies that

$$\langle Ax, x \rangle_B \leq -b \underbrace{\left(\min_{\alpha \in P} D(\alpha) \right)}_{=:D>0} ||x||_B^2.$$

As a final preparation, we note that $\mathcal{P}(\mathcal{S})$ is flow invariant for $\dot{m}(t) = f^d(m(t))$ (that is any trajectory with $m_0 \in \mathcal{P}(\mathcal{S})$ will stay in $\mathcal{P}(\mathcal{S})$ for all $t \geq 0$). This follows from the classical flow invariance theorem for ordinary differential equations (see Subsection 9.1.1) and the characterization of the tangent cone for $\mathcal{P}(\mathcal{S})$ presented in the proof of Theorem 8.4.

With these preparations we prove the theorem: We give $\mathbb{R}^S$ new coordinates via the transformation $x = m - \bar{m}$, which in particular means that we now consider $\tilde{f}^d(x) = f^d(x + \bar{m})$. By definition of the derivative, which in particular yields that $||\tilde{f}^d(x) - Ax||_B \in o(||x||_B)$ in a neighbourhood of 0, and Cauchy's inequality we have

$$0 = \lim_{x \to 0} \frac{||\tilde{f}^d(x) - Ax||_B}{||x||_B} = \lim_{x \to 0} \frac{||\tilde{f}^d(x) - Ax||_B \cdot ||x||_B}{||x||_B^2} \geq \lim_{x \to 0} \frac{\langle \tilde{f}^d(x) - Ax, x \rangle_B}{||x||_B^2} = 0.$$

If $x \in (\mathcal{P}(\mathcal{S}) - \bar{m})$, then $\langle Ax, x \rangle_B \leq -b \cdot D \cdot \underline{||x||_B^2} < 0$. This and Lemma 7.15 now imply that there is an $\delta > 0$ such that for all $x \in \overline{N_\delta(0)} \cap (\mathcal{P}(\mathcal{S}) - \bar{m})$ we have

$$\langle \tilde{f}^d(x), x \rangle_B \leq -c \cdot D \cdot ||x||_B^2 \quad \text{and} \quad \mathcal{D}(\bar{m} + x) = \{d\}.$$

Now assume that $x_0 \in N_\delta(0) \cap (\mathcal{P}(\mathcal{S}) - \bar{m})$. By Peano's existence theorem (Theorem 10.IX in Walter (1998)) there is a solution of the initial value problem $\dot{x}(t) = \tilde{f}^d(x(t))$, $x(0) = x_0$. Moreover, any solution can be extend onto $[0, \infty)$.

Let $x : [0, t_0] \to \mathbb{R}^S$ be a solution curve of the differential equation $\dot{x}(t) = \tilde{f}^d(x(t))$ in $N_\delta(0)$ and assume that $x(t) \neq 0$ for all $0 \leq t \leq t_0$. (If $x(\tilde{t}) = 0$ for some $\tilde{t} \geq 0$, then $x(t) = 0$ for all $t \geq \tilde{t}$.)

Since there is a symmetric, positive definite matrix $C \in \mathbb{R}^{S \times S}$ such that $\langle x, y \rangle_B = x^T C y$, the product rule

$$\frac{\partial}{\partial t}\langle g(t), h(t)\rangle_B = \langle \frac{\partial}{\partial t} g(t), h(t)\rangle_B + \langle g(t), \frac{\partial}{\partial t} h(t)\rangle_B$$

holds, which yields that

$$\frac{\partial}{\partial t}||x(t)||_B = \frac{\partial}{\partial t}\sqrt{\langle x(t), x(t)\rangle_B} = \frac{1}{||x(t)||_B}\langle \dot{x}(t), x(t)\rangle_B.$$

Since $\dot{x}(t) = \tilde{f}^d(x(t))$ and $\mathcal{P}(\mathcal{S}) - \bar{m}$ is flow invariant for $\dot{x}(t) = \tilde{f}^d(x(t))$, we have that $x(t) \in \mathcal{P}(\mathcal{S}) - \bar{m}$ and we obtain

$$\frac{\partial}{\partial t}||x(t)||_B = \frac{1}{||x(t)||_B}\langle f^d(x(t)), x(t)\rangle_B \leq -c \cdot D \cdot ||x(t)||_B. \tag{8.5}$$

This shows that $||x(t)||_B$ is strictly decreasing on $[0, t_0]$, which implies that $x(t) \in N_\delta(0)$ for all $t \in [t_0, t_0 + \epsilon]$. Repeating this argument, we obtain by Hirsch and Smale (1974, Section 8.5) $x(t) \in N_\delta(0)$ for all $t \geq 0$. Furthermore, the estimate (8.5) yields that

$$||x(t)||_B \leq e^{-cDt}||x(0)||_B,$$

which yields the desired exponential convergence. □

8.5. Application of the Results in the Examples of Chapter 6

For the consumer choice model with constant dynamics presented in Section 6.1 local convergence towards deterministic stationary equilibria is directly implied by Theorem 8.5, as long as the equilibrium point does not coincide with $(d_1, 1 - d_1)$ or $(d_2, 1 - d_2)$ because in the other case there would be an equilibrium point at which more than one strategy is optimal.

For the two models with non-constant dynamics, we have to check whether the Jacobian matrix $\frac{\partial}{\partial m}(Q^d(m))^T m$ given the equilibrium strategy d at the equilibrium point m has a simple 0 eigenvalue and all other eigenvalues have negative real part.

For the corruption model presented in Section 6.2 we obtain for the two candidate strategies $\{\textit{change}\} \times \{\textit{stay}\}$ and $\{\textit{stay}\} \times \{\textit{change}\}$ the following Jacobian matrices

$$\frac{\partial}{\partial m} f^{\{change\}\times\{stay\}}(m) = \begin{pmatrix} -b - q_{\text{soc}}m_H + q_{\text{inf}}m_H & -q_{\text{soc}}m_C + q_{\text{inf}}m_C & 0 \\ b - q_{\text{inf}}m_H & -q_{\text{inf}}m_C & \lambda \\ q_{\text{soc}}m_H & q_{\text{soc}}m_C & -\lambda \end{pmatrix}$$
$$\frac{\partial}{\partial m} f^{\{stay\}\times\{change\}}(m) = \begin{pmatrix} -q_{\text{soc}}m_H + q_{\text{inf}}m_H & b - q_{\text{soc}}m_C + q_{\text{inf}}m_C & 0 \\ -q_{\text{inf}}m_H & -b - q_{\text{inf}}m_C & \lambda \\ q_{\text{soc}}m_H & q_{\text{soc}}m_C & -\lambda \end{pmatrix}.$$

However, the eigenvalue structure is highly parameter dependent, we find parameters such that it is as desired and we find parameters such that it is not as desired. This behaviour is reasonable since at least the extreme stationary points $(0, 1, 0)$ and $(1, 0, 0)$ will not be sinks of the system.

For the consumer choice model presented in Section 6.3 there is the unique candidate equilibrium strategy of a stationary mean field equilibrium with a deterministic equilibrium strategy, namely $\{\textit{stay}\} \times \{\textit{change}\}$. For this strategy we obtain the following Jacobian matrix

$$\frac{\partial}{\partial m} f^{\{stay\}\times\{change\}}(m) = \begin{pmatrix} -2em_1 - \epsilon & b & \lambda \\ 0 & -b - 2em_2 - \epsilon & \lambda \\ 2em_1 + \epsilon & 2em_2 + \epsilon & -2\lambda \end{pmatrix}.$$

This matrix is a transition rate matrix and has by Asmussen (2003, Corollary II.4.9) the desired eigenvalue structure. Thus, as long as the stationary mean field equilibrium is such that the strategy $\{\textit{stay}\} \times \{\textit{change}\}$ is the unique optimal strategy for the equilibrium distribution we indeed obtain that the myopic adjustment process converges locally towards the equilibrium.

Remark 8.9. In several examples (in particular with similar congestion effects) it will be the case that the Jacobian is a transition rate matrix, which will then always guarantee local convergence. △

9. Global Convergence of the Myopic Adjustment Process towards Stationary Equilibria

In this chapter we prove that under certain conditions the myopic adjustment process converges (almost) globally towards some stationary mean field equilibrium, which means that irrespective of the initial population distribution the myopic adjustment process converges towards some mean field equilibrium. The general result will just state conditions that ensure that the process converges towards deterministic mean field equilibria or that the process will stay in a set (which usually is small) in which several strategies are optimal simultaneously. In these sets we can then often prove manually that also convergence towards stationary equilibria with a mixed equilibrium strategy hold. Additionally, nonlinear Markov chains with finite state space, which naturally arise in the investigations of global convergence, are discussed. More precisely, we prove that a Lipschitz-continuous nonlinear generator describes such a process, we illustrate through examples that the limit behaviour is more complex than in the case of standard time-homogeneous continuous time Markov chains and we provide a sufficient criterion for strong ergodicity of nonlinear Markov chains with small state spaces.

In Section 9.1 we introduce the necessary tools for the chapter. Namely, in Subsection 9.1.1 we review classical and alternative results to characterize flow invariance for ordinary differential equations and in Subsection 9.1.2 we discuss uniqueness as well as an explicit representation of all solutions to a particular type of differential inclusions. In Section 9.2 we introduce nonlinear Markov chains and consider their limit behaviour. Section 9.3 introduces the necessary assumptions as well as the main result of the chapter and Section 9.4 contains the proof of the theorem. Section 9.5 discusses in how far one of the assumptions, which is for some examples not fully satisfied, can be weakened. In Section 9.6 we propose an algorithm to verify the consistency condition of the theorem. Section 9.7 concludes the Chapter with the application of the results to the examples of Chapter 6.

9.1. Preliminaries

In this chapter we will extensively make use of the notion of flow invariance for autonomous differential equations and differential inclusions. A set D is flow invariant for an autonomous ordinary differential equation $\dot{x} = f(x)$ or an autonomous differential inclusion $\dot{x} \in F(x)$ if any trajectory that meets the set D stays in D afterwards. The first subsection 9.1.1 considers flow invariance for ordinary differential equations. More precisely, we review classical theorems and introduce, inspired by Fernandes and Zanolin (1987) and Blanchini and Miani (2015), a class of sets for which we obtain easy to handle conditions. In Subsection 9.1.2 we introduce a specific type of differential inclusions, which will occur in the discussion of the myopic adjustment process, and we review existence and uniqueness results as well as the explicit characterization of the solutions.

9.1.1. An Explicit Criterion to Check Flow Invariance

In Redheffer (1972) the classical theorems on flow invariance can be found in a unified framework. For ease of presentation we restrict ourselves to the setting of autonomous ordinary differential equations in $\mathbb{R}^S$: Let $\mathbb{R}^S$ be equipped with the dot product, let $O \subseteq \mathbb{R}^S$ be an open, connected set, let $f : O \to \mathbb{R}^S$ be a continuous function and consider the ordinary differential equation $\dot{x} = f(x)$ with initial condition $x(0) = x_0$. We say that a closed set $D \subseteq O$ is flow invariant if $x(t_0) \in D$ implies $x(t) \in D$ for all $t_0 \leq t < t_1$, where $[0, t_1)$ is the maximal interval of existence. If f is Lipschitz continuous, there are two classical sufficient conditions for flow invariance, which in our setting are even equivalent (Redheffer, 1972; Crandall, 1972):

Lemma 9.1. *Let D be a closed set, let f be a Lipschitz continuous function and assume that furthermore one of the following conditions is satisfied:*

(i) For any $x \in D$ and any outer normal vector $n(x)$ to x, which is a vector such that $N_{|n(x)|}(x + n(x)) \cap D = \emptyset$, we have $\langle n(x), f(x) \rangle \leq 0$.

(ii) For any $x \in D$ we have $f(x) \in T_D(x)$, with $T_D(x)$ being the tangent cone defined in Section 8.3.

Then D is flow invariant for the ordinary differential equation $\dot{x} = f(x)$.

However, these conditions are not helpful in all settings since the characterization of tangent cones or outer normal vectors, respectively, is complex in general. In order to achieve our goal of easy to handle conditions, we restrict our attention to sets that can be described as all those points that simultaneously satisfy several inequalities. Flow invariance for sets of a similar structure has been studied in the literature before, see Fernandes and Zanolin (1987) and Blanchini and Miani (2015). Formally, we consider sets

$$M := \{x \in O : g_i(x) \leq 0 \text{ for all } i \in \{1, \ldots, k\}\}$$

with $O \subseteq \mathbb{R}^S$ being an open set, $g_i : O \to \mathbb{R}$ ($i \in \{1, \ldots, k\}$) being continuously differentiable functions such that

(i) for all $x \in \partial M$ there is an $i \in \{1, \ldots, k\}$ such that $g_i(x) = 0$

(ii) for all $i \in \{1, \ldots, k\}$ and $x \in \{y \in O : g_i(y) = 0\} \cap M$ we have $\nabla g_i(x) \neq 0$.

For these sets we provide a much simpler sufficient condition for flow invariance:

Lemma 9.2. *Let $f : O \to \mathbb{R}^S$ be a Lipschitz continuous function and assume that for each $x \in \partial M$ and each $i \in \{1, \ldots, k\}$ such that $g_i(x) = 0$ we have*

$$\langle f(x), \nabla g_i(x) \rangle \leq 0,$$

then the set M is flow invariant with respect to the ordinary differential equation $\dot{x} = f(x)$.

Proof. Let $x_0 \in M$ and let $x(t)$ be the unique solution of $\dot{x} = f(x)$ that exists on $[0, \infty)$ (Walter, 1998, Theorem 10.VI). Assume that there is a $\tilde{t} > 0$ such that $x(t) \in M$ for all $0 \leq t \leq \tilde{t}$ and $x(t) \notin M$ for all $t \in (\tilde{t}, \tilde{t} + \epsilon_1)$ for some $\epsilon_1 > 0$. By definition of M and continuity of the functions g_i ($i \in \{1, \ldots, k\}$) we find an $\epsilon_2 > 0$ such that for some $j \in \{1, \ldots, k\}$ we have $g_j(x(\tilde{t} + s)) > 0$ for all $s \in (0, \epsilon_2)$. By definition of M it holds that $g_j(x(\tilde{t})) \leq 0$. Moreover, the continuity of $g_j(\cdot)$ implies that $g_j(x(\tilde{t})) = 0$. Therefore, the left derivative satisfies

$$0 < \frac{\partial}{\partial t} g_i(x(\tilde{t})) = \langle f(x(\tilde{t})), \nabla g_i(x(\tilde{t})) \rangle.$$

This is a contradiction to the condition that $\langle f(x), \nabla g_i(x) \rangle \leq 0$ for all $x \in \partial M$ and $i \in \{1, \ldots, k\}$ such that $g_i(x) = 0$. □

Furthermore, we can also provide a criterion that yields a sufficient condition for the flow invariance of the interior of the set M:

Lemma 9.3. *Let $f : O \to \mathbb{R}^S$ be a Lipschitz continuous function and assume that for each $x \in \partial M$ and each $i \in \{1, \ldots, k\}$ such that $g_i(x) = 0$ we have*

$$\langle f(x), \nabla g_i(x) \rangle < 0,$$

then the set int(M) is flow invariant with respect to the ordinary differential equation $\dot{x} = f(x)$.

Proof. We note that $x \in \text{int}(M)$ if and only if $g_i(x) < 0$ for all $i \in \{1, \dots, k\}$. We first show that $x \in \text{int}(M)$ implies $g_i(x) < 0$ for all $i \in \{1, \dots, k\}$: Assume that this is false, then there is a point $x \in \text{int}(M)$ such that $g_i(x) = 0$ for some $i \in \{1, \dots, k\}$. Since $\text{int}(M)$ is open, we find some $\epsilon > 0$ such that $N_\epsilon(x) \subseteq \text{int}(M) \subseteq M$. Therefore, by definition of M, it holds for all $\tilde{x} \in N_\epsilon(x)$ that $g_i(\tilde{x}) \leq 0$. Thus, x is a local maximum of $g_i(\cdot)$. However, this contradicts condition (ii) stating $\nabla g_i(x) \neq 0$. The converse implication directly follows from condition (i).

Let $x_0 \in \text{int}(M)$ and let $x(t)$ be the unique solution of $\dot{x} = f(x)$ that exists on $[0, \infty)$ (Walter, 1998, Theorem 10.VI). Assume that there is a $\tilde{t} > 0$ such that $x(t) \in \text{int}(M)$ for all $0 \leq t < \tilde{t}$ and $x(\tilde{t}) \in \partial M$. By condition (i) we find an $i \in \{1, \dots, k\}$ such that $g_i(x(\tilde{t})) = 0$ and by the previous observation regarding the interior we find that $g_i(x(t)) < 0$ for all $0 \leq t < \tilde{t}$, thus the left derivative satisfies

$$0 < \frac{\partial}{\partial t} g_i(x(\tilde{t})) = \langle f(x(\tilde{t})), \nabla g_i(x(\tilde{t})) \rangle,$$

which contradicts the condition $\langle f(x), \nabla g_i(x) \rangle < 0$. □

9.1.2. A Deeper Characterization of the Solutions of Particular Differential Inclusions

In this section we review selected results for a particular type of differential inclusions. This type of differential inclusions has been introduced in Filippov (1988, Chapter 2) in order to study discontinuous ordinary differential equations. In our case, the results will be helpful in two ways: First, they will enable us to prove flow invariance of the differential inclusion describing our myopic adjustment process in our setting. Second, they will prove helpful in the application of the main theorem by providing an explicit characterization of the solutions that stay in a set where several strategies are simultaneously optimal (see Section 9.7).

Let an open set $G \subseteq \mathbb{R}^S$ be separated into two sets G^- and G^+ by $T \subseteq \mathbb{R}^S$ being the set consisting of all points such that $\phi(x) = 0$ with $\phi(\cdot)$ being a twice continuously differentiable function such that for all $x \in T$ the gradient is non-vanishing. Furthermore, let $f^+ : G^+ \to \mathbb{R}^S$ and $f^- : G^- \to \mathbb{R}^S$ be continuous such that for each point $x^* \in \partial G^+$ and $x^* \in \partial G^-$, respectively, the limits $\lim_{x \to x^*, x \in G^+} f^+(x)$ and $\lim_{x \to x^*, x \in G^-} f^-(x)$ exists. The differential inclusion of interest is given by

$$\dot{x} \in F(x) = \begin{cases} \{f^-(x)\} & \text{if } x \in G^- \\ \{f^+(x)\} & \text{if } x \in G^+ \\ \text{conv}\{f^-(x), f^+(x)\} & \text{if } x \in T \end{cases}.$$

For this differential inclusion a solution exists (Filippov, 1988, Lemma 6.3, Theorem 7.1 and 7.2). Additionally, we obtain the following uniqueness theorem (Filippov, 1988, Theorem 10.2):

Lemma 9.4. *If at each $x \in T$ one of the following inequalities holds (possibly different inequalities for different points $x \in T$)*

$$\langle f^-(x), \nabla\phi(x)\rangle > 0, \quad \langle f^+(x), \nabla\phi(x)\rangle < 0, \tag{9.1}$$

then for any $x_0 \in G$ any two solutions of $x(t) \in F(x(t))$ satisfying $x(0) = x_0$ coincide on each interval on which they both exist and lie in the domain.

Furthermore, we can provide a deeper characterization, if we assume that the terms in (9.1) have uniform sign over T (Filippov, 1988, §4):

- $\langle f^-(x), \nabla\phi(x)\rangle > 0$ and $\langle f^+(x), \nabla\phi(x)\rangle > 0$: In this case a solution satisfying $x(t) \in T$ for some $t \geq 0$ will remain in G^+ thereafter.

- $\langle f^-(x), \nabla\phi(x)\rangle < 0$ and $\langle f^+(x), \nabla\phi(x)\rangle < 0$: In this case a solution satisfying $x(t) \in T$ for some $t \geq 0$ will remain in G^- thereafter.

- $\langle f^-(x), \nabla\phi(x)\rangle > 0$ and $\langle f^+(x), \nabla\phi(x)\rangle < 0$: In this case a solution satisfying $x(t) \in T$ for some $t \geq 0$ will remain in T afterwards. More precisely, the solution will satisfy for all $\tilde{t} \geq t$ the following ordinary differential equation

$$\dot{x} = \frac{\langle f^-(x), \nabla\phi(x)\rangle}{\underbrace{\langle f^-(x) - f^+(x), \nabla\phi(x)\rangle}_{>0}} f^+(x) + \frac{-\langle f^+(x), \nabla\phi(x)\rangle}{\underbrace{\langle f^-(x) - f^+(x), \nabla\phi(x)\rangle}_{>0}} f^-(x).$$

9.2. Nonlinear Markov Chains

This section introduces nonlinear continuous time Markov chains with finite state spaces in the spirit of Kolokoltsov (2010). The characteristic property of these processes is that the transition probabilities do not only depend on the state, but also on the distribution of the process. After the definition we will derive that given a Lipschitz continuous generator matrix there exists such a process by utilizing the standard theory for ordinary differential equations. Thereafter, we will discuss the limit behaviour of these processes. This has not been considered in the literature so far, the closest results have been obtained in Butkovsky (2014), Butkovsky (2012) and Saburov (2016), which discuss ergodicity criteria for discrete time nonlinear Markov processes relying on a generalization of Dobrushin ergodicity coefficient.

9.2.1. Definition and Characterization via Nonlinear Generators

In Kolokoltsov (2010, Section 1.1) nonlinear Markov chains with finite state space in continuous time have been introduced: For this let $(\Phi^t(\cdot))_{t\geq 0}$ be a semigroup of continuous transformations of $\mathcal{P}(\mathcal{S})$ and call it a *nonlinear Markov semigroup.* Let $P(t,m) = (P_{ij}(t,m))_{i,j\in\mathcal{S}}$ be a continuous family of stochastic matrices depending continuously on $t \geq 0$ and $m \in \mathcal{P}(\mathcal{S})$ such that the nonlinear Chapman-Kolmogorov equation

$$\sum_{i\in\mathcal{S}} m_i P_{ij}(t+s,m) = \sum_{i,k\in\mathcal{S}} m_i P_{ik}(t,m) P_{kj}\left(s, \sum_{l\in\mathcal{S}} m_l P_{l\cdot}(t,m)\right)$$

is satisfied. We will call $P(t,m)$ a *continuous family of nonlinear transition probabilities* with the interpretation that if the distribution of the process at time 0 is m, then the probability of being in state j if the initial state was i is given by $P_{ij}(t,m)$. Given a nonlinear Markov semigroup $(\Phi^t(\cdot))_{t\geq 0}$ we say that the family $P(t,m)$ is a *stochastic representation* of the semigroup $(\Phi^t)_{t\geq 0}$ if for all $m \in \mathcal{P}(\mathcal{S})$ and all $t \geq 0$ we have

$$\Phi^t_j(m) = \sum_{i\in\mathcal{S}} m_i P_{ij}(t,m). \tag{9.2}$$

Given a nonlinear Markov semigroup and its stochastic representation we then define the nonlinear Markov chain with initial distribution m_0 as the time-inhomogeneous Markov chain (in the sense of Section 2.1) with transition probabilities

$$p(s,i,t,j) = P_{ij}(t-s, \Phi^s(m_0)).$$

We remark that it is sufficient to define a nonlinear Markov chain by specifying a nonlinear Markov semigroup $(\Phi^t(\cdot))_{t\geq 0}$ of continuous transformations of $\mathcal{P}(\mathcal{S})$ since we can always construct a continuous family of nonlinear transition probabilities satisfying (9.2) from such a semigroup for example by setting $P_{ij}(t,m) = \Phi^t_j(m)$. Moreover, we remark that given a semigroup there might be several continuous families of nonlinear transition probabilities satisfying (9.2).

As in the theory of standard continuous time Markov chains the infinitesimal generator will be the cornerstone of the definition and analysis of such processes: Let $\Phi^t(m)$ be differentiable in $t = 0$ for all $m \in \mathcal{P}(\mathcal{S})$, then the *(nonlinear) infinitesimal generator* of the semigroup $(\Phi^t(\cdot))_{t\geq 0}$ is given by a transition rate matrix function $Q(\cdot)$ such that for

$$f(m) := \frac{\partial}{\partial t}\Phi^t(m)\big|_{t=0}$$

we have

$$f_j(m) = \sum_{i\in\mathcal{S}} m_i Q_{ij}(m)$$

for all $m \in \mathcal{P}(\mathcal{S})$. We remark that one can show that any differentiable nonlinear Markov semigroup has a nonlinear infinitesimal generator (Kolokoltsov, 2010, p.5).

However, in practice one is mainly interested in the converse problem: Given a transition rate matrix function (that is a function $Q : \mathcal{P}(\mathcal{S}) \to \mathbb{R}^{S \times S}$ mapping every vector m to a transition rate matrix $Q(m)$) is there a nonlinear Markov semigroup (and thus a continuous family of nonlinear transition probabilities specifying a nonlinear Markov chain) such that $Q(\cdot)$ is the nonlinear infinitesimal generator of this process. Relying on the semigroup identity $\Phi^{t+s} = \Phi^t \Phi^s$ this question can be rephrased as the following Cauchy problem: Is there for any $m_0 \in \mathcal{P}(\mathcal{S})$ a solution $(\Phi^t(m_0))_{t \geq 0}$ of

$$\frac{\partial}{\partial t}\Phi^t(m_0) = \Phi^t(m_0)Q(\Phi^t(m_0)), \quad \Phi^0(m_0) = m_0, \tag{9.3}$$

such that additionally $\Phi^t(m_0) \in \mathcal{P}(\mathcal{S})$ for all $t \geq 0$ and for any fixed $t \geq 0$ the function $\Phi^t(\cdot)$ is continuous?

This question can be answered in case of transition rate matrix functions $Q(\cdot)$ such that for all $i, j \in \mathcal{S}$ the function $Q_{ij}(\cdot)$ is Lipschitz continuous. The proof of this relies on the classical theory for ordinary differential equations (Theorem 10.VI and Theorem 12.VII in Walter (1998), Lemma 9.1):

Theorem 9.5. *Let $Q : \mathcal{P}(S) \to \mathbb{R}^{S \times S}$ and assume that $Q_{ij}(\cdot)$ is a Lipschitz continuous function for all $i, j \in \mathcal{S}$ and that $Q(m)$ is a transition rate matrix for all $m \in \mathcal{P}(\mathcal{S})$. Then there is a unique Markov semigroup $(\Phi^t(\cdot))_{t \geq 0}$ such that $Q(\cdot)$ is the infinitesimal generator for $(\Phi^t(\cdot))_{t \geq 0}$.*

Proof. We first note that $f(m) := \left(\sum_{i \in \mathcal{S}} m_i Q_{ij}(m)\right)_{j \in \mathcal{S}}$ is Lipschitz continuous on $\mathcal{P}(\mathcal{S})$. Indeed, let M be a Lipschitz constant for all functions $Q_{ij}(\cdot)$ $(i, j \in \mathcal{S})$ and L be the constant that uniformly bounds $Q_{ij}(m)$ for all $m \in \mathcal{P}(\mathcal{S})$ and $i, j \in \mathcal{S}$, then we obtain that

$$\begin{aligned}
|f(m^1) - f(m^2)|_1 &= \sum_{j \in \mathcal{S}} \left| \sum_{i \in \mathcal{S}} \left(m_i^1 Q_{ij}(m^1) - m_i^2 Q_{ij}(m^2) \right) \right| \\
&\leq \sum_{j \in \mathcal{S}} \sum_{i \in \mathcal{S}} \left(|m_i^1 Q_{ij}(m^1) - m_i^1 Q_{ij}(m^2)| + |m_i^1 Q_{ij}(m^2) - m_i^2 Q_{ij}(m^2)| \right) \\
&= \sum_{j \in \mathcal{S}} \sum_{i \in \mathcal{S}} \left(m_i^1 |Q_{ij}(m^1) - Q_{ij}(m^2)| + |Q_{ij}(m^2)| \cdot |m_i^1 - m_i^2| \right) \\
&\leq \sum_{j \in \mathcal{S}} \left(\sum_{i \in \mathcal{S}} m_i^1 M |m^1 - m^2|_1 + \sum_{i \in \mathcal{S}} L |m_i^1 - m_i^2| \right) \\
&= \sum_{j \in \mathcal{S}} \left(M |m^1 - m^2|_1 + L |m^1 - m^2|_1 \right) = (M + L) \cdot S |m^1 - m^2|_1.
\end{aligned}$$

By Corollary 1 of McShane (1934) we find a Lipschitz continuous extension $\tilde{f}(\cdot)$ of $f(\cdot)$, which means that $\tilde{f} : \mathbb{R}^S \to \mathbb{R}^S$ is Lipschitz continuous and satisfies $\tilde{f}(m) = f(m)$ for all $m \in \mathcal{P}(\mathcal{S})$.

Thus, by the classical existence and uniqueness theorem for ordinary differential equations (Walter, 1998, Theorem 10.VI), we obtain that a unique solution of $\frac{\partial}{\partial t}\Phi^t(m_0) = \tilde{f}(\Phi^t(m_0))$, $\Phi^0(m_0) = m_0$ on $[0, \infty)$ exists. Utilizing that the vectors $f(m)$ lie in $T_{\mathcal{P}(\mathcal{S})}(m)$ for all $m \in \mathcal{P}(\mathcal{S})$ (see the proof of Theorem 8.4) we obtain by Lemma 9.1 that the trajectory lies in $\mathcal{P}(\mathcal{S})$. Thus, m is also a solution of $\frac{\partial}{\partial t}\Phi^t(m_0) = f(\Phi^t(m_0))$, $\Phi^0(m_0) = m_0$. The continuity of $\Phi^t(\cdot)$ for any fixed $t \geq 0$ follow from a standard dependence theorem (Walter, 1998, Theorem 12.VII). □

9.2.2. Limit Behaviour: First Examples

We are particularly interested in the limit behaviour of these processes, especially in the following two notions: We say that a nonlinear Markov chain with semigroup $(\Phi^t(\cdot))_{t\geq 0}$ is *strongly ergodic* if there exists an $\bar{m} \in \mathcal{P}(\mathcal{S})$ such that for all $m_0 \in \mathcal{P}(\mathcal{S})$ we have $\lim_{t\to\infty} ||\Phi^t(m_0) - \bar{m}|| = 0$. For constant dynamics, that is $Q(m) = Q$ for all $m \in \mathcal{P}(\mathcal{S})$, this notion coincides with the classical notion of strong ergodicity, for which irreducibility is a sufficient criterion. Furthermore, we say that the nonlinear Markov chain *converges in the limit to some stationary distribution*, if for each $m_0 \in \mathcal{P}(\mathcal{S})$ there is a stationary point $\bar{m}(m_0)$ of $Q(\cdot)$ such that $\lim_{t\to\infty} ||\Phi^t(m_0) - \bar{m}(m_0)|| = 0$. We remark that this notion includes strong ergodicity, but covers many more cases, in particular, any Markov chain with constant dynamics (also if the generator matrix is reducible) converges in the limit to some stationary distribution.

The limit behaviour for non-constant nonlinear continuous time Markov chains cannot be simply characterized as in the case of standard continuous time Markov chains. In particular, we observe more complex phenomena for general nonlinear continuous time Markov chains, as we will discuss in the following examples: It is now possible that trajectories do not converge at all but that the ordinary equation describing the marginal distributions admits periodic solutions. Furthermore, it might happen that although $Q(m)$ is irreducible for all $m \in \mathcal{P}(\mathcal{S})$ there are several stationary distributions and that the marginal distributions converge for all initial distributions but not towards the same stationary distribution. Nonetheless, it is still possible to formulate a sufficient condition for ergodicity in the case of small state spaces, which we will present in the Subsection 9.2.3 together with a discussion of whether we can utilize the classical approaches for continuous time Markov chains to prove ergodicity.

Let us start the discussion of the limit behaviour of nonlinear Markov chains with an example of a generator matrix such that the ordinary differential equation describing the marginal distributions admits periodic solutions: Namely, we define a generator matrix on

$$B = \mathcal{P}(\{1,2,3\}) \cap \left\{m \in \mathbb{R}^3 : \min\{m_1, m_2, m_3\} \geq \tfrac{1}{10}\right\}$$

by

$$Q(m) = \begin{pmatrix} -Q_{\left\{m_2 \leq \frac{1}{3}\right\}}(m) & 0 & Q_{\left\{m_2 \leq \frac{1}{3}\right\}}(m) \\ 0 & -Q_{\left\{m_1 \geq \frac{1}{3}\right\}}(m) & Q_{\left\{m_1 \geq \frac{1}{3}\right\}}(m) \\ Q_{\left\{m_2 \geq \frac{1}{3}\right\}}(m) & Q_{\left\{m_1 \leq \frac{1}{3}\right\}}(m) & -Q_{\left\{m_2 \geq \frac{1}{3}\right\}}(m) - Q_{\left\{m_1 \leq \frac{1}{3}\right\}}(m) \end{pmatrix} \tag{9.4}$$

with

$$Q_{\left\{m_1 \leq \frac{1}{3}\right\}}(m) = \begin{cases} \frac{1}{m_3}\left(\frac{1}{3} - m_1\right) & \text{if } m_1 \leq \frac{1}{3} \\ 0 & \text{else} \end{cases}$$

$$Q_{\left\{m_1 \geq \frac{1}{3}\right\}}(m) = \begin{cases} \frac{1}{m_1}\left(m_1 - \frac{1}{3}\right) & \text{if } m_1 \geq \frac{1}{3} \\ 0 & \text{else} \end{cases}$$

$$Q_{\left\{m_2 \leq \frac{1}{3}\right\}}(m) = \begin{cases} \frac{1}{m_2}\left(\frac{1}{3} - m_2\right) & \text{if } m_2 \leq \frac{1}{3} \\ 0 & \text{else} \end{cases}$$

$$Q_{\left\{m_2 \geq \frac{1}{3}\right\}}(m) = \begin{cases} \frac{1}{m_3}\left(m_2 - \frac{1}{3}\right) & \text{if } m_2 \geq \frac{1}{3} \\ 0 & \text{else} \end{cases}.$$

We note that all transition rates are bounded Lipschitz continuous functions, which furthermore satisfy $Q_{\{ij\}}(m) \geq 0$ for all $m \in B$ and $i \neq j$. Thus, we easily find an extension of $Q_{ij}(\cdot)$ for all $i, j \in \mathcal{S}$ with $i \neq j$ such that $Q_{ij}(\cdot)$ is Lipschitz continuous and satisfies $Q_{ij}(\cdot) \geq 0$. We remark that this technique can be applied for all bounded and Lipschitz continuous functions defined on a closed subset of $\mathcal{P}(\mathcal{S})$ by a generalization of Tietze's extension theorem found in McShane (1934), which states that we can extend an Lipschitz continuous function defined on a subset of a metric space onto the whole space such that it remains Lipschitz continuous and such that the bounds are preserved.

The generator matrix has been chosen in such a way that in a neighbourhood $U \subseteq B$ of $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ the first two components of the marginals will behave as the classical harmonic oscillator, more precisely the ordinary differential equation characterizing the marginals reads

$$\dot{m}_1 = \begin{cases} m_1 \cdot \left(-\frac{1}{m_1}\left(\frac{1}{3} - m_2\right)\right) & \text{if } m_2 \leq \frac{1}{3} \\ m_3 \cdot \left(\frac{1}{m_3}\left(m_2 - \frac{1}{3}\right)\right) & \text{if } m_2 \geq \frac{1}{3} \end{cases}$$

$$= m_2 - \frac{1}{3}$$

$$\dot{m}_2 = \begin{cases} m_2 \cdot \left(-\frac{1}{m_2}\left(m_1 - \frac{1}{3}\right)\right) & \text{if } m_1 \geq \frac{1}{3} \\ m_3 \cdot \left(\frac{1}{m_3}\left(\frac{1}{3} - m_1\right)\right) & \text{if } m_1 \leq \frac{1}{3} \end{cases}$$

$$= \frac{1}{3} - m_1$$

and by symmetry

$$\dot{m}_3 = m_1 - m_2.$$

This ordinary differential equation yields periodic solutions, for example for $m_0 = (0.2, 0.4, 0.4)$ we obtain the trajectory depicted in Figure 9.1.

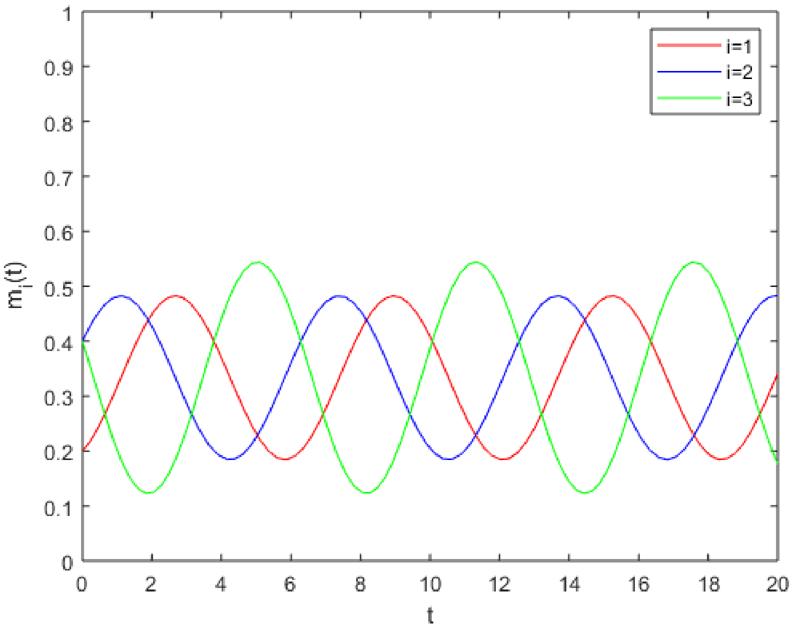

Figure 9.1.: The marginal distributions of the nonlinear continuous time Markov chain with generator (9.4) and initial distribution $(0.2, 0.4, 0.4)$.

The second example we consider is an example where for each $m \in \mathcal{P}(\mathcal{S})$ the matrix $Q(m)$ is irreducible, but still we do not observe strong ergodicity. Instead there are three stationary points and some of them are attracting, some are not:

Consider

$$Q(m) = \begin{pmatrix} -\left(\frac{29}{3}m_1^2 - 16m_1 + \frac{22}{3}\right) & \frac{29}{3}m_1^2 - 16m_1 + \frac{22}{3} \\ m_1^2 + m_1 + 1 & -(m_1^2 + m_1 + 1) \end{pmatrix}, \tag{9.5}$$

which is irreducible for all $m \in \mathcal{P}(\{1,2\})$ since $m_1^2 + m_1 + 1 \geq 1$ for all $m_1 \geq 0$ and since the minimum of the function $\frac{29}{3}m_1^2 - 16m_1 + \frac{22}{3}$ lies at $m_1 = \frac{24}{29}$ and the value at this point is $\frac{62}{87} > 0$. The dynamics are given by

$$\begin{aligned}\dot{m}_1 &= m_1 \cdot \left(-\left(\frac{29}{3}m_1^2 - 16m_1 + \frac{22}{3}\right)\right) + (1-m_1)\left(m_1^2 + m_1 + 1\right) \\ &= \underbrace{-\frac{32}{3}m_1^3 + 16m_1^2 - \frac{22}{3}m_1 + 1}_{=:f(m_1)}\,. \end{aligned} \tag{9.6}$$

We remark that this fully describes the dynamics because by symmetry $\dot{m}_2 = -\dot{m}_1$.

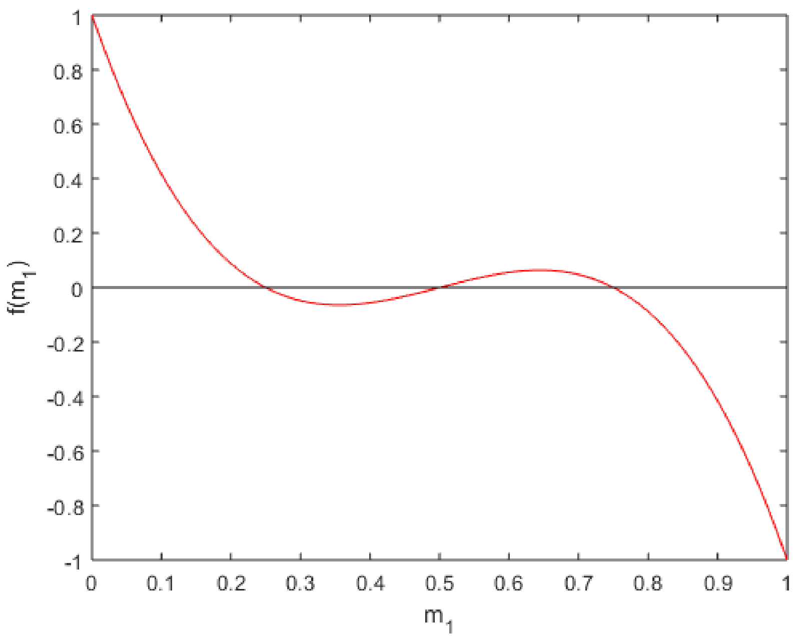

Figure 9.2.: A plot of the right-hand-side of (9.6).

In Figure 9.2 we see that there are three stationary points given $Q(m)$ at $m_1 = \frac{1}{4}$, $m_1 = \frac{1}{2}$ and $m_1 = \frac{3}{4}$. Furthermore, the plot reveals the stationary behaviour of the solutions:

- Since the function $f(\cdot)$ is strictly positive on $[0, 0.25)$, the trajectories will for all initial conditions $(m_0)_1 \in [0, 0.25)$ converge towards $m_1 = 0.25$.

- Since the function $f(\cdot)$ is strictly negative on $(0.25, 0.5)$, the trajectories will for all initial conditions $(m_0)_1 \in (0.25, 0.5)$ converge towards $m_1 = 0.25$.

- Since the function $f(\cdot)$ is strictly positive on $(0.5, 0.75)$, the trajectories will for all initial conditions $(m_0)_1 \in (0.5, 0.75)$ converge towards $m_1 = 0.75$.

- Since the function $f(\cdot)$ is strictly negative on $(0.75, 1]$, the trajectories will for all initial conditions $(m_0)_1 \in (0.75, 1]$ converge towards $m_1 = 0.75$.

In total we obtain that for all initial conditions $(m_0)_1 \in [0, 0.5)$ the trajectories converge towards $m_1 = 0.25$ and for all initial conditions $(m_0)_1 \in (0.5, 1]$ the trajectories converge towards $m_1 = 0.75$. For $(m_0)_1 = 0.5$ the trajectory stays at the stationary point $m_1 = 0.5$. This means that although $Q(m)$ is irreducible for all $m \in \mathcal{P}(\mathcal{S})$ we do not observe strong ergodicity, but only convergence towards some stationary distribution. This limit behaviour is depicted in Figure 9.3. It is an open and interesting question whether irreducibility implies at least convergence towards some stationary distribution or not (see Chapter 10).

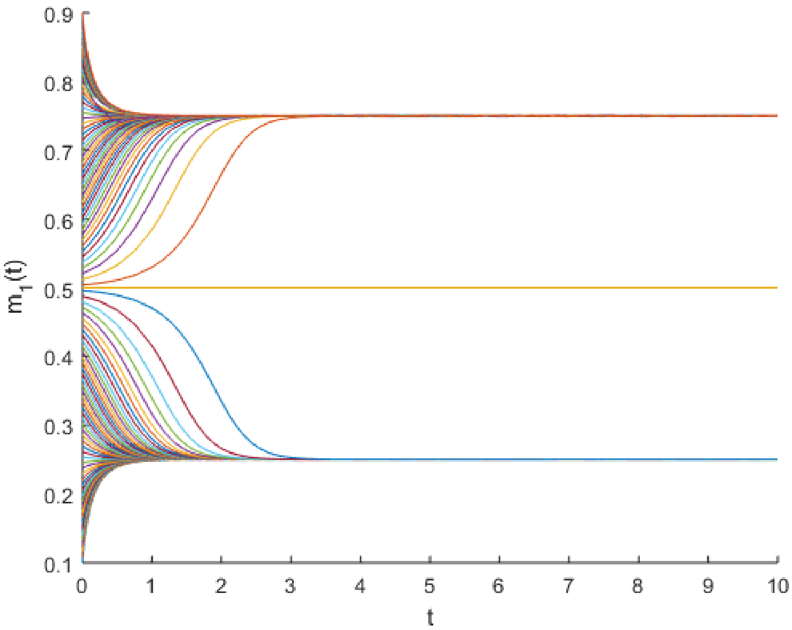

Figure 9.3.: The trajectories of the nonlinear Markov chain with generator $Q(m)$ given by (9.5) for several initial conditions.

9.2.3. Limit Behaviour: A Sufficient Condition for Strong Ergodicity

To the best of the author's knowledge, strong ergodicity as well as the convergence of the trajectories towards some stationary point for nonlinear continuous time Markov chains has not been investigated in the literature so far. The closest contributions in the literature are sufficient criteria for discrete time nonlinear Markov chains based on a generalization of Dobrushin's ergodicity coefficient: More precisely, Saburov (2016) considers nonlinear Markov chains where the Markov operator is polynomial and proves a condition for strong ergodicity relying on the notion of hypermatrices. Butkovsky (2014) proves a strong ergodicity criterion for general Markov operators and his proof relies on sensible estimates of the total variation distance between $P(m^1) \cdot m^1$ and $P(m^2) \cdot m^2$. In Butkovsky (2012) a proof of the same result relying on a modification of the coupling procedure introduced in Butkovsky and Veretennikov (2013) is sketched. However, it will not be possible to extend any of the previously discussed approaches to our setting since they all crucially rely on the sequential nature of the problem. More precisely, the first two proofs rely on the fact that we obtain the n-th distribution vector via $P(m^{n-1})m^{n-1}$ and in the third proof the coupling is constructed in a sequential fashion.

Thus, we need to search for methods that take the continuous time into account. However, although we defined our process as a time-inhomogeneous continuous time Markov chain, we cannot apply the ergodicity results known for these type of processes. The reason is that the transition probabilities depend on the initial distribution m_0, which we vary when investigating strong ergodicity. Moreover, also the approaches proving strong ergodicity of classical continuous time Markov chains cannot easily be adapted:

The method of discrete skeletons relying on Kingman's Lemma (Kingman, 1963), which yields that irreducible recurrent Markov chains that admit a stationary distribution are strongly (even exponentially) ergodic, requires to compute the transition probabilities for small time steps h. However, because of the nonlinear dynamics we cannot compute small-step transition probabilities and it is also not possible to rely on an approximation of this quantity since the approximation error is propagated in each of the small time steps.

The approach of Anderson (1991, §6.3) heavily relies on the notion of accessibility and yields convergence of the transition probabilities $P_{ij}(t)$ if and only if for one $j \in \mathcal{S}$ (and thus for all states) the supremum over all $i \in \mathcal{S} \setminus \{j\}$ of the expected value of the first hitting time given $X_0 = i$ is finite. Zeifman (1989) extends this approach to time-inhomogeneous continuous time Markov chains relying on a generalization of accessibility as well as sensible estimates of the transition probabilities. However, it is not clear how to introduce the notion of accessibility for nonlinear Markov chains since it might happen that $P_{ij}(t, m^1) \neq 0$ and $P_{ij}(t, m^2) = 0$ (consider for example the dynamics of the corruption model in Section 6.2).

Another alternative approach is the use of the coupling method, which is presented in Griffeath (1975) also for time-inhomogeneous continuous time Markov chains. With this method, it is possible to prove ergodicity given that Dobrushin's ergodicity coefficient is well behaved. However, it is not clear how to extend the coupling procedure to nonlinear Markov chains.

Johnson and Isaacson (1988) obtain ergodicity criteria (also for time-inhomogeneous Markov chains) directly linked to the time-dependent generator matrix utilizing the representation formula

$$P(s,t) = I + \int_s^t Q(u)P(u,t)\ \mathrm{d}u$$

found in Iosifescu (1980, Section 8.9), the Chapman-Kolmogorov equation as well as simple facts from linear algebra regarding the Dobrushin's ergodicity coefficient. Since we do not obtain any formula similar to the representation formula, we cannot reapply their methods.

We will follow a different approach. Namely, we directly consider the long-term behaviour of the nonlinear system (9.3) describing the marginal distributions. However, the global behaviour of nonlinear systems is complex. In $\mathbb{R}^2$ the Jordan curve theorem holds, which states that every Jordan curve (which is the image of the unit circle S^1 under a continuous injective map) dissects the plane $\mathbb{R}^2$ in two connected regions. Using this statement one obtains the Poincaré-Bendixson Theorem, which yields a description of the long-term behaviour of nonlinear systems in the plane. For higher dimension, i.e. $\mathbb{R}^n$, $n \geq 3$, this is no longer possible and indeed strange behaviour of nonlinear systems occurs. A famous example is the Lorenz equation (see Section 8.2 in Teschl (2012)). Therefore, we restrict our attention to deriving criteria for small state spaces. More precisely, we restrict on the cases $S = 2$ and $S = 3$ where we can apply, due to the flow invariance of $\mathcal{P}(\mathcal{S})$, the results for scalar and planar nonlinear systems.

For the case of two states, strong ergodicity simply follows from the existence of a unique stationary point as well as a continuity assumption:

Theorem 9.6. *Let $S = 2$ and assume that $f : [0,1] \to \mathbb{R}$ defined via*

$$f(m_1) := m_1 \cdot Q_{11}(m_1, 1 - m_1) + (1 - m_1) \cdot Q_{21}(m_1, 1 - m_1)$$

is continuous. Furthermore, assume that $(\bar{m}, 1 - \bar{m})$ is the unique stationary point given Q. Then, the nonlinear Markov chain is strongly ergodic.

Proof. We start with some preliminary facts regarding $f(\cdot)$: An equilibrium point is characterized by the property that $\dot{m} = 0$, which by flow invariance of $\mathcal{P}(\mathcal{S})$ (see Theorem 9.5) is equivalent to the fact that $\dot{m}_1 = 0$ (because $\dot{m}_1 + \dot{m}_2 = 0$). Since we have a unique equilibrium point, we obtain that $f(\bar{m}) = 0$ and $f(m_1) \neq 0$ for all $m_1 \neq \bar{m}$. Since $f(\cdot)$ is continuous, we obtain that $f(\cdot)$ is non-vanishing on $[0, \bar{m})$ and $(\bar{m}, 1]$ and has uniform sign on each of these sets. Since $Q(\cdot)$ is a conservative generator we moreover obtain that $f(0) \geq 0$ and $f(1) \leq 0$. Thus, we obtain that $f(m_1) > 0$ for all $m_1 \in [0, \bar{m})$ and $f(m_1) < 0$ for all $m_1 \in (\bar{m}, 1]$. This in turn yields that $[0,1]$ is flow invariant for $\dot{m}_1 = f(m_1)$.

Now we can prove that the systems $\dot{m} = Q(m)^T m$ and $\dot{m}_1 = f(m_1)$ are equivalent: Let $m(\cdot) = (m_1(\cdot), m_2(\cdot))$ be a solution of $\dot{m} = Q(m)^T m$ with initial condition $m(0) = m_0 \in$

$\mathcal{P}(\mathcal{S})$. By flow invariance of $\mathcal{P}(\mathcal{S})$ (see Theorem 9.5), we have $m_2(t) = 1 - m_1(t)$ for all $t \geq 0$. Thus, $\dot{m} = Q(m)^T m$ is equivalent to

$$\begin{cases} \dot{m}_1 = m_1 \cdot (Q_{11}(m_1, 1-m_1)) + (1-m_1) \cdot Q_{21}(m_1, 1-m_1) \\ -\dot{m}_1 = m_1 \cdot (-Q_{12}(m_1, 1-m_1)) + (1-m_1) \cdot Q_{22}(m_1, 1-m_1). \end{cases} \tag{9.7}$$

Therefore, $m_1(\cdot)$ is indeed a solution of $\dot{m}_1 = f(m_1)$. For the converse implication we first note that in (9.7) the last equation is the first equation multiplied by (-1) because $Q(m)$ is conservative for all $m \in \mathcal{P}(\mathcal{S})$. If $m_1(\cdot)$ satisfies $\dot{m}_1 = f(m_1)$, $m_1(0) = (m_0)_1 \in [0,1]$, then, by flow invariance, $m_1(t) \in [0,1]$ for all $t \geq 0$. Thus, the function $(m_1(\cdot), 1-m_1(\cdot))$ satisfies $\dot{m} = Q(m)^T m$.

Since $f(m_1) > 0$ for all $m_1 \in [0, \bar{m})$ and $f(m_1) < 0$ for all $m_1 \in (\bar{m}, 1]$, we obtain the desired convergence statement. □

Remark 9.7. In this two state setting we also obtain a sufficient criterion for convergence in the limit towards some stationary distribution. Indeed, requiring that the function $f(\cdot)$ is continuous is sufficient. The proof is analogous. More precisely, we again obtain that between two equilibrium points the function $f(\cdot)$ we be non-vanishing and of uniform sign. If $f(m_0) > 0$, then we observe convergence to the next equilibrium point right of m_0, and if $f(m_0) < 0$, then we observe convergence to the next equilibrium point left of m_0. △

Similarly, we obtain for systems with three states that given $m_0 \in \mathcal{P}(\mathcal{S})$ the function $m(\cdot) = (m_1(\cdot), m_2(\cdot), m_3(\cdot))$ is a solution of $\dot{m} = Q(m)^T m$, $m(0) = m_0$ if and only of $(m_1(\cdot), m_2(\cdot))$ is a solution of

$$\begin{pmatrix} \dot{m}_1 \\ \dot{m}_2 \end{pmatrix} = f \begin{pmatrix} m_1 \\ m_2 \end{pmatrix}, \quad \begin{pmatrix} m_1(0) \\ m_2(0) \end{pmatrix} = \begin{pmatrix} (m_0)_1 \\ (m_0)_2 \end{pmatrix},$$

where

$$f \begin{pmatrix} m_1 \\ m_2 \end{pmatrix} = \begin{pmatrix} Q_{31}(\hat{m}) + (Q_{11}(\hat{m}) - Q_{31}(\hat{m}))m_1 + (Q_{21}(\hat{m}) - Q_{31}(\hat{m}))m_2 \\ Q_{32}(\hat{m}) + (Q_{12}(\hat{m}) - Q_{32}(\hat{m}))m_1 + (Q_{22}(\hat{m}) - Q_{32}(\hat{m}))m_2 \end{pmatrix} \tag{9.8}$$

and $\hat{m} = (m_1, m_2, 1 - m_1 - m_2)$. Indeed, the proof is analogous to the proof for the two state case, the only step that has to be modified is the proof of the flow invariance of $\{(m_1, m_2) \in [0, \infty) : m_1 + m_2 \leq 1\}$ for $(\dot{m}_1, \dot{m}_2)^T = f(m_1, m_2)$. Nonetheless, that fact that $Q(\cdot)$ is a conservative generator and a simple application of Lemma 9.2 yield the desired statement.

Using this equivalence we obtain, relying on the Poincaré-Bendixson Theorem, a sufficient criterion for strong ergodicity. By slight abuse of notation, we will write in the following m for the three dimensional vector from $\mathcal{P}(\mathcal{S})$ as well as the two-dimensional vector $(m_1, m_2)^T$.

Theorem 9.8. *Let $O \supseteq \{(m_1, m_2) \in [0, \infty)^2 : m_1 + m_2 \leq 1\}$ be a simply connected and bounded region such that there is a continuously differentiable function $f : O \to \mathbb{R}^2$ satisfying* (9.8) *on $\mathcal{P}(\mathcal{S})$. Let $\bar{m}$ be the unique stationary point given $Q(\cdot)$. Furthermore, assume that*

(a) $\frac{\partial f_1}{\partial m_1}(m) + \frac{\partial f_2}{\partial m_2}(m)$ *is non-vanishing for all $m \in O$ and has uniform sign on O,*

(b) it holds that
$$\frac{\partial f_1}{\partial m_1}(\bar{m}) \cdot \frac{\partial f_2}{\partial m_2}(\bar{m}) - \frac{\partial f_1}{\partial m_2}(\bar{m}) \cdot \frac{\partial f_2}{\partial m_1}(\bar{m}) > 0$$
or it holds that
$$\left(\frac{\partial f_1}{\partial m_1}(\bar{m}) + \frac{\partial f_2}{\partial m_2}(\bar{m})\right)^2 - 4\left(\frac{\partial f_1}{\partial m_1}(\bar{m}) \cdot \frac{\partial f_2}{\partial m_2}(\bar{m}) - \frac{\partial f_1}{\partial m_2}(\bar{m}) \cdot \frac{\partial f_2}{\partial m_1}(\bar{m})\right) < 0.$$

Then, the nonlinear Markov chain is strongly ergodic.

In order to prove the theorem, we have to review the Poincaré-Bendixson Theorem as well as some related notions and results, for which Teschl (2012, Section 6.3 and Section 7.3) is a reference.

Let $O \subseteq \mathbb{R}^S$ be open, let $f : O \to \mathbb{R}^S$ be continuously differentiable and consider the initial value problem
$$\dot{x} = f(x), \quad x(0) = x_0. \tag{9.9}$$
Let $y \in O$ be arbitrary and denote by $(T_-(y), T_+(y))$ the maximal interval of existence of a solution of (9.9) with initial condition $x(0) = y$ and let $t \mapsto \phi(t, y)$ be the corresponding solution. The *orbit* of y is given by
$$\gamma(y) = \{\phi(t, y) : t \in (T_-(y), T_+(y))\}.$$
We remark that if $z \in \gamma(y)$, then $\gamma(y) = \gamma(z)$ and that different orbits are disjoint. Additionally, we define the *forward orbit* of y by
$$\gamma_+(y) = \{\phi(t, y) : t \in (0, T_+(y))\}$$
and the *backward orbit* by
$$\gamma_-(y) = \{\phi(t, y) : t \in (T_-(y), 0)\}.$$

We say that y is a *fixed point* if $\gamma(y) = \{y\}$. We say that $y \in O$ is a *periodic point* if there is some $T > 0$ such that $\phi(T, x) = x$ and call the infimum of such T the period. With these definitions we classify the orbits of $\dot{x} = f(x)$ as follows:

- *fixed orbits* (these correspond to a periodic point with period zero)
- *regular periodic obits* (these correspond to a periodic point with positive period)
- *non-closed orbits* (these do not correspond to periodic points).

Non-closed orbits $\gamma(y)$ that join a fixed point x to itself, formally, $\lim_{t\to\pm\infty}\phi(t,y)=x$, are often called homoclinic paths. For planar autonomous systems they can only exist if the fixed point is a saddle point, which means that the Jacobian for the fixed point has to satisfy that there are two real eigenvalues with opposite sign (Jordan and Smith, 2005, Section 3.6). Writing

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix}$$

for the Jacobian, this condition is equivalent to the requirement that $ad-bc<0$ and simultaneously $(a+d)^2-4(ad-bc)>0$.

In order to analyse the long-term behaviour of solutions, the following sets that describe all those points where an orbit accumulates of particular interest: Let $y\in O$, then the $\omega_\pm$*-limit set* is given by

$$\omega_\pm(y)=\{z\in O|\exists(t_n)_{n\in\mathbb{N}}:t_n\to\pm\infty\wedge\phi(t_n,y)\to z\}.$$

We remark that if $\hat{y}\in\gamma(y)$ then $\omega_\pm(y)=\omega_\pm(\hat{y})$. In other words, the set $\omega_\pm(y)$ depends only on the orbit $\gamma(y)$.

We are now in the position to collect the relevant results regarding the $\omega_\pm$-limit sets, which are Lemma 6.6, Lemma 6.7 and Theorem 7.16 in Teschl (2012) and Theorem 3.5 in Jordan and Smith (2005).

Lemma 9.9. *Let $\sigma\in\{+,-\}$. If $\gamma_\sigma(y)$ is contained in a compact set, then $\omega_\sigma(x)$ is non-empty, compact and connected.*

Lemma 9.10 (generalized Poincaré-Bendixson Theorem)**.** *Let $O\subseteq\mathbb{R}^2$ be open and let $f:O\to\mathbb{R}^2$ be continuously differentiable. Fix $y\in O$ and $\sigma\in\{+,-\}$. Assume that $\omega_\sigma(y)$ is non-empty, compact, connected and contains only finitely many fixed points. Then one of the following cases holds:*

(i) $\omega_\sigma(y)$ *is a fixed orbit*

(ii) $\omega_\sigma(y)$ *is a regular periodic orbit*

(iii) $\omega_\sigma(y)$ *consists of (finitely many) fixed points* $x_1,\dots,x_k$ *and non-closed orbits* $\gamma(z)$ *such that* $\omega_\pm(z)\in\{x_1,\dots,x_k\}$.

Lemma 9.11 (Bendixson's Criterion). *Let $O \subseteq \mathbb{R}^2$ be open and let $f : O \to \mathbb{R}^2$ be continuously differentiable. Let $U \subseteq O$ be a simply connected region and assume that $\frac{\partial f_1}{\partial x_1}(x) + \frac{\partial f_2}{\partial x_2}(x)$ does not vanish for $x \in U$ and has uniform sign on U. Then there are no regular periodic orbits contained (entirely) inside U.*

Lemma 9.12. *Let $\sigma \in \{+, -\}$. If $\gamma_\sigma(y)$ is contained in a compact set, then*

$$\lim_{t \to \sigma\infty} d(\phi(t, y), \omega_\sigma(y)) = 0.$$

With all these preparations we can prove the theorem:

Proof of Theorem 9.8. Since the set

$$\{(m_1, m_2) \in \mathbb{R}^2 : m_1, m_2 \geq 0 \wedge m_1 + m_2 \leq 1\}$$

is flow invariant for $(m_1, m_2)^T = f(m_1, m_2)$, we obtain, by Lemma 9.9, that $\omega_+(m_0)$ is non-empty, compact and connected. Since there is only one stationary point, we can apply the Poincaré-Bendixson Theorem (Lemma 9.10). Thus, we are in one of the following cases:

(i) $\omega_+(m_0)$ is a fixed orbit

(ii) $\omega_+(m_0)$ is a regular periodic orbit

(iii) $\omega_+(m_0)$ consists of (finitely many) fixed points $x_1, \ldots, x_k$ and non-closed orbits $\gamma(z)$ such that $\omega_\pm(z) \in \{x_1, \ldots, x_k\}$.

By condition (a) and Bendixson's criterion (Lemma 9.11), we obtain that case (ii) is not possible. Furthermore, we observe that $\bar{m}$ is the only fixed point and that it is, by condition (b), not a saddle point. Thus, also case (iii) is not possible. Therefore, $\omega_+(m_0)$ is a fixed orbit. Since $\bar{m}$ is the only fixed point, this implies that $\omega_+(m_0) = \{\bar{m}\}$. Lemma 9.12 now yields the desired result. □

Remark 9.13. We remark that the proof of the theorem crucially relies on the fact that there is a unique stationary point since otherwise it is indeed possible that non-closed orbits joining distinct fixed points exists. Thus, we cannot easily obtain a criterion for convergence towards some stationary distribution. △

Example. Let us consider the nonlinear Markov chain induced by the strategy $\{stay\} \times \{change\}$ in the consumer choice example with non-constant dynamics presented in Section 6.3. In this setting we have that

$$f(m_1, m_2) = \begin{pmatrix} \lambda - em_1^2 - \epsilon m_1 - \lambda m_1 + bm_2 - \lambda m_2 \\ \lambda - \lambda m_1 - bm_2 - em_2^2 - \epsilon m_2 - \lambda m_2 \end{pmatrix}.$$

The Jacobian matrix is given by

$$\frac{\partial f}{\partial m}(m) = \begin{pmatrix} -2em_1 - \epsilon - \lambda & b - \lambda \\ -\lambda & -b - 2em_2 - \epsilon - \lambda \end{pmatrix}$$

and it satisfies

$$\frac{\partial f_1}{\partial m_1}(m) + \frac{\partial f_2}{\partial m_2}(m) = -2em_1 - \epsilon - \lambda - b - 2em_2 - \epsilon - \lambda < 0$$

for all $m \in N_\epsilon([0,1]^2)$ as well as

$$\begin{aligned}
&\frac{\partial f_1}{\partial m_1}(m) \cdot \frac{\partial f_2}{\partial m_2}(m) - \frac{\partial f_1}{\partial m_2}(m) \cdot \frac{\partial f_2}{\partial m_1}(m) \\
&= (-2em_1 - \epsilon - \lambda)(-b - 2em_2 - \epsilon - \lambda) - (b - \lambda)(-\lambda) \\
&= (2em_1 + \epsilon)(2em_2 + \epsilon) + (2em_1 + \epsilon)(b + \lambda) + \lambda(2em_2 + \epsilon) + \lambda(b + \lambda) + \lambda(b - \lambda) \\
&= (2em_1 + \epsilon)(2em_2 + \epsilon) + (2em_1 + \epsilon)(b + \lambda) + \lambda(2em_2 + \epsilon) + 2\lambda b > 0
\end{aligned}$$

for all $m \in [0,1]^2$ and, therefore, in particular for $\bar{m}$. Thus, Theorem 9.8 implies strong ergodicity of the nonlinear Markov chain with transition rate matrix function $Q^{\{stay\}\times\{change\}}(m)$. △

9.3. The Global Convergence Theorem

This section introduces the relevant notions and the convergence result. First, we introduce several assumptions that are motivated by the fact that we have to derive that several sets are flow invariant for the differential inclusion $\dot{x} \in F(x)$. More precisely, we will assume that the process is defined on a larger set in order to allow for the application of differential calculus. Moreover, we will set up some assumptions requiring a simple description of those points for which a certain deterministic strategy is optimal in order to apply the characterization of flow invariance derived in Subsection 9.1.1. The conditions of the theorem itself are then those that are intuitively motivated and which we expect to be in some sense necessary to obtain global convergence.

The first additional assumption we set covers the necessary extension of the myopic adjustment process in order to utilize differential calculus (in particular the theorems on flow invariance presented in Section 9.1). This assumption itself is not restrictive since by classical extension theorems (see McShane (1934)) we always find such an open set and such an extension (see the proof of Theorem 9.5). However, in combination with the other assumptions it is restrictive since then not every extension is suitable.

Assumption A2. *The set O is open and satisfies $\mathcal{P}(\mathcal{S}) \subseteq O$. For all $i, j \in \mathcal{S}$ and $a \in \mathcal{A}$ the function $Q_{ija} : O \to \mathbb{R}$ is Lipschitz continuous and for all $i \in \mathcal{S}$ and $a \in \mathcal{A}$ the function $r_{ia} : O \to \mathbb{R}$ is continuous. Furthermore, the matrix $Q_{\cdot\cdot a(m)}$ is a generator matrix for all $a \in \mathcal{A}$ and $m \in O$.*

In this setting we define the extension of the myopic adjustment process as follows: For each $m \in O$ let $\mathcal{D}(m)$ be the set of all optimal deterministic stationary strategies for the continuous time Markov decision process with transition rates $Q_{ija}(m)$ and reward functions $r_{ia}(m)$ and define $F : O \to 2^{\mathbb{R}^S}$ by

$$F(m) := \operatorname{conv}\left\{ \left(\sum_{i\in\mathcal{S}} \sum_{a\in\mathcal{A}} m_i Q_{ija}(m) d_{ia} \right)_{j\in\mathcal{S}} : d \in \mathcal{D}(m) \right\}.$$

Since we want to apply the characterization of flow invariance derived in Subsection 9.1.1, we set-up additional requirements for the following sets: The set $\mathrm{Opt}^{\mathrm{unique}}(d)$ collects all $m \in O$ such that d is the unique optimal strategy for m. The set $\mathrm{Opt}^{\mathrm{some}}(d)$ collects all $m \in O$ such that d is some optimal strategy for m.

Assumption A3. *Let $\mathcal{U}$ be the set of strategies $d \in D^s$ such that there exists an $m \in O$ satisfying $\mathcal{D}(m) = \{d\}$. We assume that for each $m \in O$ there is a strategy $d \in \mathcal{U}$ such that $d \in \mathcal{D}(m)$.*

Furthermore, for each strategy $d \in \mathcal{U}$ we assume that there exists a twice continuously differentiable function $g_d : O \to \mathbb{R}$ such that

$$Opt^{unique}(d) = \{m \in O : g_d(m) < 0\}$$
$$Opt^{some}(d) = \{m \in O : g_d(m) \leq 0\}.$$

Moreover, we assume that for all $d \in \mathcal{U}$ and $m \in Opt^{some}(d) \cap \{m \in O : g_d(m) = 0\}$ we have $\nabla g_d(m) \neq 0$.

This assumption is satisfied in the case of two non-empty optimality sets $\mathrm{Opt}^{\mathrm{unique}}(d)$ whenever the transition rates and the rewards are twice continuously differentiable and the gradient of a function describing the sets $\mathrm{Opt}^{\mathrm{unique}}(d)$ is non-vanishing whenever both strategies are simultaneously optimal (see the discussion of the case of two non-empty optimality sets $\mathrm{Opt}^{\mathrm{unique}}(d)$ in the end of this section). For more than two non-empty optimality sets $\mathrm{Opt}^{\mathrm{unique}}(d)$ the case is not that clear, but we expect that economically motivated examples usually satisfy this.

The following assumption is used for two crucial steps of the proof: First, in order to prove flow invariance of $\mathcal{P}(\mathcal{S})$ and uniqueness of solutions for the differential inclusion for which we will utilize the explicit characterization results of Subsection 9.1.2. Second, in order to apply the results of Subsection 9.1.1 on the set of all points where a specific set of deterministic strategies is optimal.

Assumption A4. *For any point $m \in \mathcal{P}(\mathcal{S})$ there are at most two strategies $d \in \mathcal{U}$ such that $m \in Opt^{some}(d)$.*

This assumption is always satisfied for two non-empty optimality sets $\text{Opt}^{\text{unique}}(d)$. However, for more than two non-empty optimality sets $\text{Opt}^{\text{unique}}(d)$ this assumption is restrictive. Nonetheless, we expect that it is possible to weaken this assumption at the cost of several technicalities (see Remark 9.22 and the discussion in the outlook (Chapter 10)).

In order to systematically investigate flow invariance of the trajectories of the myopic adjustment process we will introduce a family of digraphs that will be helpful to characterize desirable properties of the trajectories of the population dynamics given the deterministic stationary strategies. In order to do this we first review the relevant definitions and results regarding digraphs (a standard reference for this is Bang-Jensen and Gutin (2010)):

A *digraph* or directed graph $D = (V, A)$ consists of a non-empty and finite set V of vertices and a set A of ordered pairs of distinct vertices from V, which are called arcs. We will often write $x \to y$ for an arc (x, y). When we work with more than one digraph, we write $V(D)$ for the vertex set of D and $A(D)$ for the arc set of D.

A *path* is a sequence $x_1 \dots x_k$ of pairwise distinct vertices such that for all $i \in \{1, \dots, k-1\}$ the arc $x_i \to x_{i+1}$ lies in D. A *cycle* is a sequence $x_1 \dots x_k$ such that $x_i \to x_{i+1} \in D$ for all $i \in \{1, \dots, k-1\}$, the vertices $x_1, \dots, x_{k-1}$ are pairwise distinct and $x_1 = x_k$. We say that a digraph D is *acyclic* if it has no cycle.

A classical way to characterize acyclic digraphs relies on the notion of acyclic orderings: Formally, we say that an ordering $\leq$ of the vertices of a digraph D is an *acyclic ordering* if for any arc $i \to j \in A(D)$ we have $i < j$. This allows us to formulate the following result, which can be found in Bang-Jensen and Gutin (2010, Proposition 2.1.3):

Lemma 9.14. *Every acyclic digraph has an acyclic ordering of its vertices.*

In order to define the family of digraphs we have to introduce one additional assumption. It basically states that the behaviour of the trajectories given any deterministic strategy is consistent within the sets $\text{Opt}^{\text{some}}(d^1) \cap \text{Opt}^{\text{some}}(d^2)$.

Assumption A5. *For all deterministic stationary strategies $d, d^1, d^2 \in \mathcal{U}$ with $d^1 \neq d^2$ and $Opt^{some}(d^1) \cap Opt^{some}(d^2) \neq \emptyset$ the term*

$$\langle (Q^d(m))^T m, \nabla g_{d^1}(m) \rangle$$

is non-vanishing for all $m \in Opt^{some}(d^1) \cap Opt^{some}(d^2)$ and has uniform sign on the set $Opt^{some}(d^1) \cap Opt^{some}(d^2)$.

Also this assumption is rather restrictive, for example the consumer choice model with non-constant dynamics presented in Section 6.3 only satisfies this assumption for some choices of parameters. We can weaken the assumption slightly by allowing that under certain conditions $\langle (Q^d(m))^T m, \nabla g_{d^1}(m)\rangle$ takes values in $[0,\infty)$ or $(-\infty, 0]$, but it again yields to several technicalities (see Section 9.5).

We define for each $d \in \mathcal{U}$ a digraph $D(d)$ with vertex set $\mathcal{U}$ and arc set defined as follows: For any two distinct deterministic strategies $d^1, d^2 \in \mathcal{U}$ it holds that $d^1 \to d^2 \in A(D(d))$ if and only if $\mathrm{Opt}^{\mathrm{some}}(d^1) \cap \mathrm{Opt}^{\mathrm{some}}(d^2) \neq \emptyset$ and

$$\langle (Q^d(m))^T m, \nabla g_{d^1}(m)\rangle > 0 \quad \text{for all} \quad m \in \mathrm{Opt}^{\mathrm{some}}(d^1) \cap \mathrm{Opt}^{\mathrm{some}}(d^2).$$

With these preparations we state the convergence result:

Theorem 9.15. *Let Assumptions A1, A2, A3, A4 and A5 hold and assume that:*

(i) For all $d \in \mathcal{U}$ the nonlinear Markov chain with transition rate matrix function $Q^d(\cdot)$ converges in the limit to some stationary distribution.

(ii) For all $d \in \mathcal{U}$ the digraph $D(d)$ is acyclic.

(iii) There exists an ordering $d^1, d^2, \dots, d^u$ of $\mathcal{U}$ such that for each $i \in \{1, \dots, u\}$ there exists an acyclic ordering $\leq_{D(d^i)}$ of $D(d^i)$ such that

$$d^i \leq_{D(d^i)} d^j \quad \text{for all} \quad j \geq i.$$

(iv) If any ordering that satisfies the conditions of (iii) has the same final vertex $\hat{d}$, then for any $d \in \mathcal{U} \setminus \{\hat{d}\}$ either $d \to \hat{d} \in D(\hat{d})$ or $Opt^{some}(d) \cap Opt^{some}(\hat{d}) = \emptyset$.

Then for any initial condition $m_0 \in \mathcal{P}(\mathcal{S})$ the myopic adjustment has a unique trajectory and the process either converges towards the distribution $\bar{m}$ of some deterministic mean field equilibrium $(\bar{m}, d)$ or there is a $T > 0$ and a strategy $d \in \mathcal{U}$ such that the trajectory stays in

$$\mathcal{P}(\mathcal{S}) \cap \left(Opt^{some}(d) \setminus Opt^{unique}(d)\right)$$

for all $t \geq T$.

This statement is a rather deep result. For any starting point $m_0 \in \mathcal{P}(\mathcal{S})$ we can describe the limit behaviour of the myopic adjustment process: Either it will converge towards some stationary mean field equilibrium or it will stay in a connected component of the set

$$\mathcal{P}(\mathcal{S}) \cap \left(\mathrm{Opt}^{\mathrm{some}}(d) \setminus \mathrm{Opt}^{\mathrm{unique}}(d)\right),$$

which is often small and for which we can often characterize the behaviour of the myopic adjustment process even further by utilizing the results presented in Subsection 9.1.2

(see Subsection 9.7.2 for an example). We remark that although the conditions (iii) and (iv) seem rather complex there is a simple (and polynomial time) algorithm to verify them (see Section 9.6).

The conditions can be intuitively justified as follows: Condition (i) guarantees that whenever we stay inside an optimality set $\text{Opt}^{\text{unique}}(d)$ for all subsequent times then we converge towards a deterministic stationary mean field equilibrium. Condition (ii) ensures that the trajectories given d_i have a suitable limit behaviour, in particular, that they do not behave like a spiral that moves through a sequence of optimality sets again and again. Condition (iii) then requires that the digraphs for individual deterministic stationary strategies are consistent such that we can describe the limit behaviour of the process. In order to ensure convergence we however need to require more, namely condition (iv) ensures that for the final vertex of the acyclic ordering we will never jump from $\text{Opt}^{\text{unique}}(d)$ to $\text{Opt}^{\text{some}}(d) \setminus \text{Opt}^{\text{unique}}(d)$ and back infinitely often.

Remark 9.16. We remark that besides assumption A4 also the assumption A5 can be weakened. Indeed, in order to establish the flow invariance of $\mathcal{P}(\mathcal{S})$ for the differential inclusion we apply the characterization result of Subsection 9.1.2 and this requires strictly positive scalar products only for all points on the boundary of $\mathcal{P}(\mathcal{S})$. For the flow invariance of the other sets (except for the vertex satisfying condition (iv)) it would actually be sufficient to require that the scalar product $\langle (Q^d(m))^T m, \nabla g_{d^1}(m)\rangle$ is either non-positive or non-negative on the set $\text{Opt}^{\text{some}}(d^1) \cap \text{Opt}^{\text{some}}(d^2)$. However, in this case the definition of the digraph, which is crucial in formalizing the relation of the trajectories given the optimal deterministic strategies, would become much more complex since it might happen that $\langle (Q^d(m))^T m, \nabla g_{d^1}(m)\rangle = 0$ on the whole set

$$\text{Opt}^{\text{some}}(d^1) \cap \text{Opt}^{\text{some}}(d^2).$$

Thus, we would not obtain conditions as simple as in Theorem 9.15. Section 9.5 discusses how Assumption A5 can be weakened in one case of particular interest, more precisely the case that Assumption A5 is not satisfied for points $m \in \mathcal{P}(\mathcal{S})$ with $m_i < 0$ for some $i \in \mathcal{S}$. Furthermore, in Subsection 9.7.2 we discuss how to weaken the assumption for a particular example. △

We conclude the section by reformulating the theorem for the simple case that there are only two strategies that are the unique optimal strategy for some $m \in O$: Let Assumption A1 and A2 hold and assume that $\mathcal{U} = \{d^1, d^2\}$. Define

$$\begin{aligned} g(m) &:= \left(V^{d^2}(m) - V^{d^1}(m)\right) \cdot (1, \dots, 1)^T \\ &= \left((\beta I - Q^{d^2}(m))^{-1} \cdot r^{d^2}(m) - (\beta I - Q^{d^1}(m))^{-1} \cdot r^{d^1}(m)\right) \cdot (1, \dots, 1)^T \end{aligned}$$

and assume that $g(\cdot)$ is twice continuously differentiable, i.e. assume that $Q_{ija}(\cdot)$ and $r(\cdot)$ are twice continuously differentiable. Then

$$\text{Opt}^{\text{unique}}(d^1) = \{m \in O : g(m) < 0\}$$

$$\begin{aligned}
\mathrm{Opt}^{\mathrm{some}}(d^1) &= \{m \in O : g(m) \leq 0\} \\
\mathrm{Opt}^{\mathrm{unique}}(d^2) &= \{m \in O : -g(m) < 0\} \\
\mathrm{Opt}^{\mathrm{some}}(d^2) &= \{m \in O : -g(m) \leq 0\}.
\end{aligned}$$

Moreover, we can rewrite the conditions (ii), (iii) and (iv) in terms of requirements for $\langle m^T Q^d(m), \nabla g(m)\rangle$:

Corollary 9.17. *Let for all $m \in O$ with $g(m) = 0$ the gradient satisfy $\nabla g(m) \neq 0$. Furthermore, assume that:*

- *The nonlinear Markov chains with transition rate matrix functions $Q^{d^1}(m)$ and $Q^{d^2}(m)$ converge in the limit towards some stationary distribution.*
- *One of the following statements holds true:*

  (a) For all $m \in Opt^{some}(d^1) \cap Opt^{some}(d^2)$ it holds that

  $$\langle (Q^{d^1}(m))^T m, \nabla g(m)\rangle > 0 \quad \text{and} \quad \langle (Q^{d^2}(m))^T m, \nabla(-g)(m)\rangle < 0$$

  (b) For all $m \in Opt^{some}(d^1) \cap Opt^{some}(d^2)$ it holds that

  $$\langle (Q^{d^1}(m))^T m, \nabla g(m)\rangle < 0 \quad \text{and} \quad \langle (Q^{d^2}(m))^T m, \nabla(-g)(m)\rangle > 0$$

  (c) For all $m \in Opt^{some}(d^1) \cap Opt^{some}(d^2)$ it holds that

  $$\langle (Q^{d^1}(m))^T m, \nabla g(m)\rangle > 0 \quad \text{and} \quad \langle (Q^{d^2}(m))^T m, \nabla(-g)(m)\rangle > 0.$$

Then for any initial condition $m_0 \in \mathcal{P}(\mathcal{S})$ the myopic adjustment process either converges towards a deterministic stationary mean field equilibrium or there is a $T > 0$ such that it remains in the set $Opt^{some}(d^1) \cap Opt^{some}(d^2)$ for all $t > T$.

Proof. It is only necessary to check whether the conditions (ii) to (iv) are satisfied:

(a) In this case, by definition of the digraphs $D(d)$, the digraph $D(d^1)$ contains the arc $d^1 \to d^2$ and the digraph $D(d^2)$ contains the arc $d^1 \to d^2$, which yields that condition (ii) is satisfied. Moreover, $d^1 \leq d^2$ is the unique ordering satisfying the consistency condition (iii) since $d^1 \leq_{D(d^1)} d^2$ and for d^2 nothing has to be checked. By assumption, we furthermore have that (iv) holds.

(b) When we exchange the roles of d^1 and d^2 we see that this case is symmetric to case (a), thus all conditions are satisfied and the ordering satisfying the consistency condition is $d^2 \leq d^1$.

(c) In this case, by definition of the digraphs $D(d)$, the digraph $D(d^1)$ contains the arc $d^1 \to d^2$ and the digraph $D(d^2)$ contains the arc $d^2 \to d^1$, which yields that condition (ii) is satisfied. We moreover obtain that $d^1 \leq d^2$ as well as $d^2 \leq d^1$ are orderings satisfying the consistency condition, which yields that also condition (iv) is satisfied.

□

9.4. Proof of the Global Convergence Theorem

In this section we will prove Theorem 9.15. For this purpose we assume throughout the whole section that Assumptions A1, A2, A3, A4 and A5 as well as the conditions (i)-(iv) of the theorem hold.

The proof relies heavily on the notion of flow invariance. More precisely, we prove that for any ordering $d^1, \ldots, d^u$ of the deterministic stationary strategies $\mathcal{U}$ that satisfy the condition (iii) and any $k < u$ the set $O \cap \left(\bigcup_{i \geq k} \mathrm{Opt}^{\mathrm{some}}(d^i)\right)$ is flow invariant for the differential inclusion $\dot{m} = F(m)$. Since proving flow invariance for differential inclusions is much more involved than proving flow invariance for ordinary differential equations, we will prove the statement indirectly. More precisely, we first prove a similar statement for acyclic orderings for $D(d)$ with respect to the ordinary differential equation $\dot{m} = (Q^d(m))^T m$. Thereafter, we will prove a technical lemma that relates flow invariance of the ordinary differential equation $\dot{x} = f(x)$ with the flow invariance of a differential inclusion, where $F(x) = \{f(x)\}$ for some set A . Using this lemma, we then derive the desired flow invariance result, which in combination with another technical lemma then allows us to prove Theorem 9.15.

The first lemma yields the desired statement on flow invariance for the ordinary differential equations $\dot{m} = (Q^d(m))^T m$. We note that the proof crucially relies on the structural assumptions regarding $\mathrm{Opt}^{\mathrm{some}}(d)$ and $\mathrm{Opt}^{\mathrm{unique}}(d)$ and in particular on the Assumption A4, which requires that at most two strategies from $\mathcal{U}$ are simultaneously optimal.

Lemma 9.18. *Let $d \in \mathcal{U}$ be a deterministic stationary strategy, let $\leq_{D(d)}$ be an acyclic ordering of the digraph $D(d)$ and enumerate the deterministic stationary strategies such that $d^1 <_{D(d)} d^2 <_{D(d)} \ldots <_{D(d)} d^u$. Furthermore, let $1 < k \leq u$. Then the set*

$$O \cap \left(\bigcup_{l \geq k} Opt^{some}(d^l) \right)$$

is flow invariant for $\dot{m} = (Q^d(m))^T m$.

Proof. We start with some observations regarding the boundary of the set

$$S_k := O \cap \left(\bigcup_{l \geq k} \mathrm{Opt}^{\mathrm{some}}(d^l) \right).$$

As a first step we note that $S_k = \{m \in O : g_{d^k}(m) \cdot \ldots \cdot g_{d^u}(m) \leq 0\}$. Indeed, let $m \in \bigcup_{l \geq k} \mathrm{Opt}^{\mathrm{some}}(d^l)$, then there is at least one $l \geq k$ such that $g_{d^l}(m) \leq 0$. If $g_{d^l}(m) = 0$, we directly have $g_{d^k}(m) \cdot \ldots g_{d^u}(m) = 0$. If $g_{d^l}(m) < 0$ then $m \in \mathrm{Opt}^{\mathrm{unique}}(d^l)$. This implies that $g_{d^j}(m) > 0$ for all $j \neq l$, thus $g_{d^k}(m) \cdot \ldots g_{d^u}(m) < 0$.

Now let $g_{d^k}(m) \cdot \ldots \cdot g_{d^u}(m) \leq 0$. If $g_{d^k}(m) \cdot \ldots g_{d^u}(m) = 0$, then there is at least one $l \geq k$ such that $g_l(m) = 0$, which implies that $m \in \mathrm{Opt}^{\mathrm{some}}(d^l)$. If $g_{d^k}(m) \cdot \ldots \cdot g_{d^u}(m) < 0$, then there is at least one $l \geq k$ such that $g_l(m) < 0$, which implies that $m \in \mathrm{Opt}^{\mathrm{unique}}(d)$.

Since $f(m) := g_{d^k}(m) \cdot \ldots \cdot g_{d^u}(m)$ is continuous and $f^{-1}((-\infty, 0)) \subseteq S_k$ as well as $S_k = f^{-1}((-\infty, 0])$ we obtain that

$$f^{-1}((-\infty, 0)) \subseteq \mathrm{int}(S_k) \quad \text{and} \quad S_k = \overline{S_k} = f^{-1}((-\infty, 0]).$$

Since additionally

$$g_{d^k}(m) \cdot \ldots \cdot g_{d^u}(m) = 0 \Leftrightarrow g_{d^l}(m) = 0 \quad \text{for some} \quad l \geq k,$$

we obtain that the boundary of the set S_k is a subset of

$$\{m \in O : g_{d^l}(m) = 0 \text{ for some } l \geq k\}.$$

As a next step, we show that the boundary is indeed a subset of

$$\bigcup_{l \geq k, j < k} \{m \in O : g_{d^l}(m) = 0 \text{ and } g_{d^j}(m) = 0\}:$$

Indeed, let $m \in O$ be such that $g_{d^l}(m) = 0$ for some $l \geq k$. Since by Assumption A3 the gradient $\nabla g_{d^l}(m)$ is non-vanishing, we obtain that m is not a local maximum of $g_{d^l}(\cdot)$. This yields that there is a sequence $m_n \to m$ such that $g_{d^l}(m_n) > 0$ for all $n \in \mathbb{N}$, which in particular implies that $m_n \notin \mathrm{Opt}^{\mathrm{some}}(d^l)$. Since by Theorem 5.2 there is an optimal deterministic stationary strategy $d^{j(n)}$ for m_n, we have $g_{d^{j(n)}}(m_n) \leq 0$ for all $n \in \mathbb{N}$. Since the set of all deterministic strategies is finite, we obtain that there is an index $\bar{j}$ occurring infinitely often in $j(n)$. Thus, there is a subsequence $(m_{n_k})_{k \in \mathbb{N}}$ such that $g_{d^{\bar{j}}}(m_{n_k}) \leq 0$. By continuity of $g_{d^{\bar{j}}}(\cdot)$, we obtain that $g_{d^{\bar{j}}}(m) \leq 0$. Since $m \in \mathrm{Opt}^{\mathrm{some}}(d^l)$, we obtain that $g_{d^{\bar{j}}}(m) = 0$.

Thus, it remains to show that there is an index $j < k$ such that $g_j(m) = 0$. Assume that this is not the case. Then by Assumption A3 we obtain that $g_{d^j}(m) > 0$ for all $j < k$. By continuity of $g_{d^j}(\cdot)$, we then find an $\epsilon > 0$ such that for all $\tilde{m} \in N_\epsilon(m)$ and $j < k$ we have $g_{d^j}(\tilde{m}) > 0$, which in particular implies that $\tilde{m} \notin \mathrm{Opt}^{\mathrm{some}}(d^j)$. Since for each point $\tilde{m} \in N_\epsilon(m)$ we find, by Theorem 5.2, an optimal deterministic strategy d^l $(l \geq k)$, we obtain that

$$N_\epsilon(m) \subseteq \bigcup_{l \geq k} \mathrm{Opt}^{\mathrm{some}}(d^l).$$

However, this contradicts the assumption that m is a boundary point of $\bigcup_{l \geq k} \mathrm{Opt}^{\mathrm{some}}(d^l)$.

With these preparations we prove the statement: For this we note that any boundary point lies in some set $\mathrm{Opt}^{\mathrm{some}}(d^j) \cap \mathrm{Opt}^{\mathrm{some}}(d^l)$ with $j < k \leq l$. By definition of the digraph $D(d)$ we have that there is no arc from l to j, which by Assumption A5 yields that

$$\langle (Q^d(m))^T m, \nabla g_{d^l}(m) \rangle < 0$$

for all $m \in \mathrm{Opt}^{\mathrm{some}}(d^j) \cap \mathrm{Opt}^{\mathrm{some}}(d^l)$. By Assumption A4 we have $g_{d^n}(m) = 0$ if and only if $n \in \{j, l\}$. Thus, it follows by the product rule that

$$\begin{aligned} \nabla g_{d^k}(m) \cdot \ldots \cdot g_{d^u}(m) &= \sum_{n=k}^{u} \left(g_{d^k}(m) \cdot \ldots g_{d^{n-1}}(m) \cdot g_{d^{n+1}}(m) \cdot \ldots g_{d^u}(m) \right) \cdot \nabla g_{d^n}(m) \\ &= \underbrace{\left(g_{d^k}(m) \cdot \ldots g_{d^{l-1}}(m) \cdot g_{d^{l+1}}(m) \cdot \ldots \cdot g_{d^u}(m) \right)}_{>0} \cdot \nabla g_{d^l}(m), \end{aligned}$$

which implies

$$\left\langle (Q^d(m))^T m, \nabla \left(g_{d^k}(m) \cdot \ldots \cdot g_{d^u}(m) \right) \right\rangle < 0.$$

Thus, Lemma 9.2 yields that S_k is flow invariant for $\dot{m} = (Q^d(m))^T m$. □

Corollary 9.19. *Let $d \in \mathcal{U}$ be a deterministic stationary strategy, let $\leq_{D(d)}$ be an acyclic ordering of the digraph $D(d)$ and enumerate the deterministic stationary strategies such that $d^1 <_{D(d)} d^2 <_{D(d)} \ldots <_{D(d)} d^u$. Furthermore, let $1 \leq k \leq u$. Then the set*

$$\mathcal{P}(\mathcal{S}) \cap \left(\bigcup_{l \geq k} Opt^{some}(d^l) \right)$$

is flow invariant for $\dot{m} = (Q^d(m))^T m$.

Proof. By Lemma 9.1 it suffices to check that $(Q^d(m))^T m \in T_{\mathcal{P}(\mathcal{S})}(m)$ in order to show that $\mathcal{P}(\mathcal{S})$ is flow invariant for $\dot{m} = (Q^d(m))^T m$. This has been done in the proof of Theorem 8.4. Thus, by Assumption A3 the statement holds for $k = 1$. For $k > 1$ note that

$$\mathcal{P}(\mathcal{S}) \cap \left(\bigcup_{l \geq k} \mathrm{Opt}^{\mathrm{some}}(d^l) \right) = \mathcal{P}(\mathcal{S}) \cap \left(O \cap \left(\bigcup_{l \geq k} \mathrm{Opt}^{\mathrm{some}}(d^l) \right) \right).$$

By Lemma 9.18, the set $O \cap \left(\bigcup_{l \geq k} \mathrm{Opt}^{\mathrm{some}}(d^l) \right)$ is flow invariant. Since $\mathcal{P}(\mathcal{S})$ is flow invariant as well, the intersection of the two sets is again flow invariant for $\dot{m} = (Q^d(m))^T m$. □

The next lemma yields the desired link between the flow invariance for the ordinary differential equation $\dot{m} = (Q^d(m))^T$ and the differential inclusion $\dot{m} \in F(m)$:

Lemma 9.20. *Let $O \subseteq \mathbb{R}^S$ be an open set and let $F : O \to 2^{\mathbb{R}^S}$ be a set-valued map. Let A, B, C be subsets of O and assume that B is closed. Let $f : O \to \mathbb{R}^S$ be a Lipschitz continuous function such that $F(x) = \{f(x)\}$ for all $x \in A \setminus B$. Furthermore, assume that $A \cup B$ is flow invariant for $\dot{x} \in F(x)$ and that $B \cup C$ is flow invariant for $\dot{x} = f(x)$. Then B is flow invariant for $\dot{x} \in F(x)$.*

Proof. Assume that there is a solution $x(\cdot)$ of $\dot{x} \in F(x)$ such that there are times $0 \leq t_0 < t_1$ such that $x(t_0) \in B$ and $x(t) \notin B$ for all $t_0 < t < t_1$. Since $A \cup B$ is flow invariant for $\dot{x} \in F(x)$, we obtain that $x(t) \in A \setminus B$ for all $t_0 < t < t_1$. Thus, the solution $x(t)$ on $[t_0, t_1)$ satisfies that $x(t_0) \in B$ and $\dot{x}(t) = f(x(t))$ for almost every $t \in [t_0, t_1)$. Therefore, $x(\cdot)$ is a Caratheodory solution of the differential equation $\dot{x} = f(x)$. By Walter (1998, Theorem 10.XVIII) there is a unique Caratheodory solution. Since $\dot{x} = f(x)$ has a unique classical solution, the unique classical solution and the unique Caratheodory solution coincide. Since the set $B \cup C$ is flow invariant for the ordinary differential equation $\dot{x} = f(x)$, we obtain a contradiction because $x(\cdot)$ would be a trajectory that leaves $B \cup C$. □

As a next step, we show that $\mathcal{P}(\mathcal{S})$ is flow invariant for the differential inclusion:

Lemma 9.21. *The set $\mathcal{P}(\mathcal{S})$ is flow invariant for $\dot{m} \in F(m)$.*

Proof. Assume that $\mathcal{P}(\mathcal{S})$ is not flow invariant for $\dot{m} \in F(m)$, then there are $\tilde{t}_2 > \tilde{t}_1 \geq 0$ such that $m(\tilde{t}_1) \in \mathcal{P}(\mathcal{S})$ and $m(t) \notin \mathcal{P}(\mathcal{S})$ for all $\tilde{t}_1 < t < \tilde{t}_2$. By Assumption A4 there are at most two strategies d^1, d^2 such that $g_{d^1}(m(\tilde{t}_1)) \leq 0$ and $g_{d^2}(m(\tilde{t}_1)) \leq 0$. In particular, there is an $\epsilon > 0$ such that $g_d(\tilde{m}) > 0$ for all $d \neq d^1, d^2$ and $\tilde{m} \in N_\epsilon(m(\tilde{t}_1))$.

Thus the differential inclusion restricted onto $N_\epsilon(m(\tilde{t}_1))$ is of the type discussed in Section 9.1.2, where $\phi = g_{d^1} = -g_{d^2}$. In particular, Assumption A5 and the consistency condition (iii) yield that one of the following three cases will hold for all $m \in \mathrm{Opt}^{\mathrm{some}}(d^1) \cap \mathrm{Opt}^{\mathrm{some}}(d^2)$ simultaneously

- $\langle (Q^{d^1}(m))^T m, \nabla\phi(m)\rangle > 0$ and $\langle (Q^{d^2}(m))^T m, \nabla - \phi(m)\rangle < 0$
- $\langle (Q^{d^1}(m))^T m, \nabla\phi(m)\rangle < 0$ and $\langle (Q^{d^2}(m))^T m, \nabla - \phi(m)\rangle > 0$
- $\langle (Q^{d^1}(m))^T m, \nabla\phi(m)\rangle > 0$ and $\langle (Q^{d^2}(m))^T m, \nabla - \phi(m)\rangle > 0$.

In particular, the results presented in Subsection 9.1.2 yield that the solution in $N_\epsilon(m(\tilde{t}_1))$ is the solution of a classical ordinary differential equation $\dot{m} = (\tilde{Q}(m))^T m$ with $\tilde{Q}(m)$

being $Q^{d^1}(m)$ in the first case, $Q^{d^2}(m)$ in the second case and

$$\frac{\langle (Q^{d^1}(m))^T m, \nabla\phi(m)\rangle}{\langle (Q^{d^1}(m))^T m, \nabla\phi(m)\rangle + (-\langle (Q^{d^2}(m))^T m, \nabla\phi(m)\rangle)} Q^{d^1}(m)$$
$$+ \frac{-\langle (Q^{d^2}(m))^T m, \nabla\phi(m)\rangle}{\langle (Q^{d^1}(m))^T m, \nabla\phi(m)\rangle + (-\langle (Q^{d^2}(m))^T m, \nabla\phi(m)\rangle)} Q^{d^2}(m)$$

in the third case. In all three cases we obtain that $\tilde{Q}(\cdot)$ is Lipschitz continuous. This implies, as in the proof of Lemma 9.19, that the set $\mathcal{P}(\mathcal{S})$ is flow invariant for these type of differential equations, which is the desired contradiction. □

Remark 9.22. In order to prove flow invariance it would be sufficient to require that for all points $m \in \mathcal{P}(\mathcal{S})$ satisfying $m_i = 0$ for some $i \in \mathcal{S}$ there are at most two strategies d^1 and d^2 such that $g_{d^1}(m) \leq 0$ and $g_{d^2}(m) \leq 0$: Indeed, we note that $\{m \in \mathbb{R}^S : \sum_{i\in\mathcal{S}} m_i = 1\}$ is flow invariant for F since $\sum_{i\in\mathcal{S}} \dot{m}_i(t) = 0$ holds almost surely because all transition rate matrices are conservative. This implies that at time $\tilde{t}_1$ where the solution leaves the set $\mathcal{P}(\mathcal{S})$ a component $m_i(\tilde{t}_1)$ has to be zero, for which we can then provide the same argument as before. △

The argument that due to Assumption A4 the solution is locally the solution of a differential inclusion of the type investigated in Subsection 9.1.2 also yields using Lemma 9.4 uniqueness of trajectories:

Lemma 9.23. *For any $m_0 \in \mathcal{P}(\mathcal{S})$ there is a unique solution of $\dot{m}(t) \in F(m(t))$.*

With all these preparations we prove the desired flow invariance theorem for the differential inclusion:

Lemma 9.24. *Let $\leq$ be an ordering of $\mathcal{U}$ satisfying (iii) and enumerate the deterministic stationary strategies such that $d^1 < d^2 < \ldots < d^u$. Furthermore, let $1 \leq k \leq u$. Then the set*

$$\mathcal{P}(\mathcal{S}) \cap \left(\bigcup_{l \geq k} Opt^{some}(d^l) \right)$$

is flow invariant for F.

Proof. We proceed by induction on k. For $k = 1$ we note that the set $\mathcal{P}(\mathcal{S})$ is flow invariant for F by Lemma 9.21.

Now let $u \geq k > 1$ and assume that we have proven that

$$\mathcal{P}(\mathcal{S}) \cap \left(\bigcup_{l \geq k-1} \mathrm{Opt}^{\mathrm{some}}(d^i) \right)$$

is flow invariant. By definition of the ordering $\leq$, we obtain that there is an acyclic ordering $\leq_{D(d^k)}$ for $D(d^k)$ such that $d^k \leq_{D(d^k)} d^j$ for all $j > k$. Setting $T_{\leq} := \{d \in \mathcal{U} : d \leq_{D(d^k)} d^k\}$ we obtain that

$$\bigcup_{d \notin T_{\leq}} \mathrm{Opt}^{\mathrm{some}}(d) \supseteq \bigcup_{i \geq k} \mathrm{Opt}^{\mathrm{some}}(d^i).$$

By Corollary 9.19, we then obtain that

$$\mathcal{P}(\mathcal{S}) \cap \left(\bigcup_{d \notin T_{\leq}} \mathrm{Opt}^{\mathrm{some}}(d) \right)$$

is flow invariant for $\dot{m} = (Q^{d^k}(m))^T m$.

As a final preparation we note that $f(m) := \sum_{i \in \mathcal{S}} m_i Q^d_{ij}(m)$ is Lipschitz continuous on $N_\epsilon(\mathcal{P}(\mathcal{S}))$ with constant $(M(1+\epsilon) + L) \cdot S$, where M is a Lipschitz constant for all function $Q_{ij}(\cdot)$ $(i, j \in \mathcal{S})$ and L is the constant that uniformly bounds $Q_{ij}(m)$ for all $m \in \mathcal{P}(\mathcal{S})$ and $i, j \in \mathcal{S}$:

$$\begin{aligned}
|f(m^1) - f(m^2)|_1 &= \sum_{j \in \mathcal{S}} \left| \sum_{i \in \mathcal{S}} \left(m_i^1 Q_{ij}(m^1) - m_i^2 Q_{ij}(m^2) \right) \right| \\
&\leq \sum_{j \in \mathcal{S}} \sum_{i \in \mathcal{S}} |m_i^1 Q_{ij}(m^1) - m_i^1 Q_{ij}(m^2)| + |m_i^1 Q_{ij}(m^2) - m_i^2 Q_{ij}(m^2)| \\
&= \sum_{j \in \mathcal{S}} \sum_{i \in \mathcal{S}} \left(|m_i^1||Q_{ij}(m^1) - Q_{ij}(m^2)| + |Q_{ij}(m^2)||m_i^1 - m_i^2| \right) \\
&\leq \sum_{j \in \mathcal{S}} \left(\sum_{i \in \mathcal{S}} |m_i^1| M |m^1 - m^2|_1 + \sum_{i \in \mathcal{S}} L |m_i^1 - m_i^2| \right) \\
&= \sum_{j \in \mathcal{S}} \left(M(1+\epsilon)|m^1 - m^2|_1 + L|m^1 - m^2|_1 \right) \\
&= (M(1+\epsilon) + L) \cdot S |m^1 - m^2|_1.
\end{aligned}$$

Thus, we apply Lemma 9.20 with

$$\begin{aligned}
A &= \mathrm{Opt}^{\mathrm{some}}(d^k) \cap \mathcal{P}(\mathcal{S}) \\
B &= \mathcal{P}(\mathcal{S}) \cap \left(\bigcup_{l \geq k+1} \mathrm{Opt}^{\mathrm{some}}(d^l) \right) = \mathcal{P}(\mathcal{S}) \cap \{ m \in O : g_{d^{k+1}}(m) \cdot \ldots \cdot g_{d^u}(m) \leq 0 \} \\
C &= \mathcal{P}(\mathcal{S}) \cap \left(\bigcup_{l < k \text{ s.t. } d^l \notin T_{\leq}} \mathrm{Opt}^{\mathrm{some}}(d^l) \right)
\end{aligned}$$

to obtain the desired claim. □

The following technical lemma deals with the special case of a vertex $d \in \mathcal{U}$ that is the final vertex of any acyclic ordering:

Lemma 9.25. *Let $d \in \mathcal{U}$ be a deterministic stationary strategy such that for any other strategy $\hat{d} \in \mathcal{U} \setminus \{d\}$ we either have $\hat{d} \to d$ or $Opt^{some}(d) \cap Opt^{some}(\hat{d}) = \emptyset$, then $\mathcal{P}(\mathcal{S}) \cap Opt^{unique}(d)$ is flow invariant for F.*

Proof. By the argument presented in the proof of Lemma 9.18 we have

$$\partial \mathrm{Opt}^{\mathrm{unique}}(d) \subseteq \bigcup_{d \in \mathcal{U} \setminus \{d\}} \left(\mathrm{Opt}^{\mathrm{some}}(d) \cap \mathrm{Opt}^{\mathrm{some}}(\hat{d}) \right).$$

By assumption we satisfy for all $m \in \mathrm{Opt}^{\mathrm{some}}(d)$, that $\langle (Q^d(m))^T m, \nabla g_d(m) \rangle < 0$, which, by Lemma 9.3, yields the flow invariance of $\mathrm{Opt}^{\mathrm{unique}}(d)$. By Lemma 9.21 the set $\mathcal{P}(\mathcal{S})$ is flow invariant for $\dot{m} \in F(m)$. Combining the two results yields the desired claim. □

With these preparations we prove the main theorem:

Proof of Theorem 9.15. By Lemma 9.23, there is a unique solution of the differential inclusion for any initial condition, and by Lemma 9.24 the set $\mathcal{P}(\mathcal{S})$ is flow invariant for the differential inclusion. Furthermore, we note that if a trajectory stays inside a set $\mathrm{Opt}^{\mathrm{unique}}(d)$ for all $t > T$, then by condition (i) the trajectory will converge towards the stationary point given $Q^d(\cdot)$.

Let $d^1, \ldots, d^u$ be a ordering that satisfies (iii) and assume that i is maximal such that $m(0) \in \mathcal{P}(\mathcal{S}) \cap \left(\bigcup_{k \geq i} \mathrm{Opt}^{\mathrm{some}}(d^k) \right)$, then $m(0) \in \mathrm{Opt}^{\mathrm{some}}(d^i)$. If the trajectory does not leave the set $\mathcal{P}(\mathcal{S}) \cap \mathrm{Opt}^{\mathrm{unique}}(d^i)$ we have by the previous observation convergence towards a deterministic stationary mean field equilibrium. Else, there is a $t_0 \geq 0$ such that the trajectory will stay in

$$\mathcal{P}(\mathcal{S}) \cap \left(\mathrm{Opt}^{\mathrm{some}}(d^i) \setminus \mathrm{Opt}^{\mathrm{unique}}(d^i) \right)$$

for all $t \geq t_0$ or there is a $t_1 > 0$ such that $m(t_1) \notin \mathcal{P}(\mathcal{S}) \cap \mathrm{Opt}^{\mathrm{some}}(d^i)$. Then we find an $\hat{i} > i$ such that $m(t_1) \in \mathcal{P}(\mathcal{S}) \cap \left(\bigcup_{k \geq \hat{i}} \mathrm{Opt}^{\mathrm{some}}(d^k) \right)$ and we can reapply the previous argument.

If we reach the final vertex d of an ordering that satisfies (iii) and there is another strategy $\hat{d}$ that is also the final vertex of an ordering, then we obtain that both $\mathcal{P}(\mathcal{S}) \cap \mathrm{Opt}^{\mathrm{some}}(d)$ and $\mathcal{P}(\mathcal{S}) \cap \mathrm{Opt}^{\mathrm{some}}(\hat{d})$ are flow invariant, which in particular yields that $\mathcal{P}(\mathcal{S}) \cap \mathrm{Opt}^{\mathrm{some}}(d) \cap \mathrm{Opt}^{\mathrm{some}}(\hat{d})$ is flow invariant, which yields that whenever a trajectory leaves $\mathcal{P}(\mathcal{S}) \cap \mathrm{Opt}^{\mathrm{unique}}(d)$, then it will remain in

$$\mathcal{P}(\mathcal{S}) \cap \mathrm{Opt}^{\mathrm{some}}(d) \cap \mathrm{Opt}^{\mathrm{some}}(\hat{d}) \subseteq \mathcal{P}(\mathcal{S}) \cap \mathrm{Opt}^{\mathrm{some}}(d) \setminus \mathrm{Opt}^{\mathrm{unique}}(d)$$

for all times, which proves the claim.

If there is a unique final vertex d then we obtain, by Lemma 9.25, that $\mathrm{Opt}^{\mathrm{unique}}(d)$ is flow invariant. Thus, if the trajectory leaves the set $\mathrm{Opt}^{\mathrm{some}}(d) \setminus \mathrm{Opt}^{\mathrm{unique}}(d)$, then we have convergence towards a deterministic stationary mean field equilibrium. □

9.5. An Alternative Formulation of the Global Convergence Theorem

It is often a problem that for $m_i < 0$ Assumption A5 does not longer hold true, an example is the case $c = 1$ of the consumer choice model presented in Section 6.3 (see Subsection 9.7.2). However, it is easily possible to prove the result given slightly different assumptions that circumvent the problem:

- Strengthen Assumption A2 by requiring that $Q(\cdot)$ and $r(\cdot)$ are defined on an open set $O \supseteq [0,\infty)^S$ and weaken Assumption A5, by requiring that
$$\langle (Q^d(m))^T m, \nabla g_{d^1}(m) \rangle$$
is non-vanishing over and has uniform over the set
$$[0,\infty)^S \cap \mathrm{Opt}^{\mathrm{some}}(d^1) \cap \mathrm{Opt}^{\mathrm{some}}(d^2)$$
instead of $\mathrm{Opt}^{\mathrm{some}}(d^1) \cap \mathrm{Opt}^{\mathrm{some}}(d^2)$ for all $d, d^1, d^2 \in \mathcal{U}$.

- Furthermore, additionally require in condition (iv) that
$$\langle Q^{\hat{d}}(m)^T m, e_i \rangle < 0$$
for all $m \in \mathrm{Opt}^{\mathrm{some}}(\hat{d}) \cap [0,\infty)^S$ and all $i \in \mathcal{S}$.

With this assumptions we can prove Theorem 9.15 with a minor adaptation, namely we have to readjust Lemma 9.18, by proving that
$$O \cap [0,\infty)^S \cap \left(\bigcup_{l \geq k} \mathrm{Opt}^{\mathrm{some}}(d^l) \right)$$
is flow invariant for $\dot{m} = (Q^d(m))^T m$. This is possible by the same techniques since the set can be described by
$$\tilde{S}_k = \{ m \in O : g_{d^k}(m) \cdot \ldots \cdot g_{d^u}(m) \leq 0 \wedge -m_i \leq 0 \quad \text{for all} \quad i \in \mathcal{S} \}$$
and the boundary of this set is a subset of $[0,\infty)^S$ since the set itself is a subset of the closed set $[0,\infty)^S$. Thus, by the same arguments as before the boundary is a subset of
$$\bigcup_{l \geq k, j < k} \{ m \in [0,\infty)^S : g_{d^l}(m) = 0 \quad \text{and} \quad g_{d^j}(m) = 0 \} \cup \{ m_i = 0 \text{ for some } i \in \mathcal{S} \}.$$

Since any boundary point lies in $[0, \infty)^S$ we obtain for all points satisfying $g_{d^k}(m) \cdot \ldots \cdot g_{d^u}(m) = 0$ that

$$\langle (Q^d(m))^T m, \nabla g_{d^k}(m) \cdot \ldots \cdot g_{d^u}(m) \rangle < 0$$

by the same arguments as in the proof of Lemma 9.18. For all points satisfying $m_i = 0$ for some $i \in \mathcal{S}$ we obtain, as in the proof of Theorem 8.4, that $\langle (Q^d(m))^T m, e_i \rangle \leq 0$. Thus, in total we can again apply Lemma 9.2, which yields the desired claim. Similarly, we can prove the accordingly modified version of Lemma 9.25 since we required that $\langle (Q^d(m))^T m, e_i \rangle < 0$ for all relevant points. All other lemmata can be proven in the same way since they do not rely on the fact that Assumption A5 holds for some open superset of $\mathcal{P}(\mathcal{S})$.

9.6. An Algorithm to Verify whether the Consistency Condition Holds

The question whether an ordering satisfying (iii) exists or not seems to be complex at first sight. One would have to check for all possible permutations whether the ordering satisfies the condition. However, due to the close connection to the notion of acyclic orderings, we can provide a polynomial algorithm that determines whether such an ordering exists or not. Moreover, relying on backtracking it is even possible, as for acyclic orderings, to provide an algorithm that yields all orderings satisfying the consistency condition (iii).

Our algorithm will be a generalization of the following simple algorithm to determine an acyclic ordering if it exists: It relies on the fact that in any acyclic digraph there is a vertex with indegree 0 and that if we remove a vertex from an acyclic graph, then it remains acyclic. More precisely, in each step we pick a vertex with indegree 0, add it to the tail of the acyclic ordering obtained so far and delete in from the digraph (Bang-Jensen and Gutin, 2010, Section 2.1), formally this yields to algorithm 1.

In our case we want to construct an ordering $\{d^1, \ldots, d^S\}$ such that $\{d^{i+1}, \ldots, d^S\}$ lies behind d^i in an acyclic ordering of $D(d^i)$ for each strategy $d^i \in \mathcal{U}$. The central modification of the algorithm is that we do not only delete arcs from the digraphs, but we also add arcs as well in order to ensure that an acyclic ordering of $D(d) \setminus \{d^1, \ldots, d^i\}$ is also an acyclic ordering of $D(d)$. More precisely, if we add d to the ordering since it had indegree 0 in $D(d)$, then we add the arcs $d^1 \to d^2$ to the graphs $D(\hat{d})$ for all $\hat{d} \in \hat{V}$ whenever $d^1 \to d$ and $d \to d^2$ are both arcs in $D(\hat{d})$. This yields that two vertices $d^1, d^2 \in \hat{V}$ are connected in the modified digraph if and only if they are connected in the original graph $D(\hat{d})$.

Algorithm 1: An Algorithm to Find an Acyclic Ordering (if it exists)

Data: A digraph $D = (V, A)$
Result: An acyclic ordering $(x^1, \dots, x^n)$ or $\emptyset$ (if no such ordering exists)

$\hat{V} \longleftarrow V$
$I \longleftarrow$ set of all vertices of D with indegree 0
$O \longleftarrow$ an empty list
while I *non-empty* **do**
  remove a vertex x from I
  remove x from $\hat{V}$
  add x to the tail of O
  for $a = (x, y) \in A$ **do**
    remove a from A
    if y *has no incomming arcs in* D **then**
      add y to I
if $\hat{V}$ *is non-empty* **then**
  return $\emptyset$
else
  return O

The following algorithm formally describes this graph modification:

Algorithm 2: GraphModification

Data: A digraph $D = (V, A)$ and a vertex x
Result: A new digraph $\hat{D} = (\hat{V}, \hat{A})$ such that $\hat{V} = V \setminus \{x\}$ and for all $x^1, x^2 \in \hat{V}$ there is a path from x^1 to x^2 in D if and only if there is a path from x^1 to x^2 in $\hat{D}$

$\hat{V} \longleftarrow V \setminus \{x\}$
$\hat{A} \longleftarrow A$
for x^1 *such that* $x^1 \to x \in A$ **do**
  for x^2 *such that* $x \to x^2 \in A$ **do**
    add $x^1 \to x^2$ to $\hat{A}$
delete all arcs containing x from $\hat{A}$

The next two results state that the graph modification has the desired properties:

Lemma 9.26. *Let D be an arbitrary digraph, let x be a vertex of D, let $\hat{D} = (\hat{V}, \hat{A})$ be the digraph resulting from the algorithm GraphModification and let $x^1, x^2 \in \hat{V}$ be two arbitrary vertices. Then there is a path from x^1 to x^2 in D if and only if there is a path from x^1 to x^2 in $\hat{D}$.*

Proof. Let $\tilde{x}^1, \dots, \tilde{x}^l$ be a path from x^1 to x^2 in D (that is $\tilde{x}^1 = x^1$ and $\tilde{x}^l = x^2$). If $x \notin \{\tilde{x}^1, \dots, \tilde{x}^l\}$ then $\tilde{x}^1, \dots, \tilde{x}^l$ is a path in $\hat{D}$ since we only deleted arcs that have x as starting point or as an end point. If $x \in \{\tilde{x}^1, \dots, \tilde{x}^l\}$ let $\tilde{x}^i = x$ then, by construction of $\hat{D}$, the arc $\tilde{x}^{i-1} \to x^{i+1}$ lies in $\hat{D}$, which yields that $\tilde{x}^1, \dots, \tilde{x}^{i-1}, \tilde{x}^{i+1}, \dots \tilde{x}^l$ is a path joining x^1 and x^2 in $\hat{D}$.

Now let $\hat{x}^1, \dots, \hat{x}^l$ be a path from x^1 to x^2 in $\hat{D}$ (that is $\hat{x}^1 = x^1$ and $\hat{x}^l = x^2$). If the arcs $\hat{x}^i \to \hat{x}^{i+1}$ ($i \in \{1, \dots, l-1\}$) all lie in D then $\hat{x}^1, \dots, \hat{x}^l$ is a path in D from x^1 to x^2. Else let i_1 be the smallest index such that $\hat{x}^{i_1} \to \hat{x}^{i_1+1} \notin D$ and i_m be the largest index such that $\hat{x}^{i_m} \to \hat{x}^{i_m+1} \notin D$. Since any arc $x^{j_1} \to x^{j_2}$ that lies in $\hat{D}$ but not in D has been constructed because $\hat{x}^{j_1} \to x$ and $x \to \hat{x}^{j_2}$ are arcs in D, we have that

$$\hat{x}^1, \dots, \hat{x}^{i_1}, x, \hat{x}^{i_m}, \dots, \hat{x}^l$$

is a path from x^1 to x^2 in D. □

Corollary 9.27. *Let D be an acyclic graph.*

(i) Let $\leq_{\hat{D}}$ be an acyclic ordering of $\hat{D}$, then there is an acyclic ordering of D such that $x^1 \leq x^2 \Leftrightarrow x^1 \leq_{\hat{D}} x^2$ for all $x^1, x^2 \in \hat{V}$.

(ii) Let $\leq_D$ be an acyclic ordering of D, then there is an acyclic ordering of $\hat{D}$ such that $x^1 \leq x^2 \Leftrightarrow x^1 \leq_{\hat{D}} x^2$ for all $x^1, x^2 \in \hat{V}$.

Proof. We start with a preliminary observation: If $\leq_D$ is an acyclic ordering of D, $x^1, x^2 \in V(\hat{D})$ and there is a path $\hat{x}^1, \dots, \hat{x}^l$ from x^1 to x^2 in $\hat{D}$ (that is $\hat{x}_1 = x^1$ and $\hat{x}_l = x^2$), then $\hat{x}^i \leq_{\hat{D}} \hat{x}^{i+1}$ for all $i \in \{1, \dots, l-1\}$. By transitivity of the acyclic ordering we obtain that $x^1 = \hat{x}^1 \leq_{\hat{D}} \hat{x}^l = x^2$.

Let $\leq_{\hat{D}}$ be an acyclic ordering, enumerate the vertices of $\hat{D}$ such that $x^1 \leq_{\hat{D}} x^2 \leq_{\hat{D}} \dots \leq_{\hat{D}} x^{u-1}$. Let $\leq_D$ be a partial ordering of $V(D)$ such that $x^1 \leq_D x^2 \leq_D \dots \leq_D x^{u-1}$: Assume that $\leq_D$ cannot be extended to an acyclic ordering. Then, by definition of acyclic orderings, there is an arc $x^j \to x^i$ with $1 \leq i < j \leq u-1$. By Lemma 9.26, this implies that there is a path from x^j to x^i in $\hat{D}$, which would yield by the preliminary observation that $x^j \leq_{\hat{D}} x^i$, which is a contradiction and thus proves (i). The proof of (ii) is analogous. □

With these preparations we formulate our algorithm:

Algorithm 3: An Algorithm to Find an Acyclic Ordering (if it exists)

Data: A family of digraphs $(D(d))_{d \in \mathcal{U}}$

Result: An ordering $(d^1, \ldots, d^u)$ of $\mathcal{U}$ satisfying (iii) or $\emptyset$ (if no such ordering exists)

$\hat{V} \longleftarrow \mathcal{U}$
$I \longleftarrow$ set of all $d \in \mathcal{U}$ with indegree 0 in $D(d)$
$O \longleftarrow$ am empty list
while I *non-empty* **do**
  remove an element d from I
  remove d from $\hat{V}$
  add d to the tail of O
  for $\hat{d} \in \hat{V}$ **do**
    $D(\hat{d}) \leftarrow \text{GraphModification}(D(\hat{d}), d))$
    if $\hat{d}$ *has no incoming arcs in* $D(\hat{d})$ **then**
      add $\hat{d}$ to I
if $\hat{V}$ *is non-empty* **then**
  return $\emptyset$
else
  return O

Theorem 9.28. *The algorithm is correct, that is whenever an ordering exists the algorithm finds one and if no ordering exists the algorithm returns* $\emptyset$.

Proof. We have to prove two things. First, that if the algorithm returns an ordering O then it satisfies condition (iii). Second, that if there is an ordering that satisfies (iii) then the algorithm will not return $\emptyset$.

If the algorithm returns an ordering $O = (d^1, \ldots, d^u)$, then in each step we find a strategy d^i such that d^i has indegree 0 in the digraph

$$\hat{D}(d^i) = \text{GraphModification}(\ldots \text{GraphModification}(D(d^i), d^1), \ldots, d^{i-1}).$$

Since d^i has indegree 0, we obtain using Algorithm 1 an acyclic ordering of $\hat{D}(d^i)$ such that $d^i \leq d^j$ for all $j > i$. By Corollary 9.27, we then also obtain an acyclic ordering of $D(d^i)$ such that $d^i \leq d^j$ for all $j \geq i$. This proves the first claim.

Now assume that $(d^1, \ldots, d^u)$ satisfies (iii). Then for each $i \in \mathcal{S}$ we find an acyclic ordering of $D(d^i)$ such that $d^i \leq d^j$ for all $j > i$. In particular, Corollary 9.27 yields that there is an acyclic ordering $\leq_{\hat{D}(d^i)}$ of

$$\hat{D}(d^i) = \text{GraphModification}(\ldots \text{GraphModification}(D(d^i), d^1), \ldots, d^{i-1})$$

such that $d^i \leq_{\hat{D}(d^i)} d^j$ for all $j > i$. However, this yields, by definition, that there is no incoming arc into d^i in $\hat{D}(d^i)$, which in particular yields that in step i the vertex d^i lies in I. Thus, choosing d^i in step i for all $1 \leq i \leq n$ yields that the algorithm outputs $(d^1, \ldots, d^u)$, which proves the second claim. □

As in the case of acyclic orderings, we can obtain all orderings satisfying the consistency condition by applying the algorithm for each possible choice of a vertex from the set I in every step. From an algorithmic viewpoint, it is important to find an efficient way to maintain the necessary information for every sub-problem. In the context of acyclic orderings such a procedure has been introduced in Knuth and Szwarcfiter (1974). A similar procedure can also be formulated for our case. However, in our case this would become even more complex due to the graph adjustments. Since the examples we consider here are small enough such that efficiency concerns do not matter so much, we omit the description of such an algorithm that outputs all acyclic orderings here.

9.7. Application of the Global Convergence Theorem in the Examples of Chapter 6

In this section we apply the global convergence theorem for the two consumer choice examples presented in Chapter 6. The theorem cannot be applied to the corruption model of Section 6.2 since the underlying nonlinear Markov chain has a reducible generator $Q(m)$ for some $m \in \mathcal{P}(\mathcal{S})$ and thus the convergence behaviour cannot be classified using the results from Section 9.2. In the consumer choice model with constant dynamics, we can explicitly compute all solutions of the myopic adjustment process. Moreover, we obtain that in some cases the theorem yields convergence towards the unique stationary mean field equilibrium. In other cases (with at least one mixed strategy equilibrium), the conditions of the theorem are violated and we indeed observe non-convergence. In the consumer choice model with non-constant dynamics, we can no longer solve the differential inclusion describing the myopic adjustment process explicitly, but we can again apply the convergence result for some cases. We thereafter analyse the special case $c = 1$ and show that we can still obtain convergence although the assumptions of the theorem are violated. Furthermore, we establish manually convergence towards the distribution of the infinitely many stationary equilibria.

9.7.1. The Consumer Choice Model of Section 6.1

We start our investigations with the consumer choice model of Section 6.1. In this setting, we can explicitly solve the differential inclusion describing the myopic adjustment process. Thus, we can discuss the limit behaviour also without applying Theorem 9.15. This discussion in particular illustrates the necessity of the consistency condition (iii). We will start by verifying that we are in the setting of this chapter by verifying that assumption A2, A3, A4 and A5 are satisfied. Then we discuss for which parameters the conditions of the theorem are satisfied and discuss the limit behaviour with and without the help of Theorem 9.15.

As a first step, we need to formulate a suitable extension of the rewards onto an open subset containing $\mathcal{P}(\mathcal{S})$ such that we satisfy Assumption A2. Let $\delta > 0$ be small enough (i.e. $\delta \ll \min\{k_1, 1-k_2\}$) and define $O_\delta := (-\delta, 1+\delta)^2$ and set

$$
\begin{aligned}
Q_{\cdot\cdot change} &: (-\delta, 1+\delta)^2 \to \mathbb{R}^{2\times 2}, m \mapsto \begin{pmatrix} -b & b \\ b & -b \end{pmatrix} \\
Q_{\cdot\cdot stay} &: (-\delta, 1+\delta)^2 \to \mathbb{R}^{2\times 2}, m \mapsto \begin{pmatrix} -\epsilon & \epsilon \\ \epsilon & -\epsilon \end{pmatrix} \\
r_{\cdot change} &: (-\delta, 1+\delta)^2 \to \mathbb{R}^{2}, m \mapsto \begin{pmatrix} \ln(f_\delta(m_1)) + s_1 - c \\ \ln(f_\delta(1-m_1)) + s_2 - c \end{pmatrix} \\
r_{\cdot stay} &: (-\delta, 1+\delta)^2 \to \mathbb{R}^{2}, m \mapsto \begin{pmatrix} \ln(f_\delta(m_1)) + s_1 \\ \ln(f_\delta(1-m_1)) + s_2 \end{pmatrix},
\end{aligned}
$$

which on $\mathcal{P}(\mathcal{S})$ coincides with the model introduced in Section 6.1 since $m_2 = 1 - m_1$. As in the discussion of optimal strategies in Section 6.1, we obtain that it is optimal to play strategy $\{change\} \times \{stay\}$ if and only if

$$\frac{f_\delta(m_1)}{f_\delta(1-m_1)} \leq \exp\left(\frac{c(\beta + 2\epsilon)}{b - \epsilon} - s_1 + s_2\right).$$

Since $f_\delta(\cdot)$ is increasing and $\delta > 0$ is small enough,this yields that it is optimal to play $\{change\}\times\{stay\}$ if and only if $m_1 \leq k_1$. For the other two strategies $\{stay\}\times\{stay\}$ and $\{stay\}\times\{change\}$ we obtain analogously that for small $\delta > 0$ the strategy $\{stay\}\times\{stay\}$ is optimal if and only if $k_1 \leq m_1 \leq k_2$ and the strategy $\{stay\} \times \{change\}$ is optimal if and only if $m_1 \geq k_2$.

With these observations, we characterize the optimality sets as follows

$$
\begin{aligned}
g_{\{change\}\times\{stay\}}(m) &= m_1 - k_1 \\
g_{\{stay\}\times\{stay\}}(m) &= -(m_1 - k_1) \cdot (k_2 - m_1) \\
g_{\{stay\}\times\{change\}}(m) &= k_2 - m_1,
\end{aligned}
$$

We remark that this is indeed a sensible characterization because $-(m_1 - k_1)(k_2 - m_1) \leq 0$ if and only if $m_1 \geq k_1$ and $m_1 \leq k_2$ or if $m_1 \leq k_1$ and $m_1 \geq k_2$ and the second case is not possible because $k_1 < k_2$.

As a next step we have to check whether the gradients are non-vanishing for all points where more than one strategy is optimal. In our case this has to be verified for the point

(k_1, m_2) for the strategies $\{change\} \times \{stay\}$ and $\{stay\} \times \{stay\}$ as well as for the point (k_2, m_2) for the strategies $\{stay\} \times \{stay\}$ and $\{stay\} \times \{change\}$:

$$\begin{aligned}
\nabla g_{\{change\}\times\{stay\}}(m) &= \begin{pmatrix} 1 \\ 0 \end{pmatrix} && \Rightarrow \nabla g_{\{change\}\times\{stay\}}(k_1, m_2) \neq 0 \\
\nabla g_{\{change\}\times\{stay\}}(m) &= \begin{pmatrix} k_2 + k_1 - 2m_1 \\ 0 \end{pmatrix} && \Rightarrow \nabla g_{\{stay\}\times\{stay\}}(k_1, m_2) \neq 0 \\
& && \quad \text{and } \nabla g_{\{stay\}\times\{stay\}}(k_2, m_2) \neq 0 \\
\nabla g_{\{stay\}\times\{change\}}(m) &= \begin{pmatrix} -1 \\ 0 \end{pmatrix} && \Rightarrow \nabla g_{\{stay\}\times\{change\}}(k_1, m_2) \neq 0,
\end{aligned}$$

where we again use for the second case that $k_1 < k_2$. This, in total, shows that Assumption A3 is satisfied. Assumption A4 is also satisfied since the three strategies $\{change\} \times \{stay\}$, $\{stay\} \times \; stay\}$ and $\{stay\} \times \{change\}$ are never simultaneously optimal.

In order to verify that Assumption A5 is satisfied it suffices by the structure of the gradients to verify that the first component of $(Q^d(m))^T m$ is non-vanishing. The first component of the term is vanishing for the strategy $\{change\} \times \{stay\}$ if m is a multiply of $\left(\frac{\epsilon}{b+\epsilon}, \frac{b}{b+\epsilon}\right)$. Similarly, it vanishes for the strategy $\{stay\} \times \{stay\}$ if m is a multiply of $\left(\frac{1}{2}, \frac{1}{2}\right)$ and for the strategy $\{stay\} \times \{change\}$ if m is a multiply of $\left(\frac{b}{b+\epsilon}, \frac{\epsilon}{b+\epsilon}\right)$. If k_1 does not coincide with $\frac{\epsilon}{b+\epsilon}$ and with $\frac{1}{2}$ and if k_2 does not coincide with $\frac{1}{2}$ and with $\frac{b}{b+\epsilon}$, then there is an open set $\mathcal{P}(\mathcal{S}) \subseteq O \subseteq O_\delta$ that does not contain a point $(k_1, m_2) = \lambda \left(\frac{\epsilon}{b+\epsilon}, \frac{b}{b+\epsilon}\right)^T$ or a point $(k_1, m_2) = \lambda \left(\frac{1}{2}, \frac{1}{2}\right)^T$ or a point $(k_2, m_2) = \lambda \left(\frac{1}{2}, \frac{1}{2}\right)^T$ or a point $(k_2, m_2) = \lambda \left(\frac{b}{b+\epsilon}, \frac{\epsilon}{b+\epsilon}\right)^T$ for some $\lambda \in \mathbb{R}$. Thus, also Assumption A5 is satisfied for a suitable chosen open set O whenever $k_1 \notin \left\{\frac{\epsilon}{b+\epsilon}, \frac{1}{2}\right\}$ and $k_2 \notin \left\{\frac{1}{2}, \frac{b}{b+\epsilon}\right\}$.

Condition (i) of Theorem 9.15 is trivially satisfied since we consider a standard continuous time Markov chain with irreducible generator. In order to verify condition (ii) we note that we can explicitly solve the differential equation $\dot{m}(t) = (Q^d(m))^T m$: Namely for the strategy $\{change\} \times \{stay\}$ we obtain that a solution for an arbitrary initial condition $m_0 = ((m_0)_1, (m_0)_2) \in O$ is

$$m(t) = ((m_0)_1 + (m_0)_2) \cdot \begin{pmatrix} \frac{\epsilon}{b+\epsilon} \\ \frac{b}{b+\epsilon} \end{pmatrix} + e^{-(b+\epsilon)t} \left(\frac{\epsilon}{b+\epsilon}(m_0)_2 - \frac{b}{b+\epsilon}(m_0)_1 \right) \begin{pmatrix} -1 \\ 1 \end{pmatrix}.$$

With this observation we obtain that given any initial condition the trajectory of $\dot{m} = (Q^d(m))^T m$ convergences towards $((m_0)_1 + (m_0)_2) \cdot (\frac{\epsilon}{b+\epsilon}, \frac{b}{b+\epsilon})^T$ and that moreover the trajectories are monotone in both components (increasing in one and decreasing in the other). For the other two strategies we obtain similar characterizations of the differential equations and the same qualitative monotonicity behaviour. Thus, condition (ii) is satisfied.

Condition (iii) holds only for certain parameter choices: Indeed, in the case (i) we obtain

$$\{change\} \times \{stay\} \leq \{stay\} \times \{stay\} \leq \{stay\} \times \{change\}$$

as the only ordering that satisfies the consistency condition, in case (iii) we obtain

$$\{change\} \times \{stay\} \leq \{stay\} \times \{change\} \leq \{stay\} \times \{stay\}$$

as the only ordering that satisfies the consistency condition and in case (viii) we obtain

$$\{stay\} \times \{change\} \leq \{stay\} \times \{stay\} \leq \{change\} \times \{stay\}$$

as the only ordering that satisfies the consistency condition. In all other cases it is not possible to find an ordering satisfying the consistency condition, in the cases (ii) an (vi) the algorithm presented in Section 9.6 cannot find a second strategy, in the cases (vi),(v) and (vii) the algorithm cannot even find a first strategy.

In the cases where the consistency condition is satisfied also condition (iv) is directly satisfied. Thus, Theorem 9.15 yields that the process converges towards the unique mean field equilibrium or stays inside some set $\mathrm{Opt}^{\mathrm{some}}(d) \setminus \mathrm{Opt}^{\mathrm{unique}}(d)$. However, these sets $\mathrm{Opt}^{\mathrm{some}}(d) \setminus \mathrm{Opt}^{\mathrm{unique}}(d)$ consist only of one point and in the three cases we are considering both strategies that are optimal for these one-point sets push the trajectory in the same direction away from that point. Thus, it is not possible to stay at such a point. Therefore, we indeed have global convergence towards the unique stationary equilibrium, which in case (i) is $\left(\frac{b}{b+\epsilon}, \frac{\epsilon}{b+\epsilon}\right)^T$ with equilibrium strategy $\{stay\} \times \{change\}$, in case (ii) it is $\left(\frac{1}{2}, \frac{1}{2}\right)^T$ with equilibrium strategy $\{stay\} \times \{stay\}$ and in case (viii) it is $\left(\frac{\epsilon}{b+\epsilon}, \frac{b}{b+\epsilon}\right)^T$ with equilibrium strategy $\{change\} \times \{stay\}$. In Figure 9.4 we illustrate how the trajectories converge towards the stationary equilibrium for some parameter set belonging to case (viii).

In the cases were the consistency condition is not satisfied we obtain that for at least one of the intersections of two optimality sets $\mathrm{Opt}^{\mathrm{some}}(d)$, which is either the point $(k_1, 1-k_1)$ or the point $(k_2, 1-k_2)$, the trajectories given the two optimal strategies point away from each other. In this case we obtain that the solution of the differential equation is no longer unique and that we therefore cannot tell for the point k_1 or k_2, respectively, what happens in the limit. For illustration, let us consider case (v). In this case $k_1 \in \left(\frac{\epsilon}{b+\epsilon}, \frac{1}{2}\right)$ and $k_2 \in \left(\frac{1}{2}, \frac{b}{b+\epsilon}\right)$.

Given the initial condition $m_0 = (k_1, 1-k_1)$ and any $T \geq 0$ the trajectory given by

$$t \mapsto \begin{cases} \begin{pmatrix} k_1 \\ 1-k_1 \end{pmatrix} & \text{if } t \leq T \\ \begin{pmatrix} \frac{\epsilon}{b+\epsilon} \\ \frac{b}{b+\epsilon} \end{pmatrix} + \left(\frac{\epsilon}{b+\epsilon} - k_1\right) e^{-(b+\epsilon)(t-T)} \begin{pmatrix} -1 \\ 1 \end{pmatrix} & \text{if } t \geq T \end{cases}$$

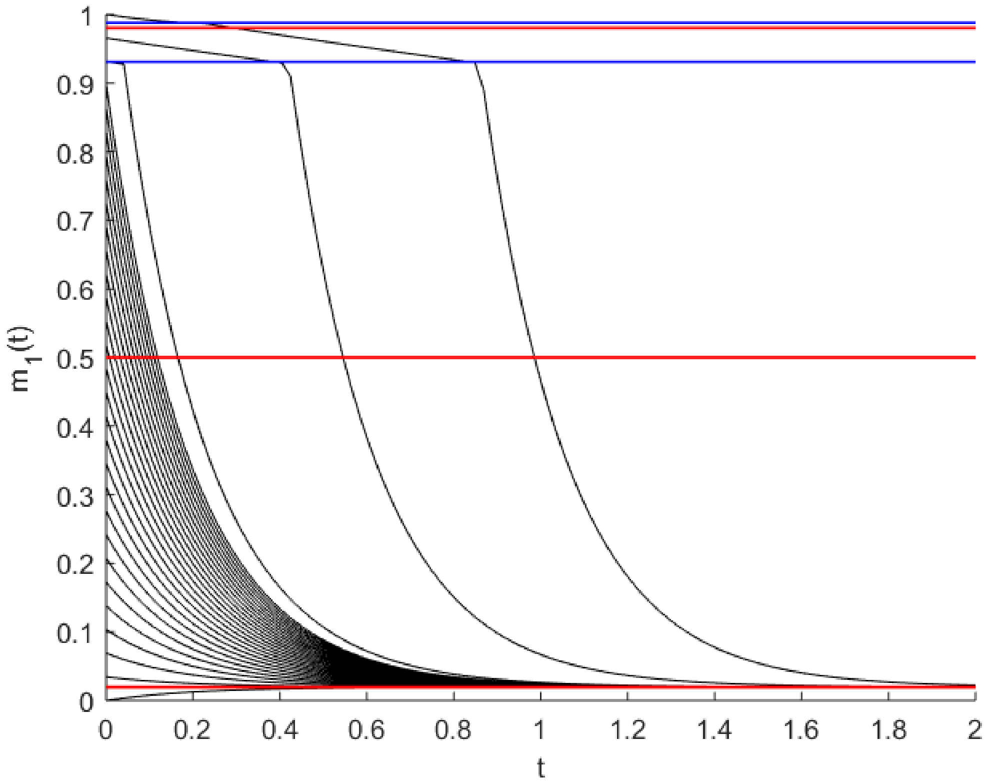

Figure 9.4.: The trajectories for different initial values $(m_0)_1$ of the myopic adjustment process for some parameter set belonging to case (viii), for which global convergence is obtained. Additionally, the red vertical lines depict the stationary points given the strategies $\{change\} \times \{stay\}$, $\{stay\} \times \{stay\}$ and $\{stay\} \times \{change\}$ (for bottom to top) and the blue vertical lines depict the crucial thresholds k_1 and k_2.

and the trajectory given by

$$t \mapsto \begin{cases} \begin{pmatrix} k_1 \\ 1-k_1 \end{pmatrix} & \text{if } t \leq T \\ \begin{pmatrix} \frac{1}{2} \\ \frac{1}{2} \end{pmatrix} + \left(\frac{1}{2} - k_1\right) e^{-2\epsilon(t-T)} \begin{pmatrix} -1 \\ 1 \end{pmatrix} & \text{if } t \geq T \end{cases}$$

are solutions of the myopic adjustment process

$$\dot{m}(t) \in \begin{cases} \begin{pmatrix} -bm_1(t) + \epsilon m_2(t) \\ bm_1(t) - \epsilon m_2(t) \end{pmatrix} & \text{if } m_1(t) < k_1 \\ \operatorname{conv}\left\{ \begin{pmatrix} -bm_1(t) + \epsilon m_2(t) \\ bm_1(t) - \epsilon m_2(t) \end{pmatrix}, \begin{pmatrix} -\epsilon m_1(t) + \epsilon m_2(t) \\ \epsilon m_1(t) - \epsilon m_2(t) \end{pmatrix} \right\} & \text{if } m_1(t) = k_1 \\ \begin{pmatrix} -\epsilon m_1(t) + \epsilon m_2(t) \\ \epsilon m_1(t) - \epsilon m_2(t) \end{pmatrix} & \text{if } k_1 < m_1(t) < k_2 \\ \operatorname{conv}\left\{ \begin{pmatrix} -\epsilon m_1(t) + \epsilon m_2(t) \\ \epsilon m_1(t) - \epsilon m_2(t) \end{pmatrix}, \begin{pmatrix} -\epsilon m_1(t) + bm_2(t) \\ \epsilon m_1(t) - bm_2(t) \end{pmatrix} \right\} & \text{if } m_1(t) = k_2 \\ \begin{pmatrix} -\epsilon m_1(t) + bm_2(t) \\ \epsilon m_1(t) - bm_2(t) \end{pmatrix} & \text{if } m_1(t) > k_2 \end{cases} .$$

Similarly, given the initial condition $(k_2, 1-k_2)$ and some $T \geq 0$ the trajectory given by

$$t \mapsto \begin{cases} \begin{pmatrix} k_2 \\ 1-k_2 \end{pmatrix} & \text{if } t \leq T \\ \begin{pmatrix} \frac{1}{2} \\ \frac{1}{2} \end{pmatrix} + \left(\frac{1}{2} - k_2\right) e^{-2\epsilon(t-T)} \begin{pmatrix} -1 \\ 1 \end{pmatrix} & \text{if } t \geq T \end{cases}$$

and the trajectory given by

$$t \mapsto \begin{cases} \begin{pmatrix} k_2 \\ 1-k_2 \end{pmatrix} & \text{if } t \leq T \\ \begin{pmatrix} \frac{b}{b+\epsilon} \\ \frac{\epsilon}{b+\epsilon} \end{pmatrix} + \left(\frac{b}{b+\epsilon} - k_2\right) e^{-(b+\epsilon)(t-T)} \begin{pmatrix} -1 \\ 1 \end{pmatrix} & \text{if } t \geq T \end{cases}$$

are solutions of the myopic adjustment process. Some of these infinitely many possible trajectories given these particular initial conditions are depicted in Figure 9.5.

These observations yield that for the initial conditions $m_0 = (k_1, 1-k_1)^T$ and $m_0 = (k_2, 1-k_2)^T$ we cannot clearly specify the limiting behaviour of the trajectory: For $m_0 = (k_1, 1-k_1)^T$ the trajectories can converge towards three equilibria, the pure strategy equilibria with equilibrium distributions $(\frac{\epsilon}{b+\epsilon}, \frac{b}{b+\epsilon})^T$ and $(\frac{1}{2}, \frac{1}{2})^T$, as well as the

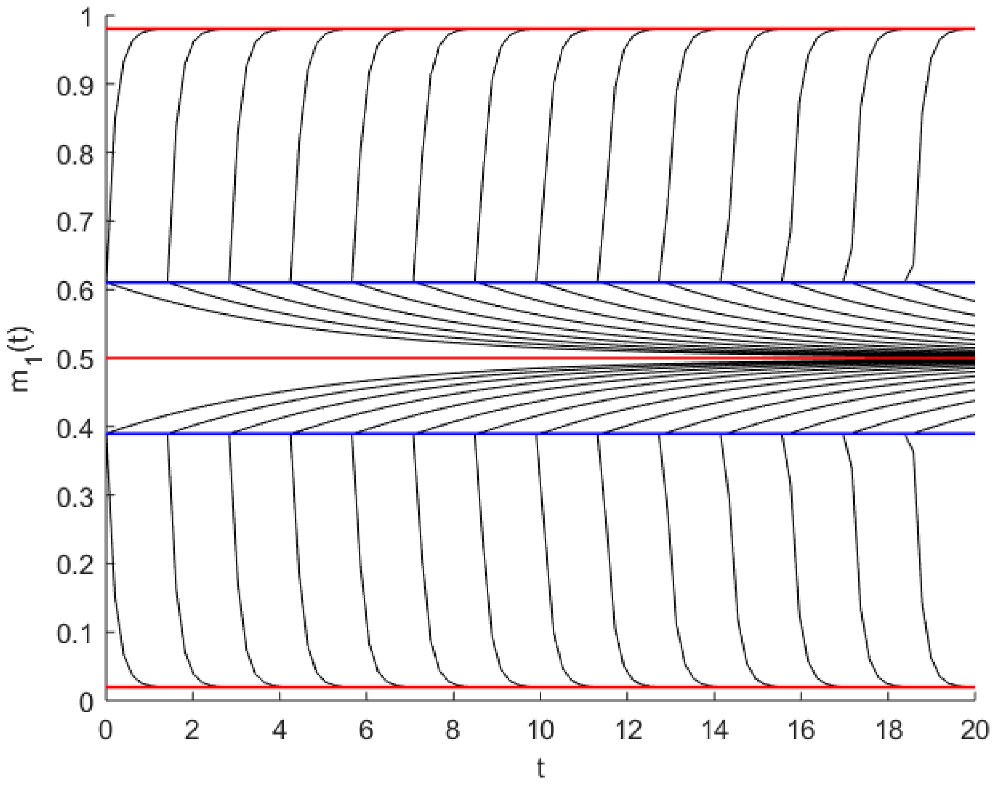

Figure 9.5.: Some of the infinitely many possible solutions with initial condition $(m_0)_1 = k_1$ and $(m_0)_1 = k_2$. Additionally, the red vertical lines depict the stationary points given the strategies $\{change\} \times \{stay\}$, $\{stay\} \times \{stay\}$ and $\{stay\} \times \{change\}$ (for bottom to top) and the blue vertical lines depict the crucial thresholds k_1 and k_2.

mixed strategy equilibrium with equilibrium distribution $(k_1, 1-k_1)^T$ and for $m_0 = (k_2, 1-k_2)$ the trajectories can either converge towards the pure strategy equilibria with equilibrium distributions $(\frac{1}{2}, \frac{1}{2})^T$ and $(\frac{b}{b+\epsilon}, \frac{\epsilon}{b+\epsilon})^T$ as well as the mixed strategy equilibrium with equilibrium distribution $(k_2, 1-k_2)^T$. For the economic application this means that we cannot make any predictions about long-term behaviour of partially rational agents in our model.

Remark 9.29. In the considered example the violated consistency condition yields non-uniqueness of trajectories, but trajectories still converge towards stationary mean field equilibria. This cannot be expected in general especially if there are more optimality sets of which some might additionally be nested in each other. Indeed, a violation of the consistency condition will most likely yield to trajectories for which we cannot make any predictions about the limit behaviour of the myopic adjustment process. △

9.7.2. The Consumer Choice Model of Section 6.3

In the second consumer choice example, we will again see that Theorem 9.15 can only be applied for some parameter sets. Furthermore, we will discuss the knife-edge case $c = 1$, where we only obtain that $\langle (Q^{sc}(m))^T m, \nabla g(m)\rangle$ is nonnegative and not that it is strictly positive. However, we can circumvent this problem and still obtain the same conclusions as Theorem 9.15 that is we obtain that after some time T the trajectory will stay in the set, where both strategies are simultaneously optimal. Using this fact, we then derive that we indeed obtain convergence towards the distribution of the stationary mixed strategy equilibria using the explicit characterization presented in Subsection 9.1.2.

We start by checking whether the assumptions of this chapter are satisfied: Assumption A2 is satisfied for any open subset of $\left(-\frac{\epsilon}{b}, \infty\right) \times \left(-\frac{\epsilon}{b}, \infty\right) \times \mathbb{R}$. Assumption A3 is satisfied since

$$\mathcal{U} = \{\{change\} \times \{stay\}, \{stay\} \times \{change\}\}$$

with $g(\cdot) = g_{\{stay\}\times\{change\}}(\cdot) = -g_{\{change\}\times\{stay\}}(\cdot)$ given by

$$g(m) = m_1 - \frac{1}{c}m_2 - \frac{1-c}{ce}(\beta + \epsilon + 2\lambda)$$

and the gradient is $\nabla g(m) = \left(1, -\frac{1}{c}, 0\right)^T \neq 0$. Assumption A4 is trivially satisfied since $\mathcal{U}$ contains only two elements.

It remains to check whether

$$\langle (Q^{cs}(m))^T m, \nabla - g(m)\rangle \quad \text{and } \langle (Q^{sc}(m))^T m, \nabla g(m)\rangle$$

have uniform sign over $\{m \in O : g(m) = 0\}$, where we again write cs for the strategy $\{change\} \times \{stay\}$ and sc for the strategy $\{stay\} \times \{change\}$: We note that

$$\begin{aligned}
&\langle (Q^{cs}(m))^T m, \nabla - g(m)\rangle \\
&= \left(-bm_1 - \epsilon m_1 - em_1^2 + \lambda m_3)\right) \cdot (-1) + \left(bm_1 - em_2^2 - \epsilon m_2 + \lambda m_3\right) \cdot \frac{1}{c} + \ldots \cdot 0 \\
&= \left(b + \epsilon + \frac{b}{c}\right) m_1 + em_1^2 + \lambda \frac{1-c}{c} m_3 - \frac{e}{c} m_2^2 - \frac{\epsilon}{c} m_2 \\
&= \left(b + \epsilon + \frac{b}{c}\right) \cdot \left(\frac{1}{c}m_2 + \frac{1-c}{ce}(\beta + \epsilon + 2\lambda)\right) + e\left(\frac{1}{c}m_2 + \frac{1-c}{ce}(\beta + \epsilon + 2\lambda)\right)^2 \\
&\quad + \lambda \frac{1-c}{c} m_3 - \frac{e}{c} m_2^2 - \frac{\epsilon}{c} m_2 \\
&= \left(\frac{b}{c} + \frac{\epsilon}{c} + \frac{b}{c^2}\right) m_2 + \frac{(1-c)}{ce}(\beta + \epsilon + 2\lambda)\left(b + \epsilon + \frac{b}{c}\right) \\
&\quad + \frac{e}{c^2} m_2^2 + 2\frac{1-c}{c^2}(\beta + \epsilon + 2\lambda) m_2 + \frac{(1-c)^2}{c^2 e}(\beta + \epsilon + 2\lambda)^2
\end{aligned}$$

$$
\begin{aligned}
&+ \lambda \frac{1-c}{c} m_3 - \frac{e}{c} m_2^2 - \frac{\epsilon}{c} m_2 \\
&= \frac{e(1-c)}{c^2} m_2^2 + \left(\frac{b}{c} + \frac{b}{c^2} + 2\frac{1-c}{c^2} (\beta + \epsilon + 2\lambda) \right) m_2 \\
&+ \frac{1-c}{ce} \left[(\beta + \epsilon + 2\lambda) \left(b + \epsilon + \frac{b}{c} \right) + \frac{1-c}{c} (\beta + \epsilon + 2\lambda)^2 \right] + \lambda \frac{1-c}{c^2} m_3.
\end{aligned}
$$

If $c < 1$ yields then $\langle (Q^{cs}(m))^T m, \nabla - g(m) \rangle$ given a fixed $m_3 \in O$ is a upwards opening parabola with a minimum at the point $\bar{m}_2$ given by

$$
\begin{aligned}
2\frac{e(1-c)}{c^2} \bar{m}_2 &= - \left(\frac{b}{c} \frac{1+c}{c} + 2\frac{1-c}{c^2} (\beta + \epsilon + 2\lambda) \right), \\
\text{i.e. } \bar{m}_2 &= - \frac{b}{2e} \frac{1+c}{1-c} - \frac{1}{e} (\beta + \epsilon + 2\lambda) < 0.
\end{aligned}
$$

Thus, since the value given $m_2 = 0$ is strictly positive, then also all values in $(-\delta, \infty)$, with $\delta > 0$ being small enough are positive since the parabola is increasing on the right of the maximum. Thus, in any case we obtain that

$$
\langle m^T Q^{cs}(m), \nabla - g(m) \rangle > 0.
$$

For the second scalar product the case is not that clear:

$$
\begin{aligned}
&\langle (Q^{sc}(m))^T m, \nabla g(m) \rangle \\
&= \left(-em_1^2 - \epsilon m_1 + bm_2 + \lambda m_3 \right) \cdot 1 + \left(-bm_2 - em_2^2 - \epsilon m_2 + \lambda m_3 \right) \cdot \left(-\frac{1}{c} \right) + \ldots \cdot 0 \\
&= -em_1^2 - \epsilon m_1 + \frac{e}{c} m_2^2 + \left(b + \frac{b}{c} + \frac{\epsilon}{c} \right) m_2 - \lambda \frac{1-c}{c} m_3 \\
&= -e \left(\frac{1}{c} m_2 + \frac{1-c}{ce} (\beta + \epsilon + 2\lambda) \right)^2 - \epsilon \left(\frac{1}{c} m_2 + \frac{1-c}{ce} (\beta + \epsilon + 2\lambda) \right) \\
&+ \frac{e}{c} m_2^2 + \left(b + \frac{b}{c} + \frac{\epsilon}{c} \right) m_2 - \lambda \frac{1-c}{c} m_3 \\
&= -\frac{e}{c^2} m_2^2 - 2\frac{1-c}{c^2} (\beta + \epsilon + 2\lambda) m_2 - \frac{(1-c)^2}{c^2 e} (\beta + \epsilon + 2\lambda)^2 - \frac{\epsilon}{c} m_2 \\
&- \epsilon \frac{1-c}{ce} (\beta + \epsilon + 2\lambda) + \frac{e}{c} m_2^2 + \left(b + \frac{b}{c} + \frac{\epsilon}{c} \right) m_2 - \lambda \frac{1-c}{c} m_3 \\
&= -\frac{e(1-c)}{c^2} m_2^2 + \left(b + \frac{b}{c} - 2\frac{1-c}{c^2} (\beta + \epsilon + 2\lambda) \right) m_2 - \frac{(1-c)^2}{c^2 e} (\beta + \epsilon + 2\lambda)^2 \\
&- \epsilon \frac{1-c}{ce} (\beta + \epsilon + 2\lambda) - \lambda \frac{1-c}{c} m_3.
\end{aligned}
$$

If m_2 is small enough this quantity will be negative and if $b + \frac{b}{c}$ is large enough it will become positive for some values $m_2 \in [0, 1]$. If this does not happen, that is if the quantity is non-zero for all $m_2 \in [0, 1]$, then it is possible to apply Theorem 9.15.

Condition (i) is met by Theorem 9.8 (see Subsection 9.2.3). Condition (ii), (iii) and (iv) hold if and only if

$$\langle (Q^{sc}(m))^T m, \nabla g(m) \rangle$$

has uniform sign over the set $\{m \in O : g(m) = 0\}$.

For the first parameter sets considered in Section 6.3 we obtain that

$$\langle (Q^{sc}(m))^T m, \nabla g(m) \rangle \approx -0.1250 m_2^2 + 0.4438 m_2 - 0.5232 - 0.1250 m_3 < 0$$

for all $m_2, m_3 \in \mathrm{Opt}^{\mathrm{some}}(\{\mathit{stay}\} \times \{\mathit{change}\}) \cap \mathrm{Opt}^{\mathrm{some}}(\{\mathit{change}\} \times \{\mathit{stay}\})$ with $\delta > 0$ being small enough. In this case, we thus obtain by Theorem 9.15 that we have convergence towards a deterministic mean field equilibrium or that we stay inside

$$\mathrm{Opt}^{\mathrm{some}}(\{\mathit{stay}\} \times \{\mathit{change}\}) \cap \mathrm{Opt}^{\mathrm{some}}(\{\mathit{change}\} \times \{\mathit{stay}\}).$$

However, by the observations made in Subsection 9.1.2 we obtain that any trajectory of the myopic adjustment process will immediately leave this intersection. Thus, we indeed obtain convergence towards the unique deterministic mean field equilibrium $(m^{cs}, \{\mathit{change}\} \times \{\mathit{stay}\})$.

In the other cases we obtain that the sign of $\langle (Q^{sc}(m))^T m, \nabla g(m) \rangle$ is not uniform on $[0, \infty)^S \cap \mathrm{Opt}^{\mathrm{some}}(\{\mathit{stay}\} \times \{\mathit{change}\}) \cap \mathrm{Opt}^{\mathrm{some}}(\{\mathit{change}\} \times \{\mathit{stay}\})$. Thus, we cannot apply Theorem 9.15. However, it would be interesting to understand how the system behaves in this case.

Let us conclude with the case $c = 1$: In this case we obtain from the previous calculations that

$$\langle (Q^{cs}(m))^T m, \nabla - g(m) \rangle = 2bm_2 \geq 0 \quad \text{and} \quad \langle (Q^{sc}(m))^T m, \nabla g(m) \rangle = 2bm_2 \geq 0.$$

However, even if we aim to use the modification of Section 9.5 we run into a problem, namely

$$\langle (Q^{cs}(m))^T, \nabla g(m) \rangle = 0 \quad \text{and} \quad \langle (Q^{sc}(m))^T, \nabla g(m) \rangle = 0$$

for all $m \in [0, \infty)^S$ with $m_2 = 0$. In general, this is problematic since to digraph, which formalized the consistent behaviour of the trajectories, has to be redefined such that it still captures the relevant relations and for some preliminary results (Lemma 9.21, Lemma 9.23 and Lemma 9.25) it is even necessary that these scalar products are non-vanishing. However, in this particular setting we can workaround the aforementioned problems: We note that we do not need Lemma 9.25. Moreover, the formulation of the digraph is not problematic since alongside the intersection of the two optimality sets $\mathrm{Opt}^{\mathrm{some}}(cs) \cap \mathrm{Opt}^{\mathrm{some}}(sc)$ the scalar product is non-zero for some values. Thus, it suffices to prove the flow invariance of $\mathcal{P}(\mathcal{S})$ as well as the uniqueness of trajectories for $\dot{m} \in F(m)$.

As in the standard case, we prove the flow invariance and remark that the uniqueness directly follows from Lemma 9.4. We note that, by Remark 9.22, the set $\{m \in \mathbb{R}^S : \sum_{i\in\mathcal{S}} m_i = 1\}$ is flow invariant for F. Assume that there are time $0 \leq \tilde{t}_1 < \tilde{t}_2$ such that $m(\tilde{t}_1) \in \mathcal{P}(\mathcal{S})$ and $m(t) \notin \mathcal{P}(\mathcal{S})$ for all $\tilde{t}_1 < t < \tilde{t}_2$. If $m(\tilde{t}_1) \neq (0,0,1)$, then we can employ the proof technique used in the classical proof of Lemma 9.21. If $m(\tilde{t}_1) = (0,0,1)$, then

$$(Q^{cs}(m))^T m = (Q^{sc}(m))^T m = \begin{pmatrix} 1 \\ 1 \\ -2 \end{pmatrix},$$

which by continuity of $Q^{cs}(m)$ and $Q^{sc}(m)$ implies for $\tilde{t}_1 < t < \tilde{t}_1 + \delta$ that

$$F(m) \subseteq N_{\frac{1}{2}}\left(\begin{pmatrix} 1 \\ 1 \\ -2 \end{pmatrix}\right).$$

However, this implies that $m_1(t)$ and $m_2(t)$ are increasing almost surely, which in particular yields that $m(t) \in \mathcal{P}(\mathcal{S})$ for $\tilde{t}_1 < t < \tilde{t}_1 + \delta$, a contradiction.

Thus, we obtain the same conclusion as with Theorem 9.15. In particular, we obtain that also in this case the trajectory will stay in

$$\mathrm{Opt}^{\mathrm{some}}(\{\mathit{stay}\} \times \{\mathit{change}\}) \cap \mathrm{Opt}^{\mathrm{some}}(\{\mathit{change}\} \times \{\mathit{stay}\})$$

or that it will converge towards a deterministic stationary mean field equilibrium. We first note that the convergence towards a stationary mean field equilibrium is not possible since for both strategies $d \in \{cs, sc\}$ the nonlinear Markov chain with nonlinear generator $Q^d(\cdot)$ is strongly ergodic and the stationary point given $Q^d(\cdot)$ lies outside $\mathrm{Opt}^{\mathrm{some}}(d)$. Thus, we will indeed stay for all $t \geq T$ in the set of all points where the two strategies are simultaneously optimal.

This insight can be used even further to characterize the long-term behaviour of the myopic adjustment process: By definition of the myopic adjustment process the unique solution $m(\cdot)$ has to satisfy $\dot{m}(t) \in F(m(t))$, which in our case means that for almost all $t \geq T$ and some $\pi_{1,change}(t), \pi_{2,change}(t) \in [0,1]$ we have

$$\begin{aligned}\dot{m}_1(t) &= -(\pi_{1,change}(t)b + em_1(t) + \epsilon)m_1(t) + \pi_{2,change}(t)bm_2(t) + \lambda m_3(t) \\ \dot{m}_2(t) &= \pi_{1,change}bm_1(t) - (\pi_{2,change}(t)b + em_2(t) + \epsilon)m_2(t) + \lambda m_3(t).\end{aligned}$$

We remark that we omit the equation for $\dot{m}_3(t)$ because the fact that $Q^d(\cdot)$ ($d \in \{cs, sc\}$) is conservative implies $\dot{m}_1(t) + \dot{m}_2(t) + \dot{m}_3(t) = 0$.

Since $m_1(t) = m_2(t)$ for all $t \geq T$ by definition of $\mathrm{Opt}^{\mathrm{some}}(\{\mathit{stay}\} \times \{\mathit{change}\}) \cap \mathrm{Opt}^{\mathrm{some}}(\{\mathit{change}\} \times \{\mathit{stay}\})$, we obtain, using that also $\dot{m}_1(t) = \dot{m}_2(t)$ has to hold for almost all $t \geq T$, that

$$-\pi_{1,change}(t)bm_1(t) - em_1(t)^2 - \epsilon m_1(t) + \pi_{2,change}(t)bm_1(t) + \lambda m_3(t)$$

$$= \pi_{1,change}(t) b m_1(t) - \pi_{2,change}(t) b m_1(t) - e m_1(t)^2 - \epsilon m_1(t) + \lambda m_3(t),$$
$$\text{i.e.} \quad - \pi_{1,change}(t) b m_1(t) + \pi_{2,change}(t) b m_1(t) = 0.$$

This yields that for almost all $t \geq T$ the trajectory of the myopic adjustment process has to satisfy

$$\dot{m}_1(t) = -e m_1(t)^2 - (\epsilon + 2\lambda) \, m_1(t) + \lambda,$$

which is a Riccati equation. Since the set $[0, 1]$ is flow invariant for the differential equation, we obtain that the right-hand side is Lipschitz continuous. Thus, by Theorem A.3 there is a unique Caratheodory solution. Since there is also a unique classical solution for any initial condition $m_0 \in [0, 1]$, this solution coincides with the Caratheodory solution. Thus, we can rely on the results regarding classical solutions of the Riccati equation stated above. However, the only known results regarding the long-term behaviour of solutions are that the solution is bounded in our setting (see Bahk et al. (2008)) and there are no general results on conditions that ensure that the solution approaches a certain limit. Nonetheless, numerical simulations indicate that in our setting with initial conditions $m_0 \in (0, 1)$ convergence towards the stationary point is likely.

10. Outlook

This chapter collects some ideas for future research for the discussed model. Additionally, several model extensions are possible directions for future research. Examples for such model extensions that have previously been discussed in examples are lump-sum state transition costs (see the corruption model of Kolokoltsov and Malafeyev (2017)), the introduction of a major player (see the inspection game of Kolokoltsov and Yang (2015)) or multi-class agent models (see the the labour market model of Guéant (2009a) and the paradigm shift model of Besancenot and Dogguy (2015)).

Numerical Methods to Obtain Stationary Mean Field Equilibria

A first central field for further research are numerical methods for finding stationary mean field equilibria. The related problem, the computation of Nash equilibria in normal form games as well as in symmetric games, is well studied: It has been shown that both problems lie in the class PPAD, which consists of all problems that can be represented by a digraph with a specific structure (which size might grow exponentially fast in the input data) and where the solutions are given by the sinks and non-standard sources of the graph (Papadimitriou, 1994, 2007). However, the problems are also PPAD-complete, which, as NP-completeness for decision problems, hints that the problem is intractable (Papadimitriou, 2007; Brandt et al., 2009). We remark, that as in the case of NP-complete problems, it might hold true that P=PPAD or even P=NP, which would in the end imply tractability.

For our problem at hand we first note that finding stationary mean field equilibria differs substantially from finding Nash equilibria: The problem is not mainly combinatorial, but of analytical nature. Thus, we expect that we can only find algorithms to obtain approximate stationary mean field equilibria. These are pairs (m, π) such that the distribution m lies close to the distribution of a stationary mean field equilibrium and the chosen strategy π of the individual agent is ϵ-optimal.

Many problems related to finding stationary mean field equilibria, like finding Nash equilibria, finding fixed points or finding competitive equilibria, also lie in PPAD (Papadimitriou, 1994, 2007). Thus, a first natural approach would be to check whether the problem itself lies in PPAD. Moreover, we expect that the problem is also PPAD-complete, which means that as a second step we would need to reduce one of the other problems onto finding stationary mean field equilibria. This is a promising approach since for arbitrary dynamics solving fixed point problems is an integral part of finding a stationary mean field equilibrium.

For constant irreducible dynamics another approach might yield a sensible algorithm, whose complexity additionally depends on certain bounds for transition rates and rewards: More precisely, the algorithm would consist of three steps. First, we compute the stationary points given each deterministic strategy, which is by Lemma 5.9 possible in polynomial time. Thereafter, we have to check for any set $\{d^1, \dots, d^k\} \subseteq D^s$ for which of the points in $\text{conv}(x^{d^1}, \dots, x^{d^k})$ all strategies $\{d^1, \dots, d^k\}$ are simultaneously optimal. Finally, for all points m obtained in the previous step we then have to solve the corresponding linear equation $(Q^\pi(m))^T m = 0$ for π, which is again possible in polynomial time. Thus, it suffices to find an sensible algorithm for the second step. A possible approach would be to define a suitable grid and use the bounds of the value function similar obtained in Section 7.3 to obtain nearly optimal strategies.

A Related Static Game

Another field of future research is the question which properties are generic, where both the topological as well as the measure-theoretic notion of genericity are of interest. Besides further investigations or generic essentiality (see Section 7.5) also the question whether the number of equilibria is generically finite is of interest. This question is most of the time answered positively in standard game theory (see Harsanyi (1973) for the case of normal form games, Haller and Lagunoff (2000) and Doraszelski and Escobar (2010) for the case of Markov perfect equilibria and Dubey (1986) as well as Anderson and Zame (2001) for the case of static games with compact, convex polyhedra as action spaces). For stationary mean field equilibria we also expect that having finitely many equilibria is a generic property.

Besides the programme sketched in Section 7.5, the following auxiliary three player static game with compact and convex action spaces leads to a promising approach for proving generic finiteness:

Player one chooses an action $m \in \mathcal{P}(\mathcal{S})$ in order to maximize his utility $-(m - x)^2$, player two chooses $\pi \in \mathcal{P}(D^s)$ such that the strategy is optimal for the Markov decision process with transition rates $Q(m)$ and rewards $r(m)$, that is to maximize

$$\sum_{d \in D^s} \pi^d (I - \beta Q^d(m))^{-1} r^d(m) \cdot 1.$$

Player 3 chooses $x \in \mathcal{P}(\mathcal{S})$ such that

$$-\sum_{j \in \mathcal{S}} \left(\sum_{i \in \mathcal{S}} x_i \sum_{d \in D^s} Q_{ij}^d(m) \pi^d \right)^2$$

is maximized. It holds that (m, π, x) is an equilibrium of this static game if and only if (m, π) is a stationary mean field equilibrium. Thus, we could follow the general programme presented in Dubey (1986) as well as Anderson and Zame (2001) to prove generic finiteness. We only have to find a different way to show that the Jacobian matrices are "generically" non-vanishing since Thom's transversality theorem is no longer directly applicable as the Jacobian matrices have a certain structure.

Global Convergence of the Myopic Adjustment Process under Weaker Assumptions

As we have seen in Section 9.5 as well as in the discussion of the examples in Section 9.7 some of the assumptions can be weakened in such a way that the theorem still holds true. In particular, we have seen that Assumption A5 can be weakened and a more general statement regarding models where $\langle (Q^d(m))^T m, \nabla g_{\hat{d}}(m) \rangle = 0$ for some points would be desirable. However, two central difficulties occur: First, we need to sensibly readjust the definition of the underlying digraphs in order to still allow them to capture the relevant information regarding the consistency of the different trajectories given the deterministic strategies. Second, we need to find new proofs for the uniqueness of trajectories and for the flow invariance of $\mathcal{P}(\mathcal{S})$ for $F(\cdot)$ since these proofs at the moment crucially rely on the fact that $\langle (Q^d(m))^T m, \nabla g_{\hat{d}}(m) \rangle$ is non-vanishing.

Furthermore, it would be desirable to discuss what happens if Assumption A4 is violated. In the current proof of the theorem, this assumption has a crucial role since if it is violated we cannot derive in the proof of Lemma 9.18 that the conditions are met to use the sufficient criterion for flow invariance presented in Lemma 9.2, which then yields that the whole proof breaks down. A first situation to investigate would be the case where two planes divide the space of population distributions into four optimality sets. This is a natural starting point for such investigations since this situation is simple, linear conditions regarding optimality naturally arose in all application and, furthermore, all other conditions except A4 are met because the distance to a convex set is continuously differentiable. In this setting, it is still possible to derive flow invariance statements for the sets where one strategy is optimal or for the unions of optimality sets that lie on one side of a plane. Due to the consistency condition we obtain for the myopic adjustment process itself even flow invariance statements for other, more general, optimality sets. The main task would be to systematically understand whether the conditions are sufficient by going through all the possible digraph configurations.

Is there a Trend towards Stationary Equilibria?

Another classical way of justifying the consideration of stationary equilibria is to prove that dynamic equilibria converge towards stationary equilibria. Similar investigations have been carried out for finite state mean field games having a unique optimizer of the Hamiltonian under strong structural assumptions (see Gomes et al. (2013) and Ferreira and Gomes (2014)). The recent preprint of Belak et al. (2019) now yields a characterization of dynamic mean field equilibria a finite time horizon version of the model discussed in this thesis. In this context it would be interesting to understand in how far stationary equilibria introduced in this thesis and dynamic equilibria in their model relate when the time horizon tends to infinity. However, for any non-trivial game with finite action spaces there will be infinitely many systems potentially describing dynamic equilibria and it is a priori not clear how they relate to each other. Thus, a first step would be to carry out such investigations for simple examples.

Limit Behaviour of Nonlinear Markov Chains

Up to the best knowledge of the author, the limit behaviour of continuous time nonlinear Markov chains has not been considered so far. In Section 9.2 we have seen that irreducibility does no longer imply strong ergodicity of the Markov chain, but it might happen that we only face convergence in the limit to some stationary distribution. In this context, we raised the question whether irreducibility implies convergence in the limit to some stationary distribution. A formal proof (or counterexample) for this statement, but also the investigation of the converse statement are of theoretical interest. However, as we have discussed in Section 9.2 classical tools from the theory of Markov chains cannot be directly applied, so one would need to find new methods. Moreover, further ergodicity conditions (in particular for larger state spaces) are desirable.

Bibliography

Sachin Adlakha and Ramesh Johari. Mean Field Equilibrium in Dynamic Games with Strategic Complementarities. *Operations Research*, 61(4):971–989, 2013. doi: 10.1287/opre.2013.1192.

Sachin Adlakha, Ramesh Johari, and Gabriel Y. Weintraub. Equilibria of dynamic games with many players: Existence, approximation, and market structure. *Journal of Economic Theory*, 156:269–316, 2015. doi: 10.1016/j.jet.2013.07.002.

Robert M. Anderson and William R. Zame. Genericity with Infinitely Many Parameters. *The B.E. Journal of Theoretical Economics*, 1(1), 2001. doi: 10.2202/1534-5963.1003.

William J. Anderson. *Continuous-Time Markov Chains: An Applications-Oriented Approach*. Springer Series in Statistics. Springer-Verlag, New York, Berlin, Heidelberg, 1991. ISBN 0-387-97369-9.

Søren Asmussen. *Applied Probability and Queues*, volume 51 of *Stochastic Modelling and Applied Probability*. Springer-Verlag, New York, 2$^{\text{nd}}$ edition, 2003. ISBN 0-387-00211-1.

Jean-Pierre Aubin and Arrigo Cellina. *Differential Inclusions: Set-Valued Maps and Viability Theory*, volume 264 of *Grundlehren der mathematischen Wissenschaften*. Springer-Verlag, Berlin, Heidelberg, 1984. doi: 10.1007/978-3-642-69512-4.

Robert J. Aumann. Markets with a Continuum of Traders. *Econometrica*, 32(1):39–50, 1964. doi: 10.2307/1913732.

Sarah Y. Bahk, Nadejda E. Dyakevich, and Stefan C. Johnson. Long term behavior of solutions for Riccati initial-value problems. *Electronic Journal of Differential Equations*, 2008(142):1–8, 2008. URL `https://ejde.math.txstate.edu/Volumes/2008/142/abstr.html`.

Jørgen Bang-Jensen and Gregory Z. Gutin. *Digraphs: Theory, Algorithms and Applications*. Springer Monographs in Mathematics. Springer-Verlag, London, 2$^{\text{nd}}$ edition, 2010. ISBN 978-0-85729-041-0.

Rani Basna, Astrid Hilbert, and Vassili N. Kolokoltsov. An epsilon-Nash equilibrium for non-linear Markov games of mean-field-type on finite spaces. *Communications on Stochastic Analysis*, 8(4):449–468, 2014. doi: 10.31390/cosa.8.4.02.

Erhan Bayraktar and Asaf Cohen. Analysis of a Finite State Many Player Game Using Its Master Equation. *SIAM Journal on Control and Optimization*, 56(5):3538–3568, 2018. doi: 10.1137/17M113887X.

Christoph Belak, Daniel Hoffmann, and Frank Thomas Seifried. Continuous-Time Mean Field Games with Finite State Space and Common Noise. Preprint, available at SSRN `https://ssrn.com/abstract=3458336`, 2019.

Alain Bensoussan, Jens Frehse, and Phillip Yam. *Mean Field Games and Mean Field Type Control Theory*. SpringerBriefs in Mathematics. Springer, New York, Heidelberg, Dordrecht, London, 2013. ISBN 978-1-4614-8507-0.

J. Bergin and D. Bernhardt. Anonymous sequential games: Existence and characterization of equilibria. *Economic Theory*, 5(3):461–489, 1995. doi: 10.1007/BF01212329.

James Bergin and Dan Bernhardt. Anonymous sequential games with aggregate uncertainty. *Journal of Mathematical Economics*, 21(6):543–562, 1992. doi: 10.1016/0304-4068(92)90026-4.

Abraham Berman and Robert J. Plemmons. *Nonnegative matrices in the mathematical sciences*. Computer science and applied mathematics. Academic Press, Inc., New York, 1979. ISBN 0-12-092250-9.

Damien Besancenot and Habib Dogguy. Paradigm Shift: A Mean Field Game Approach. *Bulletin of Economic Research*, 67(3):289–302, 2015. doi: 10.1111/boer.12024.

Franco Blanchini and Stefano Miani. *Set-Theoretic Methods in Control*. Systems & Control: Foundations & Applications. Birkhäuser, Basel, 2nd edition, 2015. doi: 10.1007/978-3-319-17933-9.

Felix Brandt, Felix Fischer, and Markus Holzer. Symmetries and the complexity of pure Nash equilibrium. *Journal of Computer and System Sciences*, 75(3):163–177, 2009. doi: 10.1016/j.jcss.2008.09.001.

Ariela Briani and Pierre Cardaliaguet. Stable solutions in potential mean field game systems. *Nonlinear Differential Equations and Applications NoDEA*, 25(1), 2018. doi: 10.1007/s00030-017-0493-3.

Arne Brøndsted. *An Introduction to Convex Polytopes*, volume 90 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1983. ISBN 0-387-90722-X.

O. A. Butkovsky. On the convergence of nonlinear Markov chains. *Doklady Mathematics*, 86(3):824–826, 2012. doi: 10.1134/S1064562412060270.

O. A. Butkovsky. On Ergodic Properties of Nonlinear Markov Chains and Stochastic McKean-Vlasov Equations. *Theory of Probability & Its Applications*, 58(4):661–674, 2014. doi: 10.1137/S0040585X97986825.

O. A. Butkovsky and A. Yu. Veretennikov. On asymptotics for Vaserstein coupling of Markov chains. *Stochastic Processes and their Applications*, 123(9):3518–3541, 2013. doi: 10.1016/j.spa.2013.04.016.

Peter E. Caines, Minyi Huang, and Roland P. Malhamé. Mean Field Games. In Tamer Basar and Georges Zaccour, editors, *Handbook of Dynamic Game Theory*. Springer, Cham, 2017. doi: 10.1007/978-3-319-27335-8_7-1.

Pierre Cardaliaguet. Notes on Mean Field Games (from P.-L. Lions' Lectures at Collège de France). 2013. URL `https://www.ceremade.dauphine.fr/~cardaliaguet/MFG20130420.pdf`.

Pierre Cardaliaguet and Saeed Hadikhanloo. Learning in mean field games: The fictitious play. *ESAIM: Control, Optimisation and Calculus of Variations*, 23(2):569–591, 2017. doi: 10.1051/cocv/2016004.

René Carmona and François Delarue. *Probabilistic Theory of Mean Field Games with Applications I: Mean Field FBSDEs, Control, and Games*, volume 83 of *Probability Theory and Stochastic Modelling*. Springer International Publishing, 2018a. doi: 10.1007/978-3-319-58920-6.

René Carmona and François Delarue. *Probabilistic Theory of Mean Field Games with Applications II: Mean Field Games with Common Noise and Master Equations*, volume 84 of *Probability Theory and Stochastic Modelling*. Springer International Publishing, 2018b. doi: 10.1007/978-3-319-56436-4.

René Carmona and Peiqi Wang. A Probabilistic Approach to Extended Finite State Mean Field Games. Preprint, available on ArXiv `https://arxiv.org/abs/1808.07635`, 2018.

René Carmona, François Delarue, and Daniel Lacker. Mean Field Games of Timing and Models for Bank Runs. *Applied Mathematics & Optimization*, 76(1):217–260, 2017. doi: 10.1007/s00245-017-9435-z.

Alekos Cecchin and Markus Fischer. Probabilistic Approach to Finite State Mean Field Games. *Applied Mathematics & Optimization*, 2018. doi: 10.1007/s00245-018-9488-7.

Alekos Cecchin and Guglielmo Pelino. Convergence, fluctuations and large deviations for finite state mean field games via the Master Equation. *Stochastic Processes and their Applications*, 129(11):4510–4555, 2019. doi: 10.1016/j.spa.2018.12.002.

Alekos Cecchin, Paolo Dai Pra, Markus Fischer, and Guglielmo Pelino. On the Convergence Problem in Mean Field Games: A Two State Model without Uniqueness. *SIAM Journal on Control and Optimization*, 57(4):2443–2466, 2019. doi: 10.1137/18M1222454.

Patrick Chan and Ronnie Sircar. Bertrand and Cournot Mean Field Games. *Applied Mathematics & Optimization*, 71(3):533–569, 2015. doi: 10.1007/s00245-014-9269-x.

Larry A. Chenault. On the uniqueness of Nash equilibria. *Economics Letters*, 20(3): 203–205, 1986. doi: 10.1016/0165-1765(86)90023-6.

Michael G. Crandall. A generalization of Peano's existence theorem and flow invariance. *Proceedings of the American Mathematical Society*, 36(1):151–155, 1972. doi: 10.2307/2039051.

Klaus Deimling. *Nonlinear Functional Analysis*. Springer-Verlag, Berlin, 1985. ISBN 3-540-13928-1.

Klaus Deimling. *Multivalued Differential Equations*, volume 1 of *de Gruyter Series in Nonlinear Analysis and Applications*. Walter de Gruyter & Co., Berlin, 1992. ISBN 3-11-013212-5.

I. Del Prete, M. Di Iorio, and S. Naimpally. Essential Fixed Points of Functions and Multifunctions. *Real Analysis Exchange*, 25(1):369–382, 1999. URL `https://projecteuclid.org/euclid.rae/1231187611`.

Josu Doncel, Nicolas Gast, and Bruno Gaujal. Mean-Field Games with Explicit Interactions. preprint, 2016a. URL `https://hal.inria.fr/hal-01277098`.

Josu Doncel, Nicolas Gast, and Bruno Gaujal. Are Mean-field Games the Limits of Finite Stochastic Games? *ACM SIGMETRICS Performance Evaluation Review*, 44 (2):18–20, 2016b. doi: 10.1145/3003977.3003984.

Josu Doncel, Nicolas Gast, and Bruno Gaujal. Discrete mean field games: Existence of equilibria and convergence. *Journal of Dynamics and Games*, 6(3):221–239, 2019. doi: 10.3934/jdg.2019016.

Ulrich Doraszelski and Juan F. Escobar. A theory of regular Markov perfect equilibria in dynamic stochastic games: Genericity, stability, and purification. *Theoretical Economics*, 5(3):369–402, 2010. doi: 10.3982/TE632.

Pradeep Dubey. Inefficiency of Nash Equilibria. *Mathematics of Operations Research*, 11(1):1–8, 1986. doi: 10.1287/moor.11.1.1.

Rick Durrett. *Essentials of Stochastic Processes*. Springer Texts in Statistics. Springer-Verlag, New York, 1999. ISBN 0-387-98836-X.

Jürgen Elstrodt. *Maß- und Integrationstheorie*. Grundwissen Mathematik. Springer-Verlag, Berlin, Heidelberg, 7$^{\text{th}}$ edition, 2011. ISBN 978-3-642-17904-4.

M. L. Fernandes and F. Zanolin. Remarks on Strongly Flow-Invariant Sets. *Journal of Mathematical Analysis and Applications*, 128(1):176–188, 1987. doi: 10.1016/0022-247X(87)90223-X.

Rita Ferreira and Diogo A. Gomes. On the convergence of finite state mean-field games through Γ-convergence. *Journal of Mathematical Analysis and Applications*, 418(1): 211–230, 2014. doi: 10.1016/j.jmaa.2014.02.044.

A. F. Filippov. *Differential Equations with Discontinuous Righthand Sides*, volume 18 of *Mathematics and its Applications*. Springer Netherlands, Dordrecht, 1988. doi: 10.1007/978-94-015-7793-9.

M. K. Fort, Jr. A unified theory of semi-continuity. *Duke Mathematical Journal*, 16(2): 237–246, 1949. doi: 10.1215/S0012-7094-49-01623-3.

M. K. Fort, Jr. Essential and Non Essential Fixed Points. *American Journal of Mathematics*, 72(2):315–322, 1950. doi: 10.2307/2372035.

Drew Fudenberg and David K. Levine. *The Theory of Learning in Games*, volume 2 of *MIT Press Series on Economic Learning and Social Evolution*. The MIT Press, Cambridge, Massachusetts, London, England, 1998. ISBN 978-0-262-06194-0.

Drew Fudenberg and David K. Levine. Whither Game Theory? Towards a Theory of Learning in Games. *Journal of Economic Perspectives*, 30(4):151–170, 2016. doi: 10.1257/jep.30.4.151.

Drew Fudenberg and Jean Tirole. *Game Theory*. Massachusetts Institute of Technology, Cambridge (Massachusetts), London (England), 1991. ISBN 0-262-06141-4.

Diogo Gomes, Robert M. Velho, and Marie-Therese Wolfram. Dual two-state mean-field games. In *53rd IEEE Conference on Decision and Control*, pages 2703–2708, 2014a. doi: 10.1109/CDC.2014.7039803.

Diogo Gomes, Roberto M. Velho, and Marie-Therese Wolfram. Socio-economic applications of finite state mean field games. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 372(2028), 2014b. doi: 10.1098/rsta.2013.0405.

Diogo A. Gomes and João Saúde. Numerical Methods for Finite-State Mean-Field Games Satisfying a Monotonicity Condition. *Applied Mathematics & Optimization*, 2018. doi: 10.1007/s00245-018-9510-0.

Diogo A. Gomes, Joana Mohr, and Rafael Rigão Souza. Discrete time, finite state space mean field games. *Journal de Mathématiques Pures et Appliquées*, 93(3):308–328, 2010. doi: 10.1016/j.matpur.2009.10.010.

Diogo A. Gomes, Joana Mohr, and Rafael Rigão Souza. Continuous Time Finite State Mean Field Games. *Applied Mathematics & Optimization*, 68(1):99–143, 2013. doi: 10.1007/s00245-013-9202-8.

Diogo A. Gomes, Levon Nurbekyan, and Edgard A. Pimentel. *Economic Models and Mean-field Games Theory*. 2015. ISBN 978-85-244-0404-7. URL `https://impa.br/wp-content/uploads/2017/04/30CBM_04.pdf`.

Andrzej Granas and James Dugundji. *Fixed Point Theory*. Springer Monographs in Mathematics. Springer-Verlag, New York, 2003. ISBN 978-0-387-00173-9.

David Griffeath. Uniform Coupling of Non-Homogeneous Markov Chains. *Journal of Applied Probability*, 12(4):753–762, 1975. doi: 10.2307/3212726.

Olivier Guéant. *Mean field games and applications to economics: Secondary topic: Discount rates and sustainable development*. PhD thesis, Université Paris-Dauphine, 2009a. URL `www.oliviergueant.com/uploads/4/3/0/9/4309511/these2.pdf`.

Olivier Guéant. A reference case for mean field games models. *Journal de Mathématiques Pures et Appliquées*, 92(3):276 – 294, 2009b. doi: 10.1016/j.matpur.2009.04.008.

Olivier Guéant. From infinity to one: The reduction of some mean field games to a global control problem. Preprint, available on ArXiv `https://arxiv.org/abs/1110.3441`, 2011.

Olivier Guéant. Existence and Uniqueness Result for Mean Field Games with Congestion Effect on Graphs. *Applied Mathematics & Optimization*, 72(2):291–303, 2015. doi: 10.1007/s00245-014-9280-2.

Olivier Guéant, Jean-Michel Lasry, and Pierre-Louis Lions. Mean Field Games and Applications. In *Paris-Princeton Lectures on Mathematical Finance 2010*, volume 2003 of *Lecture Notes in Mathematics*, pages 205–266. Springer-Verlag, Berlin, Heidelberg, 2011. doi: 10.1007/978-3-642-14660-2.

Xianping Guo and Onésimo Hernández-Lerma. *Continuous-Time Markov Decision Processes: Theory and Applications*, volume 62 of *Stochastic Modelling and Applied Probability*. Springer-Verlag, Berlin, Heidelberg, 2009. ISBN 978-3-642-26072-8.

Saeed Hadikhanloo. Learning in anonymous nonatomic games with applications to first-order mean field games. Preprint, available on ArXiv `https://arxiv.org/abs/1704.00378`, 2017.

Saeed Hadikhanloo. *Learning in Mean Field Games*. PhD thesis, Université Paris-Dauphine, Paris, France, 2018. URL `http://www.cmap.polytechnique.fr/~saeed.hadikhanloo/PhD_Thesis.pdf`.

Hans Haller and Roger Lagunoff. Genericity and Markovian Behavior in Stochastic Games. *Econometrica*, 68(5):1231–1248, 2000. doi: 10.1111/1468-0262.00156.

J. C. Harsanyi. Oddness of the Number of Equilibrium Points: A New Proof. *International Journal of Game Theory*, 2(1):235–250, 1973. doi: 10.1007/BF01737572.

Bernd Heidergott, Arie Hordijk, and Nicole Leder. Series Expansions for Continuous-Time Markov Processes. *Operations Research*, 58(3):756–767, 2010a. doi: 10.1287/opre.1090.0738.

Bernd Heidergott, Arie Hordijk, and Nicole Leder. Web Companion to: Series Expansions for Continuous-Time Markov Processes. *Operations Research*, 58(3), 2010b. URL `https://pubsonline.informs.org/doi/suppl/10.1287/opre.1090.0738`.

Harro Heuser. *Lehrbuch der Analysis Teil 2*. Mathematische Leitfäden. Vieweg + Teubner Verlag, Wiesbaden, 14^{th} edition, 2008. ISBN 978-3-8351-0208-8.

Morris W. Hirsch and Stephen Smale. *Differential Equations, Dynamical Systems, and Linear Algebra*. Pure and Applied Mathematics: A Series of Monographs and Textbooks. Academic Press, Inc., Orlando, 1974. ISBN 0-12-349550-4.

Roger A. Horn and Charles R. Johnson. *Matrix Analysis*. Cambridge University Press, New York, 2^{nd} edition, 2013. ISBN 978-0-521-54823-6.

Minyi Huang, Roland P. Malhamé, and Peter E. Caines. Large population stochastic dynamic games: closed-loop McKean-Vlasov systems and the Nash certainty equivalence principle. *Communications in Information & Systems*, 6(3):221–252, 2006. URL `https://projecteuclid.org/euclid.cis/1183728987`.

Gerhard Hübner. *Stochastik: Eine anwendungsorientierte Einführung für Informatiker, Ingenieure und Mathematiker*. Vieweg + Teubner, Wiesbaden, 5^{th} edition, 2009. ISBN 978-3-8348-0717-5.

Brian R. Hunt, Tim Sauer, and James A. Yorke. Prevalence: a translation-invariant "almost every" on infinite-dimensional spaces. *Bulletin of the American Mathematical Society*, 27(2):217–238, 1992. doi: 10.1090/S0273-0979-1992-00328-2.

Marius Iosifescu. *Finite Markov Processes and Their Applications*. Wiley series in probability and mathematical statistics. John Wiley & Sons, Ltd., Chichester, New York, Brisbane, Toronto, 1980. ISBN 0-471-27677-4.

Jiang Jia-He. Essential fixed points of multivalued mappings. *Scientia Sinica*, 11(3):293–298, 1962. URL `http://engine.scichina.com/doi/10.1360/ya1962-11-3-293`.

Jean Johnson and Dean Isaacson. Conditions for Strong Ergodicity Using Intensity Matrices. *Journal of Applied Probability*, 25(1):34–42, 1988. doi: 10.2307/3214231.

D. W. Jordan and P. Smith. *Nonlinear Ordinary Differential Equations: An Introduction to Dynamical Systems*, volume 2 of *Oxford Applied and Engineering Mathematics*. Oxford University Press, Oxford, 3^{rd} edition, 2005. ISBN 0-19-856562-3.

Boyan Jovanovic and Robert W. Rosenthal. Anonymous sequential games. *Journal of Mathematical Economics*, 17(1):77–87, 1988. doi: 10.1016/0304-4068(88)90029-8.

Prasadarao Kakumanu. Continuously Discounted Markov Decision Model with Countable State and Action Space. *The Annals of Mathematical Statistics*, 42(3):919–926, 1971. doi: 10.1214/aoms/1177693321.

Prasadarao Kakumanu. Relation between continuous and discrete time markovian decision problems. *Naval Research Logistics Quarterly*, 24(3):431–439, 1977. doi: 10.1002/nav.3800240306.

John L. Kelley. *General Topology.* The university series in higher mathematics. D. Van Nostrand Company, Inc., Princeton, New Jersey, 1955.

F. P. Kelly. *Reversibility and Stochastic Networks.* Wiley series in probability and mathematical statistics. John Wiley & Sons Ltd., Chichester, New York, Brisbane, Toronto, 1979. ISBN 0-471-27601-4.

J. F. C. Kingman. Ergodic Properties of Continuous-Time Markov Processes and Their Discrete Skeletons. *Proceedings of the London Mathematical Society*, s3–13(1):593–604, 1963. doi: 10.1112/plms/s3-13.1.593.

Donald E. Knuth and Jayme L. Szwarcfiter. A structured program to generate all topological sorting arrangements. *Information Processing Letters*, 2(6):153–157, 1974. doi: 10.1016/0020-0190(74)90001-5.

M. Kojima, A. Okada, and S. Shindoh. Strongly Stable Equilibrium Points of N-Person Noncooperative Games. *Mathematics of Operations Research*, 10(4):650–663, 1985. doi: 10.1287/moor.10.4.650.

V. N. Kolokoltsov and A. Bensoussan. Mean-Field-Game Model for Botnet Defense in Cyber-Security. *Applied Mathematics & Optimization*, 74(3):669–692, 2016. doi: 10.1007/s00245-016-9389-6.

V. N. Kolokoltsov and O. A. Malafeyev. Mean-Field-Game Model of Corruption. *Dynamic Games and Applications*, 7(1):34–47, 2017. doi: 10.1007/s13235-015-0175-x.

Vassili Kolokoltsov and Wei Yang. Inspection games in a mean field setting. Preprint, available on ArXiv `https://arxiv.org/abs/1507.08339`, 2015.

Vassili N. Kolokoltsov. *Nonlinear Markov Processes and Kinetic Equations*, volume 182 of *Cambridge Tracts in Mathematics*. Cambridge University Press, Cambridge, 2010. ISBN 978-0-521-11184-3.

Jean-Michel Lasry and Pierre-Louis Lions. Mean field games. *Japanese Journal of Mathematics*, 2(1):229–260, 2007. doi: 10.1007/s11537-007-0657-8.

Bar Light and Gabriel Y. Weintraub. Mean Field Equilibrium: Uniqueness, Existence, and Comparative Statics. Preprint, available at SSRN `https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3265048`, 2018.

Hartmut Logemann and Eugene P. Ryan. *Ordinary Differential Equations: Analysis, Qualitative Theory and Control.* Springer Undergraduate Mathematics Series. Springer-Verlag, London, 2014. ISBN 978-1-4471-6397-8.

Andreu Mas-Colell. On a theorem of Schmeidler. *Journal of Mathematical Economics*, 13(3):201–206, 1984. doi: 10.1016/0304-4068(84)90029-6.

Andreu Mas-Colell. *The theory of general economic equilibrium: A differentiable approach*, volume 9 of *Econometric Society monographs*. Cambridge University Press, Cambridge, New York, Port Chester, Melbourne, Sydney, 1985. ISBN 0-521-26514-2.

Eric Maskin and Jean Tirole. Markov Perfect Equilibrium: I. Observable Actions. *Journal of Economic Theory*, 100(2):191–219, 2001. doi: 10.1006/jeth.2000.2785.

E. J. McShane. Extension of range of functions. *Bulletin of the American Mathematical Society*, 40(12):837–842, 1934. URL `https://projecteuclid.org/euclid.bams/1183497871`.

Charafeddine Mouzouni. On Quasi-stationary Mean Field Games Models. *Applied Mathematics & Optimization*, 2018. doi: 10.1007/s00245-018-9484-y.

James Munkres. *Topology*. Pearson Education Limited, Essex, 2nd edition, 2014. ISBN 978-1-292-02362-5.

John Nachbar. Learning in Games. In Robert A. Meyers, editor, *Encyclopedia of Complexity and Systems Science*, pages 5177–5187. Springer, New York, NY, 2009. doi: 10.1007/978-0-387-30440-3_307.

J. R. Norris. *Markov Chains*. Cambridge Series on Statistical and Probabilistic Mathematics. Cambridge University Press, Cambridge, New York, Melbourne, Madrid, Cape Town, Singapore, São Paulo, 1997. ISBN 978-0-521-63396-3.

Marcel Nutz. A Mean Field Game of Optimal Stopping. *SIAM Journal on Control and Optimization*, 56(2):1206–1221, 2018. doi: 10.1137/16M1078331.

Marcel Nutz and Yuchong Zhang. A Mean Field Competition. *Mathematics of Operations Research*, 2019. doi: 10.1287/moor.2018.0966.

William Ott and James A. Yorke. Prevalence. *Bulletin of the American Mathematical Society*, 42(3):263–290, 2005. doi: 10.1090/S0273-0979-05-01060-8.

Christos H. Papadimitriou. On the Complexity of the Parity Argument and Other Inefficient Proofs of Existence. *Journal of Computer and System Sciences*, 48(3): 498–532, 1994. doi: 10.1016/S0022-0000(05)80063-7.

Christos H. Papadimitriou. The Complexity of Finding Nash Equilibria. In Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani, editors, *Algorithmic Game Theory*, pages 29–52. Cambridge University Press, Cambridge, 2007. doi: 10.1017/CBO9780511800481.004.

Martin L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley series in probability and mathematical statistics. John Wiley & Sons, Inc., New York, Chichester, Brisbane, Toronto, Singapore, 1994. ISBN 0-471-61977-9.

R. M. Redheffer. The Theorems of Bony and Brezis on Flow-Invariant Sets. *The American Mathematical Monthly*, 79(7):740–747, 1972. doi: 10.2307/2316263.

Sidney I. Resnick. *Adventures in Stochastic Processes*. Birkhäuser, Boston, 1992. ISBN 0-8176-3591-2.

Mansoor Saburov. Ergodicity of nonlinear Markov operators on the finite dimensional space. *Nonlinear Analysis: Theory, Methods & Applications*, 143:105–119, 2016. doi: 10.1016/j.na.2016.05.006.

David Schmeidler. Equilibrium Points of Nonatomic Games. *Journal of Statistical Physics*, 7(4):295–300, 1973. doi: 10.1007/BF01014905.

Horst Schubert. *Topologie: Eine Einführung*. Mathematische Leitfäden. B.G. Teubner Verlagsgesellschaft mbH, Stuttgart, 1964.

Gerald Teschl. *Ordinary Differential Equations and Dynamical Systems*, volume 140 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, Rhode Island, 2012. ISBN 978-0-8218-8328-0.

Eric van Damme. *Stability and Perfection of Nash Equilibria*. Springer-Verlag, Berlin, Heidelberg, 2nd edition, 1991. doi: 10.1007/978-3-642-58242-4.

John Von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, Princeton, 3rd edition, 1953. ISBN 0-691-00362-9.

Boto von Querenburg. *Mengentheoretische Topologie*. Springer-Lehrbuch. Springer-Verlag, Berlin, 3rd edition, 2001. ISBN 978-3-540-67790-9.

Mark Walker and John Wooders. mixed strategy equilibrium. In Steven N. Durlauf and Lawrence E. Blume, editors, *The New Palgrave Dictionary of Economics*, volume 5, pages 628–631. Macmillan Publishers Ltd., Hampshire, New York, 2nd edition, 2008. ISBN 978-0-230-22641-8.

Wolfgang Walter. *Ordinary Differential Equations*, volume 182 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1998. ISBN 0-387-98459-3.

John V. Wehausen. Transformations in metric spaces and ordinary differential equations. *Bulletin of the American Mathematical Society*, 51(2):113–119, 1945. URL `https://projecteuclid.org/euclid.bams/1183506765`.

Gabriel Y. Weintraub, C. Lanier Benkard, and Benjamin Van Roy. Markov Perfect Industry Dynamics With Many Firms. *Econometrica*, 76(6):1375–1411, 2008. doi: 10.3982/ECTA6158.

Wu Wen-Tsün and Jiang Jia-He. Essential equilibrium points of n-person non-cooperative games. *Scientia Sinica*, 11(10):1307–1322, 1962.

Holger Wendland. *Numerical Linear Algebra: An Introduction*. Cambridge Texts in Applied Mathematics. Cambridge University Press, Cambridge, 2018. ISBN 978-1-107-14713-3.

Dirk Werner. *Funktionalanalysis*. Springer-Lehrbuch. Springer-Verlag, Berlin, 7$^{\text{th}}$ edition, 2011. ISBN 978-3-642-21016-7.

A. I. Zeifman. Quasi-Ergodicity for Non-Homogeneous Continuous-Time Markov Chains. *Journal of Applied Probability*, 26(3):643–648, 1989. doi: 10.2307/3214422.

List of Symbols

The following list collects several symbols that have been used in this thesis. In the case that the notation is explained in the text the corresponding text reference is given.

2^A	power set of A
$\overline{A}$	the closure of the set A
$\langle \cdot, \cdot \rangle_B$	scalar product with respect to the basis B
$\|\| \cdot \|\|_B$	norm induced by the scalar product $\langle \cdot, \cdot \rangle_B$
a	a generic action
A	number of actions in $\mathcal{A}$
$\mathcal{A}$	action space
$\text{argmax}_{x \in X} f(x)$	the point(s) y that satisfy $f(y) = \max_{x \in X} f(x)$
β	discount factor
$\text{conv}(A)$	the convex hull of the set A
$\mathcal{C}(A, B)$	the set of all continuous functions from A to B
d	a (stationary) deterministic strategy, p. 29
$d_{i,a}(t)$	the probability to choose action a in state i and time t under strategy d, p. 29
$d_i(t)$	the action chosen by the deterministic strategy d in state i at time t, p. 29
$d^{(a_1,\ldots,a_S)}$	the stationary deterministic strategy satisfying $d(i) = a_i$ for all $i \in \mathcal{S}$, p. 58
D	the set of all deterministic strategies, p. 29
D^s	the set of all deterministic stationary strategies, p. 23

$\mathcal{D}$	the set of all optimal deterministic stationary strategies for a Markov decision process, p. 25
$\mathcal{D}(m)$	the set of all optimal deterministic stationary strategies given m, p. 51
$D(d)$	auxiliary digraph defined for the strategy $d \in \mathcal{U}$, p. 157
$d(x, y)$	distance of x and y with respect to metric d
$d(x, A)$	distance of x and the set A with respect to the metric d, that is $d(x, A) = \inf_{y \in A} d(x, y)$
$d((Q, r), (Q', r'))$	distance of the two games (Q, r) and (Q', r'), p. 93
$\hat{d}(Q, Q')$	distance of the two nonlinear generators Q and Q', p. 95
$\det(A)$	the determinant of the matrix A
δ_{ij}	Kronecker delta, that is $\delta_{ij} = 1_{\{i=j\}}$
$\mathrm{Eig}(\lambda)$	generalized eigenspace for the eigenvector λ, p. 129
$\mathrm{ext}(C)$	the set of extreme points of the closed and convex set C, p. 23
$\mathrm{FP}(f)$	the set of all fixed points of the (set-valued) map f, p. 60
$g_d(\cdot)$	function describing $\mathrm{Opt}^{\mathrm{some}}(d)$ and $\mathrm{Opt}^{\mathrm{unique}}(d)$, p. 156
$\mathcal{G}$	the set of all games, p. 93
$\mathcal{G}^{ci}$	the set of all games where $Q(\cdot)$ is constant in m and irreducible for all $d \in D^s$, p. 94
$H(A, B)$	metric on the power set of the metric space (X, d) based on the metric d, p. 94
i	a generic state
I	the identity matrix
$\mathrm{int}(A)$	the interior of the set A
m	population distribution
m_0	initial population distribution
$\dot{m}$	time derivative of m
$\mathcal{M}$	the set of all Lipschitz continuous functions ranging from $[0, \infty)$ to $\mathcal{P}(\mathcal{S})$ with Lipschitz constant L, p. 38

$N_\epsilon(x)$	ϵ-neighbourhood of x
$N_\epsilon(A)$	ϵ-neighbourhood of the set A, p. 94
∇f	gradient of f
O_i	the set of all actions maximizing the right hand side of the optimality equation of an Markov decision process for state i, p. 25
$O_i(m)$	the set of all actions that maximize the right hand side of the optimality equation given m, p. 51
$\mathrm{Opt}(A_1 \times \ldots \times A_S)$	the set of points $m \in \mathcal{P}(\mathcal{S})$ such that $O_i(m) = A_i$ for all $i \in \mathcal{S}$, p. 52
$\mathrm{Opt}^{\mathrm{some}}(d)$	the set of all those points $m \in O$ such that $d \in \mathcal{D}(m)$, p. 156
$\mathrm{Opt}^{\mathrm{unique}}(d)$	the set of all those points $m \in O$ such that $\mathcal{D}(m) = \{d\}$, p. 156
$p(s, i, t, j)$	the transition probability function of a time-inhomogeneous CTMC, p. 15
π	a (stationary) randomized strategy, p. 29
$\pi_{i,a}(t)$	the probability to choose action a in state i and time t under strategy π, p. 29
Π	the set of all strategies, p. 29
Π^s	the set of all stationary strategies, p. 23
$\mathcal{P}(A)$	the probability simplex over the set A
∂A	the boundary of the set A
$\frac{\partial}{\partial x} f(x, y)$	Jacobian of f with respect to x
$\phi(\cdot)$	best response correspondence in terms of dynamics, p. 39
$Q(t)$	transition rate matrix function for a time-inhomogeneous CTMC, p. 17
$Q_{ija}(m)$	transition rate from state i to state j using action a, when the population distribution is m, p. 29
Q^d	transition rate matrix of a continuous time Markov decision process given the deterministic stationary strategy d, p. 20
$Q^\pi(m, t)$	generator for the dynamics of the individual player playing strategy π given population distribution m at time t, p. 29

$Q^\pi(m)$	generator for the dynamics of the individual player playing the stationary strategy π given the stationary population distribution m, p. 31
$\tilde{Q}^\pi(m)$	the matrix $(Q^\pi(m))^T$, where the last row is replaced by $(1, \dots, 1)$, p. 56
Q^T	the transpose of the matrix Q
Q'_{ij}	the matrix Q, where the i-th row and j-th column is deleted
$\|\|Q\|\|$	uniform bound of $-Q_{iia}$, p. 19
(Q, r)	tuple describing a game with transition rate matrices $(Q_{\cdot\cdot a})_{a\in\mathcal{A}}$ and rewards $(r_{\cdot a})_{a\in\mathcal{A}}$, p. 93
$r_{ia}(m)$	instantaneous reward in state i, when action a is used, p. 30
r^π	reward given the strategy π, p. 19
$\mathrm{Re}(a)$	real part of a
$\rho(f, g)$	distance of the functions f and g, p. 96
S	number of states in $\mathcal{S}$
$\mathcal{S}$	state space
$\mathrm{sgn}(\cdot)$	sign function
SMFE$(\cdot)$	maps every game $(Q, r) \in \mathcal{G}$ to the set of stationary mean field equilibria (m, π), p. 94
TR	the set of all transition rate matrix functions such that $Q_{ij}(\cdot)$ is Lipschitz continuous for all $i, j \in \mathcal{S}$, p. 95
$\mathcal{U}$	the set of all strategies d such that there is an $m \in O$ such that $\{d\} = \mathcal{D}(m)$, p. 156
V^π	expected discounted reward of the strategy π in a Markov decision process, p. 19
V^*	the optimal discounted reward function for an Markov decision process, p. 20
$V_{x_0}(\pi^0, m)$	discounted reward given initial distribution x_0, flow of population distributions m and individual strategy π^0, p. 30
$x^\pi(m)$	the stationary points given $Q^\pi(m)$, p. 56

$Y(m, x)$	best possible payoff an individual can obtain given population distribution m and individual dynamics x, p. 39

A. Technical Appendix

This appendix contains a short introduction to Caratheodory solutions of ordinary differential equations and the proofs of the technical lemmata sorted by the chapters they are used in.

A.1. Caratheodory Solutions

We briefly review Caratheodory solutions of ordinary differential equations, for which Walter (1998, pp.121-125) and Filippov (1988, §1) are references.

Definition A.1. Let $I \times D \subseteq \mathbb{R}^{S+1}$ be open and let $f : I \times D \to \mathbb{R}^S$ be a function such that for every fixed $t \in I$ the function $x \mapsto f(t,x)$ is continuous and for every fixed $x \in D$ the function $t \mapsto f(t,x)$ is measurable.

A function $x : I \to \mathbb{R}^S$ is a Caratheodory solution of the initial value problem

$$\dot{x}(t) = f(t, x(t)), \quad x(0) = x_0 \tag{A.1}$$

if $x(t)$ is absolutely continuous and satisfies (A.1) for almost every $t \in I$.

We remark that classical solutions of the initial value problem (A.1) are particular Caratheodory solutions. However, not all Caratheodory solutions are also classical solutions of the initial value problem.

A classical existence theorem (Theorem 10.XX in Walter (1998)) is the following:

Lemma A.2. *Let $I \times D \subseteq \mathbb{R}^{S+1}$ be open and let $f(t,x) : D \to \mathbb{R}^S$ be a function such that for every fixed $t \in I$ the function $x \mapsto f(t,x)$ is continuous and for every fixed $x \in D$ the function $t \mapsto f(t,x)$ is measurable. Assume that for every set $J \times K \subseteq I \times D$, where $J \subseteq I$ is a finite interval and $K \subseteq D$ is a closed ball, it holds that there is a function $h(t)$ integrable over J such that $|f(t,x)| \leq h(t)$ for all $(t,x) \in J \times K$.*

Then for any initial condition $(t, x_0) \in D$ there is at least one Caratheodory solution of (A.1)*. Moreover, it can be extended to the left and right to the boundary of D.*

A classical uniqueness result is presented in Filippov (1988, Theorem 1.2):

Lemma A.3. *Let* $I \times D \subseteq \mathbb{R}^{S+1}$ *be open and let* $f(t,x) : D \to \mathbb{R}^S$ *be a function such that for every fixed* $t \in I$ *the function* $x \mapsto f(t,x)$ *is continuous and for every fixed* $x \in D$ *the function* $t \mapsto f(t,x)$ *is measurable. Let* $(t_0, x_0) \in D$ *and let* $l : \mathbb{R} \to \mathbb{R}$ *be integrable over any finite interval such that for any points* $(t,x), (t,y) \in D$

$$||f(t,x) - f(t,y)|| \leq l(t)||x - y||.$$

Then there exists at most one Caratheodory solution of (A.1) *in* D*.*

A.2. Appendix to Chapter 3

Lemma A.4. *Given a strategy* $\pi \in \Pi$ *and a Lipschitz continuous function* $m : [0,\infty) \to \mathcal{P}(\mathcal{S})$ *there is a unique transition function* $p(s,i,t,j)$ *(and thus a corresponding time-inhomogeneous continuous time Markov chain) with transition rates* $Q(t) = Q^\pi(m(t),t)$*.*

Proof. We show that $Q_{ij}(t) = Q^\pi_{ij}(m(t),t)$ satisfies the conditions from Section 2.1 that guarantee the existence of a unique, standard, stable and regular transition function (and thus the existence of a time-inhomogeneous continuous time Markov chain) with transition rates $Q^\pi(m(t),t)$.

Since the family $(Q_{\cdot\cdot a}(m))_{a\in\mathcal{A}}$ consists of conservative generators, the conditions (i) and (ii) are clearly satisfied and $Q^\pi(m(t),t)$ is conservative. Because the function $Q_{ija}(m(t))$ is a composition of Lipschitz continuous functions, it is continuous and Borel measurable. Therefore, the function $Q_{ij}(t)$ is Borel measurable because it is a sum of products of Borel measurable functions ($\pi(\cdot)$ is measurable by definition of the strategy space). Since $Q_{ija}(m(t))$ is uniformly bounded by some constant L, $Q_{ij}(t)$ is bounded the same constant and, thus, $Q_{ij}(\cdot)$ is integrable over any finite interval.

It remains to show that condition (2.2) is satisfied. For this we choose $w \equiv 1$, $L_1 = L_2 = L$ and obtain, since $Q_{ija}(m)$ is uniformly bounded by L, that for all $i \in \mathcal{S}$ and $t \geq 0$

$$\sum_{j\in\mathcal{S}} Q_{ij}(t) = \sum_{a\in\mathcal{A}} \underbrace{\left(\sum_{j\in\mathcal{S}} Q_{ija}(m(t))\right)}_{=0 \text{ as } Q_{ija}(m) \text{ is conservative}} \pi_{ia}(t) = 0 \leq L \quad \text{and} \; -Q_{ii}(t) \leq L = L \cdot 1.$$

Therefore, by the results stated in Section 2.1, we have a unique, standard, stable and regular transition function $p(s,i,t,j)$ (as well as a Markov process) with transition rates $Q_{ij}(t)$. □

Lemma A.5. *The marginal distributions $x(t)$ of the Markov process with initial distribution x^0 and infinitesimal generator $Q^\pi(m(t),t)$ satisfies*

$$\dot{x}_j(t) = \sum_{i\in\mathcal{S}} x_i(t) Q^\pi_{ij}(m(t),t), \quad \forall j \in \mathcal{S}$$

for almost all $t \geq 0$.

Proof. Let us write $(x^0_k)_{k\in\mathcal{S}}$ for the initial distribution of the Markov process. Then we have

$$x_j(t) = \sum_{k\in\mathcal{S}} x^0_k \cdot p(0,k,t,j).$$

Using the Kolmogorov forward equation (2.1) we obtain the desired characterization for almost all $t \geq 0$:

$$\begin{aligned}\dot{x}_j(t) &= \sum_{k\in\mathcal{S}} x^0_k \frac{\partial p(0,k,t,j)}{\partial t} = \sum_{k\in\mathcal{S}} x^0_k \sum_{i\in\mathcal{S}} p(0,k,t,i) Q^\pi_{ij}(m(t),t) \\ &= \sum_{i\in\mathcal{S}} \left(\sum_{k\in\mathcal{S}} x^0_k p(0,k,t,i) \right) Q^\pi_{ij}(m(t),t) = \sum_{i\in\mathcal{S}} x_i(t) Q^\pi_{ij}(m(t),t)\end{aligned}$$

□

Lemma A.6. *Given $m_0 \in \mathcal{P}(\mathcal{S})$ and $\pi \in \Pi$ there is at most one Caratheodory solution of the dynamics equation* (3.2)*.*

Proof. We apply Lemma A.3. Let M be a Lipschitz constant for all functions $Q_{ija}(\cdot)$ $(i,j \in \mathcal{S})$ with respect to the 1-norm and let L be the constant that uniformly bounds $Q_{ija}(m)$ for all $i,j \in \mathcal{S}$, $a \in \mathcal{A}$ and $m \in \mathcal{P}(\mathcal{S})$. We prove that

$$||f(t,m^1) - f(t,m^2)|| \leq (M+L) \cdot S$$

holds on $\mathcal{P}(\mathcal{S})$, where $f(t,m) = \sum_{i\in\mathcal{S}} \sum_{a\in\mathcal{A}} Q_{ija}(m)\pi_{ia}(t)m_i$:

$$\begin{aligned}&||f(t,m^1) - f(t,m^2)||_1 \\ &= \sum_{j\in\mathcal{S}} \left| \left(\sum_{i\in\mathcal{S}} \sum_{a\in\mathcal{A}} Q_{ija}(m^1)\pi_{ia}(t)m^1_i \right) - \left(\sum_{i\in\mathcal{S}} \sum_{a\in\mathcal{A}} Q_{ija}(m^2)\pi_{ia}(t)m^2_i \right) \right| \\ &= \sum_{j\in\mathcal{S}} \left| \sum_{i\in\mathcal{S}} \sum_{a\in\mathcal{A}} \left(Q_{ija}(m^1)m^1_i - Q_{ija}(m^2)m^2_i \right) \pi_{ia}(t) \right| \\ &\leq \sum_{j\in\mathcal{S}} \sum_{i\in\mathcal{S}} \sum_{a\in\mathcal{A}} \pi_{ia}(t) \left| Q_{ija}(m^1)m^1_i - Q_{ija}(m^2)m^1_i + Q_{ija}(m^2)m^1_i - Q_{ija}(m^2)m^2_i \right| \\ &\leq \sum_{j\in\mathcal{S}} \sum_{i\in\mathcal{S}} \sum_{a\in\mathcal{A}} \pi_{ia}(t) \left| Q_{ija}(m^1)m^1_i - Q_{ija}(m^2)m^1_i \right| + \left| Q_{ija}(m^2)m^1_i - Q_{ija}(m^2)m^2_i \right|\end{aligned}$$

$$
\begin{aligned}
&= \sum_{j\in\mathcal{S}}\sum_{i\in\mathcal{S}}\sum_{a\in\mathcal{A}} \pi_{ia}(t)\left(|Q_{ija}(m^1) - Q_{ija}(m^2)|m_i^1 + |Q_{ija}(m^2)|\cdot|m_i^1 - m_i^2|\right)\\
&= \sum_{j\in\mathcal{S}}\sum_{i\in\mathcal{S}}\sum_{a\in\mathcal{A}} \pi_{ia}(t)\left(M||m^1 - m^2||_1 m_i^1 + L|m_i^1 - m_i^2|\right)\\
&= \sum_{j\in\mathcal{S}}\sum_{i\in\mathcal{S}} \left(M||m^1 - m^2||_1 m_i^1 + L|m_i^1 - m_i^2|\right)\\
&= \sum_{j\in\mathcal{S}} \left(\sum_{i\in\mathcal{S}} M||m^1 - m^2||_1 m_i^1 + \sum_{i\in\mathcal{S}} L|m_i^1 - m_i^2|\right)\\
&= \sum_{j\in\mathcal{S}} \left(M||m^1 - m^2||_1 + L||m^1 - m^2||_1\right)\\
&= (M+L)\cdot S \cdot ||m^1 - m^2||_1.
\end{aligned}
$$

By Corollary 1 of McShane (1934) we find a Lipschitz continuous extension $\tilde{f} : \mathbb{R}^S \to \mathbb{R}^S$ of $f(\cdot)$. This in turn implies that at most one Caratheodory solution of (3.2) exists. □

A.3. Appendix to Chapter 4

Proof of Lemma 4.5. We set $K_n = [0, n]$ for all $n \in \mathbb{N}$. These sets are compact and furthermore we have $K_1 \subseteq K_2 \subseteq K_3 \subseteq \ldots$ as well as $\bigcup_{n\in\mathbb{N}} K_n = [0, \infty)$. Since every compact set K is bounded we find some $n \in \mathbb{N}$ such that $K \subseteq K_n$. Since the space $\mathcal{P}(\mathcal{S})$ is compact, the function $y \mapsto ||y||_\infty := \max_{i\in\mathcal{S}} |y_i|$ is a bounded. Let us denote its upper bound by M.

We now prove that the topology of uniform convergence on compacta T_C is equivalent to the topology T_d induced by the metric d. For this we have to show for all $m \in \mathcal{C}([0,\infty), \mathcal{P}(\mathcal{S}))$ and all basis elements B of one topology that there is a basis element B' of the other topology such that $m \in B' \subseteq B$ (Munkres, 2014, Lemma 13.3).

In Munkres (2014, p. 279-284) we find a suitable basis of the topology T_C of uniform convergence on compacta: We consider for all compact sets $K \subseteq [0, \infty)$, all functions $m \in \mathcal{C}([0,\infty), \mathcal{P}(\mathcal{S}))$ and all $\epsilon > 0$ the sets

$$B_K(m, \epsilon) = \left\{ m' \in \mathcal{C}([0,\infty), \mathcal{P}(\mathcal{S})) : \sup_{t\in K} ||m(t) - m'(t)||_\infty < \epsilon \right\}.$$

A basis of T_d is given by

$$B_\epsilon(m) = \{m' \in \mathcal{C}([0,\infty), \mathcal{P}(\mathcal{S})) : d(m, m') < \epsilon\}$$

for all $\epsilon > 0$ and all functions $m \in \mathcal{C}([0,\infty), \mathcal{P}(\mathcal{S}))$.

We first show that $T_C \subseteq T_d$ by proving that every basis element in T_C contains a basis element of T_d: Let $m \in \mathcal{C}([0,\infty),\mathcal{P}(\mathcal{S}))$ and $\epsilon > 0$ and choose $n \in \mathbb{N}$ such that $K_n \supseteq K$:

$$\begin{aligned}
&\left\{m' \in \mathcal{C}([0,\infty),\mathcal{P}(\mathcal{S})) : \sup_{t\in K} ||m(t) - m'(t)||_\infty < \epsilon\right\} \\
&\supseteq \left\{m' \in \mathcal{C}([0,\infty),\mathcal{P}(\mathcal{S})) : \sup_{t\in K_n, i\in\mathcal{S}} |m_i(t) - m_i'(t)| < \epsilon\right\} \\
&\supseteq \left\{m' \in \mathcal{C}([0,\infty),\mathcal{P}(\mathcal{S})) : \sup_{t\in K_n, i\in\mathcal{S}} e^{-\beta t}|m_i(t) - m_i'(t)| < e^{-\beta n}\epsilon\right\} \\
&\supseteq \left\{m' \in \mathcal{C}([0,\infty),\mathcal{P}(\mathcal{S})) : \sup_{t\geq 0, i\in\mathcal{S}} e^{-\beta t}|m_i(t) - m_i'(t)| < e^{-\beta n}\epsilon\right\} \\
&= \left\{m' \in \mathcal{C}([0,\infty),\mathcal{P}(\mathcal{S})) : d(m,m') < e^{-\beta n}\epsilon\right\}.
\end{aligned}$$

Now we show $T_d \subseteq T_C$: We are given an ϵ-ball in T_d and choose $n \in \mathbb{N}$ such that $Me^{-\beta n} < \epsilon$. This implies

$$\begin{aligned}
&\left\{m' \in \mathcal{C}([0,\infty),\mathcal{P}(\mathcal{S})) : \sup_{t\geq 0, i\in\mathcal{S}} e^{-\beta t}|m_i(t) - m_i'(t)| < \epsilon\right\} \\
&= \Bigg\{m' \in \mathcal{C}([0,\infty),\mathcal{P}(\mathcal{S})) : \sup_{t\in K_n, i\in\mathcal{S}} e^{-\beta t}|m_i(t) - m_i'(t)| < \epsilon \\
&\qquad \wedge \sup_{t\geq n, i\in\mathcal{S}} e^{-\beta t}|m_i(t) - m_i'(t)| < \epsilon\Bigg\} \\
&\supseteq \Bigg\{m' \in \mathcal{C}([0,\infty),\mathcal{P}(\mathcal{S})) : \sup_{t\in K_n} e^{-\beta t}||m(t) - m'(t)||_\infty < \epsilon \\
&\qquad \wedge \sup_{t\geq n} e^{-\beta n}||m(t) - m'(t)||_\infty < \epsilon\Bigg\} \\
&= \left\{m' \in \mathcal{C}([0,\infty),\mathcal{P}(\mathcal{S})) : \sup_{t\in K_n} e^{-\beta t}||m(t) - m'(t)||_\infty < \epsilon\right\} \\
&\supseteq \left\{m' \in \mathcal{C}([0,\infty),\mathcal{P}(\mathcal{S})) : \sup_{t\in K_n} ||m(t) - m'(t)||_\infty < \epsilon\right\}.
\end{aligned}$$

□

Lemma A.7. *Assume that X is a topological space such that any set $A \subseteq X$ is compact if and only if it is sequentially compact and, analogously, that Y is a topological space such that any set $B \subseteq Y$ is compact if and only if it is sequentially compact. Furthermore, assume that for each compact set $A \subseteq X$ and $B \subseteq Y$, respectively, the topology on A and B, respectively, is metrizable. Then the space $X \times Y$ equipped with the product topology also satisfies that a set $C \subseteq X \times Y$ is compact if and only if it is sequentially compact.*

Proof. Assume that $C \subseteq X \times Y$ is sequentially compact, that is any sequence $(x_n, y_n)_{n\in\mathbb{N}}$ has a converging subsequence with limit $(x,y) \in C$. Let $(x_n)_{n\in\mathbb{N}}$ be an arbitrary sequence in $\pi_X(C)$, where $\pi_X : X \times Y \to X$ is the projection onto X, and choose for all $n \in \mathbb{N}$ an element $y_n \in \pi_Y(C)$ such that $(x_n, y_n) \in C$. Then the sequence has a converging subsequence with limit (x, y), which especially yields that $(x_n)_{n\in\mathbb{N}}$ has a converging subsequence with limit $x \in \pi_X(C)$. Thus, the sets $\pi_X(C)$ (and analogously $\pi_Y(C)$) are sequentially compact, which by assumption means that $\pi_X(C)$ and $\pi_Y(C)$ are compact.

Let $(O_i)_{i\in I}$ be an open cover of C. Then for each $(x, y) \in C$ there is an index $i(x, y)$ such that $(x, y) \in O_{i(x,y)}$. Since $O_{i(x,y)}$ is open, we furthermore find $U_{i(x,y)} \subseteq X$ and $V_{i(x,y)} \subseteq Y$ such that

$$(x, y) \in U_{i(x,y)} \times V_{i(x,y)} \subseteq O_{i(x,y)}.$$

Now fix x. Then $(V_{i(x,y)})_{y\in\pi_Y(C)}$ is an open cover of $\pi_Y(C)$. Since $\pi_Y(C)$ is compact, we find a finite subcover $(V_{i(x,y_k(x))})_{k=1}^n$. Since the sets $U_x = \bigcap_{k=1}^n U_{i(x,y_k(x))}$ are finite intersections of open sets, they are in turn open and hence $(U_x)_{x\in\pi_X(C)}$ is an open cover of $\pi_X(C)$. Since $\pi_X(C)$ is compact, we find a finite subcover $(U_{x_j})_{j=1}^m$. Now $\big(U_{i(x_j,y_k(x_j))} \times V_{i(x_j,y_k(x_j))}\big)_{k=1,j=1}^{n,m}$ is a finite subcover of C and so is

$$\big(O_{i(x_j,y_k(x_j))}\big)_{k=1,j=1}^{n,m},$$

which proves that C is compact.

Now assume that $C \subseteq X \times Y$ is compact. Then $\pi_X(C)$ and $\pi_Y(C)$ are compact as they are the image of a compact set under a continuous map. By assumption $\pi_X(C)$ and $\pi_Y(C)$ are metrizable, thus also $\pi_X(C) \times \pi_Y(C)$ is metrizable. Since for metric space compactness and sequential compactness coincides and $C \subseteq \pi_X(C) \times \pi_Y(C)$, we obtain the desired claim. □

Lemma A.8. *Let X and Y be compact metric spaces. Assume that $f : X \times Y \to \mathbb{R}$ is a continuous function, then $g : X \to \mathbb{R}$ such that $x \mapsto g(x) := \sup_{y\in Y} f(x, y)$ is continuous.*

Proof. By compactness of Y the function $g(\cdot)$ is well-defined. In order to prove continuity we note that it is sufficient to check that $g^{-1}((-\infty, a))$ and $g^{-1}((b, \infty))$ are open since

$$\{(-\infty, a)) : a \in \mathbb{R}\} \cup \{(b, \infty) : b \in \mathbb{R}\}$$

forms a subbasis of $\mathcal{B}(\mathbb{R})$ (von Querenburg, 2001, Theorem 2.21).

We again write $\pi_Y : X \times Y \to Y, (x, y) \mapsto y$ for the projection onto Y. Using this we obtain that $g^{-1}((b, \infty)) = \pi_Y \circ f^{-1}((b, \infty))$ because $g(x) > b$ holds whenever there is some $y \in Y$ such that $f(x, y) > b$. Since $f(\cdot, \cdot)$ is continuous and $\pi_Y(\cdot)$ is open, we then obtain that $g^{-1}((b, \infty))$ is open.

To prove that $g^{-1}((-\infty, a))$ is open, we prove that for each $x \in g^{-1}((-\infty, a))$ there is an ϵ-neighbourhood such that $N_\epsilon(x) \subseteq g^{-1}((-\infty, a))$: Indeed, if $x \in g^{-1}((-\infty, a))$, then, by definition of the supremum, we have that $\{x\} \times Y \subseteq f^{-1}((-\infty, a))$. Since $f^{-1}((-\infty, a))$ is open, we obtain that for each point $y \in Y$ there are $\epsilon_1(x, y)$ and $\epsilon_2(x, y)$ such that

$$N_{\epsilon_1(x,y)}(x) \times N_{\epsilon_2(x,y)}(y) \subseteq f^{-1}((-\infty, a)).$$

Since Y is compact we furthermore have that there is a finite subcover of $\{x\} \times Y$, let this be

$$\bigcup_{i=1}^{n} \left(N_{\epsilon_1(x,y_i)}(x) \times N_{\epsilon_2(x,y_i)}(y_i)\right) \supseteq \{x\} \times Y.$$

The set

$$\bigcup_{i=1}^{n} \left(\left(\bigcap_{i=1}^{n} N_{\epsilon_1(x,y_i)}(x) \right) \times N_{\epsilon_2(x,y_i)}(y_i) \right)$$

is again a subcover of $\{x\} \times Y$, furthermore the set is open and satisfies

$$\bigcup_{i=1}^{n} \left(\left(\bigcap_{i=1}^{n} N_{\epsilon_1(x,y_i)}(x) \right) \times N_{\epsilon_2(x,y_i)}(y_i) \right) \subseteq f^{-1}((-\infty, a)).$$

This in total implies that $\bigcap_{i=1}^{n} N_{\epsilon_1(x,y_i)}(x) \in g^{-1}((-\infty, a))$, which proves the desired claim. □

A.4. Appendix to Chapter 5

Proof of Lemma 5.7. Since we face a conservative generator, we have, by definition, that $Q_{ij} \geq 0$ for all $i \neq j$ and that $\sum_{j \in \mathcal{S}} Q_{ij} = 0$ for all $i \in \mathcal{S}$. Thus, all off-diagonal entries of Q are non-negative and the row sum is always zero. Furthermore, by requiring irreducibility we do not have a row of zeros, thus the diagonal entries are strictly negative.

The matrix Q'_{SS} again has non-negative off-diagonal entries, strictly negative diagonal entries and the row sums are less or equal zero. Furthermore, the irreducibility of Q implies that Q_{iS} is non-zero for at least one $i \in \{1, \ldots, S-1\}$. Thus, the row sum for at least one i is strictly negative. With this preliminary observations we show that I_{28} is satisfied.

For this we first note that for all $k \in \{1, \ldots, S-1\}$ and all $y \in \mathbb{R}^{S-1}$

$$\begin{aligned}(-Q'_{SS}y)_k &= \sum_{l=1}^{S-1} (-Q'_{SS})_{kl} \cdot y_l = \sum_{l=1}^{S-1} (-Q)_{kl} \cdot y_l \\ &= \sum_{l \in \{1,\ldots,S-1\} \setminus \{k\}} -Q_{kl} \cdot y_l + \cdot(-Q)_{kk} \cdot y_k\end{aligned}$$

$$= \sum_{l \in \{1,\dots,S-1\} \setminus \{k\}} -Q_{kl} \cdot y_l + \sum_{l \in \mathcal{S} \setminus \{k\}} Q_{kl} \cdot y_k$$
$$= \sum_{l \in \{1,\dots,S-1\} \setminus \{k\}} Q_{kl} \cdot (y_k - y_l) + Q_{kS} \cdot y_k,$$

from which we obtain that $(-Q'_{SS}y)_k$ is increasing in y_k and decreasing in y_l since $Q_{kl} > 0$ for all $l \neq k$.

Define for all $y \in [0, \infty)^{S-1}$ the set

$$\mathcal{T}(y) = \{k \in \{1, \dots, S-1\} : (-Q'_{SS}y)_k > 0\}$$

as the set of all indices such that the k-th component of $-Q'_{SS}y$ is greater than zero. We will inductively define a sequence $(y^n)_{n \in \mathbb{N}}$ of vectors such that for some $\tilde{n} \in \mathbb{N}$ we have $\mathcal{T}(y^{\tilde{n}}) = \{1, \dots, S-1\}$. We will start with the vector y^0 with a one at every component and construct y^n in such a way that $\mathcal{T}(y^{n-1}) \subsetneq \mathcal{T}(y^n)$ for all $n \in \mathbb{N}$ and $y_i^n > 0$ for all $i \in \{1, \dots, S-1\}$ and all $n \in \mathbb{N}_0$.

Starting with y^0 being the vector of ones, we have that $\mathcal{T}(y^0) \neq \emptyset$ because we have previously seen that the row sum, which is $(-Q'_{SS}y^0)_i$, is positive for some index $i \in \{1, \dots, S-1\}$.

In the n-th step ($n \in \mathbb{N}$) we check whether $\mathcal{T}(y^{n-1}) = \{1, \dots, S-1\}$. If this is the case we have shown that I_{28} indeed holds, else there is an index $j \in \{1, \dots, S-1\} \setminus \mathcal{T}(y^{n-1})$. In this case, let $i \in \mathcal{T}(y^{n-1})$. Since the generator matrix is irreducible, the underlying transition graph is strongly connected. Thus, there is a path from j to i. Let $\tilde{k}$ be the first node on this path that lies in $\mathcal{T}(y^{n-1})$ and let k be its predecessor. We note that by definition of the transition graph $Q_{k\tilde{k}} > 0$. Now let y^n be as follows $y_l^n = y_l^{n-1}$ for all $l \in \{1, \dots, S-1\} \setminus \{\tilde{k}\}$ and $y_{\tilde{k}}^n$ such that

$$0 < y_{\tilde{k}}^n < y_{\tilde{k}}^{n-1} \quad \text{and} \quad (-Q'_{SS}y^n)_{\tilde{k}} > 0.$$

This is possible as $y_{\tilde{k}}^{n-1} > 0$ and $(-Q'_{SS}y^{n-1})_{\tilde{k}} > 0$ and $(-Q'_{SS}y)_{\tilde{k}}$ is continuous as well as increasing in $y_{\tilde{k}}$.

It is immediate that $y_l^n > 0$ for all $l \in \{1, \dots, S-1\}$. It remains to show that $p \in \mathcal{T}(y^{n-1})$ implies $p \in \mathcal{T}(y^n)$ and to show that $k \in \mathcal{T}(y^n)$ because this proves that $\mathcal{T}(y^{n-1}) \subsetneq \mathcal{T}(y^n)$.

Since $(-Q'_{SS}y^n)_{\tilde{k}} > 0$, we have by construction that $\tilde{k} \in \mathcal{T}(y^n)$. For $p \neq \tilde{k}$ we see again as $(-Q'_{SS}y)_p$ is decreasing in $y_{\tilde{k}}$ and $y_p^n = y_p^{n-1}$ that

$$(-Q'_{SS}y^n)_p = \sum_{l \in \{1,\dots,S-1\} \setminus \{p\}} Q_{pl} \cdot (y_p^n - y_l^n) + Q_{pS} \cdot y_p^n$$
$$\geq \sum_{l \in \{1,\dots,S-1\} \setminus \{p\}} Q_{pl} \cdot (y_p^{n-1} - y_l^{n-1}) + Q_{pS} \cdot y_p^{n-1} = (Q'_{SS}y^{n-1})_p.$$

Thus, if $p \in \mathcal{T}(y^{n-1})$, then $p \in \mathcal{T}(y^n)$.

For $p = k$ we have a strict inequality as $(y_k^n - y_{\tilde{k}}^n)Q_{k\tilde{k}} > (y_k^{n-1} - y_{\tilde{k}}^{n-1})Q_{k\tilde{k}}$ since $Q_{k\tilde{k}} > 0$ and $y_{\tilde{k}}^n < y_{\tilde{k}}^{n-1}$. Thus, we have $(Q'_{SS}y^n)_k > (Q'_{SS}y^{n-1})_k = 0$, which implies $k \in \mathcal{T}(y^n)$. In total we proved that $\mathcal{T}(y^{n-1}) \subsetneq \mathcal{T}(y^n)$.

Since I_{28} is satisfied and $-Q'_{SS}$ lies in $Z^{n\times n}$, we conclude that it is a non-singular M-matrix and that G_{20} holds. This implies that all eigenvalues of the matrix $-Q'_{SS}$ have positive real part and that in turn all eigenvalues of Q'_{SS} have negative real part.

Since the matrix Q'_{SS} is real, all complex eigenvalues appear in pairs of complex conjugates. Since the product of a complex number and its complex conjugate is non-negative, we obtain that the sign of the determinant, which can be computed as the product of all eigenvalues, depends only on the number of real eigenvalues or more specifically the parity of this number. As the parity of the number of real eigenvalues always equals the parity of $S - 1$, we conclude that the sign pattern of the determinant is as claimed. □

B. Addenda required by §7 of the doctoral degree regulations of the MIN Faculty

B.1. Abstract

This thesis deals with mean field games theory, a game theoretic model for dynamic games with a continuum of players. The main feature of these games is that agents do not observe the behaviour of all other agents individually, but only the distribution thereof. This simplification makes models of intertemporal interactions of many players more tractable, which leads to several economic applications. In particular, examples with finite state and action space have been considered in the literature. However, there are, except for the existence of dynamic mean field equilibria, no general results for these types of models. This thesis investigates existence, computation, stability and explanatory power of stationary equilibria in mean field games with finite state and action space. We work in the model introduced in Doncel et al. (2016a), but consider a new (and instructive) probabilistic formulation. More precisely, the individual agents control a Markov chain with transition rates that depend on the population distribution and maximize their expected discounted reward over an infinite time horizon.

In the first part of the thesis we analyse stationary mean field equilibria: We start by proving existence of stationary equilibria in mixed strategies under a continuity assumption. Thereafter, we present several tools in order to compute all stationary mean field equilibria. We show that optimal stationary strategies are convex combinations of deterministic stationary strategies choosing an action that maximizes the right-hand side of the optimality equation of an associated Markov decision process. We provide a cut criterion to reduce the set of possible stationary points of the dynamics. Moreover, under an irreducibility assumption we can reduce the problem of finding all stationary mean field equilibria to finding all fixed points of a set-valued map. We also illustrate the application of the results in three toy examples. Concluding the analysis of stationary equilibria, we investigate under which conditions stationary equilibria are stable against slight model perturbations and under which conditions this is not the case. This is a classical question in game theory and has not been considered for mean field games so

far. Since this question is closely related to the question whether fixed points are essential, we do not obtain full characterizations, but similar results as for essential fixed points. In particular, we show that the set of all games with only essential stationary equilibria is (topologically) generic.

In the second part of the thesis we consider the question in how far stationary equilibria explain dynamic behaviour of agents in our mean field game model. Indeed, the computation of dynamic mean field equilibria is for infinite time horizon not clear and for finite time horizon rather involved and, moreover, most of the time only numerically tractable. Thus, it is questionable that agents can perform these computations. Therefore, we follow the lines of learning theory in games and introduce a sensible decision mechanism for partially rational agents that under suitable conditions converges towards stationary equilibria. More precisely, we introduce a myopic adjustment process in which agents assume that the population distribution will stay constant over time and choose their best response to it, if the population distribution changes the agents re-evaluate their decision. We prove that the process converges locally to stationary equilibria with a deterministic equilibrium strategy that is the unique optimal strategy for the equilibrium distribution under certain conditions regarding the population dynamics. Furthermore, we provide a set of assumptions that ensures that the process converges (almost) globally towards stationary mean field equilibria. More precisely, we show that the process converges towards a stationary equilibrium with a deterministic equilibrium strategy or that it stays in a set, where several stationary strategies are simultaneously optimal and for which we can often prove convergence towards stationary equilibria in mixed strategies.

Additionally, this thesis contributes to the theory of continuous time nonlinear Markov chains with finite state space. We prove existence of such a process given a Lipschitz continuous nonlinear generator and provide examples illustrating that the limit behaviour is more complex than for standard time-homogeneous Markov chains. Furthermore, we provide a sufficient criterion for the existence of a unique stationary distribution and we provide for small state spaces a sufficient criterion for strong ergodicity of the process.

B.2. Zusammenfassung

Diese Arbeit beschäftigt sich mit Mean Field Games, einem spieltheoretischen Modell für zeitdynamische Spiele mit einem Kontinuum an Agenten. Das zentrale Charakteristikum dieser Spiele besteht darin, dass Agenten nicht das individuelle Verhalten der anderen Agenten beobachten können, sondern nur dessen Verteilung. Diese Vereinfachung führt dazu, dass Modelle intertemporaler Interaktionen vieler Agenten handhabbar werden, was sie für ökonomische Anwendungen nutzbar macht. Insbesondere wurden hierfür Modelle mit endlichem Zustands- und Aktionsraum genutzt. Jedoch gibt es (fast) keine allgemeinen Resultate für solche Modelle. Diese Arbeit untersucht stationäre Gleichgewichte von Mean Field Games mit endlichem Zustands- und Aktionsraum. Dabei werden folgende vier Einzelfaktoren näher betrachtet: Existenz, Berechnung, Stabilität sowie das Potential stationärer Gleichgewichte das dynamische Verhalten von Agenten zu erklären. Wir arbeiten mit dem in Doncel et al. (2016a) eingeführten Modell in einer neuen (und aufschlussreichen) probabilistischen Formulierung. In diesem Modell kontrollieren die Agenten jeweils eine Markov-Kette mit Übergangsraten, die von der Populationsverteilung abhängen, und maximieren ihren erwarteten abdiskontierten Gesamtgewinn über einen unendlichen Zeithorizont.

Im ersten Teil der Arbeit analysieren wir stationäre Mean Field Gleichgewichte: Zuerst zeigen wir die Existenz von stationären Gleichgewichten in gemischten Strategien unter einer Stetigkeitsannahme. Danach stellen wir mehrere Werkzeuge vor, um alle stationären Gleichgewichte zu berechnen. Wir zeigen dabei dass alle optimalen stationären Strategien Konvexkombinationen deterministischer stationärer Strategien sind, die eine Aktion wählen, welche die rechte Seite der Optimalitätsgleichung eines zugehörigen Markov-Entscheidungsprozesses maximiert. Außerdem leiten wir ein Schnittkriterium zur Reduktion der Menge der möglichen stationären Punkte der Dynamik her. Des Weiteren können wir unter einer Irreduzibilitätsannahme die Aufgabe, alle stationären Gleichgewichte zu finden, auf das Bestimmen aller Fixpunkte einer mengenwertigen Abbildung reduzieren. Im Anschluss veranschaulichen wir die Anwendung der Resultate in drei Beispielen. Danach untersuchen wir, unter welchen Bedingungen stationäre Gleichgewichte bezüglich leichter Modellstörungen stabil sind. Dies ist eine klassische Frage in der Spieltheorie, die für Mean Field Games noch nicht betrachtet wurde. Da diese Frage eng mit der Frage, ob Fixpunkte essentiell sind, verbunden ist, erhalten wir keine vollständigen Charakterisierungen, sondern ähnliche Resultate wie in jenem Fall. Insbesondere zeigen wir dass die Menge aller Spiele mit ausschließlich essentiellen stationären Gleichgewichten (topologisch) generisch ist.

Im zweiten Teil der Arbeit beschäftigen wir uns mit der Frage, inwiefern stationäre Gleichgewichte das dynamische Verhalten von Agenten in unserem Mean Field Game erklären. Dieses ist dadurch motiviert dass die Berechnung von dynamischen Gleichgewichten schwierig ist und es daher fraglich ist ob ein Agent dynamische Gleichgewichte berechnen kann. Analog zur der Theorie des Lernens in Spielen führen wir sodann einen realitätsnahen Entscheidungsmechanismus für partiell rationale Agenten ein, der unter geeigneten Bedingungen gegen stationäre Gleichgewichte konvergiert. Im Detail betrachten wir einen myopischen Anpassungsprozess, in dem die Agenten davon ausgehen, dass sich die Verteilung der anderen Agenten nicht ändert und daher immer die für diese Situation beste Antwort als Strategie wählen. Ändert sich die Verteilung der anderen Agenten, passen sie ihre Strategie an. Wir beweisen, unter bestimmten Voraussetzungen, dass der Prozess lokal gegen stationäre Gleichgewichte mit einer deterministischen Gleichgewichtsstrategie, die ferner die einzige optimale Strategie für die Gleichgewichtsverteilung ist, konvergiert. Darüber hinaus zeigen wir, dass der Prozess unter bestimmten Annahmen (fast) global gegen stationäre Gleichgewichte konvergiert. Genauer zeigen wir, dass der Prozess gegen ein stationäres Gleichgewicht mit einer deterministischen Gleichgewichtsstrategie konvergiert oder er in einer Menge bleibt, in der mehrere stationären Strategien gleichzeitig optimal sind und in der wir oft die Konvergenz gegen ein stationäres Gleichgewicht mit einer gemischten Gleichgewichtsstrategie zeigen können.

Ergänzend trägt die Arbeit zur Theorie der zeitstetigen nichtlinearen Markov-Ketten mit endlichem Zustandsraum bei. Im Detail zeigen wir die Existenz eines solchen Prozesses gegeben eines Lipschitz-stetigen nichtlinearen Generators und präsentieren Beispiele, die illustrieren, dass das Grenzverhalten komplexer ist als bei klassischen zeithomogenen Markov-Ketten. Ferner beweisen wir ein hinreichendes Kriterium für die Existenz einer eindeutigen stationären Verteilung sowie ein hinreichendes Kriterium für starke Ergodizität bei kleinem Zustandsraum.

B.3. Publications Derived from the Thesis up to now

- Berenice Anne Neumann. Stationary Equilibria of Mean Field Games with Finite State and Action Space. Preprint, available on ArXiv `https://arxiv.org/abs/1901.08803`, 2019.

This paper contains the stationary existence result presented in Section 4.2, the results for equilibrium computation presented in Sections 5.1, 5.2 and 5.3 as well as the equilibrium computation in the examples presented in Sections 6.1 and 6.2.

In January 2020 this preprint was published with minor revision in Dynamic Games and Applications:

- Berenice Anne Neumann. Stationary Equilibria of Mean Field Games with Finite State and Action Space, *Dynamic Games and Applications*, 2020, doi: 10.1007/s13235-019-00345-9.

B.4. Declarations on my Contributions

According to §7(3) of the doctoral degree regulations, I declare that this dissertation was written without co-authors and is therefore my own work.

B.5. Eidesstattliche Versicherung

Hiermit erkläre ich an Eides statt, dass ich die vorliegende Dissertationsschrift selbst verfasst und keine anderen als die angegebenen Quellen und Hilfsmittel benutzt habe.