

Formalization and Metastability Analysis of Agent-Based Evolutionary Models

Dissertation

zur Erlangung des Grades
eines Doktors der Naturwissenschaften (Dr. rer. nat.)
am Fachbereich Mathematik & Informatik
der Freien Universität Berlin

vorgelegt von
Mareen Hallier

Berlin, Dezember 2014

Erster Gutachter und Betreuer

Prof. Dr. Rupert Klein
Freie Universität Berlin, Fachbereich Mathematik & Informatik

Weitere Gutachter

Prof. Dr. Carlo C. Jaeger
Universität Potsdam, Fachbereich Volkswirtschaftslehre

Prof. Dr. Carsten Hartmann
Freie Universität Berlin, Fachbereich für Mathematik & Informatik

Tag der Disputation: 20. Februar 2015

Contents

Introduction	11
I Formalization	17
1 Specifications for Agent-Based Models of Exchange	19
1.1 Introduction to Agent-Based Modeling	19
1.2 Exploratory Programming and The Need for Specifications	20
1.3 Basic Notions and Notation	23
1.4 A Functional Framework for Specifying Agent-Based Models of Exchange	27
1.4.1 The Problem of Price Formation	27
1.4.2 Basic Structure of the Framework	31
1.4.3 Application	41
1.5 Conclusion	46
II Metastability Analysis	49
2 Mathematical Preliminaries	51
2.1 Probability Theory	51
2.2 Markov Chains	54
3 Stochastic Evolutionary Games	65
3.1 Definition	65
3.1.1 Games in Strategic Form	66
3.1.2 Population Games	69
3.1.3 Revision Protocol, Stochastic Evolutionary Games	70
3.2 Analysis	76
3.2.1 Nash Equilibrium and Refinements	77
3.2.2 Deterministic Approximation	78
3.2.3 Stationary Distribution and Reversibility	85
3.2.4 Stochastic Stability Analysis	86
4 Characterization of Metastability	95
4.1 Metastability based on Hitting Times	96
4.2 Metastability based on Transition Probabilities	98
4.3 Metastability in Models of Evolution with Noise	103

4.3.1	Metastability based on Hitting Times in Models of Evolution with Noise	104
4.3.2	Metastability based on Transition Probabilities in Models of Evolution with Noise	106
5	Markov State Modeling	109
5.1	Preliminaries	110
5.1.1	Transfer Operators	110
5.1.2	Orthogonal Projections	112
5.1.3	Committer Functions	112
5.2	Full Partition Markov State Models	113
5.2.1	Construction of Full Partition Markov State Models . . .	113
5.2.2	The Full Partition Markov State Model and Orthogonal Projections	117
5.3	Core Set Markov State Models	121
5.3.1	Projected Transfer Operators on General Subspaces . . .	122
5.3.2	General Markov State Models	126
5.3.3	Core Set Markov State Models: Interpretation and Estimation	127
5.4	Approximation Quality	131
5.4.1	Stationary Distribution and Stochastic Stability	131
5.4.2	Propagation of Probability Distributions	132
5.4.3	Eigenvalues	143
6	Identification of Metastable Sets for Agent-Based Evolutionary Models	147
6.1	Motivation	148
6.2	Algorithmic Strategy	151
6.3	Extended Example	155
6.4	Discussion and Outlook	161
	Summary	165
	Translation between Notations	167
	Bibliography	171
	Zusammenfassung	181

Acknowledgements

There are many people who supported me during my Phd studies. First of all, I thank my supervisor Rupert Klein for inspiring and encouraging me to explore and follow my own, interdisciplinary research interests while providing overall guidance throughout this time. I thank Carlo Jaeger for drawing my interests to economics and for seeing the potential of metastability analysis for agent-based modeling in economics, thereby laying the foundation for this thesis and allowing me to work in a captivating research field which combines my intellectual curiosity for mathematics and the study of socio-economic systems and global change. In addition, I sincerely thank Carsten Hartmann for pointing out to me and encouraging me to exploit the relationship between stochastic evolutionary games and agent-based models, for his patience, time and support what regards the detailed discussion of this thesis, and for his understanding of the not necessarily content-related questions of a PhD student.

I would also like to thank my former colleagues at the Potsdam Institute for Climate Impact Research and my colleagues at the Institute for Mathematics at the Free University of Berlin, who, taken all together, provided two amicable and stimulating working environments. In particular, I want to mention Nicola Botta, Steffen Fürst, Cezar Ionescu, Carlo Jaeger, Sophie Knell, Daniel Lincke, and Sarah Wolf as participants of the Cartesian seminar and the climate change economics discussion group – not only for the many inspiring discussions about climate change economics, game theory, computer science, and mathematics, but also about number stations, Parmesan, and other intriguing aspects of life – it was much fun and profoundly shaped my way of thinking.

I also thank Nicola Botta for intensive discussions regarding Part I of this thesis, Steffen Fürst, Steffen Galan, Sophie Knell, Ariane Papke, Sarah Wolf for reading and commenting parts of an earlier draft of this thesis, Dustin Miller for making the incredible effort of checking the whole final draft for English language errors. All remaining errors are, of course, my responsibility.

Nicht zuletzt gebührt mein großer Dank meiner Familie: ganz besonders meinen Eltern für ihre immerwährende Unterstützung, die auf vielfältige Art und Weise auch zu dieser Arbeit geführt hat; meinem Bruder Tom und meiner Schwester Bianca, als feste Bestandteile meines Lebens; meinen Schwiegereltern, dafür, dass sie so wunderbare Großeltern sind und durch ihre liebevolle Betreuung von Jakob mir so manche ruhige Stunde für diese Arbeit ermöglicht haben; meinem Sohn Jakob, einfach dafür, dass es ihn gibt und er meine Aufmerksamkeit mit Leichtigkeit auf das Jetzt und Hier lenkt und mir damit immer wieder den nötigen Abstand zu meiner Arbeit gegeben hat; meinem Mann Ole für Deine Liebe und Unterstützung in Höhen und Tiefen – es ist so schön, dass wir dieses wundersame, wundervolle Leben miteinander teilen!

This work was funded and supported by the Deutsche Telekom Stiftung and the Berlin Mathematical School.

Introduction

The advance of modern computing technologies has allowed the development of more complex and realistic models in many scientific fields than was ever possible previously. Within economics, so-called *agent-based models* study systems of boundedly rational, interacting economic agents (e.g., households, firms) by means of computer simulations. The basic idea is that the microscopic behavior of the agents is specified via individual rules of behavior and the dynamical evolution of the whole system as well as related macroscopic quantities (e.g., prices, unemployment rates) is iteratively determined by the computer (e.g., Tesfatsion and Judd, 2006). These kinds of models have many desired features that traditional economic models are lacking (e.g., Epstein and Axtell, 1996; Farmer and Foley, 2009; Gallegati and Richiardi, 2009; Tesfatsion, 2006). However, their high complexity makes them difficult to understand and analyze (e.g., Dawid and Fagiolo, 2008). Up to now, there are few approaches that go beyond traditional “look and see” analyses and it is in this context that the overall goal of this thesis, to contribute to the development of approaches to the formal description and analysis of agent-based models, is to be understood.

Specifications for Agent-Based Models of Exchange

More specifically, in the first part of this thesis, we discuss the need for formal approaches to agent-based modeling and we present a functional framework for specifying agent-based models of exchange developed together with colleagues from the Potsdam Institute for Climate Impact Research, Germany, the Centre d’Économie de la Sorbonne, Université Paris 1 – CNRS, France, and the Institute for Software Systems, Hamburg University of Technology, Germany (see Botta et al., 2013, 2011). This work builds on the observation that, from a programming perspective, agent-based models are obtained through *exploratory programming*. This means that precise specifications for the agent-based computer model, i.e., precise mathematical requirements on the accuracy of solutions delivered by the model, are not available, typically because the modeled problems are not well understood. As prototype implementations are difficult to explain, communicate and study, they are usually accompanied by other model descriptions. These additional model descriptions consist most often of narratives, solitary mathematical equations, and, rarely, pseudo-code. In practice, these additional model descriptions neither allow the identification of programming errors nor the re-implementation nor the analysis of the model since the description of a model in natural language is inherently ambiguous, i.e., it can

lead to several possibilities of interpretation and translation into natural language (e.g., Hofmann, 2007; Ionescu, 2008; Ionescu et al., 2009; Wolf, 2009; Wolf et al., 2013).

The developed functional framework for specifying agent-based models of exchange is based on the agent-based model of exchange by Herbert Gintis (2006) and represents a first attempt at introducing mathematical specifications into computational economics. We applied the framework to the Gintis model and compared the resulting reimplementations with the simulation results of the original work. This work showed that it is possible to derive consistent model reimplementations on the basis of the framework (see Botta et al., 2013).

The first part of this thesis aims to concentrate on the basic structure of the framework necessary to a) clearly formulate the problem of price formation addressed by agent-based models of exchange, b) expose in this way the relationship of agent-based models of exchange with more traditional economic theory, and c) discuss how model analysis can be supported on the basis of the framework.

Metastability in Stochastic Evolutionary Games

While the first part of this thesis concerns a formal description of agent-based models, the second part (in particular, Chapters 4 – 6) relates to their analysis. Since the model of Gintis is too complex to use as a starting point, this part focuses on stochastic evolutionary games, in which a finite number of individuals play repeatedly a game and can update their behavioral choice regularly on the basis of their observations or experiences in previous interactions with other individuals. In this sense, stochastic evolutionary games can be understood as simple agent-based models. Mathematically, these games can be described as Markov chains that fulfill certain conditions on their transition probabilities. This simple structure is exploited in the second part of the thesis.

More specifically, the second part of this thesis focuses on *metastability* in stochastic evolutionary games. This is interesting because metastability is a dynamic property that many of the stochastic evolutionary games share. It means that their sample paths exhibit long periods of stasis near one population state which are infrequently interrupted by switching events after which the sample paths stay close to a different population state, again for a long period of time. In stochastic evolutionary game theory, this property is also called *punctuated equilibrium* (Young, 1998, 2006).

Metastability has been characterized in the literature as a favorable property of stochastic evolutionary games (e.g., Young, 1998, 2006), not least because it might lead to a different perspective on modeling conventions and on the problem of equilibrium selection (Jaeger, 2008, 2012). In general, conventions are social norms or customs; that is, they constitute regularities or uniformity of behavior in populations of individuals in situations in which an individual has several behavioral alternatives to choose from. In these situations, a convention represents the traditional or ordinary way to act. Examples are rules of the road, codes of dress, greeting customs, forms of money and credit, contract standards, etc. Since Lewis (1969), conventions have traditionally been modeled as Nash equilibria of coordination games. However, since coordination games in general have multiple Nash equilibria, the question arises as to

how individuals select a certain convention. This problem of *equilibrium selection* is a central topic in game theory and numerous refinements of the Nash equilibrium solution concept have been developed to argue why a particular Nash equilibrium should be singled out by the players of the game. In fact, so many refinements have been proposed that almost any Nash equilibrium can be justified by some refinement (Binmore, 1995). This leaves the approach to equilibrium selection via refinements of the Nash equilibrium solution concept unsatisfactory. In stochastic evolutionary games with metastable dynamics, the particular population states near where the sample paths reside for a long time can be interpreted as conventions (Young, 1998, 2006). Thus, instead of fixing a single convention, conventions in stochastic evolutionary games can change – just as in the real world – and seeing conventions through metastable glasses might in this context shed new light on how to model and interpret conventions.

The methods used so far in stochastic evolutionary game theory, however, do not characterize the dynamics of the evolutionary games with respect to this property. More specifically, the analysis of evolutionary game models exclusively focuses on equilibrium selection, equilibria of mean dynamics, or the determination of stochastically stable states (Benaim and Weibull, 2003; Ellison, 2000; Foster and Young, 1990; Hofbauer and Sigmund, 2003; Kurtz, 1970; Sandholm, 2010, 2011; Weibull, 1995; Young, 1993a, 1998, 2006). All of these approaches are not able to thoroughly describe metastable dynamic behavior of the considered evolutionary processes.

This is in contrast to physics and chemistry where there has been much research in the last century on the mathematical description and analysis of metastability (for a short historical overview see, e.g., the introductory chapter of Bovier, 2009). With regard to stochastic evolutionary games, the approach to characterizing metastability via spectral properties of the transfer operator or generator of the Markov chain (e.g., Bovier et al., 2001, 2002; Davies, 1982a,b; Dellnitz and Junge, 1999; Deuffhard et al., 2000; Freidlin and Wentzell, 1984; Huisinga, 2001; Huisinga et al., 2004; Schütte and Huisinga, 2003) seems especially promising since this approach is general enough to be applicable to a wide range of models. More specifically, the approach we are pursuing in this thesis originates in the study of chemical conformations of large biomolecules. In a conformation, the large-scale geometric structure is preserved while on smaller scales the biomolecular system may very well fluctuate. The function of biomolecules depends on their dynamic characteristic to undergo statistically rare switches between these conformations. It has been shown that it is possible to interpret conformations of biomolecules as metastable subsets of state space, and function analytic approaches to the characterization of metastability as well as the numerical identification of metastable subsets of state space have been developed (e.g., Deuffhard et al., 2000; Huisinga, 2001; Huisinga and Schmidt, 2006; Schütte, 1998; Schütte et al., 1999; Schütte and Huisinga, 2003). Moreover, this approach also allows the construction of Markov models of reduced complexity and the assessment of their approximation errors (see Djurdjevac et al., 2010; Sarich, 2011; Schütte et al., 2011; Schütte and Sarich, 2013).

The central aim of the second part of this thesis is thus to build on these existing metastability approaches to present a novel approach to the analysis of stochastic evolutionary games. More specifically, we present characterizations of metastability and relate these to stochastic evolutionary games as well as to existing approaches to their analysis, in particular to the notion of stochastic

stability. We give an extensive account on the construction of Markov models of reduced complexity that approximate their essential dynamic behavior, elaborate on their relationship to the original Markov chain, and assess their approximation quality. Lastly, we present a simulation-based algorithmic strategy to identify metastable sets of population states, which is well-suited for stochastic evolutionary games.

Although the second part of this thesis focuses on stochastic evolutionary games, and in fact, often only on stochastic evolutionary games with reversible dynamics, it is especially the construction of models of reduced complexity that takes us full circle back to agent-based models as the motivation of this work. Agent-based models are in general so complex that models of reduced complexity seem to be needed – not only to see the wood for the trees but also to allow the application of advanced analytic methods (e.g., to assess the resilience of a metastable set in consideration, which was, in fact, the aim of the “PATRES: Pattern resilience” project funded by the European commission; see the monograph Deffuant and Gilbert, 2011). There is, of course, some way to go from reversible stochastic evolutionary games to agent-based models in general. But, as explained above, stochastic evolutionary games can be understood as simple agent-based models and therefore lend themselves well as a starting point to investigate approaches for the analysis of agent-based models in general. Thus, what regards the analysis of agent-based models in general, we see the contribution of this thesis in providing a proof of concept for the applicability of the presented metastability approach to agent-based models, thereby laying the foundation for future research into this field.

Road Map

- In Chapter 1, “Specifications for Agent-Based Models of Exchange”, we introduce agent-based modeling and explain the motivation as well as the expectations behind this approach to economic modeling. We discuss the need for specifications for agent-based models. We present the basic structure of the functional framework to specify agent-based models of exchange. This presentation focuses on exposing the relationship of agent-based models of exchange with evolutionary game and general equilibrium theory. More specifically, we show how agent-based models of exchange can nearly be represented as evolutionary games in the sense of Part II of this thesis. By doing so, we explicitly show that they are Markov processes – a fact that is often considered to be only theoretically possible (but not truly realizable). In addition, we use the introduced framework to formulate precisely the research question behind agent-based models of exchange and we discuss how it enables rigorous model analysis and numerical investigation.

Please note that Chapter 1 uses a different notation than the rest of the thesis. More specifically, the framework presented in Chapter 1 is written using a *functional* notation, which is better suited for writing specifications – see Section 1.3.

- Chapter 2, “Mathematical Preliminaries”, introduces the necessary mathematical concepts and notation used in the second part of this thesis. This

comprises basics of probability theory (Section 2.1) and of the theory of discrete-time Markov chains (Section 2.2).

- In Chapter 3, “Stochastic Evolutionary Games”, we introduce stochastic evolutionary games and show that they can be understood as discrete-time Markov chains. Subsequent chapters build heavily on this fact. We discuss, in addition, standard methods of their analysis, such as Nash equilibria and equilibrium selection, deterministic approximation, as well as stochastic stability analysis. We put these standard methods into the context of metastability analysis and show that they are not able to describe metastable dynamic behavior of evolutionary games.
- Chapter 4, “Characterization of Metastability”, presents two different characterizations of metastability – one being based on the notion of hitting times, the other on transition probabilities. Using the notion of *models of evolution with noise* (introduced in Ellison, 2000) to facilitate stochastic stability analysis in stochastic evolutionary games, we relate both metastability characterizations to stochastic evolutionary games.
- In Chapter 5, “Markov State Modeling”, we present the Markov state modeling approach to the approximation of Markov processes and apply it to stochastic evolutionary games. The basic idea of Markov state models is to approximate the original Markov process by a Markov chain on a small finite state space. More specifically, a Markov state model is defined as a Markov chain whose state space consists of the dominant metastable sets of the original Markov process and whose transition rates between these macrostates are given by the aggregate statistics of jumps between the metastable sets. The advantage of this approach in the context of complex models with large state spaces such as agent-based models is that the transition probabilities can be estimated on the basis of simulated short-term trajectory data. Moreover, it has been shown that Markov state models have good approximation properties if metastability is inherent in the system of interest. Specific for our focus is that we give an extensive account of the construction of these Markov state models for discrete state spaces. We elaborate on the relationship between the original Markov chain and its Markov state model approximation. Lastly, we discuss their approximation quality in terms of stochastic stability, with respect to the propagation of probability distributions in time, as well as with respect to the approximation of eigenvalues.
- In Chapter 6, “Identification of Metastable Sets for Agent-Based Evolutionary Models”, we use the approximation results of the previous chapter for an algorithmic approach to identify metastable sets of population states in stochastic evolutionary games. The algorithmic approach is well suited for stochastic evolutionary games because it exploits the fact that they depend on a noise parameter that controls the metastability of the model. Moreover, the identification algorithm preserves stochastic stability. It is based on simulated trajectory data only and seems therefore especially interesting in the context of agent-based modeling since agent-based models usually lack model specifications. One limitation of the approach, however, is that it depends on the original stochastic evolutionary game to

be reversible. We discuss this limitation and give an outlook for further research in this direction.

- The Summary briefly reviews the most important landmarks encountered on our way and closes with outlining possible future research destinations.

Part I

Formalization

Chapter 1

Specifications for Agent-Based Models of Exchange

1.1 Introduction to Agent-Based Modeling

Agent-based modeling is a relatively new modeling approach in economics which studies systems of boundedly rational, interacting economic agents (e.g., households and firms) by means of computer simulations. The basic idea is that the microscopic behavior of the agents is specified via individual rules of behaviour while the dynamical evolution of the whole system and related macroscopic quantities (e.g., prices and unemployment rates) are iteratively determined by the computer. In this sense, agent-based models (ABMs) can comply with the wish for microfoundations of macroeconomics: boundedly rational economic agents and their interactions can be modeled in an intuitive and realistic way. The resulting economic processes can be observed via simulations – does Adam Smith’s so-called “invisible hand”, that is, economic coordination, emerge in decentralized market economies? How are prices formed and how do they evolve? How do financial crises arise? These are examples of the central economic questions argued to be better addressed by agent-based models than by the mathematical models used in traditional general equilibrium theory and macroeconomics, since the latter are only mathematically tractable if the behavior of the agents is restricted in a fundamental way (see Section 1.4.1). Agent-based models, in contrast, allow the study of emerging macroscopic structures on the basis of more realistic and diverse behavioral rules (Epstein and Axtell, 1996; Farmer and Foley, 2009; Gallegati and Richiardi, 2009; Tesfatsion, 2006; Tesfatsion and Judd, 2006).

Agent-based models differ in their scientific objectives and this is reflected in their complexity. On one end of the spectrum lie those models that attempt to model economic processes as realistically as possible. Their objective is to reproduce empirical data and to derive normative statements as to the consequences of economic policies. In such models, there is a huge number and diversity of agents, such as consumers, decision-makers, bank and corporate managers,

which interact according to the rules determined by the institutions they are embedded in. Up to now, there are only a few models with this claim to reality, one of which is the EURACE model developed to provide macroeconomic policy advice for the EU (Deissenberg et al., 2008).

On the other end of the spectrum are more theoretically motivated agent-based models. Their underlying research question often relates to existing economic theory and addresses the outcomes of interaction processes with certain properties as well as of different behavioral rules assumed for the agents. In a general equilibrium model, for example, the Walrasian auctioneer is a fictive person that determines prices given the demand and supply of the agents (also see Section 1.4.1). If we do not want to rely on such an artificial construction as the Walrasian auctioneer, what kind of interactions must at least be specified in order for the resulting agent-based model to be iterable (Tsfatsion, 2006)? What prices develop? The agent-based models that can be found on this end of the spectrum are usually less complex than the ones with more claim to reality. The agent-based model of exchange by Gintis (2006), which was the point of departure for the developed functional framework for specifying agent-based models of exchange (see Section 1.4), is an example of such an agent-based model with a theoretically motivated research question.

1.2 Exploratory Programming and The Need for Specifications

While agent-based models have many desired features that traditional economic models lack, they also lead to new challenges. Indeed, there are numerous agent-based economic models (Tsfatsion, 2014; Tsfatsion and Judd, 2006) but no precise, established notion of an agent-based model so that most of the models are developed anew from scratch and assume (seemingly) arbitrary, ad hoc behavioral rules. Moreover, their inherent complexity renders them difficult to understand and analyze.

These challenges relate to the fact that agent-based models are computer-based models that are *not* developed to solve well-defined problems. In the case of scientific computing and engineering, for instance, the problems to be solved are often well understood and can be stated precisely. On this basis, computer-based algorithms can be developed and applied to deliver (usually approximate) solutions to the problems. Since the problems to be solved are well-defined, model developers and implementers are able to design crucial experiments. Negative outcomes of such crucial experiments unambiguously indicate errors in the algorithm or in its implementation. In this way, it is possible to control the accuracy of the computer-based solutions.

For example, let $r2 : \text{Real} \times \text{Real} \times \text{Real} \rightarrow \text{Complex}$ be a program that implements an algorithm for finding the roots of quadratic equations of the form $ax^2 + bx + c = 0$. In order for $r2$ to implement such an algorithm, we require that

$$\forall a, b, c \in \text{Real} : a \neq 0 \wedge y = r2(a, b, c) \Rightarrow |ay^2 + by + c| \leq \varepsilon, \quad (1.1)$$

where ε is an appropriate residuum upper bound. Now if there are $a^*, b^*, c^* \in \text{Real}$ such that $|a^*y^2 + b^*y + c^*| \geq \varepsilon$ where $y = r2(a^*, b^*, c^*)$, we unambiguously

know that $r2$ does not fulfill requirement (1.1). In this case, two interpretations are possible; either

- $r2$ is an incorrect implementation of a correct algorithm for finding the roots of a quadratic equation of the form $ax^2 + bx + c = 0$, or
- the utilized algorithm for finding the roots of a quadratic equation of the form $ax^2 + bx + c = 0$ is not correct.

In such a situation, the second possibility can be ruled out if the algorithm can be proved to be correct.

In software engineering, requirements like Eq. (1.1) are usually called *specifications*. That is, in the just given example, Eq. (1.1) is a specification for computer-based models for finding the roots of a quadratic equation of the form $ax^2 + bx + c = 0$. The example, moreover, demonstrates that

1. It is helpful to distinguish three notions: the specification, the algorithm, and the implementation.
2. Specifications can be stated in a clear mathematical notation and are often written in terms of equations. In many cases, they are called problem equations, governing equations, or governing laws.
3. Though expressed in mathematical notation, specifications necessarily depend on the computing architecture. In Eq. (1.1), for example, Real represents a computing architecture dependent set of floating point numbers and ε is a value that depends on the sets Real and Complex , and thus, depends on the computing architecture as well. It is often convenient to replace Eq. (1.1) by

$$\forall a, b, c \in \mathbb{R} : a \neq 0 \wedge y = r2(a, b, c) \Rightarrow ay^2 + by + c = 0; \quad (1.2)$$

that is, to trade precision of specifications for intuitiveness. This is the approach taken in the following sections of this chapter. Remember, however, that Eq. (1.2) is just a convenient stand-in for Eq. (1.1). All crucial experiments and proofs of correctness have to be based on the original specification.

4. Algorithms are mostly written in pseudo programming languages (Aho et al., 1982). Implementations are written in programming languages. Algorithms are more difficult to read and analyze than specifications, but still more readable than implementations.
5. For complex computer-based models it is often not feasible to derive, possibly automatically, implementations from specifications, or even to prove that implementations are correct in the sense that they fulfill the specifications. In such a situation, it is usually more realistic to implement computer-based models by combining software components that have been proved to be correct, or, at least, for which crucial experiments can be set up.

Agent-based models, in contrast, are computer-based models developed through *exploratory programming*. As an approach to software engineering, exploratory

programming is characterized by the fact that precise specifications for the computer model are not available. Typically, this results from modeled problems that are not well understood – as is often the case in agent-based modeling. Instead of formal specifications, implementations are accompanied by additional model descriptions such as narratives, solitary mathematical equations, and, rarely, pseudo-code. Note that these additional model descriptions are necessary although implementations are often considered as unambiguous, precise model descriptions themselves. This is because implementations are not very useful for explaining, communicating, and studying computer-based models. They expose too many details, lack conciseness and readability, and are thus almost inaccessible to non-programmers.

It has been argued that, in principle, these additional descriptions could be sufficient and precise enough to allow for model validation, verification as well as model reimplementation. Based on this view but aware of the inherent difficulties, there are initiatives that aim at standardizing the description of agent-based models in a way that facilitates comprehensibility as well as reimplementation, most notably the ODD protocol (Overview, Design concepts, Details; Grimm et al., 2006, 2010). Unfortunately, there have been no reported studies with regard to the reimplementation of agent-based models on the basis of, for instance, the ODD protocol. While we recognize the advantages of such a standardized description of agent-based models, we argue that in practice narrative descriptions always allow multiple possibilities of interpretation and translation into formal languages (Hofmann, 2007; Ionescu, 2008; Ionescu et al., 2009; Wolf, 2009; Wolf et al., 2013). Consequently, for consistent, independent model reimplementations as well as model analysis and dissemination, mathematical specifications are necessary and deriving such specifications from the prototype implementations and narrative descriptions represents the next step in agent-based model development.

Own work on an agent-based model of exchange put forward by Herbert Gintis called our attention to this problem setting. More specifically, Gintis (2006) proclaimed the discovery of a disequilibrium adjustment mechanism that explains price formation *without* the presence of a Walrasian auctioneer, the central authority generally assumed in mainstream economics to impose prices on players. His results are obtained by agent-based modeling and are to be seen, as Gintis (2006, p. 14) puts it, as “empirical rather than theoretical: we have created a class of economies and investigated their properties for a range of parameters”. As the understanding of price formation is a long-standing puzzle in economics, this work is relevant and intriguing. Its presentation, however, leaves a number of non-trivial questions open: Which class of economies is actually considered by the model presented by Gintis (2006)? Is it possible to express the properties investigated by Gintis (2006) formally? What exactly is the relation between the proposed agent-based model and general equilibrium models, or, in the words of Gintis (2006, Abstract), what exactly is the “underlying Walrasian model”?

The relevance of the problem together with these open questions independently motivated two groups of researchers, one at the Potsdam Institute for Climate Impact Research (see Botta et al., 2013, 2011), the other at Chalmers Technical University, Sweden (see Evensen and Mårdin, 2009), to do something which should be routine, but is hardly ever done: to reimplement the model and to try to reproduce the results reported by Gintis (2006). Although the model

seems to be simple enough to be easily understandable and reimplementable, initial attempts based only on the description given by Gintis (2006) failed. Following this, the author kindly provided the source code. Both groups were now able to discover several ways in which his implementation diverged from their interpretation of the description given in the paper and thus diverged from their reimplementations. This exemplifies the problematic ambiguity left open by the narrative provided by Gintis (2006) – an ambiguity that is often ignored in agent-based economics as scientists tend to believe that a narrative together with a few mathematical equations are sufficient descriptions for their models.

Based on these experiences, we set out to develop a formal framework to specify and study agent-based models of exchange – a framework that allows to specify agent-based models of exchange of the kind considered by Gintis (2006) in a concise and unambiguous form (see Botta et al., 2011). As formalization always involves abstraction, it is obvious that the framework we present is more general than the specific model presented by Gintis (2006). In addition to this specific model, it captures the class of models with the same structure as this specific model (for more on the process of formalization, its advantages and characteristics, see, e.g., Wolf, 2009). The framework has been applied to the model presented by Gintis (2006), see Botta et al. (2013). On the basis of the framework, it is possible to derive consistent model reimplementations, to relate the model to existing economic theory, as well as to allow rigorous model analysis and validation (see Botta et al., 2013, 2011).

It is the aim of this chapter to present the basic structure of the framework necessary to

1. clearly formulate the problem of price formation addressed by agent-based models of exchange like the one presented by Gintis (2006), and in this way to
2. expose their relationship with traditional economic theory, specifically with general equilibrium and game theory, as well as to
3. explicate how model analysis can be supported with the framework.

Thus, the focus of this chapter is to demonstrate how mathematical specifications support model analysis and problem formulation.

1.3 Basic Notions and Notation

The notation used in this chapter formulates economic notions mainly in terms of *functions* and differs in this respect from the notation in Part II of the thesis, which focuses instead on vectors. This section introduces and motivates such a functional notation and formulates the most important economic notions needed to specify models of exchange. In presenting the basic structure of the framework, we will also use basic notions from game theory and probability theory, which we are not going to introduce here. Instead we refer in the footnotes to the specific definitions given in Chapter 2 and 3.

In the following, we denote by A and G the set of agents and the set of goods, respectively. Both sets are assumed to be finite and represent the two most fundamental economic notions in the framework. Most other economic

notions which are (as just mentioned) formulated mainly as functions, take values in A or G . For example, the notion of a stock is represented as a function that associates to goods in G non-negative real numbers; that is, a stock is a function of type $G \rightarrow \mathbb{R}_{\geq 0}$ and we write

$$q : G \rightarrow \mathbb{R}_{\geq 0}.$$

to denote that q is a stock. Similarly, allocations are functions which associate stocks to the agents in A and we thus write

$$x : A \rightarrow (G \rightarrow \mathbb{R}_{\geq 0})$$

to posit that x is an allocation. Along the same lines, prices, utilities, and utility profiles are represented as functions of type $G \rightarrow \mathbb{R}_{> 0}$, $(G \rightarrow \mathbb{R}_{\geq 0}) \rightarrow \mathbb{R}$, and $A \rightarrow ((G \rightarrow \mathbb{R}_{\geq 0}) \rightarrow \mathbb{R})$, respectively. In subsequent sections, we will use the abbreviations Q , P , and U to denote in short form the set of stocks, the set of prices and the set of utilities, respectively.

Writing allocations as functions of type $A \rightarrow (G \rightarrow \mathbb{R}_{\geq 0})$ exemplifies our use of *currying* in the formulation of the framework. Currying, named after Haskell Brooks Curry, refers to the representation of functions of more than one variable as functions of one variable that return function values. In the former case, an allocation x would, for instance, be written as $x : A \times G \rightarrow \mathbb{R}_{\geq 0}$, while the curried form is $x : A \rightarrow (G \rightarrow \mathbb{R}_{\geq 0})$. Note that \rightarrow is understood to be right-associative¹, that is, a function $A \rightarrow G \rightarrow \mathbb{R}_{\geq 0}$ is understood to be of type $A \rightarrow (G \rightarrow \mathbb{R}_{\geq 0})$.

Furthermore, we denote function application by juxtaposition: $x a : G \rightarrow \mathbb{R}_{\geq 0}$ is the stock of the agent a according to the allocation x and $x a g \in \mathbb{R}_{\geq 0}$ refers to the quantity of good g according to the stock $x a$. This follows the common usage in category theory where “evaluation is a special case of composition” (Lawrence and Rosebrugh, 2003, p.6). Note that we denote multiplication of two real numbers r_1, r_2 by $r_1 * r_2$ in order to differentiate function application denoted by juxtaposition from multiplication, which is often denoted by juxtaposition as well.

Example 1.1. Let $A = \{a_1, a_2, a_3\}$, $G = \{g_1, g_2, g_3\}$. Let q_1, q_2, q_3 be defined by

$$q_i g_j = \begin{cases} 2 & \text{if } i = j \\ 0 & \text{otherwise,} \end{cases} \quad (1.3)$$

and consider the allocation x_0 in which agent a_i has stocks q_i , that is, $x_0 a_i = q_i$ for $i = 1, 2, 3$. A prominent example of a utility profile u is the one proposed by

¹Note that the notion of right-associativity we use here refers to *notational associativity*. Notational associativity is used in programming languages to declare how to interpret terms in which the same operator occurs without parentheses, e.g., $A \rightarrow G \rightarrow \mathbb{R}_{\geq 0}$. If the operations are grouped from the right, as we do in the case of \rightarrow , e.g., $A \rightarrow G \rightarrow \mathbb{R}_{\geq 0}$ is interpreted as $A \rightarrow (G \rightarrow \mathbb{R}_{\geq 0})$, the operator is called right-associative. Note that notational associativity of a specific operator may differ between programming languages. Moreover, the notion is different from the mathematical notion of associativity according to which an operator is associative if the value of an expression containing two or more occurrences of the same operator will not change if parentheses are rearranged. The standard examples for the mathematical notion of associativity are addition and multiplication. Note that an operator which is mathematically associative needs no notational associativity because it can be interpreted either way. In contrast, if the operation is not mathematically associative, its notational associativity needs to be defined in order for expressions to be unambiguously evaluable.

Scarf (1960), in which each agent has the same utility function defined by

$$u_a y = \min_{g \in G} (y_g) / (w_g), \quad (1.4)$$

where w is a function of type $G \rightarrow \mathbb{R}_{>0}$. Let us consider, for instance,

$$w = \frac{1}{|A|} * \sum_{a \in A} x_0 a, \quad (1.5)$$

where x_0 is the just defined allocation, that is, $w_{g_i} = 2/3$ for each $i = 1, 2, 3$. Then the utility of each agent given the allocation x is 0, while it is 1 for the allocation x defined by $x_a g = 2/3$ for all $a \in A, g \in G$.

We also use the notion of a *list* and the related functions `elem` as well as `fold`. Lists are recursively defined data types (Bird, 1998). The empty list is denoted by `[]`. To refer to the k th element of a list xs , we write $xs!!k$. In addition, $(x : xs)$ refers to a list with first element x and subsequent list of elements given by xs . The boolean function `elem` is used to test membership in lists:

$$\text{elem} : X \rightarrow \text{List } X \rightarrow \text{Bool}, \quad (1.6a)$$

$$\text{elem } x \text{ } xs = \begin{cases} \text{true} & \text{if } \exists k \in \mathbb{N} \text{ such that } xs!!k = x, \\ \text{false} & \text{otherwise.} \end{cases} \quad (1.6b)$$

The function `fold` is, like `elem`, a *polymorphic* function. This means that the function does not refer to specific types, but can instead be used for any type. Its type

$$\text{fold} : (X \rightarrow Y \rightarrow X) \rightarrow X \rightarrow \text{List } Y \rightarrow X \quad (1.7)$$

depends on two parameters X and Y . Like many functions that operate with lists, `fold` is defined recursively via pattern matching (Bird, 1998):

$$\text{fold } f \ x \ [] = x, \quad (1.8a)$$

$$\text{fold } f \ x \ (y : ys) = f \ (\text{fold } f \ x \ ys). \quad (1.8b)$$

Example 1.2. Let again $A = \{a_1, a_2, a_3\}, G = \{g_1, g_2, g_3\}$. Let us consider the list $ts = [(a_1, a_2), (a_3, a_1), (a_2, a_3)]$. Notice that $ts \in \text{List } (A \times A)$; in words, ts is a list of agent pairs. Moreover, the first element of the list is $ts!!1 = (a_1, a_2)$ and thus we can write $ts = ((a_1, a_2) : ts')$, where $ts' = [(a_3, a_1), (a_2, a_3)]$. We also have $\text{elem } (a_1, a_2) \ ts = \text{true}$, but $\text{elem } (a_1, a_3) \ ts = \text{false}$. Let furthermore $f : (A \rightarrow Q) \rightarrow (A \times A) \rightarrow (A \rightarrow Q)$ be defined by

$$f \ x \ (a_i, a_j) = x', \quad (1.9)$$

where x' is the resulting allocation given by $x'_a g = x_a g$ if $k \neq i, j$, and

$$x'_a g_l = \begin{cases} x_a g_l - 1 & \text{if } (l = i \wedge x_a g_i \geq 1) \\ x_a g_l + 1 & \text{if } (l = j \wedge x_a g_j \geq 1) \\ x_a g_l & \text{otherwise;} \end{cases} \quad (1.10)$$

$$x'_a g_l = \begin{cases} x_a g_l - 1 & \text{if } (l = j \wedge x_a g_j \geq 1) \\ x_a g_l + 1 & \text{if } (l = i \wedge x_a g_i \geq 1) \\ x_a g_l & \text{otherwise;} \end{cases} \quad (1.11)$$

that is, f can be interpreted as an exchange function according to which the involved agents switch one unit of “their” good, if possible, where we associated goods to agents according to their index. For the allocation x_0 defined in Example 1.1, we can now consider

$$\text{fold } f \ x_0 \ ts = \text{fold } f \ x_0 \ ((a_1, a_2) : [(a_3, a_1), (a_2, a_3)]) \quad (1.12a)$$

$$= f (\text{fold } f \ x_0 \ ((a_3, a_1) : [(a_2, a_3)])) (a_1, a_2) \quad (1.12b)$$

$$= f (f (\text{fold } f \ x_0 \ ((a_2, a_3) : [])) (a_3, a_1)) (a_1, a_2) \quad (1.12c)$$

$$= f (f (f (\text{fold } f \ x_0 \ [])) (a_2, a_3)) (a_3, a_1)) (a_1, a_2) \quad (1.12d)$$

$$= f (f (f \ x_0 \ (a_2, a_3)) (a_3, a_1)) (a_1, a_2). \quad (1.12e)$$

Notice that the list ts is executed from the right to the left. First, a_2 and a_3 exchange one unit of their good, if possible, resulting in the new allocation x' . Next, a_3 and a_1 exchange one unit of their good, if possible, resulting in the new allocation x'' . Finally, a_1 and a_2 exchange one unit of their good, if possible, resulting in x''' , which is the allocation with

$$x''' \ a_i \ g_j = \begin{cases} 0 & \text{if } i = j \\ 1 & \text{otherwise.} \end{cases} \quad (1.13)$$

This notation we just introduced is clearly influenced by category theory and functional programming languages, in particular by Haskell. This becomes especially evident in the use of currying as well as in the use of juxtaposition for function application. However, the framework and subsequent discussion can be understood without mathematical or computing science notions more advanced than those introduced up to this point.

Although the vector-based notation is more common in mathematical economics, we prefer such a functional² notation as it is clearer and closer to programming, and thus better suited for writing specifications. Stocks, for example, are usually represented in economics as vectors in $\mathbb{R}_{\geq 0}^n$ where $n = |G|$. Now, consider a set G' with $G' \neq G$ but $|G'| = |G|$. In this case, the sets $G \rightarrow \mathbb{R}_{\geq 0}$ and $G' \rightarrow \mathbb{R}_{\geq 0}$ are different sets, but their vector representations $\mathbb{R}_{\geq 0}^{|G|}$ and $\mathbb{R}_{\geq 0}^{|G'|}$ are not; that is, both sets coincide with $\mathbb{R}_{\geq 0}^n$ in the vector-based notation. This loss of information does not pose a problem if the information which is missing in the vector-based representation is clear from the context. Specifications, however, aim exactly at making this context visible, not at hiding information. Moreover, the functional notation paves the way for explicitly and concisely formulating “what” notions of interest are, i.e., how to determine and compute them, e.g., through the use of currying and by denoting function application by juxtaposition. As this is the primary goal of specifications, a functional notation as just introduced seems more appropriate than a vector-based one. Building on the formal framework, further approaches to the analysis of agent-based models of exchange can be considered in which case it is probably helpful to translate the necessary notions into the more abstract vector-based notation. In Appendix 6.4, we provide a table of possible translations between the functional notation and the vector-based notation for selective expressions used in the formal framework, which the reader unfamiliar with the functional

notation can consult if necessary.

1.4 A Functional Framework for Specifying Agent-Based Models of Exchange

The central question addressed by agent-based models of exchange such as the one proposed by Gintis (2006) is how prices emerge in a free market economy. This question has fascinated economists for a long time. In fact, it was Adam Smith (1723–1790) who in his famous work “The Wealth of Nations” (1776) was one of the first to discuss the role of prices in a free market economy. In this work, he argues that the pursuit of self-interest in a free market economy, where prices adjust themselves according to the demand and supply of goods, can indirectly promote the good of the society. It is this idea which has permeated modern economic thought and which economists have sought to refine and formalize ever since:

There is by now a long and fairly imposing line of economists from Adam Smith to the present who have sought to show that a decentralized economy motivated by self-interest and guided by price signals would be compatible with a coherent disposition of economic resources that could be regarded, in a well-defined sense, as superior to a large class of possible alternative dispositions. Moreover, the price signals would operate in a way to establish this degree of coherence. It is important to understand how surprising this claim must be to anyone not exposed to this tradition. The immediate “common sense” answer to the question “What will an economy motivated by individual greed and controlled by a very large number of different agents look like?” is probably: There will be chaos. That quite a different answer has long been claimed true and has indeed permeated the economic thinking of a large number of people who are in no way economists is itself sufficient grounds for investigating it seriously . . . it is important to know not only whether it *is* true, but also whether it *could* be true. (emphasis in original, Arrow and Hahn, 1971, p. vi-vii)

It is against this background that the motivation for and the objective of agent-based models of exchange of the kind proposed by Gintis (2006) has to be understood.

Therefore, before we lay out the basic structure of the framework for specifying agent-based models of exchange (Section 1.4.2), we firstly examine the problem of price formation and how it has been dealt with in economics (Section 1.4.1). Lastly, in Section 1.4.3 we discuss the advantages of the framework, with a focus on discussing how the framework can help to precisely formulate the research question of the model of interest and support its analysis.

1.4.1 The Problem of Price Formation

It was Léon Walras (1834–1910) who first put Smith’s idea that the pursuit of self-interest by individuals in a free market economy can lead to socially

²“Functional” here refers to “functions” and not to “functionals” in the sense used in functional analysis.

desirable outcomes in a mathematical setting. Only decades later, beginning in the 1930s, Abraham Wald, Kenneth J. Arrow, Gérard Debreu and Frank H. Hahn further developed Walras' analytical approach into the touchstone model of modern economics, the so-called *general* (or, *Walrasian*) *equilibrium model*.

As their name suggests, the *general equilibrium* is the central notion in general equilibrium models. If an economic system is in general equilibrium, so the idea behind it, demand and supply of its economic agents are compatible with each other. The Walrasian model relies on the introduction of prices in order to decentralize such an allocation problem.

In the following, we will only consider pure exchange economies, that is, economies in which the only economic agents are consumers. This is the most basic setting for general equilibrium models and it is sufficient for our purposes; that is, pure exchange economies are sufficient to formally express the notion of general equilibrium, to discuss the criticism general equilibrium models have received in the last decades, and to put the agent-based modeling approach in this context.

Definition of General Equilibrium

Now, in a pure exchange economy, we say that an allocation $x : A \rightarrow Q$ and prices $p \in P$ are *in general equilibrium* with an initial allocation $x_0 : A \rightarrow Q$ and a utility profile $u : A \rightarrow U$ if

- (optimality conditions) each agent $a \in A$ maximizes utility by choosing the stocks x_a under the prices p , that is, if

$$\forall a \in A : x_a \in \operatorname{argmax}_{z \cdot p \leq (x_0)_a \cdot p} u_a z, \quad (1.14)$$

where $z \cdot p$ denotes the scalar product between z and p ,

$$z \cdot p = \sum_{g \in G} (z \cdot g) * (p \cdot g), \quad (1.15)$$

and $\operatorname{argmax}_{z \cdot p \leq (x_0)_a \cdot p} u_a z$ denotes the set of stocks $y \in Q$ that maximize utility under the budget constraint, which means

$$y \cdot p \leq (x_0)_a \cdot p \quad (1.16a)$$

\wedge

$$u_a z \leq u_a y \quad \forall z \in Q, z \cdot p \leq (x_0)_a \cdot p; \quad (1.16b)$$

and if

- (feasibility conditions) x is a *feasible re-allocation* of x_0 , or, put differently, if total demand equals total supply:

$$\forall g \in G : \sum_{a \in A} x_a g = \sum_{a \in A} x_0_a g. \quad (1.17)$$

The notion of general equilibrium in a pure exchange economy is thus expressed as a relation \mathcal{E} between four functions and we write in short $(x, p)\mathcal{E}(x_0, u)$. The

tuple (x, p) is called a *general equilibrium* of the exchange economy with initial allocation x_0 and utility profile u (see, e.g., Varian, 1992).

Put differently, the $|A|$ equations in (1.14) express an assumption on agent behavior: agents act to maximize their utility, which derives from the consumption of goods that are offered on markets at a given price and that the individual can buy according to her budget constraint. This behavior is called *rational*.

We say that $z \cdot p$ is the *value* of the stock z according to the prices p . Note that the prices p appear in (1.14) only in the budget constraint and as linear factors on both sides of the inequality. It follows that equilibrium prices, if they exist, can only be unique up to a multiplicative factor:

$$(x, p)\mathcal{E}(x_0, u) \wedge \lambda \in \mathbb{R}_{>0} \Rightarrow (x, \lambda * p)\mathcal{E}(x_0, u). \quad (1.18)$$

In other words, what matters in the notion of general equilibrium are *not* the prices of the single goods, but instead the ratios between such prices. It is therefore common to consider *normalized* prices. These are obtained by dividing all prices by the price of a reference good, in economics called the *numéraire*.

Against the background of this assumption on agent behavior, the $|G|$ equations in (1.17) express the basic idea of the notion of general equilibrium, that supply, that is the goods provided in the initial allocation, equals demand, that is the goods in the allocation which maximizes agents' utility under the budget constraints as determined by the initial allocation and the given prices. These feasibility conditions might also be termed as law of conservation of goods.

Example 1.3. Let $A = \{a_1, a_2, a_3\}$, $G = \{g_1, g_2, g_3\}$ and let the utility profile u be defined as in Example 1.1, i.e.,

$$u a y = \min_{g \in G} (y g) / (w g), \quad (1.19)$$

where $w : G \rightarrow \mathbb{R}_{\geq 0}$. In Appendix A of Botta et al. (2011) it is proven that if the total quantity of each good in x_0 is positive and w is a multiple of these total quantities, i.e.,

$$w = \lambda * \sum_{a \in A} x_0 a \quad \text{for some } \lambda \in \mathbb{R}_{>0}, \quad (1.20)$$

any price $p \in P$ and allocation $x : A \rightarrow Q$ with

$$x a = \frac{(x_0 a) \cdot p}{w \cdot p} * w \quad (1.21)$$

are in equilibrium with (x_0, u) . Thus, with the specific x_0 and function w given in Example 1.1, and prices p defined by

$$p g_j = 1 / (w g_j), \quad (1.22)$$

we have that (x, p) is in equilibrium with (x_0, u) , where x is the allocation with $x a_i g_j = 2/3$ for every $i, j = 1, 2, 3$. The value $(x a) \cdot p$ of the stocks $x a$ in the equilibrium allocation x under the prices p is 3 for each $a \in A$

The theory of general equilibrium reformulates the conditions of general equilibrium into a fixed-point problem in order to provide sufficient conditions for its existence as well as to derive computational methods for finding an equilibrium. We refer the reader to Kehoe (1991) for an overview of these computational methods for finding equilibria.

Criticism

The general equilibrium model has been criticized on several grounds. The main arguments concern

- the behavioral assumptions: Beginning with Simon (1957), the plausibility of the rationality assumption on agent behavior has been questioned again and again. Arguments include, for example, that agents only have incomplete information about the utility of goods, and that humans only have limited computing capabilities. Therefore, they may act purposefully but not necessarily optimally in case of difficult optimization problems (Gallegati and Richiardi, 2009; Kirman, 2008).

Moreover, the careful reader will have noticed that the definition of general equilibrium does not involve explicit agent interactions. In fact, the pure exchange economy considered in this definition is rather characterized by the absence of other economic agents and processes, such as firms and production processes, than by what exchange is and how it takes place. Consumers take decisions only on the basis of prices that are externally given; that is, the model does not explain where prices actually come from and how they become established. Within the theory of general equilibrium, a fictional character, the so-called *Walrasian auctioneer*, is assumed to set prices and adjust them according to the demand and supply of the agents towards their equilibrium prices. However, besides lacking any real counterpart, the dynamical system implied by this Walrasian *tatōnnement*³ mechanism has been shown to be stable only under restricted and implausible conditions (Bala and Majumdar, 1992; Fisher, 1983; Saari, 1985). Moreover, the absence of explicit agent interactions has also been claimed as the reason for the non-uniqueness of the general equilibrium (see the next point).

- The (non-)uniqueness of the general equilibrium: While existence of a general equilibrium is granted under fairly general conditions (the first proof for an economy with production given in Arrow and Debreu, 1954), uniqueness requires very restricting conditions on agent behavior (Debreu, 1974; Mantel, 1976; Sonnenschein, 1972). Put differently, the usual assumptions on utility functions (e.g., continuity, monotonicity, strict quasi-concavity) are not sufficient to guarantee the uniqueness of the general equilibrium and in this sense do not determine the macroscopic behavior of the economy. A prominent example of a utility profile which supports multiple equilibria is the one proposed by Scarf (1960), see Example 1.3.

This has far-reaching consequences, especially for macroeconomic modeling: uniqueness of the equilibrium is a necessary premise for the application of macroeconomic methods. The restricting conditions under which uniqueness is guaranteed have led to the critical and artificial construction of the *representative agent*, a construction that is formally unjustifiable (Kirman, 1992) but is ubiquitously used in mainstream macroeconomics because of its analytical simplicity (Stern, 2006) and the lack of alternatives (Edenhofer, 2010).

³French for “groping”.

It is in response to this criticism that the agent-based models of exchange have to be understood. More specifically, the aim is to lay out plausible rules of trading interactions that might

1. explain how equilibrium prices become established, and
2. ensure that equilibrium allocations are reached under these equilibrium prices.

This research objective itself is not quite new; the agent-based modeling approach, however, is. In fact, there are game-theoretic approaches to general equilibrium dating already from the 1980s and early 1990s (e.g., Gale, 1986a,b, 2000; Osborne and Rubinstein, 1990; Rubinstein and Wolinsky, 1985). This comes as no surprise as game theory is the economic theory of interactions (Osborne and Rubinstein, 1994). Gintis himself sees agent-based modeling in this game-theoretic tradition when he states: “An agent-based model is a computer simulation of the repeated play of a game” (Gintis, 2006, p. 1). Game-theoretic approaches, however, often focus on computing for a given game the Nash equilibria or some refinement thereof and on showing that these equilibria correspond to the Walrasian ones. This focus has the limitation that games have to stay quite simple in order to be analyzable. Moreover, no dynamic mechanism explaining price formation has yet been proposed. The aim of agent-based modeling, in contrast, is not to determine any Nash equilibrium or the like, but to simulate the outcome of interactions and numerically investigate whether the proposed interactions might lead to a general equilibrium. This approach has the advantage that interactions need not be restricted to be analytically simple. This deserves mentioning at this point as, firstly, the model proposed by Gintis (2006) can be better understood in this context. Secondly, it points to a relationship between agent-based modeling and economic theory, which can be exploited as a starting point for the development of a theoretical basis of agent-based modeling (e.g., Izquierdo, 2008; Littman, 1994).

1.4.2 Basic Structure of the Framework

In this section, we introduce the functional framework for specifying agent-based models of exchange (Botta et al., 2013, 2011). As the framework itself is quite comprehensive with lots of details and its application to the specific agent-based model of exchange proposed by Gintis (2006) adds even more to these details, only a basic structure is going to be presented. The aim is to see the forest for the trees. The focus that guided the selection of notions to be included in this basic structure has been to show how the research problem described in 1. and 2. above can be formulated precisely on the basis of the framework. In addition, for the reasons given in the last paragraph, that is, to expound on the relationship between game theory and agent-based modeling in the way used by Gintis (2006), the presentation given here differs from the one given by Botta et al. (2011) and Botta et al. (2013). In particular, we show how the agent-based models of exchange can be understood as evolutionary games in the sense used in Part II of the thesis. That way, we establish a link between the two parts of this thesis. It is for this reason that the terminology differs from the one used in Botta et al. (2013, 2011) as it uses more terms from game theory. It should not, however, pose any difficulties to relate the basic structure of the framework as

presented here to the functional framework in Botta et al. (2011) as well as to its application in Botta et al. (2013) to the model presented by Gintis (2006), and to fill in any details of personal interest. Moreover, we refer in the footnotes to specific probability and game theoretic definitions provided in Part II.

Overview and Preliminaries

The agent-based models of exchange considered here are based on the one proposed by Gintis (2006). They consist of:

1. a time-discrete model for the *evolution of prices*, which models in general a learning process and depends at each iteration on the outcome of
2. a *trading game*, in which the agent-specific prices determine the trading strategies of the agents. The outcome of the trading game under these strategies determines the *fitness*⁴ of the prices, on which in turn the prices of the next time step depend.

In what follows, we formulate what this precisely means. In particular, we will see that in terms of evolutionary games (Chapter 3), the learning process in the model of evolution of prices is a kind of revision protocol⁵ and the trading game can be understood as a population game⁶. We will often exemplify the given specifications on the agent-based model of exchange presented by Gintis (2006), which we will also refer to as the *Gintis model*⁷.

Model of Price Formation

In agent-based models of exchange, the evolution of prices is time-dependent and usually non-deterministic. In general, time-dependent, non-deterministic states of a system of interest, such as prices, can be understood as a stochastic process⁸. More specifically, the evolution of states can be represented by a discrete-time Markov process $(X_t)_{t \in \mathbb{N}}$ with sample space Θ and state space Z . It is defined in terms of an initial state $z_0 \in Z$ and a transition kernel

$$tr : Z \rightarrow \mathcal{B}(Z) \rightarrow [0, 1], \quad (1.23)$$

where $\mathcal{B}(Z)$ denotes a sigma-algebra of Z , which means that 1) $tr z : \mathcal{B}(Z) \rightarrow [0, 1]$ is a probability measure for every $z \in Z$, and 2) for every $A \in \mathcal{B}(Z)$ the

⁴Fitness and utility are related notions. While fitness is mostly used in computing or biology-related disciplines, utility is more common in economics. We use the term fitness here to differentiate between the notion of utility as used in the trading game.

⁵Def. 3.4.

⁶Def. 3.3.

⁷The careful reader will notice at some points in the following presentation that the model as presented in the paper (Gintis, 2006) differs from the implementation as provided by the author. This raises, of course, the question of what do we actually refer to by “the” Gintis model – the implementation or its natural language description given in Gintis (2006). However, as the formal framework captures both and the differences are small details, it should not pose a problem to refer to both, the implementation and the description, as “the” Gintis model. If necessary, we explicitly state which of the two we consider.

⁸Def. 2.1.

map

$$Z \rightarrow [0, 1] \quad (1.24a)$$

$$z \mapsto tr\ z\ A \quad (1.24b)$$

is measurable⁹. The value of $tr\ z\ A$ can be interpreted as the probability of being in subset A of state space after one time step given the current state is z . We can think of such a stochastic process as a random variable of type $\Theta \rightarrow \mathbb{N} \rightarrow Z$. This way, we write (with obvious abuse of notation)

$$z : \Theta \rightarrow \mathbb{N} \rightarrow Z \quad (1.26)$$

to denote that z is a time-dependent, non-deterministic state, defined by

$$z\ \theta\ t = X_t\ \theta \quad (1.27)$$

for each $t \in \mathbb{N}$. Accordingly, $z\ \theta$ for $\theta \in \Theta$ refers to the trajectory of z given θ and $z\ \theta\ t$ is the state of the system at time $t \in \mathbb{N}$ under θ .

In general, it is not possible to generate a complete realization of the infinite trajectory $z\ \theta$ by simulations. What can only be done, instead, is to sample sequences of finite length, that is, a random variable z_J with

$$z_J : \Theta \rightarrow J \rightarrow Z \quad (1.28a)$$

$$z_J\ \theta\ j = z\ \theta\ j, \quad \forall j \in J, \quad (1.28b)$$

where $J \subset \mathbb{N}$ is a finite subset of \mathbb{N} . Realizations of this random variable are generated iteratively on the basis of tr . This reflects the simulation-based approach of agent-based modeling, where trajectories of finite length are sampled and subsequently analysed.

Now, this general case is introduced since the notion of price will vary from model to model. The agent-based model by Gintis (2006), for instance, considers agent-specific *private* prices. The idea behind agent-specific private prices is that instead of being “price-takers” as in the Walrasian model, agents are “price-setters”. Each agent has, through repeated trading encounters, a mental picture of the current prices and engages in trades accordingly. As a result, prices might vary from trading encounter to trading encounter. In the long run, so the hope, “quasi-public” prices evolve in the sense that all agents use more or less the same prices and these prices reflect the equilibrium prices of the underlying Walrasian model (see Section 1.4.3 for a precise formulation). Note that in this respect agent-specific private prices differ from the agent-independent *public* prices used

⁹For a definition of a measurable function and of a probability measure see Section 2.1. Moreover, note that the notion of transition kernel is the generalization of the notion of a transition matrix for discrete-time, discrete state space Markov processes (see Def. 2.2) to the general state space case. We consider here the general state space case since prices are usually thought of as elements in \mathbb{R} . Given a transition kernel tr , a sequence of random variables $(X_t)_{t \in \mathbb{N}}$ with state space Z is called a discrete-time Markov process if for any $t \in \mathbb{N}$ and any sequence of states x_0, x_1, \dots, x_t the conditional distribution of X_{t+1} given $X_t = x_t, \dots, X_0 = x_0$ is

$$\mathbf{P}(X_{t+1} \in A \mid X_t = x_t, \dots, X_0 = x_0) = \mathbf{P}(X_{t+1} \in A \mid X_t = x_t) = tr\ x_t\ A. \quad (1.25)$$

For a textbook on Markov processes on general state spaces, see, e.g., Meyn and Tweedie (2009)

in the Walrasian model of exchange. Besides agent-specific private and agent-independent public prices, price systems that are constant on certain groups of people or that depend on tuples of agents are yet other possibilities of how the notion of price could be modeled.

Formally, private prices can be represented as functions $A \rightarrow P$ and comprise the set of states in the model of price formation used in the Gintis model. That is, in this specific example, $Z = A \rightarrow P$. We write $pp \theta : \mathbb{N} \rightarrow A \rightarrow P$ to denote time-dependent private prices for a given random event $\theta \in \Theta$.

The transition function tr in agent-based modeling generally derives from the specification of some learning process (see, e.g., Brenner, 2006), or, in terms of evolutionary game theory, from the specification of a revision protocol. In the Gintis model, for instance, the transition kernel tr represents an evolutionary algorithm. Since its details are not important to the subsequent presentation and application of the framework, we refer the reader to Botta et al. (2013) for its complete specification. We just note here that tr generally depends not only on the current state but also on some assessment of its fitness (see Section 3.1.3). This fitness is calculated on the basis of the outcome of a population game in which the current state determines the strategy profile. In the case of agent-based models of exchange, the population game is the trading game specified further below. We discuss there that in contrast to the evolutionary games we consider in Part II, the fitness function used in agent-based models is in general not a deterministic function.

Bilateral exchanges and bilateral trades

The trading games considered in the following are *decentralized* in that they focus only on bilateral exchanges and bilateral trades. This is motivated by the observation that, often, trades and resulting exchanges of goods take place only between two parties.

More specifically, we say that two allocations $x, x' : A \rightarrow Q$ are related through a *bilateral exchange* between two agents $a_1, a_2 \in A, a_1 \neq a_2$, if

$$x' a \neq x a \Rightarrow a = a_1 \vee a = a_2 \quad (1.29a)$$

$$x' a_1 - x a_1 = x a_2 - x' a_2. \quad (1.29b)$$

In this case, we write $x(\mathcal{X} a_1 a_2)x'$. In specification (1.29b), $(x' a_1 - x a_1) g = x' a_1 g - x a_1 g$ is the amount of good g exchanged between a_1 and a_2 . Obviously, allocations which are related by a bilateral exchange are re-allocations of each other; that is,

$$x(\mathcal{X} a_1 a_2)x' \Rightarrow \forall g \in G : \sum_{a \in A} x' a g = \sum_{a \in A} x a g. \quad (1.30)$$

If, moreover, x' is a re-allocation of x which preserves the values of the stocks of the agents; that is,

$$\forall p \in P, a \in A : (x' a) \cdot p = (x a) \cdot p, \quad (1.31)$$

we call the re-allocation *value-preserving*.

In addition, the framework considers only exchanges of exactly two goods, which are called *elementary* bilateral exchanges in the functional framework. These goods are determined by the *sectors* of the two interacting agents. More specifically, sectors partition the set of agents such that all agents of a given sector have in common a sector-specific good. Sectors are represented through a surjective function

$$sct : A \rightarrow G \quad (1.32a)$$

$$\forall g \in G : sct^{-1} g \neq \emptyset, \quad (1.32b)$$

where sct^{-1} denotes the preimage, i.e., $sct^{-1} g = \{a \in A : sct a = g\}$. The idea behind sectors is that an agent can only offer and give away a good g if he belongs to the sector represented by g (for more on sectors, their interpretation in terms of, e.g., networks, or with respect to time, see Botta et al., 2011, p. 4032f.).

Thus, given a sector function sct , we say that two allocations $x, x' : A \rightarrow Q$ are related through an elementary bilateral exchange of two goods $g_1, g_2 \in G$ between two agents $a_1, a_2 \in A$ if

$$x(\mathcal{X}_{a_1} a_2)x', \text{ and} \quad (1.33a)$$

$$(x' - x)_{a_1} g \neq 0 \Rightarrow g = sct a_1 \vee g = sct a_2. \quad (1.33b)$$

We write in short $x(\mathcal{X}_{a_1} a_2)x'$ (leaving the sectors of the agents implicit).

Example 1.4. Let again $A = \{a_1, a_2, a_3\}$ and $G = \{g_1, g_2, g_3\}$. Consider the sector function $sct : A \rightarrow G$ with $sct a_i = g_i$. Then, for example, $x'(\mathcal{X}_{a_1} a_2)x$ where $x' = f x(a_1, a_2)$ with the exchange function f defined in Example 1.2. In contrast, the equilibrium allocation of Example 1.3 is not related by an elementary bilateral exchange with the initial allocation x_0 from Example 1.1, because in order to get from x_0 to x we need elementary bilateral exchanges between more than just two agents, i.e., we need a sequence of elementary bilateral exchanges.

Elementary bilateral trades in turn are characterized as bilateral interactions which lead to elementary bilateral exchanges as specified by (1.33a) and (1.33b) above. The definition of an elementary bilateral trade depends on two *policies*, an agent-specific *offer-and-demand* policy and an agent-independent *trade-resolving* policy. The term *policy* is used here in the sense of control theory, where policies are functions that map *states* into *actions* or *controls*.

In the case of the offer-and-demand policy, the actions consist of the agent's offer and demand for two specific goods and are represented as an element of $\mathbb{R}_{\geq 0}^2$. The state is model-dependent and agent-specific; that is, it would not depend on, e.g., the stocks of other agents. It depends, however, on the sectors of the agents engaged in trade; thus, the offer and demand of an agent $a \in A$ might depend not only on her own sector but also on the sector of the other agent she is engaged in trade with. In the case of the Gintis model, the state of an agent $a \in A$ depends on model setup values and functions such as its given initial stocks $x_0 a$, its utility function $u a \in U$, its sector $sct a$, as well as on its agent-specific current prices $pp \theta t a$, the sector represented by $g \in G$ of the agent it is engaged in trade with, and on its current stock $x a$. Leaving

model setup functions implicit in the notation, the offer-and-demand policy of the Gintis model is a function of type $A \rightarrow P \rightarrow G \rightarrow Q \rightarrow \mathbb{R}_{\geq 0}^2$.

In order to formulate minimal specifications for the offer-and-demand policy of an agent, let (o, d) be an outcome of the application of agent a 's policy. Here, o is the amount of good $g = \text{sct } a$ the agent offers in exchange for d units of some other good g' . It is then required that

$$x \ a \ g - o \geq 0. \quad (1.34)$$

In words, agent a needs to have at least o units of good g “in stock”. The idea behind this requirement is that if the agent’s offer and demand are matched by another agent’s demand and offer at most o units of g will be taken out from his stock. Specification (1.34) ensures that the resulting stocks will be non-negative, as required by the type $Q = G \rightarrow \mathbb{R}_{\geq 0}$ of stocks.

The states of the trade-resolving policy, in turn, consist of the offers and demands of the two interacting agents as determined by their offer-and-demand policies as well as in other model-specific parameters. The trade-resolving policy is agent-independent if it does not explicitly depend on the interacting agents. Put differently, while the trade-resolving policy certainly depends implicitly on the interacting agents through their offers and demands, it is not the agents’ “identities” that are of interest; every pair of agents is treated as every other pair¹⁰.

When applied to the offers and demands of two interacting agents, the resulting action is an element of $\mathbb{R}_{\geq 0}^2$ that corresponds to the amounts of the goods to be exchanged. More explicitly, let $(o_1, g_1), (d_1, g'_1)$ and $(o_2, g_2), (d_2, g'_2)$ denote the offers and demands of two interacting agents a_1, a_2 , where o_1 refers to the amount of good g_1 offered by agent a_1 to be exchanged in return of d_1 units of good g'_1 . Similarly, o_2, d_2 are the amounts of g_2, g'_2 offered and asked for by a_2 . In addition, let (δ_1, δ_2) be the amounts of goods to be exchanged as determined by the application of the trade-resolving policy. That is, agent a_1 gives δ_1 units of g_1 in return for δ_2 units of g_2 (and vice versa in case of agent a_2). Minimal requirements for the trade-resolving policy could thus read:

$$g_1 \neq g'_2 \vee g'_1 \neq g_2 \Rightarrow \delta_1 = \delta_2 = 0 \quad (1.35a)$$

$$g_1 = g'_2 \wedge g'_1 = g_2 \wedge d_1 \leq o_2 \wedge d_2 \leq o_1 \Rightarrow \delta_1 = d_1 \wedge \delta_2 = d_2. \quad (1.35b)$$

The interpretation of this specification is that for a non-zero exchange to take place, the goods demanded (offered) by a_1 have to coincide with the goods offered (demanded) by a_2 . If this is the case and the demands of the two agents are matched by the corresponding offers, then the amounts of g_1, g_2 to be exchanged shall coincide with the demands of the agents. Several trade-resolving policies are conceivable that fulfill this specification. An example is a simple “fill and kill” rule that returns zero exchanges whenever the left-hand side of Eq. (1.35b) evaluates to false. Note at this point that the description of the trading mechanism in the Gintis model does not yield an unambiguous specification for a trade-resolving policy. However, its implementation, provided by the author, is consistent with a specification that does not fulfill Eq. (1.35b). It is a function of type $P \rightarrow P \rightarrow (\mathbb{R}_{\geq 0} \times G) \rightarrow (\mathbb{R}_{\geq 0} \times G) \rightarrow (\mathbb{R}_{\geq 0} \times G) \rightarrow (\mathbb{R}_{\geq 0} \times G) \rightarrow \mathbb{R}_{\geq 0}^2$.

¹⁰It is possible to imagine agent-specific trade-resolving policies, e.g., to model “privileged” agents. This, however, is not the focus of the agent-based models of exchange considered here.

Its complete specification is given in Eqs. (21) - (24) of Botta et al. (2013). Informally, it includes *rationing constraints* that ensure, for instance, that the amount of good g_1 exchanged does not exceed the demand of agent a_2 . As it turns out, these rationing constraints are vital for the price dynamics reported by Gintis (2006). We discuss this fact in more detail in Section 1.4.3.

Taken together, an elementary bilateral trade between two agents a_1, a_2 can be described as a three-step process:

1. Two agents apply their offer and demand policies yielding (o_1, d_1) and (o_2, d_2) .
2. The trade-resolving policy is applied to the offer and demands of the two agents yielding (δ_1, δ_2) .
3. The two agents exchange the amounts of goods as determined by the trade-resolving policy.

This means formally that an elementary bilateral trade ebt is a function of type $(A \rightarrow Q) \rightarrow A \times A \rightarrow (A \rightarrow Q)$ with minimal specification

$$\begin{aligned} ebt \ x \ (a_1, a_2) &= x' \\ \Rightarrow \\ x(\mathcal{X}_e a_1 a_2) x' & \end{aligned} \tag{1.36a}$$

$$(x' \ a_1 - x \ a_1) \ (sct \ a_2) = \delta_2 \tag{1.36b}$$

$$(x' \ a_2 - x \ a_2) \ (sct \ a_1) = \delta_1. \tag{1.36c}$$

The complete specification is, of course, model-specific. The following is the specification of ebt as used in the Gintis model, which we present as an example of general value of how the offer-and-demand policies and the trade-resolving policy together determine the elementary bilateral trades. More specifically, the function ebt as used in the Gintis model fulfills in addition to Eqs. (1.36a)-(1.36c) the following:

$$(\delta_1, \delta_2) = trp \ p_1 \ p_2 \ (o_1, g_1) \ (d_1, g_2) \ (o_2, g_2) \ (d_2, g_1) \tag{1.37a}$$

$$(o_1, d_1) = odp \ a_1 \ p_1 \ g_2 \ (x \ a_1) \tag{1.37b}$$

$$(o_2, d_2) = odp \ a_2 \ p_2 \ g_1 \ (x \ a_2) \tag{1.37c}$$

$$g_1 = sct \ a_1 \tag{1.37d}$$

$$g_2 = sct \ a_2 \tag{1.37e}$$

$$p_1 = pp \ \theta \ t \ a_1 \tag{1.37f}$$

$$p_2 = pp \ \theta \ t \ a_2, \tag{1.37g}$$

where trp denotes the trade-resolving policy and $pp \ \theta \ t$ refers to the current prices under $\theta \in \Theta$.

Example 1.5. As a simple example, consider again $A = \{a_1, a_2, a_3\}$ and $G = \{g_1, g_2, g_3\}$. Let the offer-and-demand policy odp be a function of the same type as the one used in the Gintis model, i.e., $odp : A \rightarrow P \rightarrow G \rightarrow Q \rightarrow \mathbb{R}_{\geq 0}^2$, and define it by

$$(o, d) = odp \ a_i \ p \ g_j \ q, \tag{1.38}$$

where $d = 1$ and $o = 1$ if $q g_j \geq 1$ and $o = 0$ otherwise. Moreover, let the trade-resolving policy trp be of the same type $P \rightarrow P \rightarrow (\mathbb{R}_{\geq 0} \times G) \rightarrow (\mathbb{R}_{\geq 0} \times G) \rightarrow (\mathbb{R}_{\geq 0} \times G) \rightarrow (\mathbb{R}_{\geq 0} \times G) \rightarrow \mathbb{R}_{\geq 0}^2$ as used in the Gintis model, and define it by

$$(\delta_1, \delta_2) = trp p_1 p_2 (o_1, g_1) (d_1, g_2) (o_2, g_2) (d_2, g_1) \quad (1.39)$$

with $\delta_i = 1$ if both $o_1 \geq 1$ and $o_2 \geq 1$, otherwise $(\delta_1, \delta_2) = (0, 0)$. If we now consider the elementary bilateral trade function ebt as defined according to Eq. (1.36a) - (1.37g), then

$$ebt x (a_i, a_j) = f x (a_1, a_2), \quad (1.40)$$

where f is the exchange function defined in Example 1.2.

The Trading Game

The trading game in the agent-based models of exchange we consider can be understood as a population game¹¹ – with one difference being how utilities are defined. More specifically, a population game¹² is defined by

- a population A of agents;
- a set of *strategies* S available to each agent in the population, and the corresponding set of *strategy profiles* $\mathcal{S} = A \rightarrow S$;
- a consequence function of type $\mathcal{S} \rightarrow (\Omega \rightarrow C)$, where C denotes the set of consequences and Ω refers to some probability space $(\Omega, \mathcal{F}, \mathbb{P})$. The consequence function associates to each strategy profile in \mathcal{S} a probability distribution over consequences;
- for each agent, a preference relation \preceq over the set $\Omega \rightarrow C$; that is, a relation on $\Omega \rightarrow C$ which is complete, reflexive, and transitive.

Usually in game theory, it is assumed that the considered preference relations over $\Omega \rightarrow C$ fulfill the so-called *von-Neumann-Morgenstern* axioms (e.g., Osborne and Rubinstein, 1994). These ensure the existence of a utility function on the set of consequences; that is a function $u : C \rightarrow \mathbb{R}$, that represents preferences on C and that can be extended to a utility function $\bar{u} : (\Omega \rightarrow C) \rightarrow \mathbb{R}$, called the von-Neumann-Morgenstern utility (see Kreps, 1988, 1990). \bar{u} represents preferences on $\Omega \rightarrow C$ and is computed on the basis of u by taking the expected value. That is, the von-Neumann-Morgenstern-utility \bar{u} *rvC* of the random variable $rvC : \Omega \rightarrow C$ is given by

$$\bar{u} \text{ rvC} = \mathbb{E} u.\text{rvC}, \quad (1.41)$$

¹¹Def. 3.3.

¹²The definition of a population game considered here is slightly different from Def. 3.3. It is adapted to our functional framework and is thus more explicit. For example, the set of population states here is the set of strategy profiles $A \rightarrow S$, in contrast to the set of strategy distributions in the population as in Def. 3.3. Moreover, we use the term utility (instead of payoff) as this reflects the context of the Walrasian general equilibrium model, and we consider explicitly random outcomes as an intermediate step between states and their utilities. All in all, it should pose no difficulties to translate from one definition to the other.

where $u.rvC$ denotes the composition of rvC and u ; thus, $u.rvC : \Omega \rightarrow \mathbb{R}$. We show now how the trading game can be understood as such a population game – except for the definition of utilities on random outcomes.

In the following, let, as before, A be the set of agents, $x_0 \in A \rightarrow Q$ a fixed initial allocation, and $u : A \rightarrow U$ a given utility profile. As we are going to show, the set of strategies in the Gintis model corresponds to the set of prices P . Firstly, however, we are going to take the set of strategy profiles as a black box and specify a consequence function for a general trading game; that is, we specify the outcome of a specific play of the trading game.

In order to do so, we need the notion of a *trading schedule* as well as that of a *trading round*. A trading schedule determines which pairs of agents engage in an elementary bilateral trade as well as the order in which these trades take place. This order determines in general the outcome of a trading round; in a trading round, the pairs given by the trading schedule engage in elementary bilateral trades one after another in the order given by the trading schedule. Therefore, a trading schedule ts is represented by a list of agent pairs. That is, $ts \in List(A \times A)$. A minimal specification for ts is:

$$\forall (a, a') \in A \times A : \text{elem } (a, a') \text{ } ts \Rightarrow \text{sct } a \neq \text{sct } a'. \quad (1.42)$$

The outcome of a trading round with trading schedule ts on an allocation x can now be expressed by folding ebt on ts from x ,

$$\text{fold } ebt \ x \ ts. \quad (1.43)$$

Thus, in Eq. (1.7), $X = A \rightarrow Q$ and $Y = A \times A$, in agreement with the type of $ebt : (A \rightarrow Q) \rightarrow A \times A \rightarrow (A \rightarrow Q)$.

Example 1.6. Let ebt be the elementary bilateral trade function defined in Example 1.5, let x_0 be the allocation defined in Example 1.1, and let ts be the list of agent pairs defined in Example 1.2. Then, the outcome of the trading round with trading schedule ts on the allocation x_0 is

$$x''' = \text{fold } ebt \ x_0 \ ts, \quad (1.44)$$

with x''' defined in Example 1.2.

Note that trading rounds as specified by Eq. (1.43) inherit the properties of elementary bilateral trades. For instance,

- if $x' = \text{fold } ebt \ x \ ts$, then x' is a re-allocation of x ;
- for any $a \in A$, $x' \ a \ (\text{sct } a) \leq x \ a \ (\text{sct } a)$; and
- for any $g \in G$ with $g \neq \text{sct } a$, we have $x' \ a \ g \geq x \ a \ g$.

Still leaving the set of strategies a black box, we are now in a position to define the consequence function $c : \mathcal{S} \rightarrow (\Omega \rightarrow C)$. In the agent-based models of exchange we consider, Ω refers to a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ on a set Ω of finite trading lists, and the set of consequences corresponds to the set of allocations; that is, $C = A \rightarrow Q$. Of course, the particular probability space is model-specific. See Botta et al. (2013, Section 3.4.1) for the specific case of the

Gintis model. Given a strategy profile $s \in \mathcal{S}$, and a trading schedule $ts \in \Omega$, the resulting allocation $c s ts$ is given by the outcome of the corresponding trading round with initial allocation x_0 as defined in Eq. (1.43); that is,

$$c s ts = \text{fold } ebt \ x_0 \ ts, \quad (1.45)$$

where the strategy profile s is implicitly used in the function ebt .

As for the set of strategies an agent can choose from in a trading game, consider the following question: What is agent-specific about a trading round as specified in Eq. (1.43)? It is only the agent-specific states of the offer-and-demand policy. Moreover, we can differentiate between the components of such an agent-specific state in the following way: 1) there are those components that are fixed for all times, e.g., model input values and functions such as initial stocks of an agent, its utility function etc.; 2) there are values or functions that stay fixed during a trading round, an example being the agent-specific prices in the Gintis model; 3) there are values or functions that change within a trading round, such as the current stocks of an agent. It is the second of these kinds of components that constitute the strategies of an agent in a trading game. It is in this sense that prices in the Gintis model can be understood, as the author explicitly points out, as the strategies of the trading game.

Lastly, we explain how agent-based models of exchange use the utility functions on stocks to derive utilities on $\Omega \rightarrow C$, where $C = A \rightarrow Q$ as shown above. The underlying idea is that, in complex games, it might be costly or simply not possible to compute the expected value for each strategy profile. In addition, it is often argued that in real-world situations, people do not choose their strategy with respect to an expected utility. Instead, a few sample consequences must suffice as an indicator of the expected utility on which to base a decision.

To be more precise in the case of agent-based models of exchange, let $u : A \rightarrow (Q \rightarrow \mathbb{R})$ denote the given utility profile and consider its extension to a utility profile on allocations¹³ $u_{ext} : A \rightarrow (A \rightarrow Q) \rightarrow \mathbb{R}$ defined by

$$u_{ext} \ a \ x = u \ a \ (x \ a); \quad (1.46)$$

that is, the utility u_{ext} of an allocation x to an agent a is simply the original utility u of agent a 's stock in the allocation x . Furthermore, let $rvC : \Omega \rightarrow A \rightarrow Q$ represent a random consequence in a specific play of the game. Now, agent-based models of exchange use a random sample¹⁴ of rvC in order to compute an indicator of $\mathbb{E} (u_{ext} \ a).rvC$. To be more precise, let rvC_j for $j = 1, \dots, n$ denote the random variables of the sample of length n . Then a utility profile $\bar{u} : A \rightarrow (\Omega \rightarrow A \rightarrow Q) \rightarrow (\Omega \rightarrow \mathbb{R})$ is defined by

$$\bar{u} \ a \ rvC = \frac{1}{n} \sum_{j=1}^n (u_{ext} \ a).rvC_j. \quad (1.47)$$

Notice that the thus defined utility profile is not deterministic anymore: for each agent, the utility of a random consequence is again a random variable, a

¹³This is just one way of dealing with the fact that the utility functions we have considered so far are defined as functions on stocks while the consequences of a trading round are allocations.

¹⁴For a definition, see p. 53.

random utility.

The computation of \bar{u} brings us back to the model of price formation as outlined above. In evolutionary game theory or in evolutionary algorithms, a fitness function f is in general a deterministic function of type $\mathcal{S} \rightarrow A \rightarrow \mathbb{R}$ associating to each agent's strategy in the strategy profile $s \in \mathcal{S} = A \rightarrow S$ a real number, its fitness. Put differently, $f s a \in \mathbb{R}$ denotes the fitness of agent a 's strategy in the strategy profile s . As in the case of \bar{u} , the fitness function in agent-based models is often non-deterministic. In fact, \bar{u} and f are related through the consequence function: given the consequence function $c : \mathcal{S} \rightarrow \Omega \rightarrow A \rightarrow Q$ and $\bar{u} : A \rightarrow (\Omega \rightarrow A \rightarrow Q) \rightarrow \Omega \rightarrow \mathbb{R}$, the fitness function f is defined by

$$f : \mathcal{S} \rightarrow A \rightarrow \Omega \rightarrow \mathbb{R} \quad (1.48a)$$

$$f s a = \bar{u} a (c s). \quad (1.48b)$$

It is such a fitness function that the learning process in the model of price formation uses to calculate the prices in the next time period.

Although the decision to consider the average of a certain number of samples of $(u_{ext} a).rvC : \Omega \rightarrow \mathbb{R}$ instead of its expected value is, in fact, a modeling issue, it raises the question of what difference it makes for the dynamic evolution of the model. Of course, in the limit of large sample sizes, the average approaches the expected value by the strong law of large numbers¹⁵. But what happens in the case of small sample sizes? How does it effect the learning process via the fitness function? We are not going to discuss these question further here. We only point out that it illustrates the usefulness of the framework; by formulating precisely what the model is supposed to do, mathematical questions are raised that need to be investigated in order to gain a thorough understanding of the model (see also Section 1.4.3).

1.4.3 Application

As we have emphasized from the beginning, the functional framework for agent-based models of exchange as presented by Botta et al. (2013, 2011) does not only provide complete specifications of an agent-based model of exchange and thus allows its independent reimplementaion, but can be applied in a number of ways. In the last section (Section 1.4.2), for instance, we introduced and used the basic structure of the framework so as to make the relationship (in terms of similarities and differences) of agent-based modeling and game theory more explicit. In this section, we apply the framework to 1) explicate how to formulate the research question behind agent-based models of exchange precisely, and to 2) discuss how it enables rigorous model analysis and numerical investigations. Lastly, we refer to Botta et al. (2011, Section 4.1) for how the framework directly supports model implementation and documentation by, e.g., using the

¹⁵The assumptions commonly made on utility functions, in particular, continuity, as well as the fact that the resources are finite; that is, $\sum_{a \in A} x_0 a g < \infty$ for all $g \in G$ implies $(u_{ext} a).rvC$ is integrable and of finite variance and thus the strong law of large numbers is applicable since the $(u_{ext} a).rvC_j$'s constitute a sample of $(u_{ext} a).rvC$ and are thus independent.

specifications to formulate pre- and post-conditions for specific function interfaces.

Problem formulation

As we have discussed in Section 1.4.1, agent-based models of exchange have to be understood as a response to the criticism on general equilibrium models of exchange and price formation. There we informally described their research objective as to “lay out plausible rules of trading interactions that

1. might explain how equilibrium prices become established, and
2. ensure equilibrium allocations are reached under these equilibrium prices.”

We are now in a position to consider this objective again, this time in a more formal way. In the following, let x_0 be an initial allocation and $u : A \rightarrow U$ a utility profile such that there exists at least one pair (x_e, p_e) in general equilibrium with (x_0, u) according to Eqs. (1.14) and (1.17). In addition, we require that x_e is unique for p_e ; that is, (x, p_e) is *not* in general equilibrium with (x_0, u) for $x \neq x_e$.

We will focus on point 2. of the above formulated research objective first. The question is: what does it precisely mean to find plausible trading rules that under given equilibrium prices ensure equilibrium allocations are reached? This question involves several aspects:

1. We have to define what “plausible trading rules” are. In terms of the introduced framework, trading depends on the offer-and-demand policy odp as well as on the trade-resolving policy trp . To find plausible trading rules thus translates into finding plausible offer-and-demand as well as trade-resolving policies. In this context, defining what plausibility of these policies means entails two things. The first one – what does it mean for policies to be “economically plausible”? – is a modeling issue that we are not concerned with here. The second concerns “how to formulate” whatever it means for a trading rule to be (economically) plausible. The answer is, of course, to translate the economic idea of a plausible trading policy into specifications for it. Examples of such translations were given by the minimal specifications for odp and trp in Eqs. (1.34) and (1.35b), respectively.
2. We have to specify what equilibrium prices p_e exactly mean in the agent-based model. As we have pointed out in Section 1.4.2, what constitutes prices differs from model to model. The Walrasian model considers agent-independent public prices $P = G \rightarrow \mathbb{R}$, while the Gintis model, for instance, uses agent-specific private prices of type $A \rightarrow P$. In this latter case, we could say that private prices $pp \in A \rightarrow P$ correspond to the equilibrium prices p_e if each agent’s prices are the equilibrium prices; that is, if $pp \ a = p_e$.

Moreover, we have to specify how the thus determined prices translate into strategies in the trading game since it is the trading game in which allocations are changed and reached. As we have discussed in the previous section, private prices in the Gintis model translate directly into strategy profiles for the trading game.

3. Given the strategy profile s_e that corresponds to equilibrium prices p_e , we can consider the associated consequences in terms of allocations $c_{s_e} : \Omega \rightarrow (A \rightarrow Q)$. Remember that $\Omega = (\Omega, \mathcal{F}, \mathbb{P})$ denotes a model-specific probability space on a set of trading schedules. For $ts \in \Omega$, $c_{s_e} ts$ denotes the end-allocation resulting from the elementary bilateral trades given by the trading schedule ts . In this context, i.e., as c_{s_e} is not a deterministic function, we have to define what it means to “ensure equilibrium allocations are reached”. Do we ask for all trading schedules $ts \in \Omega$ to end in the equilibrium allocation x_e , i.e.,

$$c_{s_e} ts = x_e \quad (1.49)$$

for (almost) all $ts \in \Omega$? Yet another possibility is to require the distribution of the random variables $\tilde{c}_{s_e} a g : \Omega \rightarrow \mathbb{R}$ to have mean $x_e a g$ and small variance for each $a \in A, g \in G$, where \tilde{c} is the function defined from c by switching some of the input arguments:

$$\tilde{c} : \mathcal{S} \rightarrow A \rightarrow G \rightarrow \Omega \rightarrow \mathbb{R} \quad (1.50a)$$

$$\tilde{c} s a g ts = c s ts a g. \quad (1.50b)$$

These examples should suffice to demonstrate that it is not per se clear what it means that “equilibrium allocations are reached”. It is a notion that needs explicit definition.

As for point 1. of the research objective formulated above – what does it mean for “equilibrium prices to become established”? – this question relates to the evolution of prices which we represented as a stochastic process

$$z : \Theta \rightarrow \mathbb{N} \rightarrow Z, \quad (1.51)$$

where Z is a set of states that represents prices. Again, we have to define how the equilibrium prices p_e translate into prices of the agent-based model (see 2. above). Moreover, we have to state more precisely what “become established” means. This probably involves some notion of convergence of stochastic processes, for example, almost sure convergence, convergence in probability, or convergence in distribution. In addition, since the analysis of agent-based models depends on simulation data, it has to be made precise how to measure convergence in the output data. In Gintis (2006), for instance, the author measures “short-term convergence” of prices in terms of three standard deviations: a “cross-period standard deviation of mean private prices” and “inter-agent standard deviations of private producer and consumer prices”. Unfortunately, the author did not specify how to actually compute these quantities, and the vertical scale in Figure 1 of Gintis (2006) suggests that either these quantities are not standard deviations or that the probability distribution for the initial prices is not uniform. We mention this example to emphasize the need for formal descriptions of agent-based models and their analysis¹⁶.

¹⁶In order to make the discussion not too complex, we provide another possible interpretation of “prices becoming established” only here on the side. This interpretation relates to the fact that prices in the Walrasian model correspond to the exchange ratios of goods (Hildenbrand and Kirman, 1988, p. 77). Following this understanding of prices in the Walrasian model, we could decide to look at the exchange ratios according to which bilateral trades take

At this point, a last remark seems to be in order. We arranged the just given discussion of the formal interpretation of 1) and 2) above such that it corresponds to the intentions and specific setup of the Gintis model. It is, however, limited in several ways. Firstly, for simplicity we restricted the above discussion to the case that the equilibrium allocation x_e is unique for p_e . In the case of more than one equilibrium allocation, it seems rather odd to ask for all trading schedules to reach the equilibrium allocation x_e . Instead, of course, we would ask for the trading schedules to reach one of the equilibrium allocations corresponding to the initial allocation x_0 ; which equilibrium allocation is reached will certainly depend on the specific trading schedule. Secondly, and more importantly, the interpretation of 1) and 2) above is naturally linked to the question of which timescales we consider. As we will see in Part II of the thesis, it could very well be that, e.g., some equilibrium prices are reached on a short time scale in the sense that the price process stays close to these equilibrium prices on this short time scale, while at much longer time scales the stochastic price process switches to other equilibrium prices, again in the sense that the price process stays close to these other equilibrium prices, see Chapter 4.

Model analysis and numerical investigation

Specifications do not only allow to formulate precisely the problem investigated by a model, but are a necessary basis for reasoning rigorously about the model¹⁷. In particular, specifications allow us to deduce model properties.

As an interesting example that relates to the question of whether equilibrium allocations are reached under equilibrium prices we come back to the offer-and-demand and trade-resolving policies used in the Gintis model, see Section 1.4.2. Again, let x_0 denote an initial allocation and $u : A \rightarrow U$ a utility profile such that there is at least one pair (x_e, p_e) in general equilibrium with (x_0, u) . Let x_e be unique for p_e . In addition, p denotes (public) prices $p \in P$. We have already pointed out that the trade-resolving policy trp of the Gintis model does not fulfill Eq. (1.35b). It instead satisfies additional rationing constraints, which result in trp being *value-preserving* and *demand-limited*:

$$\begin{aligned} (\delta_1, \delta_2) &= trp \ p \ p \ (o_1, g_1) \ (d_1, g_2) \ (o_2, g_2) \ (d_2, g_1) \\ &\Rightarrow \\ \delta_1 * (p \ g_2) &= \delta_2 * (p \ g_1) \wedge (\delta_1 \leq d_1 \wedge \delta_2 \leq d_2). \end{aligned} \quad (1.52)$$

Value-preserving trade-resolving policies ensure that all bilateral exchanges in a trading round yield value-preserving re-allocations. This guarantees, for every possible trading schedule, that if (x_e, p) is in equilibrium with (x_0, u) then it is in equilibrium with all allocations that occur along a trading round under

place in the trading game, interpret these exchange ratios as the prices “on the markets”, and ask whether these exchange ratios “correspond to equilibrium prices p_e ” (of course, this phrase needs a definition, too). This is different from the approach via the stochastic price process as just outlined since it is not obvious how these prices (i.e., those that determine the strategies in the trading game) actually determine exchange ratios in the trading game.

¹⁷Or, more precisely, model “classes” since specifications usually do not define a single model but a class of models.

the given trading schedule. In addition, the offer-and-demand policy odp of the Gintis model satisfies:

$$x_e a g - x a g = 0 \Rightarrow odp a p_e g = (0, 0), \quad (1.53)$$

where x denotes the current allocation, see Botta et al. (2013, 2011). These constraints on odp and trp are necessary for x_e to be reachable from x_0 through a sequence of elementary bilateral trades. The reason for this is obvious; if the amount of good g that an agent a receives in a trade exceeds $x_e a g - x a g$, then this agent is going to be left with an amount of good g in excess of its optimal value. Now, since every agent can only offer the good of his sector, agent a cannot give away this excess amount of good g and thus will not be able to reach its optimal stock $x_e a$.

This is interesting because the numerical investigations of the Gintis model presented by Botta et al. (2013) as well as the analysis of this model presented by Mandel and Botta' (2009) suggest that the rationing constraints which ensure that trp is demand-limited (see specification (1.52)), prevent the convergence of prices to certain special prices $\hat{p} \in P$ defined by

$$\hat{p} g = \frac{1}{w g}, \quad (1.54)$$

where it is assumed that the utility function of each agent is the Scarf utility function defined in Example 1.3 with $\lambda = 1/|A|$. That is, for each agent $a \in A$,

$$u a g = \min_{g \in G} (y g) / (w g), \quad (1.55)$$

where

$$w = \frac{1}{|A|} * \sum_{a \in A} x_0 a. \quad (1.56)$$

Such a convergence to special prices had been reported by Gintis (2006) as a property of the model which presupposes rationing, and thus this appears to be in opposition to the numerical investigations of the same model as presented by Botta et al. (2013). Since this observation leads to questions about the correctness of the implementation by Botta et al. (2013) and by Gintis (2006) as well as to the question of whether the specification given by Botta et al. (2013) is actually a good description of the Gintis model, it motivated further numerical experiments. In particular, the model was run with a different trade-resolving policy which did not fulfill the rationing constraints (the rest of the setup stayed the same)¹⁸. If only this trade-resolving policy is changed, agent-specific prices appear to converge at large times towards the reported special prices. Thus, while our theoretical reasoning above had led to the conclusion that rationing constraints in the trade-resolving policy are a necessary condition for equilibrium allocation to be reached at given equilibrium prices, they seem to prevent the convergence of prices towards the special prices. Note, however, that all numerical results reported by Gintis (2006) as well as by Botta et al. (2013) have been obtained by assuming a very special utility profile and thus might

¹⁸See Eq. (34) and (35) of Botta et al. (2013) for a specification of the other trade-resolving policy considered.

not be representative for other utility profiles.

Lastly, besides using specifications to deduce properties of theoretical interest, such as in the example just given, they are also useful for deducing properties that model implementations have to display if they are consistent with the specification. Such properties are a prerequisite for designing crucial experiments in order to validate model implementations. Negative outcomes of such crucial experiments unambiguously indicate errors in the implementation (see Section 1.2). Take, for example, again an offer-and-demand policy which fulfills Eq. (1.53), a trade-resolving policy which is value-preserving and demand-limited (Eq. (1.52)), constant agent-specific prices $pp : A \rightarrow P$, i.e., $pp a = p$ for some $p \in P$, as well as the existence of an allocation x_e such that (x_e, p) are in equilibrium with given initial allocation $x_0 : A \rightarrow Q$ and utility profile $u : A \rightarrow U$. From these assumptions, it follows directly from the definition of a general equilibrium (see Eqs. (1.14) and (1.17)) that

$$(x_e, p)\mathcal{E}(x_0, u) \Rightarrow (x_e, p)\mathcal{E}(x, u) \quad (1.57)$$

for all allocations x that result from x_0 via sequences of elementary bilateral trades and thus for all allocations x that occur in every possible trading round. We can furthermore deduce that

$$\sum_{a \in A} \sum_{g \in G} (x_e a g - x a g) \quad (1.58)$$

is non-increasing in time. Thus, if a model implementation yields at one point in time a trading schedule an allocation x for which the sum in (1.58) is smaller than for an allocation at some later point in that trading schedule, it cannot be correct.

Such crucial experiments are essential for validating model implementations. In particular, in case model implementations cannot be automatically derived from specifications and proven to be correct, systematically exercising a model implementation on carefully chosen sets of crucial experiments is the only and best thing one can do. Crucial experiments, however, can only be set up on the basis of formal model specifications. If such specifications are not available, actually very little can be said about a model: implementations can still be used to obtain numerical results, but the interpretation of these results becomes problematic not least because the correctness of the implementation cannot be assessed.

1.5 Conclusion

In this chapter, we discussed the need for specifications for agent-based models. Specifications provide an intermediate layer between implementations and narrative descriptions of a model – an intermediate layer which is necessary but often not provided in the context of agent-based modeling. We introduced the basic structure of the functional framework for specifying agent-based models of exchange (Botta et al., 2013, 2011), which is based on the agent-based model presented by Gintis (2006). The presentation in this chapter exposed the relationship of agent-based models of exchange with general equilibrium

theory as well as game theory. The motivation behind this focus was that 1) agent-based models of exchange of the kind proposed by Gintis (2006) can be better understood in this context, and that 2) it points to a relationship between agent-based modeling and economic theory, which can potentially be used as a starting point for the development of a theoretical basis of agent-based modeling. Subsequently, we applied the framework to demonstrate how the framework helps to formulate the research question behind agent-based models of exchange precisely. Moreover we discussed how it constitutes a starting point for rigorous model analysis and numerical investigations. On the basis of the given formal representation of the Gintis model as an evolutionary game, next steps of analysis could include the investigation of Nash equilibria, the deterministic approximation of the model, the analysis of stochastically stable states, or a possible application of methods to analysis metastability and build models of reduced complexity such as the one presented in Part II of this thesis or other approaches to analyze the model output (e.g., Horenko, 2011).

Part II

Metastability Analysis

Chapter 2

Mathematical Preliminaries

This chapter introduces the basic notions and notation used in subsequent chapters. We start with the probabilistic framework within which everything will be situated (Section 2.1). Next, Markov chains as stochastic processes on finite state spaces and in discrete time will be defined and their characteristics will be described (Section 2.2).

2.1 Probability Theory

In Chapter 3, evolutionary games in finite populations are formalized as certain *stochastic processes*. This means that the state of the population changes in time, but that we are not able to specify deterministic laws according to which these state changes take place. Instead, we only consider how probable certain states and state changes are. The following aims at specifying the underlying probabilistic framework within which we can define the stochastic processes that formalize these ideas. Note that we focus only on finite sample and state spaces since this context suffices for our purpose.

Probability Space

A *probability space* consists of a triple $(\Omega, \mathcal{A}, \mathbb{P})$, where

- Ω is a set;
- \mathcal{A} is a σ -algebra on Ω ; that is, \mathcal{A} is a collection of subsets of Ω with
 1. $\emptyset, \Omega \in \mathcal{A}$,
 2. $A \in \mathcal{A} \Rightarrow \Omega \setminus A \in \mathcal{A}$, and
 3. $A_1, A_2, \dots \in \mathcal{A} \Rightarrow \bigcup_{i=1}^{\infty} A_i \in \mathcal{A}$;
- \mathbb{P} is a *probability measure* on \mathcal{A} , i.e., $\mathbb{P} : \mathcal{A} \rightarrow [0, 1]$ such that
 1. $\mathbb{P}(\Omega) = 1$, and

2. $\mathbb{P}(\cup_{i=1}^{\infty} A_i) = \sum_{i=1}^{\infty} \mathbb{P}(A_i)$ for every sequence A_1, A_2, \dots of pairwise disjoint elements of \mathcal{A} , i.e., for every $A_1, A_2, \dots \in \mathcal{A}$ with $A_i \cap A_j = \emptyset$, $i \neq j$.

Ω represents the set of possible “outcomes” in the random experiment or situation we consider and is called *sample space*. \mathcal{A} specifies the “events” we are interested in, i.e., the properties which can be observed, or, in other words, the questions that can be answered after the experiment is done (e.g., is the outcome $\omega \in A$ for $A \in \mathcal{A}$?). Note that the tuple (Ω, \mathcal{A}) is called a *measurable space*.

Now, given (Ω, \mathcal{A}) , the probability measure assigns to each such subset A of Ω a number between 0 and 1 with the interpretation that the higher this number, the higher the probability that the event will happen. Thus, the outcome $\omega \in A$, where $\mathbb{P}(A) = 1$ means the event happens for certain and $\mathbb{P}(A) = 0$ means the event does not happen for certain. A function $\mu : \mathcal{A} \rightarrow [0, \infty]$ that satisfies the 2 above but that is not necessarily normed to $\mu(\Omega) = 1$ is called a *measure*. Probability measures are also called *probability distributions*¹.

By condition 2 above, probability measures on finite sample spaces are characterized by their values on the elementary outcomes $\mathbb{P}(\{\omega\})$, $\omega \in \Omega$. The σ -algebra of Ω is the power set of Ω , i.e., the set of all subsets of Ω , and we can define the probability of any subset A of Ω by the sum of the probabilities of its elements: $\mathbb{P}(A) = \sum_{\omega \in A} \mathbb{P}(\{\omega\})$. The resulting function \mathbb{P} is a well-defined probability measure. Thus, in the case of a finite sample space, we can represent probability measures as functions $p : \Omega \rightarrow [0, 1]$ with the property that $\sum_{\omega \in \Omega} p(\omega) = 1$. We call such a function p probability measure or probability distribution, too. We furthermore note that we can also represent probability measures on finite sample spaces by *probability vectors*, i.e., by a vector $(p_i : i = 1, \dots, n) \in \Delta^{n-1}$, where $n = |\Omega|$ and $\Delta^{n-1} = \{(x_1, \dots, x_n) \in \mathbb{R}^n : \sum_{i=1}^n x_i = 1 \text{ and } x_i \geq 0 \text{ for all } i\}$ is the unit simplex.

Random Variables

Random variables formalize the idea that we are often not interested in the outcome of a random experiment, but in properties derived from the nature of the outcomes. The terminology is old and somewhat unfortunate, because random variables are neither “random” nor “variables” in the sense commonly used in mathematics².

A *random variable* on a sample space $(\Omega, \mathcal{A}, \mathbb{P})$ with *state space* (Z, \mathcal{B}) , where

¹Distribution functions, however, are a somewhat different albeit related concept.

²Probability theory textbooks sometimes provide lengthy explanations to get used to the term. One of these, for example, reads (Jacod and Protter, 2004, p.27)

A random variable represents an unknown quantity (hence the term variable) that varies not as a variable in an algebraic relation (such as $x^2 - 9 = 0$), but rather varies with the outcome of a random event. Before the random event, we know which values X could possibly assume, but we do not know which one it will take until the random event happens. This is analogous to algebra when we know that x can take on *a priori* any real value, but we do not know which one (or ones) it will take on until we solve the equation $x^2 - 9 = 0$ (for example).

\mathcal{B} is a σ -algebra on a set Z , is a function $X : \Omega \rightarrow Z$ for which

$$X^{-1}(B) := \{\omega \in \Omega : X(\omega) \in B\} \in \mathcal{A} \quad \text{for every } B \in \mathcal{B}. \quad (2.1)$$

The random variable X thus *induces* a probability measure $\mathbb{P}_X : \mathcal{B} \rightarrow [0, 1]$ on (Z, \mathcal{B}) by

$$\mathbb{P}_X(B) = \mathbb{P}(X^{-1}(B)). \quad (2.2)$$

The probability measure \mathbb{P}_X is called the *distribution of* or the *law of* X . In order to define such a probability measure on \mathcal{B} , we need condition (2.1) above, i.e., we need to be able to measure the probability of every pre-image of subsets of \mathcal{B} . The collection of all pre-images of sets in \mathcal{B} is called the σ -algebra *generated* by X . To denote that a random variable X has distribution λ we write $X \sim \lambda$. Given a random variable $X \sim \lambda$, a random *sample* of length $n \in \mathbb{N}$ is a set of n independent random variables with identical distribution λ .

The *expected value* of a random variable X with discrete state space $Z \subset \mathbb{R}$ is defined by

$$\mathbb{E}(X) = \sum_{z \in Z} z \mathbb{P}(X = z). \quad (2.3)$$

Conditional Probability and Independence

In Section 3 we are going to represent evolutionary games as stochastic processes with a specific dependency structure. In order to define this dependency structure we need the notion of conditional probability.

Given a probability space $(\Omega, \mathcal{A}, \mathbb{P})$ and two events $A, B \in \mathcal{A}$ with $\mathbb{P}(B) > 0$, we define the *conditional probability* $\mathbb{P}(A | B)$ of A given B by

$$\mathbb{P}(A | B) = \frac{\mathbb{P}(A \cap B)}{\mathbb{P}(B)}. \quad (2.4)$$

The underlying idea is that the information that B happened changes the probability measurement of A ; if we know that B happened, the possible outcomes that lead to A lie in $A \cap B$. We divide by $\mathbb{P}(B)$ in order to get a probability measure $\mathbb{P}^B : \mathcal{A} \rightarrow [0, 1]$ by $\mathbb{P}^B(A) = \mathbb{P}(A | B)$ for $A \in \mathcal{A}$.

Now, what can happen is that the information that B occurred does not change the probability that A takes place, which is, according to our definition of conditional probability in Eq. (2.4) above, the case if $\mathbb{P}(A | B) = \mathbb{P}(A)$, which is equivalent to $\mathbb{P}(A \cap B) = \mathbb{P}(A)\mathbb{P}(B)$. The latter formulation has the advantage of being defined even in the case of $\mathbb{P}(B) = 0$ and being easily extendible to the case of a finite number of events; we therefore define: two events A and B are independent if $\mathbb{P}(A \cap B) = \mathbb{P}(A)\mathbb{P}(B)$. Moreover, we call two random variables $X : \Omega \rightarrow Z_1$ and $Y : \Omega \rightarrow Z_2$ independent if every pair of events $X^{-1}(B_1), Y^{-1}(B_2)$, $B_1 \in \mathcal{B}_1, B_2 \in \mathcal{B}_2$ is independent.

Extending this definition to the case of collections of events, we say that a collection of events $(A_k)_{k \in K}$, where K is some index set, is *mutually independent* or an *independent collection* if for every finite subset $L \subset K$

$$\mathbb{P}(\cap_{k \in L} A_k) = \prod_{k \in L} \mathbb{P}(A_k). \quad (2.5)$$

A family $(\mathcal{A}_k)_{k \in K}$ of sub σ -algebras of \mathcal{A} is independent if for every finite subset L of K

$$\mathbb{P}(\cap_{k \in L} A_k) = \prod_{k \in L} \mathbb{P}(A_k) \text{ for all } A_k \in \mathcal{A}_k. \quad (2.6)$$

Furthermore, we call a collection of random variables $(X_k)_{k \in K} : (\Omega, \mathcal{A}, \mathbb{P}) \rightarrow (Z_k, \mathcal{B}_k)$ independent if the generated σ -algebras $X_k^{-1}(\mathcal{B}_k)$ are independent. In the special case of two random variables, this definition is equivalent to the one given above.

Stochastic Processes

Let $(\Omega, \mathcal{A}, \mathbb{P})$ be a probability space with Ω being the sample space, \mathcal{A} a σ -algebra on Ω , and \mathbb{P} a probability measure on Ω . A random variable $X : \Omega \rightarrow Z$, where Z is a set endowed with a σ -algebra \mathcal{B} , can then be interpreted as a random state of a system. Stochastic processes add to this construction a time parameter and thus formalize the fundamental idea of state changes in time which are affected by chance. More formally, a stochastic process assigns to every point in time $t \in T$ a random variable $X_t : \Omega \rightarrow Z$, where $T \subseteq [0, \infty)$:

Definition 2.1 (Stochastic process). Let $(\Omega, \mathcal{A}, \mathbb{P})$ be a probability space and (Z, \mathcal{B}) a measurable space. A *stochastic process* on Ω with state space Z is a family $X = (X_t)_{t \in T}$ of random variables $X_t : \Omega \rightarrow Z$, where T is an index set with $T \subseteq [0, \infty)$.

Given a stochastic process $X = (X_t)_{t \in T}$ and $\omega \in \Omega$, we call the mapping $X(\omega) : T \rightarrow Z$ defined by $X(\omega)(t) = X_t(\omega)$ the *path* of ω or the *trajectory* (or, *realization*) of the stochastic process X associated with ω .

In this thesis, we concentrate on discrete-time stochastic processes on a finite state space. That is, the index set T is discrete.

2.2 Markov Chains

In this section, we introduce discrete-time Markov chains on a finite state space. These are stochastic processes with a specific dependency structure between the random variables:

Definition 2.2 (Markov chain). Let $(\Omega, \mathcal{A}, \mathbb{P})$ be a probability space, (Z, \mathcal{B}) a finite state space with \mathcal{B} being its power set, and let $X = (X_k)_{k \in \mathbb{N}}$ be a discrete-time stochastic process. X is called a *Markov chain* if

$$\mathbb{P}(X_{k+1} = i_{k+1} \mid X_k = i_k, \dots, X_0 = i_0) = \mathbb{P}(X_{k+1} = i_{k+1} \mid X_k = i_k) \quad (2.7)$$

for all $i_{k+1}, \dots, i_0 \in Z, k \in \mathbb{N}$.

This means that X is a Markov chain if the conditional probability distributions of the process at time $(k+1)$ given the complete history of the process and the previous state only are indistinguishable. In other words: conditional on

$X_k = i_k$ the distribution of X_{k+1} is independent of X_0, \dots, X_{k-1} . The Markov chain is *time-homogeneous* if the probability

$$\mathbb{P}(X_{k+1} = i_{k+1} \mid X_k = i_k)$$

does not depend on the actual time step k ; that is,

$$\mathbb{P}(X_{k+1} = j \mid X_k = i) = \mathbb{P}(X_{n+1} = j \mid X_n = i) \quad \text{for all } k, n \in \mathbb{N}. \quad (2.8)$$

In this case, we can define a *stochastic matrix* $P = (p_{ij})_{i,j \in Z}$; that is, a function

$$P : Z \times Z \rightarrow [0, 1], \quad (i, j) \mapsto p_{ij}$$

with the following properties:

1. $0 \leq p_{ij} \leq 1$ for all $i, j \in Z$, and
2. $\sum_{j \in Z} p_{ij} = 1$ for all $i \in Z$.

This means that each row of P represents a conditional probability distribution on Z . The matrix P is defined by the *transition probabilities* of X :

$$p_{ij} = \mathbb{P}(X_{k+1} = j \mid X_k = i) \quad (2.9)$$

for any $k \in \mathbb{N}$. The thus defined matrix P is called the *transition matrix* of the Markov chain.

Given an *initial distribution* λ , i.e., $\mathbb{P}[X_0 = i] = \lambda_i$ for all $i \in Z$, and the transition matrix P , the finite-dimensional distributions of the Markov chain are completely determined:

$$\mathbb{P}(X_0 = i_0, \dots, X_k = i_k) = \lambda_{i_0} p_{i_0 i_1} \cdots p_{i_{k-1} i_k}. \quad (2.10)$$

Moreover, given a discrete-time stochastic process whose finite-dimensional distributions are defined as in Eq. (2.10) for all $k \in \mathbb{N}$, the process is a Markov chain with initial distribution λ and transition probabilities given by the stochastic matrix P . Thus, the stochastic matrix P together with an initial distribution λ uniquely characterize a Markov chain. The following theorem guarantees its existence (for a proof see, e.g., Meintrup and Schäffler, 2005, p. 554):

Theorem 2.1 (Existence of Markov Chains). *Let Z be a finite state space, P a stochastic matrix and λ a probability distribution on Z . There exists a Markov chain $X = (X_k)_{k \in \mathbb{N}}$ with*

$$\mathbb{P}(X_{k+1} = j \mid X_k = i) = p_{ij}, \quad (2.11a)$$

$$\mathbb{P}(X_0 = i) = \lambda_i \quad (2.11b)$$

for all $i, j \in Z$.

We write \mathbb{P}_i to denote the probability measure $\mathcal{A} \rightarrow [0, 1]$ defined by

$$\mathbb{P}_i(A) = \mathbb{P}(A \mid X_0 = i). \quad (2.12)$$

That is, \mathbb{P}_i is the probability measure on (Ω, \mathcal{A}) conditional on $X_0 = i$. Equivalently, \mathbb{P}_λ is the probability measure conditional on $X_0 \sim \lambda$.

Transfer of Probability Distributions

We have seen that a Markov chain is uniquely determined by the transition matrix P and the initial distribution λ . Since we consider only finite state spaces Z , we may write $Z = \{1, \dots, l\}$ for some $l \in \mathbb{N}$. The transition matrix then defines a linear mapping $\mathcal{P} : \mathbb{R}^l \rightarrow \mathbb{R}^l$, which is defined by

$$\mathcal{P}(x) = x^T P = \left(\sum_{i=1}^l p_{ij} x_i : j = 1, \dots, l \right) \quad (2.13)$$

Restricted to probability distributions, i.e., $x \in \mathbb{R}^l$ with $x_i \geq 0$ and $\sum_{i=1}^l x_i = 1$, the result of the application of \mathcal{P} is again a probability distribution. Note that in Eq. (2.13), we did not differentiate between the probability distribution and its vector representation with respect to the canonical basis $\{e_1, \dots, e_l\}$ of \mathbb{R}^l , which is common in the case of the canonical basis. In Chapter 5, we are going to consider other bases as well and thus a differentiation will become necessary.

Given the initial distribution λ represented as a vector in \mathbb{R}^l , we can use the law of total probability to compute

$$(\mathcal{P}(\lambda))_j = \sum_{i=1}^l p_{ij} \lambda_i \quad (2.14a)$$

$$= \sum_{i=1}^l \mathbb{P}(X_1 = j \mid X_0 = i) \mathbb{P}(X_0 = i) \quad (2.14b)$$

$$= \mathbb{P}_\lambda(X_1 = j). \quad (2.14c)$$

In the same manner, using the homogeneity property of the Markov chain, we can compute

$$(\mathcal{P}^k(\lambda))_j = (\lambda^T P^k)_j \quad (2.15a)$$

$$= \mathbb{P}_\lambda(X_k = j), \quad (2.15b)$$

where \mathcal{P}^k refers to the k -fold application of \mathcal{P} and we write $P^k = (p_{ij}^{(k)})_{i,j \in Z}$ to denote the k -step transition matrix and its entries. Thus, given an initial distribution, the linear mapping \mathcal{P} can be used to determine the probability distribution of states at later points in time. In this way, we can interpret the working of the linear mapping \mathcal{P} as a transfer of probability distributions in time.

Invariant Distributions

We can interpret \mathcal{P} (restricted to probability distributions) as the transition function of a discrete dynamical system. A fixed point of this system, i.e., a probability distribution μ that satisfies

$$\mathcal{P}(\mu) = \mu \quad (2.16)$$

is called a *stationary* or *invariant* distribution of the Markov chain. Thus, from the definition of \mathcal{P} it follows that invariant distributions are left eigenvectors of the transition matrix P associated with the eigenvalue 1. We know that there

exists at least one such distribution since $P\mathbf{1} = \mathbf{1}$, where $\mathbf{1} = (1, \dots, 1) \in \mathbb{R}^l$ and thus there is a left eigenvector associated with the eigenvalue 1 which can be normalized to a probability distribution.

The invariant distribution is an important notion in the study of Markov chains. In order to characterize the long-term behavior of Markov chains with the help of invariant distributions, we first need to introduce the following notions:

1. *Class structure and irreducibility:*

Communicating classes can be used to break down a Markov chain into smaller pieces in a such a way that allows us to study the dynamical behavior of each piece separately and put together the chain again to understand the behavior of the whole process. More specifically, we say that a state $i \in Z$ leads to state $j \in Z$ (or: j is accessible from i) if

$$\mathbb{P}_i(X_k = j \text{ for some } k \in \mathbb{N}) > 0 \quad (2.17)$$

We write $i \rightarrow j$. We furthermore say that the states i, j communicate if $i \rightarrow j$ and $j \rightarrow i$ and write $i \leftrightarrow j$. The relation \leftrightarrow on Z defines an equivalence relation whose equivalence classes partition Z . We call these equivalence classes *communicating classes*. A Markov chain $(X_n)_{n \in \mathbb{N}}$ is called *irreducible* if Z is a single communicating class.

A communicating class C is called *closed* if

$$i \in C, i \rightarrow j \Rightarrow j \in C. \quad (2.18)$$

2. *Aperiodicity:*

We say that $i \in Z$ is *aperiodic* if the greatest common divisor of the set $\{n \geq 0 : p_{ii}^{(n)} > 0\}$ is 1. Aperiodicity is a class property. Moreover, an irreducible Markov chain $(X_k)_{k \in \mathbb{N}}$ can be shown to be aperiodic if for some $i \in Z$ we have $\mathbb{P}(X_1 = i \mid X_0 = i) > 0$.

We are now in a position to characterize the long-term behavior of Markov chains on a finite state space by the following theorem (e.g., Norris, 1998):

Theorem 2.2 (Unique invariant distribution; ergodic theorem; convergence to the invariant distribution). *Let $(X_k)_{k \in \mathbb{N}}$ be an irreducible Markov chain on a finite state space Z with initial distribution λ . Then there exists a unique invariant distribution μ and, moreover, we have*

$$\mathbb{P}\left(\frac{v_i(k)}{k} \rightarrow \mu_i \text{ as } k \rightarrow \infty\right) = 1 \quad (2.19)$$

where $v_i(k) = \sum_{n=0}^{k-1} \mathbf{1}_{X_n=i}$ denotes the number of visits to i before k . If the Markov chain is in addition aperiodic, we have

$$\mathbb{P}_\lambda(X_k = j) \rightarrow \mu_j \quad \text{as } k \rightarrow \infty. \quad (2.20)$$

Thus, in particular,

$$p_{ij}^{(k)} \rightarrow \mu_j \quad \text{as } k \rightarrow \infty. \quad (2.21)$$

The interpretation of Eq. (2.19) is that if the Markov chain is irreducible, the empirical frequency with which states $i \in Z$ are visited during the first n periods (the probability of being in state i up through time n) converges to the unique invariant distribution μ with probability 1, independent of the initial conditions. If the process is also aperiodic, Eq. (2.19) together with Eq. (2.20) means that for $k \in \mathbb{N}$ sufficiently large, the probability of being in a certain state $i \in Z$ at time $k \in \mathbb{N}$ is essentially the same as the probability of being in state i through time k since both converge to the unique invariant distribution independently of the initial state.

Example 2.1 (Stationary Distribution for Birth-and-Death Chains). In Section 3.1, we show that, in a special case, stochastic evolutionary games can be represented as *birth-and-death* chains on a finite state space $Z = \{0, \dots, l\}$. These are discrete-time Markov chains $(X_k)_{k \in \mathbb{N}}$ on Z with a transition matrix P of the following form:

$$P = \begin{pmatrix} 1 - \alpha_0 & \alpha_0 & 0 & \dots & & \\ \beta_1 & 1 - (\beta_1 + \alpha_1) & \alpha_1 & 0 & \dots & \\ 0 & \beta_2 & 1 - (\beta_2 + \alpha_2) & \alpha_2 & 0 & \dots \\ & & \vdots & & & \\ 0 & \dots & & & 0 & \beta_l & 1 - \beta_l \end{pmatrix}, \quad (2.22)$$

that is,

$$p_{i(i+1)} = \mathbb{P}(X_{n+1} = i + 1 \mid X_n = i) = \alpha_i \text{ for } i = 0, \dots, l - 1, \quad (2.23a)$$

$$p_{j(j-1)} = \mathbb{P}(X_{n+1} = j - 1 \mid X_n = j) = \beta_j \text{ for } j = 1, \dots, l, \quad (2.23b)$$

$$p_{ij} = \mathbb{P}(X_{n+1} = j \mid X_n = i) = 0 \text{ if } j \notin \{i - 1, i, i + 1\}, \quad (2.23c)$$

where we assume $0 < \alpha_i, \beta_j < 1$ for each $i = 0, \dots, l - 1, j = 1, \dots, l$, and $\alpha_m + \beta_m < 1$ for $m = 1, \dots, l - 1$. The parameters α_i and β_j are called birth and death parameters.

The thus defined birth-and-death chain is clearly aperiodic and irreducible. The previous theorem thus implies the existence of a unique stationary distribution that we can compute by solving Equation (2.16):

$$\mu^T P = \mu^T \quad (2.24)$$

which is entrywise

$$\mu(i)(1 - p_{ii}) = \sum_{j \in Z, j \neq i} \mu(j)p_{ji}. \quad (2.25)$$

A solution can be determined recursively and is given by

$$\mu(i) = \mu(0) \frac{\alpha_0 \cdots \alpha_{i-1}}{\beta_1 \cdots \beta_i} \text{ for } i = 1, \dots, l. \quad (2.26)$$

The value of $\mu(0)$ is determined by the requirement that

$$\sum_{i \in Z} \mu(i) = 1 \quad (2.27)$$

which leads to

$$\mu(0) \left(1 + \sum_{i=1}^l \frac{\alpha_0 \cdots \alpha_{i-1}}{\beta_1 \cdots \beta_i} \right) = 1. \quad (2.28)$$

Invariant Distributions and Spectral Properties of the Transition Matrix

We have noted above that we can represent a stationary distribution of a Markov chain as a left eigenvector of the transition matrix P with respect to the eigenvalue 1. The Perron-Frobenius theorem for non-negative matrices tells us even more about the eigenstructure of transition matrices (see, e.g., Berman and Plemmons, 1979; Seneta, 1981):

Theorem 2.3 (Perron-Frobenius theorem). *Let $(X_k)_{k \in \mathbb{N}}$ be an irreducible and aperiodic Markov chain on a finite state space Z and let P denote its transition matrix. Then*

1. $\lambda = 1$ is a simple eigenvalue of P . The associated right eigenvector is $\mathbf{1} = (1, \dots, 1)^T$, i.e., $P\mathbf{1} = \mathbf{1}$. The associated left eigenvector μ , i.e., $\mu^T P = \mu^T$ can be chosen to be positive.
2. Any other eigenvalue λ' of P is strictly smaller (in modulus) than λ , i.e., $|\lambda'| < 1$ for every $\lambda' \in \sigma(P)$ with $\lambda' \neq \lambda$ and where $\sigma(P)$ denotes the spectrum of P , i.e., the set of all eigenvalues of P .
3. If $\lambda_1, \lambda_2, \dots, \lambda_r$ with $r \leq l$ are the eigenvalues of P which are ordered such that

$$1 = \lambda_1 > |\lambda_2| \geq |\lambda_3| \geq \dots \geq |\lambda_r| \quad (2.29)$$

and such that if $|\lambda_2| = |\lambda_j|$ for any $2 < j \leq r$ the algebraic multiplicity of λ_2 is greater or equal to the algebraic multiplicity of λ_j , then

$$P^n = \mathbf{1}\mu^T + \mathcal{O}(n^{m-1}|\lambda_2|^n) \quad (2.30)$$

where m denotes the algebraic multiplicity of λ_2 .

The interpretation of the first statement above tells us that if the Markov chain is irreducible and aperiodic, then there is a unique stationary distribution with positive probability weight on every state. The associated eigenvalue is 1 and all other eigenvalues of the transition matrix are strictly smaller in absolute value (second statement). The third statement tells us that the rate of convergence to the unique stationary distribution μ is geometric and dictated by the second largest eigenvalue.

Reversible Markov Chains

We can say even more for the class of Markov chains for which

$$\mu(i)p_{ij} = \mu(j)p_{ji} \quad (2.31)$$

where μ denotes a strictly positive distribution on the state space Z of the Markov chain $(X_k)_{k \in \mathbb{N}}$ with transition matrix P . In this case, the Markov chain is called *reversible* and P and μ are said to be in *detailed balance*. Eq. (2.31) implies that μ is a stationary distribution of (X_k) . It is unique if (X_k) is irreducible.

Example 2.2 (Reversibility of Birth-and-Death Chains). Consider again a birth-and-death chain as introduced in Example 2.1. Instead of solving the global balance equation in order to determine the unique stationary distribution μ of the chain, we now use the detailed balance equation (2.31). Since $p_{ij} = 0$ for $j \notin \{i-1, i, i+1\}$, we only need to solve for

$$\mu(i)p_{i(i+1)} = \mu(i+1)p_{(i+1)i} \text{ for } i = 0, \dots, n-1. \quad (2.32)$$

Via recursion, we get again

$$\mu(i) = \mu(0) \frac{\alpha_0 \cdots \alpha_{i-1}}{\beta_1 \cdots \beta_i} \quad \text{for } i = 1, \dots, l \quad (2.33)$$

and $\mu(0)$ determined again by the requirement that the entries of μ sum up to one, in order for μ to be a probability distribution, i.e.,

$$\mu(0) \left(1 + \sum_{i=1}^l \frac{\alpha_0 \cdots \alpha_{i-1}}{\beta_1 \cdots \beta_i} \right) = 1. \quad (2.34)$$

This implies that birth-and-death chains are reversible.

An important characteristic of an irreducible Markov chain $(X_k)_{k \in \mathbb{N}}$ with transition matrix P is that reversibility of (X_k) with respect to the distribution μ is equivalent to the transfer operator \mathcal{P} being self-adjoint in $l^2(\frac{1}{\mu})$; that is,

$$\langle \mathcal{P}(w_1), w_2 \rangle_{\frac{1}{\mu}} = \langle w_1, \mathcal{P}(w_2) \rangle_{\frac{1}{\mu}}, \quad (2.35)$$

where $l^2(\frac{1}{\mu})$ denotes the vector space \mathbb{R}^l equipped with the $\frac{1}{\mu}$ -weighted scalar product $\langle \cdot, \cdot \rangle_{\frac{1}{\mu}}$:

$$\langle w_1, w_2 \rangle_{\frac{1}{\mu}} = \sum_{i=1}^l w_1(i)w_2(i) \frac{1}{\mu(i)}. \quad (2.36)$$

See, e.g., Brémaud (1999, Ch. 6.2) for details. In the following, we will consider the associated norms

$$\|w\|^2 = \langle w, w \rangle_{\frac{1}{\mu}}, \text{ and } \|w\|_1 = \sum_{z \in Z} |w(z)| \frac{1}{\mu(z)}. \quad (2.37)$$

From \mathcal{P} being self-adjoint in $l^2(\frac{1}{\mu})$ follows that \mathcal{P} has only real eigenvalues and the associated eigenvectors can be chosen to form an orthonormal basis of $l^2(\frac{1}{\mu})$.

Using the last observation, we can derive a spectral decomposition of \mathcal{P} and demonstrate the relationship between the eigenstructure of \mathcal{P} and the speed of convergence to the unique invariant distribution in a straightforward way. To do so, let $(X_k)_{k \in \mathbb{N}}$ denote a reversible Markov chain on a finite state space Z with transition matrix P and let μ denote its unique invariant distribution. Furthermore, let the distinct real eigenvalues be denoted by $1 = \lambda_1 > |\lambda_2| \geq \dots \geq |\lambda_l|$, where $l = |Z|$, and let $\mu = u_1, u_2, \dots, u_l$ and $\mathbf{1} = v_1, \dots, v_l$ be the associated l linear independent, orthonormal left and right eigenvectors of P , respectively. It thus holds that $\langle u_i, u_j \rangle_{\frac{1}{\mu}} = 0$ for $i \neq j$ and $\langle u_i, u_i \rangle_{\frac{1}{\mu}} = 1$ for all

$i = 1, \dots, l$, similarly for the right eigenvectors v_1, \dots, v_j . The matrix P can thus be rewritten as

$$P = V\Lambda U^T \quad (2.38a)$$

$$= \sum_{i=1}^l \lambda_i v_i u_i^T, \quad (2.38b)$$

where $V = (\mathbf{1}, v_2, \dots, v_l)$ is the matrix consisting of the right eigenvectors and $U = (\mu, u_2, \dots, u_l)$ is the matrix consisting of the left eigenvectors. Since $U^T V = Id$, we get the following *spectral decomposition*

$$P^k = \sum_{i=1}^l \lambda_i^k v_i u_i^T. \quad (2.39)$$

Now, let v be an initial probability distribution on Z represented as a vector $v \in \mathbb{R}_{\geq 0}^l$ with $\sum_{j=1}^l v_j = 1$. Since μ, u_2, \dots, u_l are a basis of \mathbb{R}^l , we may write v as a linear combination of these vectors:

$$v = c_1 \mu + c_2 u_2 + \dots + c_l u_l. \quad (2.40)$$

Let $v^{(k)}$ denote the probability distribution of X_k , i.e., $v_j^{(k)} = \mathbb{P}_v(X_k = j)$. Then,

$$v^{(k)T} = v^T P^k \quad (2.41a)$$

$$= (c_1 \mu^T + c_2 u_2^T + \dots + c_l u_l^T) P^k \quad (2.41b)$$

$$= \{ \text{since } \langle u_i, v_j \rangle_{\frac{1}{\mu}} = 0 \text{ if } i \neq j \} \quad (2.41c)$$

$$c_1 \mu^T + c_2 \lambda_2^k u_2^T + \dots + c_l \lambda_l^k u_l^T, \quad (2.41d)$$

and thus,

$$\|v^{(k)} - c_1 \mu\|_1 \leq \|v^{(k)} - c_1 \mu\| \quad (2.42a)$$

$$\leq |c_2| |\lambda_2|^k \|u_2\| + \dots + |c_l| |\lambda_l|^k \|u_l\| \quad (2.42b)$$

$$\rightarrow 0 \quad \text{for } k \rightarrow \infty. \quad (2.42c)$$

It follows that $c_1 = 1$ since both $v^{(k)}$ and μ are probability distributions. The convergence results from the fact that $|\lambda_i| < 1$ for $i = 2, \dots, l$.

This illustrates how the eigenvalues of the transition matrix P relate to its *relaxation timescales*, i.e., to the timescales on which initial probability distributions “relax” to the stationary distribution (see, e.g., Levin et al., 2008). These are defined as $T_j = 1/|\log |\lambda_j||$. More specifically, the relaxation timescale linearly approximates the time to decay of the influence of the j -th eigenvector on the evolution of $v^{(k)}$ in the following sense: let $\nu_j = -\log |\lambda_j|$, i.e., $|\lambda_j| = e^{-\nu_j}$. We can interpolate the sequence $|\lambda_j|, |\lambda_j|^2, \dots$ by $e^{-t\nu_j}$ with $t \in [0, \infty)$, and interpret ν_j as the decay rate associated with the eigenvector u_j . Its *mean lifetime* is given by $\tau_j = 1/\nu_j$. τ_j can be interpreted as the amount of time needed for the influence of u_j as measured by $\|u_j\|$ to be reduced by a factor of $1/e$. In terms of the discrete time sequence $|\lambda_j|, |\lambda_j|^2, \dots$ this implies a timescale of $T_j = 1/|\log |\lambda_j||$.

Put differently, we have

$$\|v^{(k)} - \mu\|^2 = \|c_2 \lambda_2^k u_2 + \cdots + c_l \lambda_l^k u_l\|^2 \quad (2.43a)$$

$$= \sum_{z \in Z} \frac{1}{\mu(z)} \left(\sum_{j=2}^l c_j \lambda_j^k u_j(z) \right)^2 \quad (2.43b)$$

$$\leq \{1 > |\lambda_2| \geq |\lambda_j| \text{ for each } j = 3, \dots, l\} \quad (2.43c)$$

$$|\lambda_2|^{2k} \sum_{z \in Z} \frac{1}{\mu(z)} \sum_{j=2}^l (c_j u_j(z))^2 \quad (2.43d)$$

$$= |\lambda_2|^{2k} \|v - \mu\|^2, \quad (2.43e)$$

which implies

$$\|v^{(k)} - \mu\| \leq |\lambda_2|^k \|v - \mu\|. \quad (2.44)$$

Thus, because

$$|\lambda_2|^k = e^{k \log |\lambda_2|}, \quad (2.45)$$

we get for $k \geq \frac{1}{|\log |\lambda_2||} \approx \frac{1}{1-|\lambda_2|}$ that the distance between the initial distribution v and the stationary distribution μ has reduced by more than $\frac{1}{e}$. Moreover, for a probability distribution v defined by

$$v = \frac{1}{m_2} u_2 + \mu, \quad (2.46)$$

where

$$m_2 = -\min_{i \in Z} \frac{u_2(i)}{\mu(i)} > 0, \quad (2.47)$$

we get equality in (2.44), which shows that the upper bound in this equation is sharp. v is a probability distribution because m_2 has been chosen such that $v(i) \geq 0$, and because $u_2^T \mathbf{1} = 0$, which implies $\sum_{i \in Z} v(i) = 1$. As λ_2 is the second largest eigenvalue in absolute value, there are no processes that are slower to converge to the stationary distribution than the v defined in Eq. (2.46). If we start the Markov chain with initial distribution v , there are a few trajectories starting near the minimum of u_2 weighted with μ and many trajectories near the maximum of u_2 weighted with μ . As time advances, this imbalance gets equilibrated³. The larger $|\lambda_2| < 1$, the slower this process. Similar considerations for the other eigenvalues show that the relaxation timescales can be interpreted as convergence rates. In this way, we can speak of *slow* and *fast* processes associated with large and small eigenvalues (in absolute value) and relaxation timescales, respectively. Thus, $T_1 = \infty$ is the trivial timescale on which the process relaxes to its invariant measure, T_2 is the slowest non-trivial

³Why the term “equilibrated” fits well here might be easier to see in terms of probability densities with respect to μ . We say that ν is a probability distribution with respect to μ if $\nu(i) \geq 0$ for all $i \in Z$ and $\sum_{i \in Z} \nu(i) \mu(i) = 1$. Then, $v = (\frac{1}{m_2} u_2 / \mu + \mathbf{1}) \mu$ and thus its probability density with respect to μ , which is $\frac{1}{m_2} u_2 / \mu + \mathbf{1}$, gets equilibrated to the probability density $\mathbf{1}$ with respect to μ , i.e., to the density which is uniform on all states.

timescale, etc. If some eigenvalues are particularly close to 1, the corresponding timescales are very large and significantly longer than the other timescales. These eigenvalues are called *dominant*.

Thus, the observation that dominant eigenvalues and the corresponding eigenvectors relate to the long-time behavior of the Markov chain suggests a relationship between the spectral properties of the transition matrix of a Markov chain and metastability in the sense that metastability corresponds to the presence of dominant eigenvalues. The characterization of metastability in terms of spectral properties of the transition matrix of a Markov chain has been initiated by Davies (1982a,b, 1983) and has been pursued more recently by, e.g., Bovier et al. (2001, 2002); Gaveau and Schulman (1998); Huisinga and Schmidt (2006); Schütte and Huisinga (2003), see Chapter 4 for a detailed discussion and further references. In Chapter 5, we study the Markov state modelling approach to construct models of reduced complexity. The indicated relationship between the spectral properties of the transition matrix of the Markov chain and metastability implies that we would like a small approximation error in the dominant eigenvalues in order to capture the essential dynamics of the original Markov chain in terms of metastability. We remark at this point, however, that although eigenvalues near -1 are related to the long-time behavior of the Markov chain, they correspond rather to quasi-periodic behavior than to metastable dynamic behavior. In other words, for a Markov chain with metastable dynamic behavior we expect *positive* eigenvalues near 1.

Lastly, we note that the given relationship between the largest eigenvalues in absolute value and the rate of convergence to stationarity depends on the special l^2 context we chose. There are, of course, other notions that measure distances between probability distributions such as the total variance distance. It is therefore not hard to imagine that there are different related notions of mixing times, i.e., measurements of the time it takes to be “close” to the stationary distribution (see, e.g., Levin et al., 2008), some of which might be in some cases considerably smaller than the above given relaxation times (e.g., Diaconis and Saloff-Coste, 1996).

Hitting Times

Hitting times are fundamental notions to the study of metastability (see Chapter 4).

In the following, let $(X_k)_{k \in \mathbb{N}}$ be an irreducible time-discrete Markov chain on a finite state space Z with transition matrix P . Let μ be the unique stationary distribution.

Definition 2.3 ((First) hitting time). Let $A \subset Z$. The *(first) hitting time* (or, *first passage time*) τ_A of the Markov chain into the set A measures the first non-zero time that the Markov chain enters the set A :

$$\tau_A = \inf\{k > 0 : X_k \in A\}. \quad (2.48)$$

We are particularly interested in the mean (first) hitting time $\mathbb{E}_x \tau_A$ of the Markov chain with initial state $x \in Z$ to reach a set $A \subseteq Z$. The mean hitting

times of the Markov chain can be computed in the following way. Let $w_A : Z \rightarrow \mathbb{R}_{\geq 0}$ be defined by

$$w_A(x) = \begin{cases} \mathbb{E}_x \tau_A & \text{if } x \notin A \\ 0 & \text{if } x \in A. \end{cases} \quad (2.49)$$

We can compute w_A as the unique solution of the inhomogeneous Dirichlet problem (Norris, 1998):

$$(P - Id)w_A(x) = -1, \quad x \in Z \setminus A \quad (2.50a)$$

$$w_A(x) = 0, \quad x \in A, \quad (2.50b)$$

where Id is the identity matrix on Z . Then

$$\mathbb{E}_x \tau_A = \begin{cases} w_A(x) & \text{if } x \notin A \\ Pw_A(x) + 1 & \text{if } x \in A. \end{cases} \quad (2.51)$$

Finally, we define the notion of *recurrence* based on hitting times:

Definition 2.4. A state $x \in Z$ is *recurrent* if $\mathbb{P}_x(\tau_x < \infty) = 1$. Otherwise, it is called *transient*.

Put differently, a state $x \in Z$ is recurrent if the chain hits x in finite time with probability one conditional on having started in x . Note that recurrence is a class property. Moreover, in the finite state space case, which we consider throughout, we have that a communication class $A \subset Z$ is recurrent if and only if it is closed.

Chapter 3

Stochastic Evolutionary Games

In this chapter, we define stochastic evolutionary games and show that they can be represented as discrete-time Markov chains (Section 3.1). Subsequently, we discuss common approaches to their analysis (Section 3.2). For a comprehensive introduction to game theory and evolutionary game theory, respectively, we refer the reader to, e.g., Osborne and Rubinstein (1994) and Sandholm (2010).

3.1 Definition

Evolutionary game theory studies the dynamic process of interactive decision-making in a population of agents. More specifically, it studies how frequencies of strategies within a population change in time according to the decision-making of the agents, which is “interactive” in the sense of being interdependent through the current state of the population as well as through the strategies’ success. We represent evolutionary games in terms of two notions. The first one is the notion of a *population game*, which describes the basic structure of a game played in a population of agents. The second, in turn, is called *revision protocol* and models the decision-making of the agents given a population game. We show that evolutionary games in finite populations lead to stochastic dynamics that can be modeled as discrete-time Markov chains. Moreover, they can be seen as simple, agent-based models¹ in the sense used in Chapter 1. Note that in all that follows we do *not* model the actual play of the game and it is in this respect of no importance when and how often the game is played. We are only interested in aggregate behavior changes over time.

Since evolutionary game theory builds on classical game theory and lots of examples of evolutionary games are classical games played in a population of players, we first introduce the classical notion of a game before that of an evolutionary one.

¹ Or rather, computer-based implementations that simulate the dynamic behavior of the evolutionary game would be called agent-based model.

3.1.1 Games in Strategic Form

Classical game theory differentiates between the *representation* of a game and *solutions* for the game. A representation specifies the rules of the game – the players, the order of moves and possible actions of the players, etc. – but does not include any assumptions about the behavior of the agents or conclusions from these assumptions about how agents should reasonably behave. Solutions of a game suggest reasonable choices of the participants of the game under certain behavioral assumptions. *The* fundamental solution concept is the notion of *Nash equilibrium*. In risk of oversimplification, one can say that classical game theory thus differs from evolutionary game theory in that it focuses on how players should reasonably choose their strategies, e.g., in order to maximize their payoff which depends on the strategies chosen by the co-players (who, in turn, strive to maximize their payoff).

A standard representation of a game is the *strategic form* as given in the following definition:

Definition 3.1 (Game in strategic form). A game G in strategic (or normal) form consists of

- a finite set $N = \{1, \dots, n\}$ of *players*,
- a set $S = \times_{j \in N} S_j$ of *strategy profiles*, where for each player $i \in N$ the set S_i is a nonempty set of *strategies* available to player i ,
- for each player $i \in N$ a function $\pi_i : S \rightarrow \mathbb{R}$, called the payoff function of player i , which specifies for each player $i \in N$ the payoff $\pi_i(s)$ of a strategy profile $s \in S$.

We write in short $G = (N, S, (\pi_i)_{i \in N})$.

As such, games in strategic form only describe the *structure* of game-like situations, but *not* the actual play. Simple examples are board games, how to greet (shake hands or bow), as well as paying or evading taxes, etc. Games in strategic form cover the class of games in which the players' moves are sequentially taken such as in the games tic-tac-toe or chess².

In the following, we often consider games with two players in which each player can choose from two strategies. We denote these games as 2x2-games. A typical graphical representation of a 2x2 game is a table (see Figure 3.1). The *payoff matrices* $A = (a_{ij})$ and $B = (b_{ij})$ for $i \in S_1$ and $j \in S_2$ are defined by

$$a_{ij} = \pi_1(i, j) \tag{3.1a}$$

$$b_{ij} = \pi_2(i, j). \tag{3.1b}$$

Moreover, our focus are *symmetric* 2x2-games. These are games in which, according to the notation used in Figure 3.1, $\pi_1(T, L) = \pi_2(T, L)$, $\pi_1(B, R) = \pi_2(B, R)$, $\pi_1(T, R) = \pi_2(B, L)$, and $\pi_1(B, L) = \pi_2(T, R)$. Thus, in symmetric 2x2-games, if we identify T with L and B with R (e.g., T and L are the first strategy available, and R and B are the second strategy available), it does not

²Actually, the formal representation that can be seen as the direct translation of the sequential structure of these games is called the *extensive form*. There is, however, a strategic form game for every extensive form one.

matter in which player position the players play the game. Instead of T or L , and B or R , we write 1 and 2 for “first” and “second” strategy, respectively (see Figure 3.2). Note that in symmetric games, we have $A = B^T$ for the payoff matrices.

		Player 2	
		L	R
Player 1	T	$\pi_1(T, L), \pi_2(T, L)$	$\pi_1(T, R), \pi_2(T, R)$
	B	$\pi_1(B, L), \pi_2(B, L)$	$\pi_1(B, R), \pi_2(B, R)$

Figure 3.1: Graphical representation of a two-player game in strategic form in which each player has two possible strategies.

		1	2
	1	a, a	b, c
	2	c, b	d, d

Figure 3.2: Graphical representation of a symmetric 2x2-game.

In the following, let $S_{-i} = \times_{j \in N, j \neq i} S_j$ denote the strategy profile of all players except player i and let $(s_i, s_{-i}) \in S_i \times S_{-i}$ denote the strategy profile $(s_1, \dots, s_{i-1}, s_i, s_{i+1}, \dots, s_n) \in S$. That is, (s_i, s_{-i}) is the profile of strategies $s \in S$ according to which the i -th player plays strategy s_i and the rest of the players play according to s_{-i} .

The fundamental solution concept in game theory, the Nash equilibrium, has been introduced by John Nash in his PhD dissertation (Nash, 1950). In a Nash equilibrium, every players choice is optimal given the choices of the others.

Definition 3.2 (Nash equilibrium). A *Nash equilibrium* of the strategic game $(N, S, (\pi_i))$ is a strategy profile $s^* \in S$ for which

$$\pi_i(s^*) \geq \pi_i(s_i, s_{-i}^*) \quad \text{for all } s_i \in S_i. \tag{3.2}$$

It is called *strict* if the above inequality holds strictly.

One traditional interpretation of the Nash equilibrium solution concept – that is, an explanation of when and why we might expect players to actually play a Nash equilibrium – assumes that a Nash equilibrium

“results from analysis and introspection by the players in a situation where the rules of the game, the rationality of the players, and the players’ payoff functions are all common knowledge” (Fudenberg and Levine, 1998, p. 1).

According to this interpretation, a Nash equilibrium is the only “viable outcome of careful reasoning by ideally rational players” (Binmore, 1995, p. x).

This interpretation has received similar criticism as the general equilibrium model (see Section 1.4.1)³. In short, it has been argued that this interpretation is

- unreasonable because of the underlying behavioral assumptions: real-world players are often not fully rational and do not have common knowledge and information about the game structure and the other players’ rationality (Simon, 1957). Moreover, they often do not have the computational capabilities of calculating the optimal play (as is the case for, e.g., chess, and other complex decision situations; Gallegati and Richiardi, 2009; Kirman, 1992); and
- uninformative because most games possess multiple equilibria: if a game has multiple Nash equilibria, it is not possible to single out a unique play of the game on the basis of this notion.

Example 3.1. A simple example of a game with multiple equilibria, which we will often use in the subsequent sections and chapters, is a 2x2 *pure coordination* game, see Figure 3.3 for its graphical representation. The payoffs in a pure coordination game are such that (1,1) and (2,2) are strict Nash equilibria, payoffs of the remaining strategy profiles (1,2) and (2,1) are 0. The question thus is on which of the Nash equilibria do the players coordinate?

	1	2
1	a, a	$0, 0$
2	$0, 0$	b, b

Figure 3.3: A pure coordination game.

An interpretation of such a pure coordination game can be phrased in terms of currencies. Each of the strategies is a currency, e.g., strategy 1 is “gold” and strategy 2 is “silver”. Players can only exchange goods and thus obtain a positive payoff from the exchange if they use the same currency. If players hold different mediums of exchange, they are not able to exchange goods and thus the payoff from the encounter is zero.

The problem setting described in the second point above has been termed the problem of *equilibrium selection*. *Refinements* of the Nash equilibrium notion, such as payoff or risk dominance (for a definition of these notions see, e.g., Osborne and Rubinstein, 1994), constitute one approach to the problem of equilibrium selection – an approach that is, however, still unsatisfactory since the behavioral assumptions are even stronger and because, as Binmore (p. ix 1995) puts it,

³This is not a coincidence but results from the fact that the general equilibrium model is an *abstract* game. Arrow and Debreu (1954) introduced abstract games as a generalization of a game and used a corresponding generalization of the Nash equilibrium notion in order to prove the existence of a general equilibrium in economies with production.

“different game theorists proposed so many different rationality definitions that the available set of refinements of Nash equilibrium became embarrassingly large. Eventually, almost any Nash equilibrium could be justified in terms of someone or other’s refinement”.

This discontent with traditional game theory has motivated the study of evolutionary games, which we pursue in the following.

3.1.2 Population Games

In contrast to games in strategic form, population games put emphasis not on each player’s strategy choices and the resulting strategy profile but on the distribution of strategies in the population.

In the following, we consider only games played by a single, finite population.

Definition 3.3 (Population game). A *population game* consists of

- a population of n agents,
- a strategy set $S = \{1, \dots, m\}$ available to each agent in the population, i.e., each player faces the same set of strategies,
- the set of population states is $\Delta_n^{m-1} = \Delta^{m-1} \cap \frac{1}{n}\mathbb{Z}^m = \{x \in \Delta^{m-1} : nx \in \mathbb{Z}^m\}$, where $\Delta^{m-1} = \{x \in \mathbb{R}_{\geq 0}^m : \sum_{j \in S} x_j = 1\}$ is the unit $(m-1)$ -simplex and the j -th component of $x \in \Delta_n^{m-1}$ represents the proportion of agents choosing the strategy j in the population; thus, Δ_n^{m-1} is the grid in the simplex Δ^{m-1} with spacing $\frac{1}{n}$, which has $\binom{n+m-1}{m-1} = \binom{n+m-1}{n}$ states (Scheffé, 1958), and
- a payoff function $F : \Delta_n^{m-1} \rightarrow \mathbb{R}^m$, where $F_i(x)$ represents the payoff to playing strategy i when the population state is x .

Note that we restrict ourselves to payoff functions which are independent of the population size n ; this is sufficient for our purposes. See (Sandholm, 2010) for population games with population size dependent payoff functions. The following example draws a relationship between population games and games in strategic form:

Example 3.2 (Complete matching in symmetric two-player strategic form games). Let $A \in \mathbb{R}^{m \times m}$ be the payoff matrix in a symmetric, two-player strategic form game with a strategy set of size m . We define a payoff function $F : \Delta_n^{m-1} \rightarrow \mathbb{R}^m$ for a population of n agents by $F(x) = Ax$. These payoffs could, for instance, be interpreted as resulting from a complete matching of the agents in the population to play the given game in strategic form.

As a concrete example, consider agents in a population of size n are randomly matched to play the 2x2 pure coordination game with parameters a and b from Example 3.1. For $x \in \Delta_n^1$, let x_1 denote the proportion of agents in the population playing strategy 1 (“gold”) and $x_2 = 1 - x_1$ the proportion of agents playing strategy 2 (“silver”). The payoff an agent playing strategy 1 or 2 obtains is thus $F_1(x_1, x_2) = ax_1$ or $F_2(x_1, x_2) = bx_2$, respectively. Because of the interpretation of the strategies in terms of currencies (see Example 3.1), we will call this specific population game with parameters a and b the *currency*

game.

Further prominent examples of population games are congestion games and potential games (see, e.g., Sandholm, 2010, and Chapter 6.3).

3.1.3 Revision Protocol, Stochastic Evolutionary Games

Given a population game, we now model the strategy updating process of the agents. The idea is the following: at every moment in time, each agent has chosen a strategy in the strategy set S . At times $t = k\delta$, where $\delta = 1/n, k \in \mathbb{N}$, exactly one agent is randomly drawn (with equal probability for all agents) to reconsider her strategy choice. We assume statistical independence between successive draws. Thus, the expected time interval an individual has to wait between opportunities to change the strategy is 1.

The following definition of a revision protocol⁴ formulates how agents choose a strategy given a revision opportunity. In what follows, we assume for simplicity that all agents display the same strategy-updating behavior, i.e., act according to the same revision protocol.

Definition 3.4 (Revision Protocol). A revision protocol is a function $\rho : \Delta^{m-1} \times \mathbb{R}^m \rightarrow \mathbb{R}_{\geq 0}^{m \times m}$ with $\sum_{j=1}^m \rho_{ij}(x, \pi) = 1$ for each $i \in S$, all population states $x \in \Delta^{m-1}$, and all possible payoffs $\pi \in \mathbb{R}^m$.

The revision protocol thus associates to each population state $x \in \Delta^{m-1}$ and payoffs $\pi \in \mathbb{R}^m$ a matrix of transition probabilities⁵ $\rho(x, \pi) = (\rho_{ij}(x, \pi))_{i,j=1,\dots,m}$ where $\rho_{ij}(x, \pi)$ represents the probability of the agent to switch from strategy i to strategy j given the current population state x and payoffs π .

Remark 3.1. Note that we modeled a revision protocol as a map with domain $\Delta^{m-1} \times \mathbb{R}^m$. Of course, if we focus on a specific population of size n , we could instead focus on maps with domain $\Delta_n^{m-1} \times \mathbb{R}^m$. Revision protocols are, however, often defined independent of a specific population and thus independent of a specific population size. For that reason, we prefer the domain to be $\Delta^{m-1} \times \mathbb{R}^m$.

We now give prominent examples of revision protocols.

Example 3.3 (Best response). The *best response* (BR) strategy updating protocol with its resulting mean dynamics (see Section 3.2.2) is one of the leading dynamics studied in evolutionary game theory. It is defined as follows: Let $B : \Delta^{m-1} \rightarrow \mathfrak{P}(S)$, where $\mathfrak{P}(S)$ is the power set of S , be the *best response function* defined by

$$B(x) = \{i \in S : F_i(x) \geq F_j(x) \text{ for all } j \in S\} \quad (3.3a)$$

$$= \operatorname{argmax}_{i \in S} F_i(x), \quad (3.3b)$$

i.e., $B(x)$ represents the set of strategies that optimize payoff given the current population state x . We shall write (with obvious abuse of notation) $j = B(x)$

⁴The definition of a revision protocol given here differs from the one given in (Sandholm, 2010) in that we consider time-discrete updating processes instead of time-continuous ones.

⁵The revision protocol is not to be confused with the transition matrix of the aggregate strategy updating process, see below.

in the case of a unique best response j to the population state $x \in \Delta_n^{m-1}$. Under the best response revision protocol, updating players always choose a best response to the current population state (Gilboa and Matsui, 1991). More formally, for $j \neq i$,

$$\rho_{ij}(x, F(x)) = 0 \text{ if } j \notin B(x), \quad (3.4a)$$

$$\rho_{ij}(x, F(x)) \in [0, 1] \text{ if } j \in B(x), \text{ such that} \quad (3.4b)$$

$$\begin{cases} \sum_{j \in B(x), j \neq i} \rho_{ij}(x, F(x)) = 1 & \text{if } i \notin B(x) \\ \sum_{j \in B(x), j \neq i} \rho_{ij}(x, F(x)) \leq 1 & \text{if } i \in B(x). \end{cases} \quad (3.4c)$$

In words, the revising player chooses a probability distribution on the set of best replies to the current population state⁶. The resulting revision protocol is therefore not uniquely defined.

One often-used possibility to circumvent this is to focus on population games where the best responses are uniquely determined for each population state. In the case of the currency game (Example 3.2), this could be established by slightly changing the payoff parameters. More specifically, the best response is unique for all population states $x \in \Delta_n^1$ with $x_1 \neq \frac{b}{a+b}$. Thus, if $x_1 \neq \frac{b}{a+b}$ for all $x \in \Delta_n^1$, the best response revision protocol is well-defined. To ensure that $x_1 \neq \frac{b}{a+b}$ for all $x \in \Delta_n^1$, we could assume, e.g., that a is irrational. The resulting revision protocol is provided by

$$\rho_{21}(x, F(x)) = \begin{cases} 1 & \text{if } x_1 > \frac{b}{a+b} \\ 0 & \text{if } x_1 < \frac{b}{a+b} \end{cases} \quad (3.5)$$

$$\rho_{12}(x, F(x)) = \begin{cases} 0 & \text{if } x_1 > \frac{b}{a+b} \\ 1 & \text{if } x_1 < \frac{b}{a+b}. \end{cases} \quad (3.6)$$

Example 3.4 (Best response with mutations). A revising agent using the *best response with mutations* (BRM) revision protocol at *mutation rate* $0 \leq \varepsilon < 1$ updates his strategy choice as follows: with probability $(1 - \varepsilon)$ he chooses a best response $b \in B(x)$ to the current population state, while with probability ε he chooses a strategy $s \in S$ at random (uniform distribution; Kandori et al., 1993; Young, 1993a). Note that for $\varepsilon = 0$, the BRM revision protocol reduces to the best response revision protocol (Example 3.3).

More formally, this means for $j \neq i$,

$$\rho_{ij}(x, F(x)) = \frac{\varepsilon}{m} \text{ if } j \notin B(x) \quad (3.7a)$$

$$\rho_{ij}(x, F(x)) \in \left[\frac{\varepsilon}{m}, 1 - \frac{m-1}{m}\varepsilon \right] \text{ if } j \in B(x) \text{ and such that} \\ \begin{cases} \sum_{j \in B(x), j \neq i} (\rho_{ij}(x, F(x)) - \frac{\varepsilon}{m}) = 1 - \varepsilon & \text{if } i \notin B(x), \\ \sum_{j \in B(x), j \neq i} (\rho_{ij}(x, F(x)) - \frac{\varepsilon}{m}) \leq 1 - \varepsilon & \text{if } i \in B(x). \end{cases} \quad (3.7b)$$

Note that Eq. (3.7b) translates in the case of a unique optimal strategy given the population state $x \in \Delta_n^{m-1}$ to the following

$$\rho_{ij}(x, F(x)) = \begin{cases} 1 - \frac{m-1}{m}\varepsilon & \text{if } j = B(x) \\ \frac{m-1}{m}\varepsilon & \text{otherwise.} \end{cases} \quad (3.8)$$

⁶This is called a *mixed strategy* in game theory.

Under the BRM revision protocol, the probability of choosing a suboptimal strategy is the same for each suboptimal strategy, e.g., it is independent of the payoffs associated with each suboptimal strategy. This means that the mutations do not favor alternatives with higher payoffs over those with lower payoffs. As $\varepsilon \rightarrow 0$, the protocol chooses the best replies with probability one.

More specifically, in the case of the currency game as defined in Example 3.2, the best response with mutation revision protocol is given by (assuming, as in the case of the best response revision protocol that $x_1 \neq \frac{b}{a+b}$ for all $x \in \Delta_n^1$):

$$\rho_{21}(x, F(x)) = \begin{cases} 1 - \frac{\varepsilon}{2} & \text{if } x_1 > \frac{b}{a+b} \\ \frac{\varepsilon}{2} & \text{if } x_1 < \frac{b}{a+b} \end{cases} \quad (3.9)$$

$$\rho_{12}(x, F(x)) = \begin{cases} \frac{\varepsilon}{2} & \text{if } x_1 > \frac{b}{a+b} \\ 1 - \frac{\varepsilon}{2} & \text{if } x_1 < \frac{b}{a+b} \end{cases} \quad (3.10)$$

Notice that $\rho_{12}(x, F(x))$ and $\rho_{21}(x, F(x))$ are not continuous in x . Moreover, it is only the ratio $\frac{b}{a+b}$ that matters for the BRM revision protocol and not the specific parameters a and b . We can therefore restrict the analysis of the currency game to the $b = 1$ case⁷.


Figure 3.4a displays ρ_{12} as a function of the proportion x_1 of strategy 1 players in the population for various values of ε with $a = 1, b = 1$. Notice that in the $m = 2$ case we consider here, the proportion x_1 uniquely determines the values of the revision protocol since given the proportion x_1 of strategy 1 players, the proportion x_2 of strategy 2 players and thus the population state $x = (x_1, x_2) \in \Delta_n^1$ is given by $1 - x_1$ and $(x_1, 1 - x_1)$, respectively. Moreover, notice that the corresponding ρ_{21} is given by $\rho_{21} = 1 - \rho_{12}$. It shows that for $\varepsilon \rightarrow 0$, the functions approach those of the best response (without mutations) revision protocol. For $\varepsilon \rightarrow 1$, the functions approach those of a uniform choice protocol. Figure 3.4b displays the function for various values of the payoff parameter a with $\varepsilon = .3, b = 1$. It shows how the value $\rho_{12}(x, F(x))$ (and thus also $\rho_{21}(x, F(x))$) depends on whether $x_1 < \frac{b}{a+b}$ or $x_1 > \frac{b}{a+b}$. For sake of clarity, we do not consider a value of $x = (b/(a+b), 1 - b/(a+b))$ in the figures since the value of these functions at these points is not uniquely defined. It might assume any value in $[\frac{\varepsilon}{2}, 1 - \frac{\varepsilon}{2}]$.

Example 3.5 (Logit choice). The logit choice revision protocol is one of the most widely used models of discrete choice making (McFadden, 2001; Train, 2003) and has a long history in economics and psychology dating back to Thurstone (1927). In the context of stochastic evolutionary games it has been introduced by Blume (1993). It is defined for $j \neq i$ by


$$\rho_{ij}(x, F(x)) = \frac{\exp(\sigma F_j(x))}{\sum_{i \in S} \exp(\sigma F_i(x))} \quad (3.11)$$

where $F : \Delta_n^{m-1} \rightarrow \mathbb{R}_{\geq 0}^m$ denotes the payoff function of the population game. The parameter $\sigma \geq 0$ is called the *noise level*. If $\sigma = 0$, the choice probabilities are uniform; if $\sigma \rightarrow \infty$, then the revision protocol approximates the best response revision protocol. In contrast to the best response revision protocol, the logit choice is continuous in x and is therefore also called *smoothed best response* (see Figure 3.5).

⁷If $b \neq 1$, we can find a' such that $\frac{1}{a'+1} = \frac{b}{a+b}$. Since the ratios coincide, so do the corresponding revision protocols.




(a)




(b)

Figure 3.4: Best response with mutations revision protocol: ρ_{12} as a function of the proportion x_1 of strategy 1 players: (a) shows the functions for various values of ϵ , where the payoff parameter a is set to 1; (b) shows the functions for various values of the payoff parameter a and $\epsilon = .3$.

Rewriting Eq. (3.11) using payoff differences to an optimal strategy k^* , i.e.,



(a)



(b)

Figure 3.5: Logit choice revision protocol: $\rho_{12}(x)$ as a function of the proportion x_1 of strategy 1 players. (a) shows the function for various values of the noise parameter σ and payoff parameter $a = 1$; (b) shows the same for $a = 10$. As before, the graph of the function ρ_{21} can be deduced from the one of ρ_{12} via $\rho_{21} = 1 - \rho_{12}$.

by multiplying both numerator and denominator by $\exp(-\sigma F_{k^*}(x))$, we get

$$\rho_{ij}(x, F(x)) = \frac{\exp(\sigma(F_j(x) - F_{k^*}(x)))}{\sum_{i \in S} \exp(\sigma(F_i(x) - F_{k^*}(x)))} \quad (3.12)$$

and see that in contrast to the BRM revision protocol, the logit choice revision protocol puts a higher probability on choosing a strategy with a higher payoff, i.e., better strategies are more likely to be chosen. In other words, more costly mistakes are less likely to be made. In the context of the currency game (Example 3.2), this is also visible by comparing Figure 3.4 with Figure 3.5. The values of $\rho_{12}(x, F(x))$ and $\rho_{21}(x, F(x))$ for the BRM revision protocol only depend on whether $x_1 < \frac{b}{a+b}$, $x_1 = \frac{b}{a+b}$, or $x_1 > \frac{b}{a+b}$ (and not on the specific x). In contrast, in the case of the logit choice revision protocol, $\rho_{12}(x, F(x))$ and $\rho_{21}(x, F(x))$ depend on the specific values of x and $F(x)$. More specifically, for small σ , the functions are roughly constant (and thus reflect a uniform choice distribution); for medium size values of σ (e.g., $\sigma = 5$ in Figure 3.5), the functions reflect that the revision protocol puts a higher probability on choosing a strategy with a higher payoff, or in other words, the smaller the difference in payoffs between the two strategies (this is the case for population states x with x_1 approaching $\frac{b}{a+b}$), the more the values of the revision protocol reflects a uniform choice distribution on the two strategies (this is, $\rho_{12}(x, F(x)) \rightarrow .5$ and $\rho_{21}(x, F(x)) \rightarrow .5$ for $x \rightarrow (\frac{b}{a+b}, 1 - \frac{b}{a+b})$). For large values of σ , the functions approach, however, the best response revision protocol ($\varepsilon = 0$) and thus the less the payoff difference matters (see also Figure 3.5). More specifically, the logit choice revision protocol pointwise converges to the best response protocol for $\sigma \rightarrow \infty$ for all $x \in \Delta_n^{m-1}$ for which the best response is unique (since $\exp(\sigma(F_j(x) - F_{k^*}(x))) \rightarrow 1$ for $\sigma \rightarrow \infty$ if $j \in B(x)$ and $\exp(\sigma(F_j(x) - F_{k^*}(x))) \rightarrow 0$ for $\sigma \rightarrow \infty$ if $j \notin B(x)$).

Notice that our running example in the following chapters builds on the BRM revision protocol. In Chapter 6.3, we consider an extended example in which we concentrate on the logit choice revision protocol. For more examples of revision protocols we refer the reader to Sandholm (2010).

We now turn to the aggregate strategy updating process; given a given population game with payoff function F and revision protocol ρ , only one agent is drawn from the whole population (with uniform distribution) to reconsider its strategy choice. This means, on the aggregate level, that transitions between population states are only possible, i.e., have a probability greater 0, if they differ in at most one component by at most $1/n$. Moreover, the probability of drawing an agent that currently holds strategy $i \in S$ corresponds to the share x_i of strategy i in the current population state x , and the probability that an agent holding strategy i changes to strategy j when given the chance to reconsider the strategy choice is given by $\rho_{ij}(x, F(x))$. Assuming statistical independence, the strategy updating process on the population level is thus a time-discrete Markov chain $X = (X_t)_{t \in \mathbb{T}}$ on the set of population states Δ_n^{m-1} , where $\mathbb{T} = \{k\delta \mid k \in \mathbb{N}, \delta = 1/n\}$. Its transition matrix $P = (p_{xy})_{x, y \in \Delta_n^{m-1}}$ is

given by

$$p_{xy} = \begin{cases} x_i \rho_{ij}(x, F(x)) & \text{if } y = x + \frac{1}{n}(e_j - e_i), i, j \in S, i \neq j, \\ 1 - \sum_{i \in S} \sum_{j \neq i} x_i \rho_{ij}(x, F(x)) & \text{if } x = y, \\ 0 & \text{otherwise.} \end{cases}$$

In the examples that follow, we will mostly concentrate on games with two strategies only. In this case, we can identify the state $x \in \Delta_n^1 \subset \mathbb{R}^2$ with $\chi = x_1$ (since $x_2 = 1 - x_1$). We can thus restrict our analysis of the chain $(X_t)_{t \in \mathbb{T}}$ to the state space $Z = \{0, \frac{1}{n}, \dots, 1\}$ and we will write (with abuse of notation) $F(\chi)$ for $F(x)$ (i.e., $F : Z \rightarrow \mathbb{R}^2$) and $\rho(\chi, F(\chi))$ for $\rho(x, F(x))$ (i.e., $\rho : Z \times \mathbb{R}^2 \rightarrow \mathbb{R}_{\geq 0}^{2 \times 2}$). Moreover, the aforementioned considerations about representing general evolutionary games as discrete-time Markov chains imply that stochastic evolutionary games with two strategies are birth-and-death chains (see Examples 2.1 and 2.2) with transition matrix

$$P = \begin{pmatrix} 1 - \alpha_0 & \alpha_0 & 0 & \dots & & & \\ \beta_1 & 1 - (\beta_1 + \alpha_1) & \alpha_1 & 0 & \dots & & \\ 0 & \beta_2 & 1 - (\beta_2 + \alpha_2) & \alpha_2 & 0 & \dots & \\ & & \vdots & & & & \\ 0 & \dots & & & 0 & \beta_n & 1 - \beta_n \end{pmatrix}, \quad (3.13)$$

where the parameters are

$$\alpha_j = \left(1 - \frac{j}{n}\right) \rho_{21}\left(\frac{j}{n}, F\left(\frac{j}{n}\right)\right) \quad \text{for } j = 0, \dots, n-1 \quad (3.14)$$

$$\beta_j = \frac{j}{n} \rho_{12}\left(\frac{j}{n}, F\left(\frac{j}{n}\right)\right) \quad \text{for } j = 1, \dots, n. \quad (3.15)$$

Example 3.6. Consider the currency game introduced in Example 3.2, and let $\rho : \Delta_n^1 \times \mathbb{R}^2 \rightarrow \mathbb{R}_{\geq 0}^{2 \times 2}$ be the BRM revision protocol with noise parameter ε . The resulting evolutionary game represented by the discrete-time Markov chain $(X_t)_{t \in \mathbb{T}}$ with state space $Z = \{0, \frac{1}{n}, \dots, 1\}$ is going to be *the* running example in the following sections and chapters and will be referred to as *our (running) example*. Figure 3.6 gives an impression of characteristic sample paths for a resulting evolutionary process with parameters $a = b = 1, n = 11, \varepsilon = .3$. It shows the characteristic metastable behavior that we are going to focus on in the subsequent chapters.

3.2 Analysis

This section considers existing approaches to the analysis of evolutionary games and puts them in the context of metastability.




Figure 3.6: Typical sample path of the number of agents holding strategy 1 in the evolutionary game defined by the currency game and the BRM revision protocol ($a = b = 1$, $n = 11$, $\varepsilon = .3$).

3.2.1 Nash Equilibrium and Refinements

The evolutionary game theory program is a different approach to the problem of equilibrium selection, an approach that is viewed as more realistic than Nash equilibrium refinements since it does not rely on overly idealistic requirements on the behavior of the agents (see Section 3.1.1). The body of literature on evolutionary game theory and its approach to the problem of equilibrium selection is large (for a comprehensive introduction and references, see, e.g., Sandholm, 2010; Weibull, 1995). Here, we can only give a rough overview necessary for an understanding of the subsequent sections.

The aim of the evolutionary game theory approach to the problem of equilibrium selection is an actual population foundation that emphasizes the individual decision-making of boundedly rational players and wants to observe the resulting aggregate dynamics. Typical questions in this context are:

- Can a Nash equilibrium of the underlying game in strategic form be implemented in the population of agents in the sense that to each of these Nash equilibria there is a distribution of strategies in the population?
- Does a Nash equilibrium arise as the long-run outcome of an evolutionary game?
- Does a Nash equilibrium emerge in the limit of large population sizes $n \rightarrow \infty$?

As we will see in the following sections, the second of these questions relates to the analysis of the stochastically stable states of an evolutionary game (Section 3.2.4) while the third is linked to its deterministic approximation (Section 3.2.2). With respect to the first question, consider the following definition:

Definition 3.5 (Nash Equilibrium of Finite-Population Games). The population state $x \in \Delta_n^{m-1}$ is a *Nash equilibrium of the finite population game* with payoff function $F : \Delta_n^{m-1} \rightarrow \mathbb{R}^m$ if no individual player receives a higher payoff by changing his strategy. More formally, this means that $x \in \Delta_n^{m-1}$ is a Nash equilibrium for the finite population game if

$$x_i > 0 \Rightarrow F_i(x) \geq F_j(x) \text{ for all } i, j \in S. \quad (3.16)$$

We denote the set of Nash equilibria of a finite population game with payoff function F by $\text{NE}(F)$.

Let F be the payoff function of a finite population game which is defined by a complete matching of agents to play a strategic form game (see Example 3.2). The definition then implies that every Nash equilibrium of the population game is a Nash equilibrium of the strategic game. In this case, we can view the Nash equilibrium of the population game as *implementing* the Nash equilibrium of the underlying strategic form game. We only mention that the other way around needs not hold, which is different to the case of continuous population games, and refer to Sandholm (2010) and Bernergård and Wärneryd (2011) for examples and discussion of further particularities of finite population games.

Example. In the currency game with parameters a, b , and a population of size n , the set of Nash equilibria of the population game resulting from a complete matching of the agents is

$$\text{NE}(F) = \begin{cases} \{(0, 1), (1, 0), (\frac{b}{b+a}, 1 - \frac{b}{a+b})\} & \text{if } (\frac{b}{a+b}, 1 - \frac{b}{a+b}) \in \Delta_n^1 \\ \{(0, 1), (1, 0)\} & \text{otherwise.} \end{cases} \quad (3.17)$$

The Nash equilibria of the population game thus relate to the Nash equilibria of the underlying pure coordination game (Example 3.1). If $(\frac{b}{a+b}, 1 - \frac{b}{a+b})$ is a population state, i.e., $(\frac{b}{a+b}, 1 - \frac{b}{a+b}) \in \Delta_n^1$ it corresponds to a *mixed* strategy Nash equilibrium, i.e., to a probability distribution over strategies which is a Nash equilibrium, for a definition see Sandholm (2010).

Lastly, with a view on metastability, we note that the notion of Nash equilibrium is obviously not a dynamic notion. If the evolutionary game has metastable dynamics between different Nash equilibria, an analysis of the dynamics from the viewpoint of metastability might enhance the understanding of the evolutionary process of interest.

3.2.2 Deterministic Approximation

This section addresses the question of how well stochastic evolutionary games can be approximated by deterministic, mean dynamics. There are two reasons why this question is of interest. First, from a mathematical and computational point of view, the analysis of the differential equations that describe the mean dynamics of the stochastic evolutionary game under consideration is often more

feasible and efficient. Second, major results in evolutionary game theory assume a continuum of agents, where the dynamics are considered deterministic and are given in terms of ordinary differential equations that are viewed to describe the changes in population shares over a large number of individual strategy switches (e.g., Weibull, 1995). In order for these results to be of relevance, the solutions to the differential equations describing the mean dynamics need to be good approximations of the discrete-time Markov chains that represent the stochastic evolutionary games in finite but large populations.

We first consider the case of mean dynamics with a Lipschitz continuous vector-field. After this, we turn to best response dynamics. Lastly, we interpret the results with a view on metastability and equilibrium selection.

In the following, let $(X^n(t))_{t \in \mathbb{T}}$, where $\mathbb{T} = \{\delta_n k \mid k \in \mathbb{N}, \delta_n = 1/n\}$, be the Markov chain defined by the population game $(n, S, \Delta_n^{m-1}, F)$ and the revision protocol ρ . Since we are going to consider different population sizes, it is important for us to note that this notation makes the dependence of the Markov chain on the population size explicit; while the state space Δ_n^{m-1} as well as the time steps δ_n depend on n , we assume for simplicity that both the payoff function F as well as the revision protocol ρ do not depend on the population size n . Under certain conditions, however, the following considerations and results hold also in the case of population size dependent payoff functions or revision protocols (Kurtz, 1970; Sandholm, 2010). Note, moreover, that in the following, equilibria (without the ‘‘Nash’’ prefix) always refer to equilibria in the dynamical systems, not to Nash equilibria. If we speak of Nash equilibria, we will always use the prefix ‘‘Nash’’.

Lipschitz Continuous Vector-Field

As shown in Section 3.1.3, $(X^n(t))_{t \in \mathbb{T}}$ is a Markov chain with transition matrix $P = (p_{xy})_{x, y \in \Delta_n^{m-1}}$ given by

$$p_{xy} = \begin{cases} x_i \rho_{ij}(x, F(x)) & \text{if } y = x + \frac{1}{n}(e_j - e_i), \ i, j \in S, \ i \neq j, \\ 1 - \sum_{i \in S} \sum_{j \neq i} x_i \rho_{ij}(x, F(x)) & \text{if } x = y, \\ 0 & \text{otherwise.} \end{cases}$$

Now we are in a position to calculate the *expected net increase per time unit* $V_i(x)$ associated with strategy $i \in S$ in the proportion of players conditional on the current population state $x \in \Delta_n^{m-1}$; that is, V_i is a function of type

$\Delta_n^{m-1} \rightarrow \mathbb{R}$ defined by:

$$V_i(x) = \frac{1}{\delta_n} \mathbb{E} \left(X_i^n((k+1)\delta_n) - X_i^n(k\delta_n) \mid X^n(k\delta_n) = x \right) \quad (3.18a)$$

$$= \frac{1}{\delta_n} \sum_{y \in \Delta_n^{m-1}} (y_i - x_i) \mathbb{P} \left(X_i^n((k+1)\delta_n) = y_i \mid X^n(k\delta_n) = x \right) \quad (3.18b)$$

$$= \frac{1}{\delta_n} \sum_{y \in \Delta_n^{m-1}, y_i \neq x_i} (y_i - x_i) p_{xy} \quad (3.18c)$$

$$= \frac{1}{\delta_n} \left(\sum_{j \neq i} \frac{1}{n} x_j \rho_{ji}(x, F(x)) - \sum_{j \neq i} \frac{1}{n} x_i \rho_{ij}(x, F(x)) \right) \quad (3.18d)$$

$$= \sum_{j \neq i} x_j \rho_{ji}(x, F(x)) - x_i \sum_{j \neq i} \rho_{ij}(x, F(x)), \quad (3.18e)$$

where the second to last equals sign follows from the fact that if $y_i \neq x_i$ and neither $y = x + \frac{1}{n}(e_i - e_j)$ nor $y = x + \frac{1}{n}(e_j - e_i)$ for some $j \in S, j \neq i$, then $p_{xy}^n = 0$.

Notice that since we assume that neither F nor ρ depends on n , the function V_i as defined by Eq. (3.18e) does also not depend on n . This means that it is uniquely defined for all $x \in \Delta^{m-1}$ and thus we can view it as a function $V_i : \Delta^{m-1} \rightarrow \mathbb{R}$. Moreover, we have

$$\sum_{i \in S} V_i(x) = 0 \quad (3.19)$$

for every $x \in \Delta^{m-1}$ and thus, the V_i 's together yield a function $V : \Delta^{m-1} \rightarrow \mathbb{T}\Delta^{m-1}$, where $\mathbb{T}\Delta^{m-1} = \{z \in \mathbb{R}^m \mid \sum_{i \in S} z_i = 0\}$ is the tangent space of Δ^{m-1} , and $V(x) = (V_1(x), \dots, V_m(x))$. Notice that if F is a continuous payoff function and ρ a continuous revision protocol, then V is continuous and bounded as well.

In what follows, we will identify Δ^{m-1} with $\mathbb{R}_{\geq 0}^{m-1}$ and $\mathbb{T}\Delta^{m-1}$ with \mathbb{R}^{m-1} . This is possible because in both cases the spaces are isomorph. Moreover, by Tietze's extension theorem, we can extend a continuous and bounded function $\mathbb{R}_{\geq 0}^{m-1} \rightarrow \mathbb{R}^{m-1}$ to a function $\mathbb{R}^{m-1} \rightarrow \mathbb{R}^{m-1}$ with the same properties. It is in this way that we regard V in the following as a function $\mathbb{R}^{m-1} \rightarrow \mathbb{R}^{m-1}$. The function $V : \mathbb{R}^{m-1} \rightarrow \mathbb{R}^{m-1}$ is called the *vector-field associated with the Markov chain* $(X^n(t))_{t \in \mathbb{T}}$. It gives for large populations and short time intervals the *expected net increase* in each population *share* during the time interval, *per time unit*.

Now, consider the associated *mean-field equations*:

$$\dot{x}_i = V_i(x) = \sum_{j \neq i} x_j \rho_{ji}(x, F(x)) - x_i \sum_{j \neq i} \rho_{ij}(x, F(x)). \quad (3.20)$$

Assume that V is locally Lipschitz continuous; that is, V is continuous and for each $x_0 \in \Delta^{m-1}$ exists an open neighborhood $U(x_0) \subseteq \Delta^{m-1}$ such that there is a constant $L \in \mathbb{R}_{\geq 0}$ with

$$\|V(x) - V(y)\| \leq L\|x - y\| \quad \text{for each } x, y \in U(x_0). \quad (3.21)$$

Then we know from the theorem of Picard and Lindelöf that the initial value problem

$$\dot{x} = V(x), \quad x(0) = x_0$$

has a unique solution for every $x_0 \in \Delta^{m-1}$. Furthermore, the solutions leave the set of population states Δ^{m-1} *forward invariant* since $\sum_{i=1}^m V_i(x) = 0$, i.e., the sum of population shares are left constant, and $V_i(x) \geq 0$ if $x_i = 0$, i.e., population shares cannot turn negative. This just means that if $x(t)$ is a solution of

$$\dot{x} = V(x), \quad x(0) \in \Delta^{m-1},$$

then $x(t) \in \Delta^{m-1}$ for all $t \geq 0$. An example of a revision protocol which leads to a Lipschitz continuous vector field is the logit choice revision protocol, see Example 3.5. We will not, however, consider this example in depth since deterministic approximations are not the main focus of this thesis. We instead give an example in the case of the best response with mutations revision protocol (see below) and refer to, e.g., Sandholm (2010) for more on the logit choice revision protocol and its deterministic approximation.

The following result is due to Kurtz (1970). Stronger approximation results have recently been established by Benaim and Weibull (2003).

Theorem 3.1. *Let $(X^n(t))_{t \in \mathbb{T}}$ be the discrete-time Markov chain associated with the population game $(n, S, \Delta_n^{m-1}, F)$ and revision protocol ρ . Let V denote the associated vector-field and suppose V is locally Lipschitz continuous. Let $x(\cdot; x_0)$ denote the solution to the initial value problem*

$$\dot{x} = V(x), \quad x(0) = x_0.$$

Suppose that⁸ $X^n(0) = x_0^n \rightarrow x_0 \in \Delta^{m-1}$ for $n \rightarrow \infty$. Then for any $T > 0$ and $\varepsilon > 0$,

$$\lim_{n \rightarrow \infty} \mathbb{P} \left(\sup_{\tau \in [0, T]} \|X^n(\lfloor \tau \rfloor) - x(\tau; x_0)\|_\infty < \varepsilon \mid X^n(0) = x_0^n \right) = 1. \quad (3.22)$$

where $\|X^n(\lfloor \tau \rfloor) - x(\tau; x_0)\|_\infty = \max_{i \in S} |X_i^n(\lfloor \tau \rfloor) - x_i(\tau; x_0)|$ and $\lfloor \cdot \rfloor$ denotes the floor function, i.e., $\lfloor \tau \rfloor$ is the largest integer not greater than τ .

Thus, for any finite time horizon, the solution to the mean-field equations (uniformly) approximates the evolutionary games arbitrarily well as the population size goes to plus infinity. Put differently, given a finite time span T and $\varepsilon, \eta > 0$, then for large enough population sizes n , the probability that $(X^n(t))$ stays within ε of $(x(t; x_0))$ through time T is at least $1 - \eta$.

Lastly, we want to point out that since V is in general nonlinear, we have in the finite population case

$$\mathbb{E}(V(X^n(t))) \neq V(\mathbb{E}(X^n(t))), \quad (3.23)$$

and thus, the solution to the mean-field equations do not coincide with the mean of the sample paths in the finite population case, as one might expect from their name. It is only in the limit $n \rightarrow \infty$ that this is true.

⁸Notice that since the state space Δ_n^{m-1} changes with changing population size n , the initial state $X^n(0) \in \Delta_n^{m-1}$ of the Markov chain $(X^n(t))_{t \in \mathbb{T}}$ changes as well.

Best Response with Mutations Dynamics

The mean dynamics of the best response with mutations revision protocol can be derived similar to Eq. (3.18e), for details see Sandholm (2010). They are in general described by the differential inclusion

$$\dot{x} \in \left(1 - \frac{\varepsilon}{m}\right) BM(x) - x, \quad (3.24)$$

where $BM : \Delta^{m-1} \rightrightarrows \Delta^{m-1}$ is the (mixed strategies') *best response correspondence*⁹, i.e., $BM(x) \subseteq \Delta^{m-1}$ with $y \in BM(x) \equiv y \in \operatorname{argmax}_{z \in \Delta^{m-1}} z^T F(x)$. In other words, $BM(x)$ is the set of probability distributions over strategies that are a best reply to the current population state x . Eq. (3.24) is a differential inclusion. It can be shown that for continuous payoff functions F , the right-hand side of Eq. (3.24) is non-empty, convex and bounded as well as upper semicontinuous. Thus, there is at least one solution $x : \mathbb{R}_{\geq 0} \rightarrow \Delta^{m-1}$ with $x(0) = x_0$, which is absolutely continuous and satisfies Eq. (3.24) for almost every $\tau \in \mathbb{R}_{\geq 0}$ (Aubin and Cellina, 1984).

Note, that if $x(\tau)$ is a solution of Eq. (3.24) such that for all τ in the time interval $[0, T]$ the best response to the states $x(\tau)$ is unique, e.g., $i \in S$, then during this time interval, the evolution of the mean dynamics is described by the differential equation

$$\dot{x} = e_i \left(1 - \frac{\varepsilon}{m}\right) - x, \quad (3.25)$$

where e_i denotes the vector with 1 in the i th coordinate and 0's elsewhere. Thus, it is clear that the solution moves straight to $e_i(1 - \frac{\varepsilon}{m})$. It is explicitly given by

$$x(\tau) = (1 - e^{-\tau}) \left(1 - \frac{\varepsilon}{m}\right) e_i + e^{-\tau} x_0 \quad \text{for all } \tau \in [0, T]. \quad (3.26)$$

More specifically, in the case of our running example (see Example 3.6), the mean dynamics are given as follows:

$$\dot{x} = \begin{cases} (1 - \frac{\varepsilon}{2}) - x & \text{if } x > \frac{b}{a+b} \\ \frac{\varepsilon}{2} - x & \text{if } x < \frac{b}{a+b} \end{cases} \quad (3.27a)$$

$$\dot{x} \in \left(1 - \frac{\varepsilon}{2}\right)[0, 1] - x \quad \text{if } x = \frac{b}{a+b}. \quad (3.27b)$$

If $\frac{\varepsilon}{2} < \frac{b}{a+b}$ and $1 - \frac{\varepsilon}{2} > \frac{b}{a+b}$, the equilibria are the population states $x = \frac{\varepsilon}{2}$ and $x = 1 - \frac{\varepsilon}{2}$, respectively. Both are asymptotically stable. If $\frac{\varepsilon}{2} > \frac{b}{a+b}$ or $1 - \frac{\varepsilon}{2} < \frac{b}{a+b}$, then the equilibrium $x = \frac{\varepsilon}{2}$ or $x = 1 - \frac{\varepsilon}{2}$ disappears, respectively (see Figure 3.7).

Solutions to this differential equation with initial population density $x(0) = x_0 \neq \frac{b}{a+b}$ are

$$x(\tau) = \begin{cases} (1 - \frac{\varepsilon}{2}) - e^{-\tau} (1 - \frac{\varepsilon}{2} - x_0) & \text{if } x > \frac{b}{a+b} \\ \frac{\varepsilon}{2} - e^{-\tau} (\frac{\varepsilon}{2} - x_0) & \text{if } x < \frac{b}{a+b}. \end{cases} \quad (3.28)$$

⁹Note that BM differs from the function B introduced in Ex. 3.3 in that B considers "strategies" which are a best reply whereas BM chooses "probability distributions over strategies" which are a best reply. Probability distributions over strategies are called "mixed strategies" in game theory.




Figure 3.7: Equilibria of the best response with mutations ($\varepsilon = .3$) mean dynamics in our running example (Example 3.6) as a function of the payoff parameter a (assuming $b = 1$). The payoff ratio $1/(a + 1)$ divides the state space into two basins of attraction: all solutions of initial points $x_0 > 1/(a + 1)$ approach $x = 1 - \frac{\varepsilon}{2}$ provided that $1 - \frac{\varepsilon}{2} > \frac{1}{a+1}$, whereas all solutions with initial values $x_0 < 1/(a + 1)$ approach $x = \frac{\varepsilon}{2}$ provided $\frac{\varepsilon}{2} < \frac{1}{a+1}$.

As the mean dynamics of the best response with mutations revision protocol are given by a differential inclusion, whose vector-field is not Lipschitz continuous even when the mean dynamics reduce to differential equations, Theorem 3.1 does not apply. There is, however, an analogue for mean dynamics of best response with mutations revision protocols (Roth and Sandholm, 2013):

Theorem 3.2. *Let $(X^n(t))_{t \in \mathbb{T}}$ be the continuous-time Markov chain associated with the population game $(n, S, \Delta_n^{m-1}, F)$ and a best response with mutations protocol ρ , i.e., Eq. (3.7b) is satisfied for some $0 < \varepsilon < 1$. Let Φ denote the set of all solutions to the differential inclusion associated with the mean dynamics of the Markov chain for all possible initial states $x(0) = x \in \Delta^{m-1}$, i.e., Φ represents the set of all solutions to Eq. (3.24). Then for any $T > 0$ and $\alpha > 0$,*

$$\lim_{n \rightarrow \infty} \mathbb{P} \left(\inf_{z \in \Phi} \sup_{0 \leq s \leq T} \|X^n(\lfloor s \rfloor) - z(s)\|_\infty \geq \alpha \mid X^n(0) = x \right) = 0 \quad (3.29)$$

uniformly in x .

In words, if ρ is a best response with mutations revision protocol and $X^n = (X^n(t))_{t \in \mathbb{T}}$ is the associated discrete-time Markov chain, then for any $T >$

$0, \alpha > 0, \eta > 0$ there is at least one solution to the associated differential inclusion (not necessarily with the same initial condition as the one of X^n) such that the probability that sample paths of population shares of the evolutionary game X^n deviate more than α from the population shares determined by that solution in the time interval $[0, T]$ is less than η for all initial states of the Markov chain $X^n(0) = x \in \Delta_n^{m-1}$, provided n is large enough. See Figure 3.8 for a comparison of mean dynamics with sample paths in case of our example (Example 3.6).




Figure 3.8: Comparison of a single sample paths with the corresponding solution of the mean dynamics in our running example (see Example 3.6; $a = b = 1, n = 11, \varepsilon = .3$).

Deterministic Approximation in the Context of Metastability and Equilibrium Selection

We now put the results of this section into the context of Part II of this thesis. It should be clear that, as metastability is an inherently probabilistic notion, deterministic approximations of stochastic evolutionary games are in general not able to capture any metastable dynamic behavior, see, e.g., Figure 3.8.

The deterministic approximation results necessarily rely on bounded time intervals. This is obvious since, for instance, for fixed population size n , the evolutionary games that result from full support revision protocols (see Def. 3.6) are irreducible discrete-time Markov chains. Thus, every population state is visited infinitely often with probability one, which would not be possible if sample paths were to stay near the deterministic solution of the mean dynamics for all times.

However, on short time scales, the approximation results might give a hint on where to find metastable subsets of population state space. For certain

classes of stochastic evolutionary games, global convergence to Nash equilibria or perturbed¹⁰ Nash equilibria has been established (see, e.g., Sandholm, 2010; Weibull, 1995). As, for instance, in our running example (Example 3.6), the population states $1 - \frac{\varepsilon}{2}$ and $\frac{\varepsilon}{2}$ are asymptotically stable equilibria of the associated mean dynamics (provided $1 - \frac{\varepsilon}{2} > \frac{b}{a+b}$ and $\frac{\varepsilon}{2} < \frac{b}{a+b}$, respectively) and their basins of attraction almost cover the whole state space (except the population state $\frac{b}{a+b}$). We therefore know that if the deterministic dynamics start in a population state $x \neq \frac{b}{a+b}$, it will converge to either equilibrium. Both of these states can be understood as perturbed Nash equilibria. Now, if the convergence rate is not too slow, as is the case in our example where convergence is exponentially fast, we might thus expect the metastable subsets of state space around the perturbed Nash equilibria.

3.2.3 Stationary Distribution and Reversibility

Both the BRM and the logit choice revision protocol are so called *full revision protocols* that lead to an irreducible and aperiodic Markov chain $X = (X_t)_{t \in \mathbb{T}}$ at the population level:

Definition 3.6 (Full support revision protocol). A revision protocol $\rho : \Delta^{m-1} \times \mathbb{R}^m \rightarrow \mathbb{R}_{\geq 0}^{m \times m}$ is a *full support revision protocol* if there is a constant $C > 0$ such that

$$\rho_{ij}(x, \pi) \geq C \quad \text{for all } i, j \in S, x \in \Delta^{m-1}, \pi \in \mathbb{R}^m. \quad (3.30)$$

Thus, in the case of a full support revision protocol, the probability that a revising agent switches to the strategy j is always strictly positive. At the population level, it follows that there is a strictly positive possibility to move between two arbitrary population states $x_1, x_2 \in \Delta_n^{m-1}$. This means that the Markov chain at the population level is irreducible. It is also aperiodic since the probability of staying in the same population state is strictly positive.

Thus, as we have outlined in Section 2.2, Theorem 2.2, irreducibility and aperiodicity of the Markov chain implies

1. the existence of a unique invariant distribution μ ;
2. the convergence of the time t probability distributions to the unique invariant distribution as $t \rightarrow \infty$, no matter the initial distribution:

$$\lim_{t \rightarrow \infty} \mathbb{P}(X_t = j) = \mu_j \quad \text{for all } j \in S \quad (3.31)$$

(from an *ex-ante* point of view, this means that, in the sufficiently distant future, we would expect to find the chain to be in locations determined by μ);

¹⁰For instance, in the case of perturbed best response dynamics, it cannot be expected that the dynamics lead to Nash equilibria. The appropriate notion instead is that of *perturbed* Nash equilibria, which are population states “near” the Nash equilibria. How far the perturbed Nash equilibria are from the Nash equilibria usually depends on the perturbation parameter, ε and σ in our examples of the best response with mutations revision protocol, Ex. 3.4, and the logit choice revision protocol, Ex. 3.5, respectively (see, e.g., Sandholm, 2010).

3. that the proportion of time spent at each state $j \in S$ in the long run and along almost every sample path is again determined by μ (*ex-post* point of view):

$$\mathbb{P}_\nu \left(\lim_{t \rightarrow \infty} \int_0^t \mathbf{1}_{\{X_t=j\}} dt = \mu_j \right) = 1 \text{ for all } j \in S, \quad (3.32)$$

where ν denotes the initial distribution.

In the case of population games with two strategies, we can easily determine its unique stationary distribution (see Examples 2.1 and 2.2); for $\chi \in \{\frac{1}{n}, \dots, \frac{n-1}{n}, 1\}$, the weights $\mu(\chi)$ of the stationary distribution are determined by

$$\mu(\chi) = \mu(0) \frac{\alpha_0 \alpha_1 \cdots \alpha_{n\chi-1}}{\beta_1 \beta_2 \cdots \beta_{n\chi}} \quad (3.33a)$$

$$= \mu(0) \prod_{j=1}^{n\chi} \frac{\alpha_{(j-1)}}{\beta_j} \quad (3.33b)$$

$$= \mu(0) \prod_{j=1}^{n\chi} \frac{(1 - \frac{j-1}{n})}{\frac{j}{n}} \cdot \frac{\rho_{21}(\frac{j-1}{n}, F(\frac{j-1}{n}))}{\rho_{12}(\frac{j}{n}, F(\frac{j}{n}))} \quad (3.33c)$$

$$= \mu(0) \prod_{j=1}^{n\chi} \frac{n-j+1}{j} \cdot \frac{\rho_{21}(\frac{j-1}{n}, F(\frac{j-1}{n}))}{\rho_{12}(\frac{j}{n}, F(\frac{j}{n}))} \quad (3.33d)$$

with $\mu(0)$ determined by the restriction that $\sum_{\chi \in \Delta_1^n} \mu(\chi) = 1$.

The example 2.2 also showed that birth-and-death chains are reversible. This is an important property, on which most of the metastability methods introduced in the subsequent chapters rely. We only note here that population games with two strategies under full support revision protocols do not constitute the only class of examples with reversible dynamics. Another large class consists of finite-population potential games under the logit choice revision protocol (see Sandholm, 2010, and Chapter 6.3).

Example. Figure 3.9 shows the weights of the stationary distribution in our example (Example 3.6) with parameters $a = b = 1, n = 11$. This stationary distribution μ can be computed according to Eq. (3.33d). The example demonstrates that the stationary distribution can predict the empirical frequency of population states, but it does not tell us that typical sample paths stay in one subset of population state space for a long time and then switch to the other, as shown by the typical sample path in Figure 3.6.

3.2.4 Stochastic Stability Analysis

Introduction and Definition

The concept of stochastic stability is related to the infinite-horizon analysis of Markov chains. In Section 3.2.2, we summarized results on the deterministic approximation of stochastic evolutionary games and showed that over finite-time horizons, the mean dynamics provide a good prediction of the dynamic




Figure 3.9: Stationary distribution μ for our running example for $a = b = 1$, $n = 11$, $\epsilon = 0.3$ (see Example 3.6).

behavior of the Markov chain conditional on the initial state. However, if we are interested in the infinite-horizon dynamic behavior, we need to consider the stationary distribution, see Section 3.2.3. For full support revision protocols, the stationary distribution is unique, has positive mass on each population state, and is independent of the initial state – in contrast to the respective mean dynamics, which is initial state dependent. Consider, for example, a stochastic evolutionary game with multiple locally stable equilibria in the mean dynamics. Whereas the mean dynamics converge to a locally stable equilibrium, the stochastic dynamics of the evolutionary game can drive the system into a different basin of attraction and thus to a different locally stable equilibrium. Thus, the central motivation for studying the stationary distribution is that it might provide predictions about a population’s very long-run behavior independent of the initial state.

More specifically, the hope is that the limiting stationary distribution for the population size $n \rightarrow \infty$ or the noise parameter $\epsilon, \sigma \rightarrow 0$ converges to a point mass at a single population state; in the best case, this population state is a state in which all players use the same strategy, a Nash equilibrium of the underlying strategic game (see Figure 3.10 for the case of our running example, Example 3.6). In such a case, this population state is seen as the *conventional* way of playing the game and it is said that the agents *coordinate* on this Nash equilibrium (Young, 1993a). Moreover, in this case, the sketched approach can single out a Nash equilibrium as more plausible than the other Nash equilibria and thus provides an approach to the problem of equilibrium selection. Besides this game theoretical interest to study the limiting stationary distribution, there is also the practical motivation that the limiting distribution might be easier to

calculate than the stationary distribution for a specific ε .

Here, we focus on stochastic stability as the noise parameter approaches 0, which is what most of the literature does and which suits well the context of this thesis. We refer the reader interested in the limiting distribution when the population size n approaches ∞ to Sandholm (2010, Chapter 12 and corresponding chapter notes), which gives an overview of the small literature on this topic. In the following, we will always denote the noise parameter by ε , but the reader should keep in mind that in doing so we do not always refer to the best response with mutation revision protocol, but instead to the general case of a noise parameter dependent revision protocol.

In order to formulate a precise definition of stochastically stable states, let $(X^\varepsilon(t))_{t \in \mathbb{T}}$ denote the discrete-time Markov chain defined by the population game $(n, S, \Delta_n^{m-1}, F)$ and full support revision protocol ρ^ε . Note that this notation, in contrast to the one used in Section 3.2.2, makes the dependence on the noise parameter of the revision protocol as well as the resulting Markov chain explicit. The population size n is fixed. Let μ^ε denote its unique stationary distribution. If $\lim_{\varepsilon \rightarrow 0} \mu^\varepsilon$ exists, we can define:

Definition 3.7 (Stochastically Stable States). A population state $x \in \Delta_n^{m-1}$ is called *stochastically stable* if $\mu^*(x) = \lim_{\varepsilon \rightarrow 0} \mu^\varepsilon(x) > 0$.

This notion of stochastically stable states was introduced into evolutionary game theory by Foster and Young (1990), Young (1993a), and Kandori et al. (1993). Building largely on work by Freidlin and Wentzell (1984) on perturbed dynamical systems, they show how to calculate the stochastically stable states of Markov chains that comply with certain regularity conditions by solving a series of shortest path problems in graphs¹¹. Building on this approach, a large literature has developed with the goal of determining the stochastically stable states for different types of games and adjustment processes (e.g., Ben-Shoham et al., 2004; Binmore and Samuelson, 1997; Binmore et al., 2003; Ellison, 1993; Kandori and Rob, 1995, 1998; Nöldeke and Samuelson, 1993, 1997; Robson and Vega-Redondo, 1996; Serrano and Volij, 2008; Young, 1993b, 1998).

The algorithm by Foster and Young (1990), Kandori et al. (1993), and Young (1993a) has the advantage of being always applicable. However, it is difficult and not so intuitive to apply in the case of complex evolutionary models. It involves, for example, determining all recurrent classes of the underlying “unperturbed” process and the computation of a certain “least-cost tree” – which is a complicated graph-theoretic problem.

Here, we will focus on a different approach to the identification of stochastically stable states, which in contrast to the algorithm by Foster and Young (1990), Kandori et al. (1993), and Young (1993a), does not provide a necessary but only a sufficient condition for population states to be stochastically stable. For our purposes, it has the advantage of suggesting a relationship to metastability, which we will exploit in more detail in Section 4.3.

In short, the approach by Ellison (2000) is based on two measures, the *radius* and (*modified*) *coradius* describing the basins of attraction of the underlying unperturbed process. These two measures give bounds on the expected time to stay in a basin of attraction of a recurrence class and on the expected time to

¹¹In fact, the results of Young (1993a) can be considered a finite version of Freidlin and Wentzell’s results on continuous diffusion processes (see Young, 1993a, Appendix).

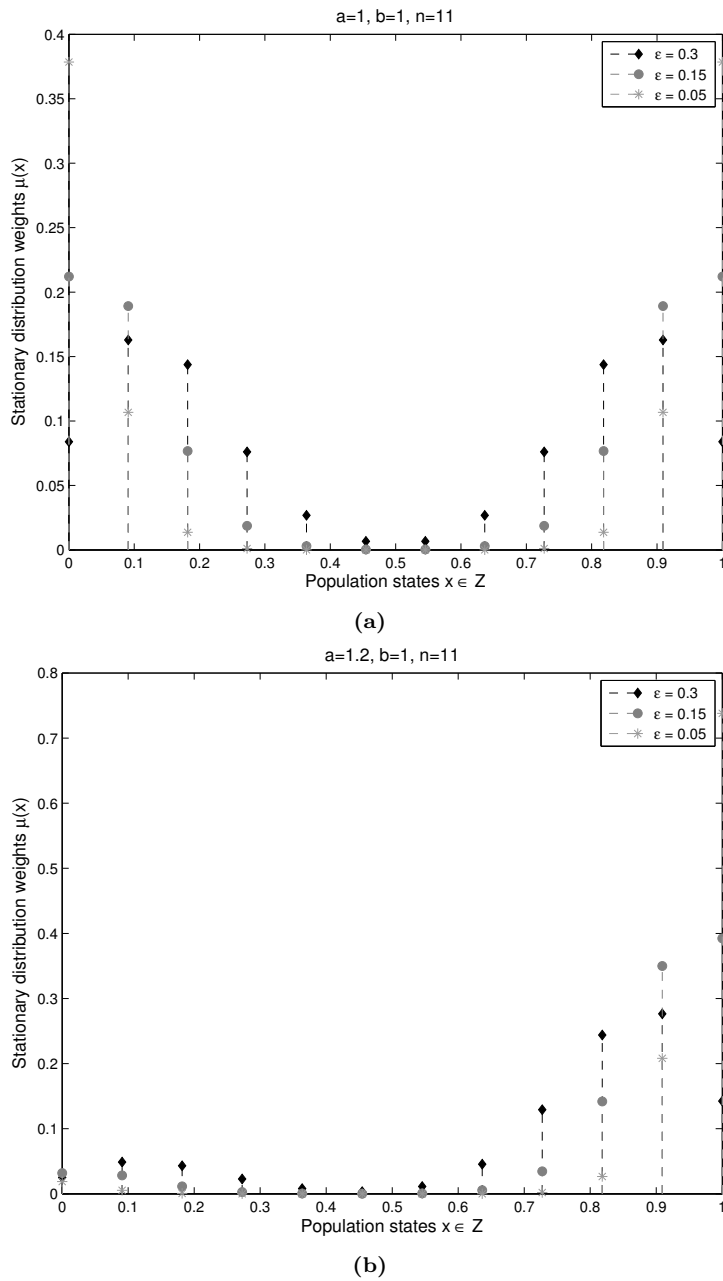


Figure 3.10: Stationary distribution weights of the evolutionary game in our running example (see Example 3.6) for $\epsilon = .3, .15, .05$. (3.10a): Parameters are $a = b = 1, n = 11$; for $\epsilon \rightarrow 0$, the weights approach .5 for the population states $x = 0, 1$ and 0 otherwise. (3.10b): Parameters are $a = 1.2, b = 1, n = 11$; for $\epsilon \rightarrow 0$, the stationary distribution seems to approach a point mass at $x = 1$.

return to a basin of attraction from the basin of attraction of another recurrence class, respectively. A recurrence class then is stochastically stable if the

expected time to return is smaller than the expected time to stay in that class. This approach thus not only determines the stochastically stable population states but gives additionally bounds on expected hitting times.

Identification of Stochastically Stable States

Ellison (2000) uses the following definition of a *model of evolution with noise*¹² as a framework for his approach:

Definition 3.8 (Model of Evolution with Noise). Let $(X^\varepsilon)_{\varepsilon \geq 0}$, be a family of discrete-time Markov chains $X^\varepsilon = (X^\varepsilon(k))_{k \in \mathbb{N}}$ on a finite state space Z . Let P^ε denote the respective transition matrices. We call $(X^\varepsilon)_{\varepsilon \geq 0}$ a *model of evolution with noise* if there exists an $\bar{\varepsilon} > 0$ such that

1. P^ε is ergodic for every $\varepsilon \in (0, \bar{\varepsilon})$,
2. P^ε is continuous in ε , and
3. there exists a (possibly asymmetric) cost function $c : Z \times Z \rightarrow \mathbb{R}_{\geq 0} \cup \{\infty\}$ such that for all pairs of states $z, z' \in Z$, $0 < \lim_{\varepsilon \rightarrow 0} P_{zz'}^\varepsilon / \varepsilon^{c(z, z')} < \infty$ exists if $c(z, z') < \infty$ and $P_{zz'}^\varepsilon = 0$ for sufficiently small ε if $c(z, z') = \infty$.

We call $X^0 = (X^0(k))_{k \in \mathbb{N}}$ the *unperturbed* Markov chain.

This definition captures many of the models studied in the literature. It implies the existence of a unique stationary distribution μ^ε for each $\varepsilon \in (0, \bar{\varepsilon})$ as well as of the limiting distribution $\mu^* = \lim_{\varepsilon \rightarrow 0} \mu^\varepsilon$ (see Ellison, 2000, p. 22, footnote 5). The cost function can be interpreted as indicating how unlikely it is for transitions between states to take place.

Example. For each $\varepsilon \in [0, 1)$, let $(X^\varepsilon(t))_{t \in \mathbb{T}}$ be the Markov chain of the evolutionary game defined by our running example (Example 3.6). We assume that $\frac{b}{a+b} \notin Z = \{0, \frac{1}{n}, \dots, 1\}$. Let P^ε denote the respective transition matrices. Then this family of Markov chains is a model of evolution with noise with cost function $c : Z \times Z \rightarrow \mathbb{R}_{\geq 0} \cup \{\infty\}$ defined by

$$c(x, y) = \begin{cases} 0 & \text{if } P_{xy}^0 > 0, \\ 1 & \text{if } (y < \frac{b}{a+b} \wedge y = x + \frac{1}{n}) \vee (y > \frac{b}{a+b} \wedge y = x - \frac{b}{a+b}) \vee \\ & (x = y \neq 0, 1), \\ \infty & \text{otherwise.} \end{cases}$$

In order to state the main results by Ellison (2000), we need the following definitions:

Definition 3.9 (Path, Cost of a Path). Let $A, B \subset Z$. A *path*¹³ from A to B is a finite sequence of distinct states (z_1, \dots, z_n) with $z_1 \in A, z_k \notin B$ for $k < n$,

¹²This definition differs from the one given by Ellison (2000) in that it is given in terms of the Markov chains and not just in terms of transition matrices.

¹³Note that the notion of a path given in this definition is not related to and should not be confused with the notion of a (sample) path of a stochastic process.

and $z_n \in B$. The *cost* of a path (z_1, \dots, z_n) is

$$c(z_1, \dots, z_n) = \sum_{k=1}^{n-1} c(z_k, z_{k+1}). \quad (3.34)$$

In what follows, we denote by $S(A, B)$ the set of all paths from A to B . Moreover, we extend the cost function on paths to a set-to-set cost function C by

$$C(A, B) = \inf_{(z_1, \dots, z_n) \in S(A, B)} c(z_1, \dots, z_n). \quad (3.35)$$

From Young (1993a) we know that the stochastically stable states of a model of evolution with noise are contained in the recurrent classes of the unperturbed chain. In the literature on evolutionary models with noise the recurrent sets are also called *limit sets*.

Definition 3.10 (Basin of Attraction). Let L be a limit set of the unperturbed Markov chain X^0 . The *basin of attraction* $\mathfrak{B}(L)$ of L is the set of states from which the unperturbed chain reaches L with probability one. More formally,

$$\mathfrak{B}(L) = \{z \in Z : \mathbb{P}_z(X^0(k) \in L) = 1 \text{ for some } k \in \mathbb{N}\}. \quad (3.36)$$

We are now in a position to introduce the radius and coradius of a basin of attraction:

Definition 3.11 (Radius of a Basin of Attraction of a Limit Set). The *radius* $\mathfrak{r}(L)$ of a basin of attraction of a limit set L is the minimum cost of any path starting in L and leaving its basin of attraction $\mathfrak{B}(L)$; that is,

$$\mathfrak{r}(L) = C(L, Z \setminus \mathfrak{B}(L)). \quad (3.37)$$

In essence, the radius counts the minimum number of “mutations” (ε - probability events) that are necessary to escape the basin of attraction of a limit set. Ellison (2000) shows that the ε -perturbed chain X^ε spends at least $k\varepsilon^{-\mathfrak{r}(L)}$ number of periods in $\mathfrak{B}(L)$ before leaving, where k is some constant. Put differently,

$$\frac{1}{\mathbb{E}_x \tau_y^\varepsilon} = \mathcal{O}(\varepsilon^{\mathfrak{r}(L)}) \quad \text{for all } x \in L, y \notin L, \quad (3.38)$$

where τ^ε denotes the first hitting time of the ε -perturbed chain X^ε (see Definition 2.3).

Definition 3.12 (Coradius of a Basin of Attraction of a Limit Set). The *coradius* $\mathfrak{cr}(L)$ of the basin of attraction of L is defined by

$$\mathfrak{cr}(L) = \sup_{y \notin L} C(y, L). \quad (3.39)$$

Note that $\sup_{y \notin L} C(y, L) = \sup_{y \notin \mathfrak{B}(L)} C(y, \mathfrak{B}(L))$ since the cost of getting to L from any element of $\mathfrak{B}(L)$ is zero (Ellison, 2000, p. 23). The coradius provides a measure of the maximum of the minimum number of mutations necessary to reach the limit set L from outside.

Ellison (2000) shows that the expected number of periods the chain X^ε spends outside of the limit set L is bounded from above by $k'\varepsilon^{-\mathbf{cr}(L)}$ for some constant k' , i.e., $\mathbb{E}_y\tau_L^\varepsilon = \mathcal{O}(\varepsilon^{-\mathbf{cr}(L)})$ for $y \notin L$. In this sense, the coradius presents a measure of the attractiveness of the limit set. The following modified coradius can be used to obtain a tighter bound.

Definition 3.13 (Modified Cost Function, Modified Coradius of a Basin of Attraction). Let L be a union of limit sets and let (z_1, \dots, z_n) be a path from $z_1 = y$ to L . Suppose that A_1, \dots, A_k is a sequence of limit sets the path hits consecutively with $A_i \not\subset L$ for $i < k$ and $A_k \subset L$, and with the convention that a limit set can appear on the list multiple times but not successively. The *modified cost function* c^* for such a path is defined by

$$c^*(z_1, \dots, z_n) = c(z_1, \dots, z_n) - \sum_{i=2}^{k-1} \mathbf{r}(A_i). \quad (3.40)$$

Extending this definition to a point-to-set concept by

$$c^*(y, L) = \inf_{(z_1, \dots, z_n) \in \mathcal{S}(\{y\}, L)} c^*(z_1, \dots, z_n), \quad (3.41)$$

the *modified coradius* is defined by

$$\mathbf{cr}^*(L) = \sup_{y \notin L} c^*(y, L). \quad (3.42)$$

Note that $\mathbf{cr}^*(L) \leq \mathbf{cr}(L)$, and $\mathbb{E}_y\tau_L^\varepsilon = \mathcal{O}(\varepsilon^{-\mathbf{cr}^*(L)})$ for $y \notin L$. The idea behind the modified coradius is that larger state changes happen more rapidly if they can be accomplished by passing through a sequence of limit sets.

Using the results on the radius and coradius of a basin of attraction, Ellison (2000) shows that if we have $\mathbf{r}(L) > \mathbf{cr}^*(L)$ for a limit set L , then this set L contains all stochastically stable states of the model of evolution with noise. We summarize the results in the following

Theorem 3.3. *Let $(X^\varepsilon)_{\varepsilon \geq 0}$ be a model of evolution of noise. Let L be a union of limit sets of X^0 . Then,*

$$\frac{1}{\mathbb{E}_x\tau_y^\varepsilon} = \mathcal{O}(\varepsilon^{\mathbf{r}(L)}) \quad \text{for all } x \in L, y \notin L, \text{ and} \quad (3.43)$$

$$\mathbb{E}_y\tau_L^\varepsilon = \mathcal{O}(\varepsilon^{-\mathbf{cr}^*(L)}) \quad \text{for } y \notin L. \quad (3.44)$$

Moreover, $\mathbf{r}(L) > \mathbf{cr}^*(L)$ implies that L contains all stochastically stable states.

Example 3.7. In our running example (Example 3.6), the recurrent sets of the unperturbed chain X^0 are the sets $L_0 = \{0\}$, $L_1 = \{1\}$. Their radii and coradii are given by

$$\mathbf{r}(L_0) = \left\lfloor \frac{b}{a+b}n \right\rfloor + 1 \quad (3.45)$$

$$\mathbf{r}(L_1) = n - \left\lfloor \frac{b}{a+b}n \right\rfloor \quad (3.46)$$

$$\mathbf{cr}^*(L_0) = \mathbf{cr}(L_0) = n + 1 - \mathbf{r}(L_0) = \mathbf{r}(L_1) \quad (3.47)$$

$$\mathbf{cr}^*(L_1) = \mathbf{cr}(L_1) = n + 1 - \mathbf{r}(L_1) = \mathbf{r}(L_0). \quad (3.48)$$

It thus follows from Theorem 3.3 that the stochastically stable population states are

- 0 and 1 if $a = b = 1$;
- 1 if $a > b$; and
- 0 if $b > a$.

This example reflects the impression we got from Figure 3.10. In addition, it points to the fact that stochastically stable states per se are not able to capture metastable dynamics. Firstly, in the case of more than one stochastically stable state, as is the case in our running example for the parameters $a = b = 1$, applying the notion of stochastic stability might tell us to expect these states in the limit $\varepsilon \rightarrow 0$, but it does not inform us about the dynamics between those states. Secondly, the notion of stochastic stability depends on the limit $\varepsilon \rightarrow 0$. We are, however, often interested in the behavior of a model of evolution with noise for fixed $\varepsilon > 0$. In this case, there might be other metastable states which are visited often, but less and less for smaller and smaller ε . In our running example, for instance, the state 0 is such a metastable state that is not stochastically stable for $a > b$, but for $\varepsilon > 0$ is still visited often. And again, the notion of stochastic stability does not inform us about the essential dynamics for $\varepsilon > 0$. In Chapter 5, we will introduce the Markov state modeling approach, which can be used to fill exactly this gap. Moreover, in Chapter 4.3, we study metastability in models of evolution with noise using the notions of radius and coradius introduced by Ellison (2000) in order to identify stochastically stable states.

Chapter 4

Characterization of Metastability

Metastability is a characteristic of dynamical systems exhibiting different time scales; on a short time scale, the system stays within a confined subset of the state space and thus appears to be in equilibrium, while at much longer time scales, transitions to other subsets of state space are possible, in which the system again stays for a long period of time. These areas of state space where the systems stays for a long period of time are called *metastable*. Often, metastability is not built into the model as such but results as a macroscopic feature from the microscopic specification of the system. In this case, the length and time scales of the microscopic specification can differ by many orders of magnitude from the resulting macroscopic effect.

Metastability is of interest in many different fields. Examples range from the classical example of phase transitions in physics, to conformational changes in large bio-molecules, long-term changes in the climate system, insect or pathogen outbreaks in population biology and epidemics, and business cycles and price fluctuations in economics and finance – just to name a few.

Here, we focus on metastability in stochastic evolutionary games. This is interesting because metastability is a dynamic property that many stochastic evolutionary games share. In the evolutionary game literature, this property is also called *punctuated equilibrium* and is considered a favorable property of stochastic evolutionary games (e.g., Young, 1998, 2006), not least because considering metastability in stochastic evolutionary games might lead to a different perspective on modeling conventions and the problem of equilibrium selection (Jaeger, 2008, 2012, see also the Introduction and Chapter 3). As we have outlined in the previous chapter (Section 3.2), however, the methods used so far in stochastic evolutionary game theory do not characterize the dynamics of the evolutionary games with respect to this property. This is in contrast to physics and chemistry where there has been much research in the last century on the mathematical description and analysis of metastability (for a short historical overview see, e.g., the introductory chapter of Bovier, 2009). It is this fact that we are going to exploit in this and the following chapters when we introduce methods to analyze metastability in stochastic evolutionary games that stem for the most part from the study of large bio-molecules (e.g., Deuffhard et al.,

2000; Huisinga, 2001; Huisinga and Schmidt, 2006; Schütte, 1998; Schütte et al., 1999; Schütte and Huisinga, 2003).

Against this background, it is the aim of this chapter to characterize the given intuitive notion of metastability in a more formal way. There are several characterizations of metastability (see, e.g., Bovier et al., 2001; Davies, 1982a,b; Schütte and Huisinga, 2003). We will present two different characterizations and relate these to the study of stochastic evolutionary games. Section 4.1 characterizes metastability in terms of the time the system spends within a subset of state space. According to this characterization, the subset is metastable if the system stays with high probability within the subset longer than some reference macroscopic time scale. Section 4.2, on the other hand, characterizes metastability in terms of the transition probability between sets. Here, a subset is called metastable if the probability of exiting this subset in the next time step is significantly small. We are going to relate these characterizations to the spectral properties of the transition matrix of the Markov chain under investigation. Moreover, Section 4.3 uses the approach of Ellison (2000) to relate both characterizations to the notion of stochastic stability as introduced in the previous chapter (Section 3.2.4).

In the subsequent chapter, we present different ways of approximating stochastic evolutionary games and assess to what extent these approximations capture the dynamic behavior on the metastable sets. The motivation stems from the fact that, often, one is interested in resolving not all the microscopic details of the dynamics, but to take advantage of the metastable property and to derive approximative models of reduced complexity that still capture the dynamical behavior on the length- and time-scales of interest.

One final comment seems to be in order at this point. As the introductory words to this chapter indicate, the concept of metastability is inherently an imprecise notion, as Davies (1982a) puts it:

“qualitative properties [of metastable states] are really not exact, but only appear so because various errors are small, and certain instabilities only appear after very long times” (p. 133).

This is going to be reflected in the formal characterizations of metastability in that they involve certain quantities to be *sufficiently* small or large.

4.1 Metastability based on Hitting Times

The intuitive characterization given at the beginning of this chapter leads to the following possible definition of a metastable Markov chain (also compare Bovier, 2006):

Definition 4.1 (Metastable Markov chain based on Mean Hitting Times). Let T denote some discrete time index set. The Markov chain $(X_t)_{t \in T}$ on the finite state space Z is called *metastable* if there are $k \geq 2$ disjoint subsets $C_1, \dots, C_k \subset Z$, such that

$$\frac{R_C}{\min_{i=1, \dots, k} W_{C_i}} \ll 1, \quad (4.1)$$

where $C = \cup_{i=1}^k C_i$ denotes the union of the subsets C_1, \dots, C_k ; R_C for $C \subsetneq Z$ is the *return time to C*, i.e.,

$$R_C = \sup_{x \notin C} \mathbb{E}_x \tau_C, \quad (4.2)$$

and in case $C = Z$, R_C denotes a reference time scale τ_{ref} on which single step transitions take place; and W_{C_i} for $i = 1, \dots, k$ is the (mean) *residence time* (or, *waiting time*) of the process in the subset C_i ; that is,

$$W_{C_i} = \frac{1}{\mu(C_i)} \sum_{x \in C_i} \mu(x) \mathbb{E}_x \tau_{C \setminus C_i}. \quad (4.3)$$

In this definition, the reference time scale τ_{ref} represents the time scale on which single step transitions take place and is built into the definition to capture the case of a full decomposition of state space for which the return time $R_C = -\infty$. In the case of the stochastic evolutionary games given by a population game $(n, S, \Delta_n^{m-1}, F)$ and revision protocol ρ , the reference time scale is $\tau_{\text{ref}} = 1/n$. Furthermore, R_C can be interpreted as the longest expected time the evolutionary process takes to return to C from outside of C , and W_{C_i} represents the μ -weighted mean time to exit the subset C_i .

Def. 4.1 of metastable discrete-time Markov chains is thus a straightforward translation of the intuitive notion of metastability given at the beginning of this chapter; a discrete-time Markov chain is metastable according to this definition if there are disjoint subsets of state space such that the residence times in the subsets are much longer than the return time to the union of the subsets. In other words, the process stays in general for a long time in one of the subsets and when it leaves a specific subset, it returns quickly to one of the subsets again, see Example 4.1 for the application of the definition to our example (Example 3.6).

Remark 4.1. In the case of a noise parameter dependent family of Markov chains such as models of evolutions with noise (Def. 3.8), the definition might be interpreted to mean that the left hand side of Eq. (4.1) gets small as $\varepsilon \rightarrow 0$, i.e.,

$$\frac{R_C}{\min_{i=1, \dots, k} W_{C_i}} = o(1) \quad \text{for } \varepsilon \rightarrow 0. \quad (4.4)$$

If $C \subsetneq Z$, it is possible to relate the return time R_C to spectral characteristics of the operator \mathcal{P} of the Markov chain under investigation. In order to do so, let \mathcal{P}_C denote the *projected* operator defined by $\mathcal{P}_C = S\mathcal{P}S$, where S is the orthogonal projection of \mathbb{R}^l onto $Z \setminus C$ (with respect to the standard scalar product $\langle \cdot, \cdot \rangle$). \mathcal{P}_C is thus an operator on the space \mathbb{R}^{l_C} , $l_C = l - |C|$. The matrix P_C representing the linear operator \mathcal{P}_C results from deleting all rows and columns associated with states in C from the transition matrix P , which represents the operator \mathcal{P} , see Eq. (2.13). Thus,

$$P_C(x, y) = P(x, y) \quad \text{for } x, y \in Z \setminus C. \quad (4.5)$$

Theorem 4.1 (Bovier (2006)). *Let λ_C be the largest eigenvalue of \mathcal{P}_C . Then,*

$$\frac{1}{1 - \lambda_C} \leq R_C. \quad (4.6)$$

This theorem provides a lower bound on the return time R_C to the set C from states outside of C . It implies that a necessary condition for quick returns to C is that the largest eigenvalue of \mathcal{P}_C is not near 1; otherwise the fraction on the left side of Eq. (4.6), and thus also the return time R_C , gets large. The projection of \mathcal{P} onto \mathbb{R}^l defined above must therefore remove all dominant timescales of the Markov chain under investigation (see Section 2.2, “Invariant Distribution and Spectral Properties”).

4.2 Metastability based on Transition Probabilities

The characterization of metastability in terms of transition probabilities is based on the idea that a metastable subset of state space of a system under investigation is *almost invariant* in the sense that the system exits this subset only with small probability. The presentation in this section is based on Huisinga (2001), Schütte and Huisinga (2003), and Huisinga and Schmidt (2006).

In the following, let $(X_k)_{k \in \mathbb{N}}$ denote a discrete-time, reversible Markov chain on a finite state space Z with $|Z| = l$ with transition matrix P and unique stationary distribution μ .

The characterization of metastability in terms of transition probabilities relies on the following notion of transition probabilities between subsets of state space:

Definition 4.2. Let $A, B \subseteq Z$. The *transition probability from A to B* is the conditional probability

$$p(A, B) = \mathbb{P}_\mu[X_1 \in B | X_0 \in A] = \frac{1}{\mu(A)} \sum_{x \in A, y \in B} \mu(x) P(x, y) \quad (4.7)$$

if $\mu(A) > 0$ and 0 otherwise, where $\mu(A) = \sum_{x \in A} \mu(x)$.

In words, $p(A, B)$ gives the probability of the system having started in A and distributed according to the stationary distribution to be found in set B after one time step and thus measures the dynamical fluctuations within the stationary distribution μ .

Definition 4.3 (Invariant Subset, Metastable Subset, Joint Metastability). We call $A \subseteq Z$ *invariant* if $p(A, A) = 1$ and *metastable* if $p(A, A) \approx 1$. The *joint metastability* of a partition $D = \{A_i\}_{i=1, \dots, m}$ of state space is defined by

$$M(D) = \sum_{i=1}^m p(A_i, A_i). \quad (4.8)$$

We call D metastable if $M(D) \approx m$.

Given an arbitrary decomposition $D = \{A_i\}_{i=1}^m$ of state space, we give in Theorem 4.2 sharp upper and lower bounds on its joint metastability which are related to the spectral properties of the transition matrix. To do so, let

$$\lambda_l \leq \dots \leq \lambda_m \leq \dots \leq \lambda_2 < \lambda_1 = 1$$

be the l real eigenvalues of \mathcal{P} counted according to their (algebraic) multiplicity and let v_1, \dots, v_l denote the corresponding orthonormal (right) eigenvectors. Let furthermore $Q : l^2(\mu) \rightarrow l^2(\mu)$ be the orthogonal projection onto $\text{span}\{\mathbf{1}_{A_1}, \dots, \mathbf{1}_{A_m}\}$ with respect to the μ -weighted scalar product $\langle \cdot, \cdot \rangle_\mu$; that is,

$$Q(v) = \sum_{i=1}^m \frac{\langle v, \mathbf{1}_{A_i} \rangle_\mu}{\langle \mathbf{1}_{A_i}, \mathbf{1}_{A_i} \rangle_\mu} \mathbf{1}_{A_i}. \quad (4.9)$$

Theorem 4.2 (Huisinga and Schmidt (2006)). *The metastability of an arbitrary partition of state space $D = \{A_1, \dots, A_m\}$ is bounded from above and from below by*


$$1 + \rho_2 \lambda_2 + \dots + \rho_m \lambda_m + c \leq m(D) \leq 1 + \lambda_2 + \dots + \lambda_m \quad (4.10)$$

where $\rho_j = \|Q(v_j)\|^2 = \langle Q(v_j), Q(v_j) \rangle_\mu \in [0, 1]$ and $c = \lambda_l(1 - \rho_2 + \dots + 1 - \rho_m)$.


Thus, the joint metastability of a partition of state space $\{A_1, \dots, A_m\}$ cannot be larger than the sum of the first m eigenvalues, $1 + \lambda_2 + \dots + \lambda_m$. This upper bound is thus independent of the actual decomposition considered. The lower bound is close to this upper bound if $\rho_j \approx 1$. For each j , the value ρ_j measures how constant the eigenvector v_j is on the subsets of the decomposition $\{A_1, \dots, A_m\}$; equality holds if and only if v_j restricted to each of the subsets of the decomposition is constant. Thus, in order to maximize the lower bound, one should look for a decomposition on which the eigenvectors are as constant as possible. This is the basis for algorithmic strategies of metastable state space decompositions (Deuffhard et al., 2000; Deuffhard and Weber, 2005). The overall quality of the decomposition with respect to its joint metastability can be judged by comparing the upper and the lower bound.

Example 4.1. We now apply both characterizations to our example (Example 3.6). The simulations (see Figure 3.6) suggest that there are two metastable subsets $A_1 = \{0, 1/n, \dots, r/n\}$ and $A_2 = \{(r+1)/n, \dots, 1\}$ for some $r \in \{0, \dots, n\}$. Figures 4.1–4.3 show for different sets of parameters the calculated metastability $M(C)$ of the decomposition $C = (A_1, A_2)$ for $r = 0, \dots, n$; the lower bound as provided by Theorem 4.2; as well as $\min\{W_{A_1}, W_{A_2}\}$ which determines the metastability in terms of mean hitting times (see Def. (4.1)).

The figures show that both the decomposition C with highest joint metastability $M(C)$ in terms of transition probabilities as well as the decomposition with largest minimal waiting times depends on the ratio of the payoff parameters $b/(a+b)$. It is given by $r = \lfloor bn/(a+b) \rfloor$. These metastable decompositions thus correspond (as expected) to the basins of attraction of the deterministic approximation (see Section 3.2.2, pp. 82–83). The smaller the perturbation parameter ε , the higher the value $M(C)$ for all decompositions. The lower bound seems to be a good approximation for the partitions with highest joint metastability.



(a)



(b)

Figure 4.1: Metastability in terms of (a) transition probabilities and in terms of (b) minimum waiting times in the case of our example (Example 3.6) for different decompositions $C = (A_1, A_2)$, where $A_1 = \{0, 1/n, \dots, r/n\}$ and $A_2 = \{(r+1)/n, \dots, 1\}$ for $r \in \{0, \dots, n\}$. The star (*) on the x-axis indicates the value of $\gamma = \frac{b}{a+b} = \frac{1}{a+1}$. The parameters are $a = 1, b = 1, \varepsilon = .3, n = 11$. The upper bound $1 + \lambda_2$ of the joint metastability $M(C)$ for these parameters is 1.9986.







Figure 4.2: Metastability in terms of (a) transition probabilities and in terms of (b) minimum waiting times in the case of our example (Example 3.6) for different decompositions $C = (A_1, A_2)$, where $A_1 = \{0, 1/n, \dots, r/n\}$ and $A_2 = \{(r + 1)/n, \dots, 1\}$ for $r \in \{0, \dots, n\}$. The star (*) on the x-axis indicates the value of $\gamma = \frac{b}{a+b} = \frac{1}{a+1}$. The parameters are $a = 2, b = 1, \epsilon = .3, n = 11$. The upper bound $1 + \lambda_2$ of the joint metastability $M(C)$ for these parameters is 1.9890.



(a)



(b)

Figure 4.3: Metastability in terms of (a) transition probabilities and in terms of (b) minimum waiting times in the case of our example (Example 3.6) for different decompositions $C = (A_1, A_2)$, where $A_1 = \{0, 1/n, \dots, r/n\}$ and $A_2 = \{(r+1)/n, \dots, 1\}$ for $r \in \{0, \dots, n\}$. The star (*) on the x-axis indicates the value of $\gamma = \frac{b}{a+b} = \frac{1}{a+1}$. The parameters are $a = 2, b = 1, \varepsilon = .1, n = 11$. The upper bound $1 + \lambda_2$ of the joint metastability $M(C)$ for these parameters is 1.9996.

Table 4.1 compares for different values of a and ε (assuming $b = 1, n = 11$) the weight in the stationary distribution, the transition probabilities as well as the waiting times for the sets A_1, A_2 in the decomposition $C = (A_1, A_2)$ with $r = \lfloor bn/(a+b) \rfloor$; that is, for the decomposition with the maximal joint metastability as well as maximal minimal waiting times. The table shows that while a decomposition might be metastable according to both definitions, it can happen that the statistical weight of one of the subsets is close to 0. The empirical frequency with which sample paths visit such a subset is thus close to 0, which means that in such a case we are not going to observe very often sample paths that visit this subset (see Theorem 2.2). Put differently, it is unlikely to observe metastability in the sample paths.

subset	$\varepsilon = .4$		$\varepsilon = .3$		$\varepsilon = .1$	
	A_1	A_2	A_1	A_2	A_1	A_2
$a = 1$						
$\mu(A_i)$.5	.5	.5	.5	.5	.5
$p(A_i, A_i)$.996	.996	.999	.999	1	1
W_{A_i}	331	331	1188	1188	$3.7 * 10^5$	$3.7 * 10^5$
$a = 1.2$						
$\mu(A_i)$.204	.796	.151	.849	.05	.95
$p(A_i, A_i)$.983	.996	.994	.999	1	1
W_{A_i}	155	331	428	1188	$4.2 * 10^4$	$3.7 * 10^5$
$a = 2$						
$\mu(A_i)$.003	.997	.001	.999	0	1
$p(A_i, A_i)$.962	1	.982	1	1	1
W_{A_i}	39	9404	82	$7 * 10^4$	2346	$2.7 * 10^8$

Table 4.1: Statistical weight $\mu(A_i)$, transition probability $p(A_i, A_i)$, as well as waiting times W_{A_i} for the sets $A_i, i = 1, 2$, in the decomposition $C = (A_1, A_2)$ with $r = \lfloor \frac{b}{a+b}n \rfloor$. Note that all values are approximative.

4.3 Metastability in Models of Evolution with Noise

In this section, we study the two different characterizations of metastability in models of evolution with noise. In order to differentiate the two characterizations of metastability given in the previous sections, we write in short *metastable-H* to mean to “metastable based on hitting times” and *metastable-T* to mean “metastable in terms of transition probabilities”; similarly for *metastability-H* and *metastability-T*.

In the following, let $X^\varepsilon = (X^\varepsilon(k))_{k \in \mathbb{N}}$ for $\varepsilon \geq 0$ be a model of evolution with noise with state space Z (Def. 3.8). Let P^ε denote for each $\varepsilon \geq 0$ the transition matrix of X^ε , and let μ^ε denote the corresponding stationary distribution provided $\varepsilon > 0$. For each $\varepsilon \geq 0$, let τ_A^ε denote the ε -dependent hitting time of

the set A of the Markov chain X^ε (Def. 2.3).

4.3.1 Metastability based on Hitting Times in Models of Evolution with Noise

As remarked in Section 3.2.4, it is the measures of radius and coradius that turn out to be useful to study metastability-H in models of evolution with noise:

Proposition 4.3. *Let $L_1, \dots, L_k, k \geq 2$, be disjoint limit sets of the unperturbed Markov chain X^0 . Let $L = \cup_i L_i$. Then the sets L_1, \dots, L_k are metastable-H if*

$$\min_{i=1, \dots, k} \mathfrak{r}(L_i) > \mathfrak{c}\mathfrak{r}^*(L). \quad (4.11)$$

Proof. For each $i = 1, \dots, k$ we have by Theorem 3.3:

$$\frac{1}{\mathbb{E}_x \tau_y} = \mathcal{O}(\varepsilon^{\mathfrak{r}(L_i)}) \quad \text{for } x \in L_i, y \notin L, \varepsilon \rightarrow 0, \quad (4.12)$$

which implies

$$\frac{1}{W_{L_i}} = \mathcal{O}(\varepsilon^{\mathfrak{r}(L_i)}) \quad \text{for } \varepsilon \rightarrow 0, \quad (4.13)$$

since by definition of W_{L_i} we have

$$W_{L_i} = \frac{1}{\mu^\varepsilon(L_i)} \sum_{w \in L_i} \mu^\varepsilon(w) \mathbb{E}_w \tau_{L \setminus L_i}^\varepsilon \quad (4.14a)$$

$$\geq \frac{1}{\mu^\varepsilon(L_i)} \sum_{w \in L_i} \mu^\varepsilon(w) \min_{x \in L_i} \mathbb{E}_x \tau_{Z \setminus L_i}^\varepsilon \quad (4.14b)$$

$$\geq \min_{x \in L_i, y \in Z \setminus L_i} \mathbb{E}_x \tau_y^\varepsilon. \quad (4.14c)$$

Moreover, again by Theorem 3.3,

$$R_L = \sup_{y \notin L} \mathbb{E}_y \tau_L^\varepsilon = \mathcal{O}(\varepsilon^{-\mathfrak{c}\mathfrak{r}^*(L)}). \quad (4.15)$$

It thus follows that

$$\frac{R_L}{\min_{i=1, \dots, k} W_{L_i}} = \mathcal{O}(\varepsilon^{\min_{i=1, \dots, k} \mathfrak{r}(L_i) - \mathfrak{c}\mathfrak{r}^*(L)}), \quad (4.16)$$

which implies the assertion. \square

Since $\mathfrak{c}\mathfrak{r}^*(L) = 0$ if $\mathfrak{B}(L) = Z$ while always $\min_i \mathfrak{r}(L_i) > 0$, we immediately get:

Corollary 4.4. *Let $L_1, \dots, L_k, k \geq 2$, be disjoint limit sets of the unperturbed Markov chain X^0 such that $\mathfrak{B}(L) = Z$ for $L = \cup_i L_i$. Then L_1, \dots, L_k are metastable-H.*

It does not suffice to replace the condition (4.11) in Proposition 4.3 by the sufficient condition for stochastic stability $\mathfrak{r}(L) > \mathfrak{c}\mathfrak{r}^*(L)$, as this example demonstrates:




Figure 4.4

Example 4.2. Consider a model of evolution with noise with limit sets L_1, L_2, L_3 and transition costs between these limit sets as depicted in Figure 4.4.

Then,

$$\tau(L_2) = 1 \tag{4.17}$$

$$\tau(L_3) = 1 \tag{4.18}$$

$$\tau(L) = +\infty \tag{4.19}$$

$$c\tau^*(L) = 9, \tag{4.20}$$

where $L = L_2 \cup L_3$. It follows that $\tau(L) > c\tau^*(L)$ (and thus L contains all stochastically stable states by Theorem 3.3). However, L_2, L_3 are not metastable-H since

$$\frac{R_L}{\min\{W_{L_2}, W_{L_3}\}} \rightarrow +\infty \text{ for } \varepsilon \rightarrow 0. \tag{4.21}$$

Thus, L contains all stochastically stable states but L_2, L_3 is not metastable-H.

Furthermore, the following example shows that the converse direction of Proposition 4.3 need not hold.

Example 4.3. Consider a model of evolution with noise with limit sets L_1, L_2, L_3, L_4 and transition costs between these limit sets as depicted in Figure 4.5.




Figure 4.5

Then,

$$\tau(L_2) = 10 \tag{4.22}$$

$$\tau(L_3) = 4 \tag{4.23}$$

$$c\tau^*(L) = 9, \tag{4.24}$$

where $L = L_2 \cup L_3$. Thus, $c\tau^*(L) = 9 > 4 = \min\{\tau(L_2), \tau(L_3)\}$, but $(X^\varepsilon(t))$ is metastable-H since for the sets L_2, L_3 and $L = L_2 \cup L_3$ we have

$$\frac{R_L}{\min\{W_{L_2}, W_{L_3}\}} = \mathcal{O}(\varepsilon^{10-9}) = \mathcal{O}(\varepsilon) \text{ for } \varepsilon \rightarrow 0. \tag{4.25}$$

This follows from the observation that the waiting times W_{L_2}, W_{L_3} consider only transitions from states in L_2 to states in L_3 and from states in L_3 to states in L_2 , respectively, while the radii $\mathfrak{r}(L_2)$ and $\mathfrak{r}(L_3)$ consider all transitions out of L_2 and L_3 , respectively (and thus also transitions from L_2 or L_3 to L_4), compare Definition 4.1.

4.3.2 Metastability based on Transition Probabilities in Models of Evolution with Noise

There is a direct relationship of limit sets of the unperturbed Markov chain X^0 as well as of their basins of attraction (see Def. 3.10) with metastable-T, i.e., almost invariant, subsets of state space:

Proposition 4.5. *Let L be a recurrent set of the unperturbed Markov chain X^0 and let $\mathfrak{B}(L)$ denote its basin of attraction. Then L as well as $\mathfrak{B}(L)$ are almost invariant under P^ε .*

Proof. First, notice that since L is a recurrent set and thus a closed communication class of X^0 , we have for each $x \in L$:

$$\sum_{y \in L} P^0(x, y) = 1. \quad (4.26)$$

Moreover, the same holds for $x, y \in \mathfrak{B}(L)$: If not, there is a $z \in Z \setminus \mathfrak{B}(L)$ such that $P^0(x, z) > 0$. But this implies that $x \notin \mathfrak{B}(L)$ since in this case, $\mathbb{P}_x(X^0(k) \in L) < 1$ for all $k \in \mathbb{N}$. Thus, we have proved by contraposition that for each $x \in \mathfrak{B}(L)$ we have

$$\sum_{y \in \mathfrak{B}(L)} P^0(x, y) = 1. \quad (4.27)$$

Now, since $(X^\varepsilon)_{\varepsilon \geq 0}$ is a model of evolution with noise, the transition probabilities $P^\varepsilon(x, y)$ are continuous in ε with limit $P^0(x, y)$. This implies

$$p^\varepsilon(L, L) = \frac{1}{\mu^\varepsilon(L)} \sum_{x \in L} \mu^\varepsilon(x) \sum_{y \in L} P^\varepsilon(x, y) \quad (4.28a)$$

$$\rightarrow 1 \quad \text{for } \varepsilon \rightarrow 0. \quad (4.28b)$$

Similarly for $\mathfrak{B}(L)$. □

It follows directly that a decomposition of state space into disjoint limit sets is metastable-T:

Corollary 4.6. Let (L_1, \dots, L_k) be a decomposition of state space Z into disjoint limit sets of the unperturbed Markov chain X^0 . Then, (L_1, \dots, L_k) represents a metastable-T decomposition of Z .

The following example demonstrates that almost invariant subsets of state space might not coincide with limit sets:

Example 4.4. Consider a model of evolution with noise $(X^\varepsilon)_{\varepsilon \geq 0}$ on the state space $Z = \{0, 1, 2\}$ with transition matrix P^0 of the unperturbed Markov chain X^0 given by

$$P^0 = \begin{pmatrix} 1 - \delta & \delta & 0 \\ \frac{1}{2} & 0 & \frac{1}{2} \\ 0 & \delta & 1 - \delta \end{pmatrix} \quad (4.29)$$

where $1 > \delta > 0$ is another parameter. Then, P^0 (for fixed δ) has a single limit set which consists of the whole state space Z . For $\delta \ll 1$, the sets $\{0\}$ and $\{1\}$ are, however, almost invariant, as well.

Summarizing the results of this and the previous subsection, we get

Proposition 4.7. *Let $(X^\varepsilon)_{\varepsilon \geq 0}$ be a model of evolution with noise with state space Z . Let L_1, \dots, L_k , $k \geq 2$, be disjoint limit sets of the unperturbed Markov chain X^0 with $\mathfrak{B}(L) = Z$, where $L = \cup_i L_i$. Then (L_1, \dots, L_k) is both metastable-H and metastable-T.*

Proof. By Corollary 4.4 and Corollary 4.6. □

In contrast to the case of metastability-H, where a union L of limit sets which fulfills the condition $\tau(L) > \mathfrak{c}^*(L)$ for containing all stochastically stable states (see Theorem 3.3) need not be metastable-H, see Example 4.2, all limit sets and all unions of limit sets are metastable-T in the sense of being almost invariant. Thus, also a union of limit sets L for which $\tau(L) > \mathfrak{c}^*(L)$ is metastable-T.

Lastly, we point out that the results of this and the previous subsection present an approach to the identification of metastable subsets of state space for models of evolution with noise – provided that the identification of the limit sets of the unperturbed Markov chain as well as their radii and coradii is not too involved.

Example. In our running example (Example 3.6), the state space $Z = \{0, \frac{1}{n}, \dots, 1\}$ can be decomposed into the basins of attraction of the two limit sets $L_0 = \{0\}$ and $L_1 = \{1\}$ of the unperturbed Markov chain. Thus, by Proposition 4.7, these limit sets are metastable-H as well as metastable-T. This confirms the numerical results given in Example 4.1.

Chapter 5

Markov State Modeling

Markov state modelling is a modelling approach for the approximation of complex dynamical processes that are characterized by metastable behavior. It has been developed primarily for the application context of molecular dynamics (see Deuffhard et al., 2000; Djurdjevac et al., 2010; Noé et al., 2007a, 2009; Schütte, 1998; Schütte et al., 1999; Schütte and Huisinga, 2003; Schütte et al., 2011, and references therein). The motivation for such an approximation approach is two-fold. First, models of large biomolecules like most models of complex systems are usually too complicated to be analyzed analytically. Thus, computer simulations are an important tool to explore their dynamic behavior. Moreover, macromolecules are characterized by metastable dynamic behavior, which is essential for understanding their biochemical functioning¹. In this case of metastable dynamic behavior on large state spaces, however, direct numerical simulation becomes very costly or even infeasible from a computational point of view since the transitions between metastable subsets of state space are rare events. Second, although modern computing technologies make it more and more possible to simulate large scale systems and to provide insights into complex dynamical systems, the amount of data generated is too large to be understood without rigorous methods for the analysis of their essential structure and dynamics (e.g., Sarich et al., 2014; Schütte et al., 2011). Against this background, Markov state models have been developed as models of reduced complexity whose construction takes only short trajectory data obtained by simulation into account. As agent-based models are also computational models of large scale and high complexity (see Chapter 1), it seems worthwhile to use the Markov state modelling approach as a tool for their analysis in case metastability is present.

The basic idea of Markov state models is to approximate the original Markov process by a Markov chain on a small finite state space. More specifically, a Markov state model is defined as a Markov chain whose state space consists of the dominant metastable sets of the original Markov process and whose transition rates between these macrostates are given by the aggregate statistics of

¹ More specifically, macromolecules conserve their large scale geometric structure on long time-scales – the so-called *conformations* – while on smaller time-scales the system may well change. Switches between conformations are rare. Thus, conformations can be understood as metastable subsets of state space.

jumps between the metastable sets. The advantage of this approach in the context of complex models with large state spaces is that the transition probabilities can be estimated on the basis of short-term trajectory data. Moreover, it has been shown that Markov state models have good approximation properties if metastability is inherent in the system of interest (Sarich, 2011; Sarich et al., 2010; Schütte and Sarich, 2013).

In Section 5.1, we first introduce further notions necessary for the understanding of the rest of the chapter such as the transfer operator \mathcal{T} , orthogonal projections, as well as committor functions. In Section 5.2, we introduce full partition Markov state models as the Markov chain that approximates the jump process between the subsets of a given partition. This construction allows us to estimate the transition matrix of the full partition Markov state model from trajectory data. We show how this construction relates to orthogonal projections; more precisely, we show that the transition matrix of the full partition Markov state models corresponds to the matrix representation of the projected transfer operators onto the subspace of \mathbb{R}^l , $l = |Z|$, spanned by the indicator functions on the subsets of the partition. Subsequently, in Section 5.3, we use this last observation as motivation to define general Markov state models as the Markov chains with the transition matrix which corresponds to the matrix representation of the projected transfer operators. The core set Markov state models are a special case of this construction in which the subspace of \mathbb{R}^l considered for the projection is spanned by the committor functions on disjoint, nonempty subsets of \mathbb{R}^l that do not necessarily constitute a full partition. Moreover, there is an interpretation of the transition matrix of core set Markov state models that allows as well its estimation from trajectory data. Finally, in Section 5.4 we discuss the approximation quality of Markov state models with respect to stochastic stability and the propagation of probability distributions as well as with respect to the approximation of eigenvalues. In Chapter 6, we relate these approximation results to the identification of metastable subsets of state space.

Throughout the chapter, let Z denote the state space with $|Z| = l$. Let $(X_k)_{k \in \mathbb{N}}$ be an irreducible, reversible Markov chain on Z with transition matrix P and let μ denote its unique stationary distribution.

5.1 Preliminaries

5.1.1 Transfer Operators

In Chapter 2.2, we introduced the transfer operator \mathcal{P} that transfers probability distributions in time. We can similarly consider a transfer operator \mathcal{T} that transfers probability *densities* with respect to the stationary distribution μ ; that is, functions $v : Z \rightarrow \mathbb{R}_{\geq 0}$ with $\sum_{i \in Z} v(i)\mu(i) = 1$. It is defined by

$$\mathcal{T} : \mathbb{R}^l \rightarrow \mathbb{R}^l \tag{5.1}$$

with

$$\mathcal{T}(v)(j) = \frac{1}{\mu(j)} \sum_{i=1}^l v(i)\mu(i)P(i, j). \tag{5.2}$$

The operator \mathcal{T} maps probability densities with respect to μ to probability densities with respect to μ . Note that we did not differentiate between the probability density with respect to μ and its representation as a vector since we considered again the canonical basis $\{e_1, \dots, e_l\}$ of R^l .

In matrix notation, we may write

$$\mathcal{T}(v) = Qv, \quad (5.3)$$

where $Q = (MPM^{-1})^T$ is the transition matrix of the *time-reversed* Markov chain, $M = \text{diag}(\mu(1), \mu(2), \dots, \mu(l))$. Written differently, $Q(i, j) = \frac{\mu(j)}{\mu(i)}P(j, i)$. In general, the two operators are related via

$$\mu\mathcal{T}(v) = \mathcal{P}(\mu v), \quad (5.4)$$

where for a density v with respect to μ we denote by μv the probability distribution w defined by $w(j) = v(j)\mu(j)$. Notice that (X_k) is reversible if and only if $Q = P$ and thus $\mathcal{T}(v) = Pv$. Moreover, we have noted in Chapter 2.2 that the reversibility of (X_k) is equivalent to \mathcal{P} being self-adjoint in $l^2(\frac{1}{\mu})$. In the case of the transfer operator \mathcal{T} this translates to (X_k) being reversible with respect to μ if and only if \mathcal{T} is self-adjoint in $l^2(\mu)$; that is,

$$\langle \mathcal{T}(v_1), v_2 \rangle_\mu = \langle v_1, \mathcal{T}(v_2) \rangle_\mu, \quad (5.5)$$

where $l^2(\mu)$ is the vector space \mathbb{R}^l equipped with the scalar product $\langle \cdot, \cdot \rangle_\mu$,

$$\langle v_1, v_2 \rangle_\mu = \sum_{i=1}^l v_1(i)v_2(i)\mu(i). \quad (5.6)$$

Note that

$$\langle v_1, v_2 \rangle_\mu = \langle \mu v_1, \mu v_2 \rangle_{\frac{1}{\mu}}. \quad (5.7)$$

In what follows, we see \mathcal{T} as an operator on $l^2(\mu)$ and \mathcal{P} as an operator on $l^2(\frac{1}{\mu})$. Doing so allows to take advantage of the Hilbert space framework. In the finite state space case we consider in this thesis, this does not represent a restriction, but in the continuous state space it does. We note that in this setting, Equation (5.4) can be depicted as a commutative diagram, see Figure 5.1. The arrows $l^2(\frac{1}{\mu}) \xrightarrow{1/\mu} l^2(\mu)$ in the diagram denote the reweighting of probability distributions $w \in l^2(\frac{1}{\mu})$ with $1/\mu$ which yields a probability density $v = w/\mu \in l^2(\mu)$ with respect to μ defined entrywise by $v(j) = w(j)/\mu(j)$. Similarly, the arrows $l^2(\frac{1}{\mu}) \xleftarrow{\mu} l^2(\mu)$ denote the entrywise reweighting with μ of probability densities $v \in l^2(\mu)$ with respect to μ to yield probability distributions $w = \mu v \in l^2(\frac{1}{\mu})$.

We mention these relationships between the spaces $l^2(\frac{1}{\mu})$ and $l^2(\mu)$, and between the transfer operators \mathcal{P} and \mathcal{T} , because in the discrete state space case we consider here it is common to think and work in terms of probability distributions, but less so in probability densities with respect to the stationary distribution. For this reason, we will state the main results and their interpretation in the following in terms of probability distributions, and thus, in $l^2(\frac{1}{\mu})$ and using the operator \mathcal{P} . Using Eq. (5.4), we can relate the results to the case of $l^2(\mu)$ and the transfer operator \mathcal{T} , and thereby to the existing literature.

$$\begin{array}{ccc}
l^2(\frac{1}{\mu}) & \begin{array}{c} \xrightarrow{\frac{1}{\mu}} \\ \xleftarrow{\mu} \end{array} & l^2(\mu) \\
\mathcal{P} \downarrow & & \downarrow \mathcal{T} \\
l^2(\frac{1}{\mu}) & \begin{array}{c} \xrightarrow{\frac{1}{\mu}} \\ \xleftarrow{\mu} \end{array} & l^2(\mu)
\end{array}$$

Figure 5.1: Commutative diagram for Equation (5.4).

5.1.2 Orthogonal Projections

The principal idea behind Markov state modelling is the best approximation and thus orthogonal projection of the dynamics on a lower dimensional subspace.

In the case of a given orthogonal basis $\{u_1, \dots, u_m\}$ of the considered subspace $D \subset \mathbb{R}^l$, we can compute the result of the orthogonal projection of a vector $v \in \mathbb{R}^l$ onto D via

$$\mathcal{Q}_D(v) = \sum_{j=1}^m \frac{\langle v, u_j \rangle}{\langle u_j, u_j \rangle} u_j. \quad (5.8)$$

In general, given a (not necessarily orthogonal) basis $\{q_1, \dots, q_m\}$, we can write

$$\mathcal{Q}_D(v) = \sum_{i,j=1}^m S^{-1}(i,j) \langle v, q_i \rangle q_j, \quad (5.9)$$

where S is the matrix with entries $S(i,j) = \langle q_i, q_j \rangle$. The matrix S is invertible since the q_i 's are assumed linearly independent.

Note that in the above equations we have not specified the scalar product considered. In the following, we will denote by \mathcal{Q}_D^μ the orthogonal projection onto D with respect to $\langle \cdot, \cdot \rangle_\mu$ and by $\mathcal{Q}_D^{1/\mu}$ the orthogonal projection onto D with respect to $\langle \cdot, \cdot \rangle_{\frac{1}{\mu}}$.

5.1.3 Committed Functions

Let C_1, \dots, C_m be nonempty, disjoint sets that do not constitute a full partition of state space Z . Let $C = \cup_{i=1}^m C_i$. The *forward committor* $q_i^+ : Z \rightarrow [0, 1]$ is the function that gives for a state $z \in Z$ the probability that the Markov chain (X_k) will visit the set C_i first rather than $C \setminus C_i$. More formally,

$$q_i^+(z) = \mathbb{P}[\tau_{C_i}^0 < \tau_{C \setminus C_i}^0 \mid X_0 = z], \quad (5.10)$$

where $\tau_A^k = \inf\{k' \geq k \mid X_{k'} \in A\}$ for $k \geq 0$. Note that in this notation the hitting time τ_A introduced in Def. 2.3 is just τ_A^1 . The forward committor q_i^+ is the solution to the following linear system with boundary conditions (see

Metzner et al., 2009b):

$$(P - Id)q_i^+(z) = 0 \text{ if } z \in Z \setminus C, \quad (5.11a)$$

$$q_i^+(z) = 1 \text{ if } z \in C_i, \quad (5.11b)$$

$$q_i^+(z) = 0 \text{ if } z \in C \setminus C_i. \quad (5.11c)$$

Similarly, the *backward committor* $q_i^- : Z \rightarrow [0, 1]$ is defined as the probability that (X_k) came last from C_i rather than from $C \setminus C_i$, conditional on being in state $z \in Z$. The function q_i^- solves

$$(Q - Id)q_i^-(z) = 0 \text{ if } z \in Z \setminus C, \quad (5.12a)$$

$$q_i^-(z) = 1 \text{ if } z \in C_i, \quad (5.12b)$$

$$q_i^-(z) = 0 \text{ if } z \in C \setminus C_i, \quad (5.12c)$$

where Q refers to the transition matrix of the time-reversed Markov chain, see Section 5.1.1. Now, if (X_k) has a unique stationary distribution which is non-vanishing on all of the sets C_1, \dots, C_m then Eqs. (5.11) and (5.12) have a unique solution. Moreover, for reversible Markov chains (X_k) we know that $P = Q$ and thus $q_i^+ = q_i^-$. In this case, we write in short q_i and refer to the committor function.

Example 5.1. Let $(X_t)_{t \in \mathbb{T}}$ be the Markov chain on state space $Z = \{0, \frac{1}{n}, \dots, 1\}$ with transition matrix P of our example (see Example 3.6). Let $C_1 = \{0\} \subset Z$ and $C_2 = \{1\} \subset Z$. Figure 5.2 displays the committor function q_1 as well as the eigenvector u'_2 associated with the eigenvalue of second largest absolute value normalized such that $u'_2(0) = 1$. The figure hints at the observation that, for $\varepsilon \rightarrow 0$, committor functions get close to eigenvectors (Bovier et al., 2002).

5.2 Full Partition Markov State Models


In this section, we are going to focus on full partition Markov state models. These models consider a full partition of state space and the aggregate jump statistics between the sets in the partition. As we show in the next section, the full partition Markov state models are a special case of core set Markov state models.

5.2.1 Construction of Full Partition Markov State Models


Let $(X_k)_{k \in \mathbb{N}}$ be a discrete-time Markov chain on a finite state space $Z = \{1, \dots, l\}$. Our goal is to construct a Markov chain $(\hat{X}_k)_{k \in \mathbb{N}}$ on the state space $\hat{Z} = \{1, \dots, m\}$ with m considerably smaller than l such that (\hat{X}_k) captures the essential dynamics of the original Markov chain (X_k) .

For the construction of a Markov state model based on a *full partition* of state space, we consider subsets A_1, \dots, A_m of Z that partition Z ; that is,

$$\bigcup_{j=1}^m A_j = Z \quad \text{and} \quad A_i \cap A_j = \emptyset \text{ for } i \neq j. \quad (5.13)$$



(a)



(b)

Figure 5.2: Comparison of the committor function q_1 with the eigenvector u'_2 associated with the eigenvalue of second largest absolute value normalized such that $u'_2(0) = 1$ for the parameter values $a = 2, b = 1, n = 11$ and (a) $\varepsilon = .3$, (b) $\varepsilon = .1$.

The reduced chain (\hat{X}_k) on $\hat{Z} = \{1, \dots, m\}$ with transition matrix $\hat{P} = (\hat{p}_{ij})$ is then defined by setting

$$\hat{p}_{ij} = \mathbb{P}[\tilde{X}_1 = j \mid \tilde{X}_0 = i], \quad (5.14)$$

where $(\tilde{X}_k)_{k \in \mathbb{N}}$ is the discrete-time process on \hat{Z} that describes the dynamics of (X_k) between the sets A_1, \dots, A_m , i.e.,

$$\tilde{X}_k = i \Leftrightarrow X_k \in A_i. \quad (5.15)$$

Note that we differentiate between (\hat{X}_k) and (\tilde{X}_k) because (\tilde{X}_k) is in general not Markovian.

Example 5.2. Consider the Markov chain $(X^\varepsilon(t))_{t \in \mathbb{T}}$ of our running example (Example 3.6). Let (A, B) be the partition of its state space $Z = \{0, \frac{1}{n}, \dots, 1\}$ into the basins of attraction of the limit sets of the unperturbed Markov chain $(X^0(t))$; that is, $A = \{0, \dots, r/n\}$ and $B = \{r+1/n, \dots, 1\}$, where $r = \lfloor bn/(a+b) \rfloor$. Then,

$$\mathbb{P}_\mu(X^\varepsilon(t+1/n) \in B \mid X^\varepsilon(t) \in A) = \frac{1}{\mu(A)} \sum_{i \in A} \mu(i) \sum_{j \in B} p(i, j) \quad (5.16a)$$

$$= \frac{\mu(r)}{\mu(A)} p(r, r+1), \quad (5.16b)$$

where the last equality follows from the fact that the only transition possible from A to B is from r to $r+1$. Furthermore,

$$\mathbb{P}_\mu(X^\varepsilon(t+1/n) \in B \mid X^\varepsilon(t) \in A, X^\varepsilon(t-1/n) \in B) \quad (5.17a)$$

$$= \mathbb{P}_\mu(X^\varepsilon(t+1/n) \in B \mid X^\varepsilon(t) = r) \quad (5.17b)$$

$$= p(r, r+1) \quad (5.17c)$$

$$\neq \mathbb{P}_\mu(X^\varepsilon(t+1/n) \in B \mid X^\varepsilon(t) \in A), \quad (5.17d)$$

since $\mu(A) > \mu(r)$ if $r \neq 0$ because in this case the set A contains more than one population state each of which has positive weight in the stationary distribution. Thus, the process $(\tilde{X}^\varepsilon(k))$ defined according to Eq. (5.15) with respect to the partition (A, B) is not memoryless and thus does not have the Markov property.

The reason why (\tilde{X}_k) is not Markovian is called the *recrossing problem*. This name refers to the issue that transitions between the subsets of state space are much more likely at the boundaries of the sets. However, we still want to approximate (X_k) by a Markov chain which is why we consider (\hat{X}_k) . In Section 5.4, we give an upper bound on the error of approximating (X_k) by (\hat{X}_k) and discuss under which conditions the approximation is nevertheless acceptable. Apparent advantages of a reduced state space Markov chain are that it is easier to compute eigenvalues and eigenvectors as well as other properties such as waiting times.

Example (Continuation). In our example (Example 3.6) with parameters $a = b = 1, n = 11, \varepsilon = .3$, we have $r = 5/11$ and $A = \{0, \dots, 5/11\}$, $B =$

$\{6/11, \dots, 1\}$. The resulting matrix \hat{P} is given by (rounded to four decimal places)

$$\hat{P} = \begin{pmatrix} .9989 & .0011 \\ .0011 & .9989 \end{pmatrix}. \quad (5.18)$$

From this we can easily calculate $\hat{W}_1 = \mathbb{E}_1 \hat{\tau}_2 = 1/p(A, B)$ ($= \hat{W}_2 = \mathbb{E}_2 \hat{\tau}_1 = 1/p(B, A)$, since \hat{P} is symmetric), where \hat{W}_1, \hat{W}_2 and $\hat{\tau}$ refer to the waiting times in state 1, 2, and to the hitting time of the chain $(\hat{X}_t)_{t \in \mathbb{T}}$, respectively. See Table 5.1 for a comparison of W_A and \hat{W}_1 as well as of the two dominant eigenvalues of P with the two eigenvalues of \hat{P} . It shows that the approximation underestimates the waiting times.

(X_t)	(\hat{X}_t)
$W_A = 1188$	$\hat{W}_1 = 909$
$\lambda_1 = 1$	$\hat{\lambda}_1 = 1$
$\lambda_2 = .9986$	$\hat{\lambda}_2 = .9978$

Table 5.1: Waiting times and dominant eigenvalues of (X_t) and (\hat{X}_t) , respectively.

The following proposition shows that (\hat{X}_k) is reversible if (X_k) is:

Proposition 5.1 (Reversibility of (\hat{X}_k)). *Let $(X_k)_{k \in \mathbb{N}}$ be a reversible discrete-time Markov chain on a finite state space Z with transition matrix P and unique stationary distribution μ . Let A_1, A_2, \dots, A_m be a full partition of Z and let $(\hat{X}_k)_{k \in \mathbb{N}}$ be the Markov chain on $\hat{Z} = \{1, \dots, m\}$ whose transition matrix $\hat{P} = (p_{ij})_{i,j=1,\dots,m}$ is defined by Eq. (5.14). Then, $\{\hat{X}_k\}$ has the stationary distribution $\hat{\mu}$ defined on \hat{Z} by*

$$\hat{\mu}(i) = \mu(A_i) \quad \text{for all } i = 1, \dots, m. \quad (5.19)$$

Moreover, \hat{P} and $\hat{\mu}$ are in detailed balance:

$$\hat{\mu}(i)\hat{P}(i, j) = \hat{\mu}(j)\hat{P}(j, i) \quad \text{for all } i, j = 1, \dots, m. \quad (5.20)$$

Thus, (\hat{X}_k) is reversible as well.

Proof. It follows straightforwardly from the reversibility of (X_k) with respect to μ that

$$\hat{\mu}(i)\hat{P}(i, j) = \mu(A_i) \frac{1}{\mu(A_i)} \sum_{k \in A_i} \mu(k) \sum_{l \in A_j} P(k, l) \quad (5.21a)$$

$$= \sum_{k \in A_i} \sum_{l \in A_j} \mu(l) P(l, k) \quad (5.21b)$$

$$= \mu(A_j) \frac{1}{\mu(A_j)} \sum_{l \in A_j} \mu(l) \sum_{k \in A_i} P(l, k) \quad (5.21c)$$

$$= \hat{\mu}(j)\hat{P}(j, i). \quad (5.21d)$$

□

Last but not least, another important feature of the transition matrix \hat{P} is that it can be estimated from trajectory data of the original process (X_k) only. Given a trajectory of M datapoints $x_k, k = 1, \dots, M$ where x_k refers to the realization of the random variable X_k . Then, \hat{P} can be estimated by

$$\hat{p}_{ij}^* = \frac{N_{ij}}{n_i}, \quad (5.22)$$

where n_i is the number of times the process spent in A_i , i.e., $n_i = |\{x_k : x_k \in A_i\}|$, and N_{ij} is the number of transitions made from the set A_i to the set A_j , i.e., $N_{ij} = |\{(x_k, x_{k+1}) : x_k \in A_i \text{ and } x_{k+1} \in A_j\}|$. \hat{P}^* is a maximum likelihood estimator for \hat{P} . The resulting statistical error $\|\hat{P} - \hat{P}^*\|$ can be further analyzed (Metzner et al., 2009a; Noé, 2008; Röblitz, 2008; Singhal and Pande, 2005).

5.2.2 The Full Partition Markov State Model and Orthogonal Projections

In this section, we are going to give another interpretation of the reduced chain (\hat{X}_k) and its transition matrix \hat{P} in terms of projected transfer operators. As we said before, one way to arrive at an approximation of the chain (X_k) is via orthogonal projection of its dynamics onto a lower dimensional subspace $D \subset \mathbb{R}^l$. As a matter of fact, the transition matrix \hat{P} is the matrix representation of the transfer operators \mathcal{T} and \mathcal{P} projected onto the subspace $D = \text{span}\{\mathbb{1}_{A_1}, \dots, \mathbb{1}_{A_m}\}$ and μD , respectively, where $\mu D = \text{span}\{\mu \mathbb{1}_{A_1}, \dots, \mu \mathbb{1}_{A_m}\}$ (just as P is the matrix representation of the transfer operators \mathcal{P} and \mathcal{T}). See Figure 5.3 for a commutative diagram that serves as an orientation for the relationships shown in this section.

$$\begin{array}{ccccc}
 l^2(\frac{1}{\mu}) \supset \mu D & \xleftarrow{Q_{\mu D}^{1/\mu}} & l^2(\frac{1}{\mu}) & \xrightleftharpoons[\mu]{\frac{1}{\mu}} & l^2(\mu) & \xrightarrow{Q_D^\mu} & D \subset l^2(\mu) \\
 \downarrow Q_{\mu D}^{1/\mu} \mathcal{P} & & & & & & \downarrow Q_D^\mu \mathcal{T} \\
 l^2(\frac{1}{\mu}) \supset \mu D & \xleftarrow[\mu]{\frac{1}{\mu}} & & \xrightarrow[\mu]{\frac{1}{\mu}} & & & D \subset l^2(\mu)
 \end{array}$$

Figure 5.3: Relationships between the projected transfer operators $Q_{\mu D}^{1/D} \mathcal{P} Q_{\mu D}^{1/\mu}$ and $Q_D^\mu \mathcal{T} Q_D^\mu$.

Now, let $D = \text{span}\{\mathbb{1}_{A_1}, \dots, \mathbb{1}_{A_m}\}$. Note that $\mathbb{1} \in D$ and thus D as a subspace of $l^2(\mu)$ can be considered the space of probability densities with respect to μ that are still representable in our reduced model. In other words, only probability densities with respect to μ that are constant on the subsets A_1, \dots, A_m of the considered partition are elements of D .

In this way, μv for a $v \in D$ is a probability distribution and we can consider the subspace $\mu D = \text{span}\{\varphi_1, \varphi_2, \dots, \varphi_m\}$ of $l^2(\frac{1}{\mu})$ where $\varphi_j = \frac{\mu \mathbf{1}_{A_j}}{\hat{\mu}(j)}$ for the subspace of probability distributions which are still representable in the reduced model. In other words, μD can be interpreted as the subspace of probability distributions that are considered for the approximation. A distribution w is an element of μD if its probability density w/μ with respect to μ defined entrywise by $w(j)/\mu(j)$ is an element of D and thus if w/μ is constant on the subsets of the partition. Put differently, $w \in \mu D$ if it is distributed within the subsets of the partition according to the stationary distribution μ . We consider the φ_j as a basis of this subspace (instead of, e.g., $\mu \mathbf{1}_{A_j}$) since if $\mu D \ni w = \sum_{j=1}^m \hat{w}(j) \varphi_j$ is a vector representation of w with respect to the basis $\{\varphi_1, \dots, \varphi_m\}$ then

$$1 = \sum_{i=1}^l w(i) \quad (5.23a)$$

$$= \sum_{i=1}^l \left(\sum_{j=1}^m \hat{w}(j) \varphi_j \right)(i) \quad (5.23b)$$

$$= \sum_{j=1}^m \frac{\hat{w}(j)}{\hat{\mu}(j)} \sum_{i \in A_j} \mu(i) \quad (5.23c)$$

$$= \sum_{j=1}^m \hat{w}(j). \quad (5.23d)$$

It follows that \hat{w} is a probability distribution on $\hat{Z} = \{1, \dots, m\}$ if and only if w is a probability distribution on Z (the non-negativity of \hat{w} follows from the non-negativity of w).

Now, for $w' \in l^2(\frac{1}{\mu})$, we have that its projection $w = \mathcal{Q}_{\mu D}^{1/\mu}(w')$ can be represented in terms of the basis $\{\varphi_j\}$ as $w = \sum_{j=1}^m \hat{w}(j) \varphi_j$. Moreover, since by straightforward calculation we get

$$\langle \varphi_j, \varphi_j \rangle_{\frac{1}{\mu}} = \frac{1}{\hat{\mu}(j)} \quad (5.24)$$

as well as

$$\langle \mathcal{P}(\varphi_k), \varphi_j \rangle_{\frac{1}{\mu}} = \frac{1}{\hat{\mu}(k)\hat{\mu}(j)} \sum_{z \in Z} (\mu \mathbf{1}_{A_k})^T P(z) \mu \mathbf{1}_{A_j}(z) \frac{1}{\mu(z)} \quad (5.25a)$$

$$= \frac{1}{\hat{\mu}(k)\hat{\mu}(j)} \sum_{z \in A_j} (\mu \mathbf{1}_{A_k})^T P(z) \quad (5.25b)$$

$$= \frac{1}{\hat{\mu}(k)\hat{\mu}(j)} \sum_{z \in A_j} \sum_{x \in Z} (\mu \mathbf{1}_{A_k})(x) P(x, z) \quad (5.25c)$$

$$= \frac{1}{\hat{\mu}(k)\hat{\mu}(j)} \sum_{x \in A_k} \mu(x) \sum_{z \in A_j} \mathbb{P}(X_1 = z \mid X_0 = x) \quad (5.25d)$$

$$= \frac{1}{\hat{\mu}(k)\hat{\mu}(j)} \sum_{x \in A_k} \mu(x) \mathbb{P}(X_1 \in A_j \mid X_0 = x) \quad (5.25e)$$

$$= \frac{1}{\hat{\mu}(j)} \hat{P}(k, j), \quad (5.25f)$$

it follows by definition of the orthogonal projection $\mathcal{Q}_{\mu D}^{1/\mu}$ that

$$\mathcal{Q}_{\mu D}^{1/\mu} \mathcal{P}(\varphi_k) = \sum_{j=1}^m \frac{\langle \mathcal{P}(\varphi_k), \varphi_j \rangle_{\frac{1}{\mu}}}{\langle \varphi_j, \varphi_j \rangle_{\frac{1}{\mu}}} \varphi_j \quad (5.26a)$$

$$= \sum_{j=1}^m \hat{P}(k, j) \varphi_j. \quad (5.26b)$$

Thus,

$$\mathcal{Q}_{\mu D}^{1/\mu} \mathcal{P}(w) = \sum_{j=1}^m (\hat{w}^T \hat{P})(j) \varphi_j, \quad (5.27)$$

which means that $\hat{w} \hat{P}$ is a vector representation of $\mathcal{Q}_{\mu D}^{1/\mu} \mathcal{P}(w)$ with respect to the basis $\{\varphi_j\}$ and thus \hat{P} is a matrix representation of the projected transfer operator $\mathcal{Q}_{\mu D}^{1/\mu} \mathcal{P}|_{\mu D}$.

Now, given a probability density $v' \in l^2(\mu)$ with respect to μ . Since $\{\mathbf{1}_{A_1}, \dots, \mathbf{1}_{A_m}\}$ is a basis of D , its projection $v = \mathcal{Q}_D^\mu(v')$ can be represented by

$$v = \sum_{j=1}^m \hat{v}(j) \mathbf{1}_{A_j}. \quad (5.28)$$

Notice at this point that \hat{v} is a probability density on $\hat{Z} = \{1, \dots, m\}$ with respect to $\hat{\mu}$ if and only if v is a probability density with respect to μ :

$$1 = \sum_{i=1}^l v(i) \mu(i) \quad (5.29a)$$

$$= \sum_{i=1}^l \mu(i) \left(\sum_{j=1}^m \hat{v}(j) \mathbf{1}_{A_j} \right)(i) \quad (5.29b)$$

$$= \sum_{j=1}^m \hat{v}(j) \sum_{i \in A_j} \mu(i) \quad (5.29c)$$

$$= \sum_{j=1}^m \hat{v}(j) \hat{\mu}(j). \quad (5.29d)$$

This is why we consider the basis $\{\mathbf{1}_{A_i}\}$ for D (the non-negativity of \hat{v} follows from the non-negativity of v).

Moreover, we have by the definition of the orthogonal projection \mathcal{Q}_D^μ that

$$\mathcal{Q}_D^\mu \mathcal{T} \mathcal{Q}_D^\mu (\mathbf{1}_{A_i}) = \sum_{j=1}^m \frac{\langle \mathcal{T}(\mathbf{1}_{A_i}), \mathbf{1}_{A_j} \rangle_\mu}{\langle \mathbf{1}_{A_j}, \mathbf{1}_{A_j} \rangle_\mu} \mathbf{1}_{A_j} \quad (5.30a)$$

$$= \sum_{j=1}^m \frac{\langle \mathcal{P}(\mu \mathbf{1}_{A_i}), \mu \mathbf{1}_{A_j} \rangle_{\frac{1}{\mu}}}{\mu(A_j)} \mathbf{1}_{A_j} \quad (5.30b)$$

$$= \sum_{j=1}^m \hat{\mu}(i) \langle \mathcal{P}(\varphi_i), \varphi_j \rangle_{\frac{1}{\mu}} \mathbf{1}_{A_j} \quad (5.30c)$$

$$= \sum_{j=1}^m \frac{\hat{\mu}(i)}{\hat{\mu}(j)} \hat{P}(i, j) \mathbf{1}_{A_j}, \quad (5.30d)$$

which implies

$$\mathcal{Q}_D^\mu \mathcal{T} \mathcal{Q}_D^\mu (\mathbf{1}_{A_i}) = \sum_{j=1}^m \hat{P}(j, i) \mathbf{1}_{A_j}, \quad (5.31)$$

since $(\hat{P}, \hat{\mu})$ are in detailed balance, see Proposition 5.1.

From this follows for $v = \sum_{j=1}^m \hat{v}(j) \mathbf{1}_{A_j}$ that

$$\mathcal{Q}_D^\mu \mathcal{T} \mathcal{Q}_D^\mu (v) = \sum_{j=1}^m (\hat{P} \hat{v})(j) \mathbf{1}_{A_j}, \quad (5.32)$$

which means that $\hat{P} \hat{v}$ is a vector representation of $\mathcal{Q}_D^\mu \mathcal{T} \mathcal{Q}_D^\mu (v)$ with respect to the basis $\{\mathbf{1}_{A_j}\}$ and \hat{P} is a matrix representation of the projected transfer operator $\mathcal{Q}_D^\mu \mathcal{T}|_D$ with respect to the basis $\{\mathbf{1}_{A_j}\}$.

Lastly, to complete the sketched relationship between the projected transfer operators, note that if $w' = \mu v'$, then $\hat{w} = \hat{\mu} \hat{v}$ where \hat{w} and \hat{v} are the vector representations of $\mathcal{Q}_{\mu D}^{1/\mu}(w')$ and $\mathcal{Q}_D^\mu(v')$ with respect to the bases $\{\varphi_j\}$ and $\{\mathbf{1}_{A_j}\}$, respectively. It then follows that

$$\mathcal{Q}_{\mu D}^{1/\mu} \mathcal{P} \mathcal{Q}_{\mu D}^{1/\mu}(w') = \sum_{j=1}^m \hat{w}^T \hat{P}(j) \varphi_j \quad (5.33a)$$

$$= \sum_{j=1}^m ((\hat{\mu} \hat{v})^T \hat{P})(j) \varphi_j \quad (5.33b)$$

$$= \{ \text{detailed balance of } (\hat{P}, \hat{\mu}) \} \quad (5.33c)$$

$$\sum_{j=1}^m \hat{P} \hat{v}(j) (\mu \mathbf{1}_{A_j}) \quad (5.33d)$$

$$= \mu \sum_{j=1}^m \hat{P} \hat{v}(j) \mathbf{1}_{A_j} \quad (5.33e)$$

$$= \mu \mathcal{Q}_D^\mu \mathcal{T} \mathcal{Q}_D^\mu (v'). \quad (5.33f)$$

Summarizing, we get

Proposition 5.2. *Let $(X_k)_{k \in \mathbb{N}}$ be a reversible Markov chain with transition matrix P and stationary distribution μ . Let (A_1, \dots, A_m) be a partition of state space Z and let $(\hat{X}_k)_{k \in \mathbb{N}}$ denote the reduced Markov chain defined in Eqs. (5.14) and (5.15) with state space $\hat{Z} = \{1, \dots, m\}$, transition matrix \hat{P} , and stationary distribution $\hat{\mu}$. Let furthermore D and μD be the subspaces of \mathbb{R}^l defined by $D = \text{span}\{\mathbb{1}_{A_1}, \dots, \mathbb{1}_{A_m}\}$ and $\mu D = \text{span}\{\varphi_1, \dots, \varphi_m\}$, where $\varphi_j = \mu \mathbb{1}_{A_j} / \hat{\mu}(j)$. Then \hat{P} is the matrix representation of the projected transfer operators $\mathcal{Q}_{\mu D}^{1/\mu} \mathcal{P}_{\mu D}$ and $\mathcal{Q}_D^\mu \mathcal{T}_D$ with respect to the multiplication of vectors to the matrix from the left and right, respectively:*

$$\begin{aligned} \mathcal{Q}_{\mu D}^{1/\mu} \mathcal{P}_{\mu D} : \mu D &\rightarrow \mu D \\ w &\mapsto \sum_{j=1}^m (\hat{w}^T \hat{P})(j) \varphi_j \quad \text{where } w = \sum_{j=1}^m \hat{w}(j) \varphi_j, \end{aligned}$$

and

$$\begin{aligned} \mathcal{Q}_D^\mu \mathcal{T}_D : D &\rightarrow D \\ v &\mapsto \sum_{j=1}^m (\hat{P} \hat{v})(j) \mathbb{1}_{A_j} \quad \text{where } v = \sum_{j=1}^m \hat{v}(j) \mathbb{1}_{A_j}. \end{aligned}$$

Moreover, we have the relationship

$$\mathcal{Q}_{\mu D}^{1/\mu} \mathcal{P}_{\mu D} \mathcal{Q}_{\mu D}^{1/\mu} (w') = \mu \mathcal{Q}_D^\mu \mathcal{T}_D \mathcal{Q}_D^\mu (v') \quad (5.34)$$

for all $v', w' \in \mathbb{R}^l$ such that $w' = \mu v'$.

5.3 Core Set Markov State Models

In the previous section, we defined Markov state models for a full partition of state space and showed that there is a one-to-one relationship between the Markov chain (\hat{X}_k) and the matrix representation of the projected transfer operators. In this section, we generalize the projected transfer operator approach to subspaces D that are not necessarily spanned by the indicator functions on the sets of a partition of state space. We derive the matrix representation for the projected transfer operators and show that the same relationships between the projected transfer operators hold (see again Figure 5.3). On this basis, we define general Markov state models as the Markov chain whose transition matrix corresponds to the matrix representation of the projected transfer operators. Subsequently, we show that there is an interpretation of the matrix representation which allows to estimate it from trajectory data in the case the subspace D is spanned by committor functions with respect to so-called *core sets*; that is, some nonempty, disjoint sets which do not necessarily form a full partition of state space. It is this essential feature that makes core set Markov state models so interesting.

5.3.1 Projected Transfer Operators on General Subspaces

In what follows, let $\{q_1, \dots, q_m\}$ be a basis of D . The only assumption we will make on D is that $\mathbf{1} \in D$, which means that the probability density with respect to μ is an element of D . As in the case of full partition models, we interpret D as the subspace of \mathbb{R}^l which represents the probability densities with respect to μ that are representable in the reduced model, and we interpret the associated subspace μD as the subspace of \mathbb{R}^l which comprises the probability distributions representable in the reduced model.

Note that if we define $\hat{\mu}$ by

$$\hat{\mu}(j) = \sum_{i \in Z} \mu(i) q_j(i), \quad (5.35)$$

then, as before, we have that $v \in D$ is a probability density with respect to μ if and only if \hat{v} is a probability density with respect to $\hat{\mu}$:

$$1 = \sum_{i \in Z} v(i) \mu(i) \quad (5.36a)$$

$$= \sum_{i \in Z} \mu(i) \left(\sum_{j=1}^m \hat{v}(j) q_j \right) (i) \quad (5.36b)$$

$$= \sum_{j=1}^m \hat{v}(j) \sum_{i \in Z} \mu(i) q_j(i) \quad (5.36c)$$

$$= \sum_{j=1}^m \hat{v}(j) \hat{\mu}(j). \quad (5.36d)$$

Now, let $\{\varphi_1, \dots, \varphi_m\}$ be a basis of μD , where $\varphi_j = \frac{\mu q_j}{\hat{\mu}(j)}$. This basis seems natural for μD since if $w \in \mu D$ is a probability distribution, then its vector representation with respect to the φ_j 's, i.e., $w = \sum_{j=1}^m \hat{w}(j) \varphi_j$, is again a probability distribution:

$$1 = \sum_{i \in Z} w(i) \quad (5.37a)$$

$$= \sum_{i \in Z} \left(\sum_{j=1}^m \hat{w}(j) \varphi_j \right) (i) \quad (5.37b)$$

$$= \sum_{j=1}^m \hat{w}(j) \sum_{i \in Z} \varphi_j(i) \quad (5.37c)$$

$$= \sum_{j=1}^m \hat{w}(j) \sum_{i \in Z} \frac{\mu q_j(i)}{\hat{\mu}(j)} \quad (5.37d)$$

$$= \sum_{j=1}^m \hat{w}(j) \frac{1}{\hat{\mu}(j)} \sum_{i \in Z} \mu(i) q_j(i) \quad (5.37e)$$

$$= \sum_{j=1}^m \hat{w}(j). \quad (5.37f)$$

$$(5.37g)$$

We define the matrices \hat{P} and W by

$$\begin{aligned}\hat{P}(j, k) &= \hat{\mu}(k) \langle \mathcal{P}(\varphi_j), \varphi_k \rangle_{\frac{1}{\mu}} \\ &= \frac{\langle \mathcal{T}(q_j), q_k \rangle_{\mu}}{\hat{\mu}(j)},\end{aligned}\quad (5.38)$$

and

$$\begin{aligned}W(i, j) &= \hat{\mu}(j) \langle \varphi_i, \varphi_j \rangle_{\frac{1}{\mu}} \\ &= \frac{\langle q_i, q_j \rangle_{\mu}}{\hat{\mu}(i)}.\end{aligned}\quad (5.39)$$

At this point, notice that $(\hat{P}, \hat{\mu})$ are in detailed balance:

$$\hat{\mu}(j) \hat{P}(j, k) = \langle \mathcal{T}(q_j), q_k \rangle_{\mu} \quad (5.40a)$$

$$= \langle q_j, \mathcal{T}(q_k) \rangle_{\mu} \quad (5.40b)$$

$$= \hat{\mu}(k) \hat{P}(k, j) \quad (5.40c)$$

by self-adjointness of \mathcal{T} with respect to $\langle \cdot, \cdot \rangle_{\mu}$.

Furthermore, we get by straightforward calculation,

$$\langle \varphi_i, \varphi_j \rangle_{\frac{1}{\mu}} = \frac{S(i, j)}{\hat{\mu}(i) \hat{\mu}(j)}, \quad (5.41)$$

where $S(i, j) = \langle q_i, q_j \rangle_{\mu}$. Thus if $S(\varphi)$ is the matrix defined by $S_{\varphi}(i, j) = \langle \varphi_i, \varphi_j \rangle_{\frac{1}{\mu}}$, then

$$S_{\varphi}^{-1}(i, j) = \hat{\mu}(i) \hat{\mu}(j) S^{-1}(i, j). \quad (5.42)$$

With this we get by the definition of the orthogonal projection $\mathcal{Q}_{\mu D}^{1/\mu}$

$$\mathcal{Q}_{\mu D}^{1/\mu} \mathcal{P}(\varphi_k) = \sum_{i, j=1}^m S_{\varphi}^{-1}(i, j) \langle \mathcal{P}(\varphi_k), \varphi_i \rangle_{\frac{1}{\mu}} \varphi_j \quad (5.43a)$$

$$= \sum_{j=1}^m \varphi_j \sum_{i=1}^m S^{-1}(i, j) \hat{\mu}(i) \hat{\mu}(j) \frac{\hat{P}(k, i)}{\hat{\mu}(i)} \quad (5.43b)$$

$$= \sum_{j=1}^m \varphi_j \sum_{i=1}^m \hat{P}(k, i) W^{-1}(i, j) \quad (5.43c)$$

$$= \sum_{j=1}^m (\hat{P} W^{-1})(k, j) \varphi_j. \quad (5.43d)$$

This means that $\hat{P} W^{-1}$ is a matrix representation of $\mathcal{Q}_{\mu D}^{1/\mu} \mathcal{P}|_{\mu D}$ with respect to the basis $\{\varphi_j\}$. If \hat{w} is a vector representation of $w \in \mu D$ with respect to $\{\varphi_j\}$, i.e., $w = \sum_{j=1}^m \hat{w}(j) \varphi_j$, then $\hat{w}^T \hat{P} W^{-1}$ is a vector representation of $\mathcal{Q}_{\mu D}^{1/\mu} \mathcal{P}(w)$.

Similarly in the case of the transfer operator \mathcal{T} , we get by the definition of \mathcal{Q}_D^μ that

$$\mathcal{Q}_D^\mu \mathcal{T}(q_k) = \sum_{i,j=1}^m S^{-1}(i,j) \langle \mathcal{T}(q_k), q_i \rangle_\mu q_j \quad (5.44a)$$

$$= \{W(i,j) = \frac{S(i,j)}{\hat{\mu}(i)} \Rightarrow W^{-1}(j,i) = \hat{\mu}(i)S^{-1}(j,i) = \hat{\mu}(i)S^{-1}(j,i)\} \quad (5.44b)$$

$$\sum_{j=1}^m q_j \sum_{i=1}^m \frac{W^{-1}(j,i)}{\hat{\mu}(i)} \hat{\mu}(k) \hat{P}(k,i) \quad (5.44c)$$

$$= \{ \text{detailed balance of } (\hat{P}, \hat{\mu}) \} \quad (5.44d)$$

$$\sum_{j=1}^m q_j \sum_{i=1}^m W^{-1}(j,i) \hat{P}(i,k) \quad (5.44e)$$

$$= \sum_{j=1}^m (W^{-1} \hat{P})(j,k) q_j. \quad (5.44f)$$

It follows that if $v \in D$ has the vector representation \hat{v} with respect to the basis $\{q_i\}$, i.e., $v = \sum_{j=1}^m \hat{v}(j)q_j$, then

$$\mathcal{Q}_D^\mu \mathcal{T}(v) = \sum_{j=1}^m (W^{-1} \hat{P} \hat{v})(j) q_j, \quad (5.45)$$

and thus, $W^{-1} \hat{P} \hat{v}$ is a vector representation of $\mathcal{Q}_D^\mu \mathcal{T}(v)$ with respect to the basis $\{q_i\}$ and $\mathcal{Q}_D^\mu \mathcal{T}|_D$ has the matrix representation $W^{-1} \hat{P}$ with respect to the basis $\{q_i\}$.

Lastly, we show that

$$\mathcal{Q}_{\mu D}^{1/\mu} \mathcal{P} \mathcal{Q}_{\mu D}^{1/\mu}(w') = \mu \mathcal{Q}_D^\mu \mathcal{T} \mathcal{Q}_D^\mu(v'), \quad (5.46)$$

where $w' = \mu v'$. In order to show this, notice that if $w = \mathcal{Q}_{\mu D}^{1/\mu}(w') = \sum_{j=1}^m \hat{w}(j)q_j$ and $v = \mathcal{Q}_D^\mu(v') = \sum_{j=1}^m \hat{v}(j)q_j$, then by straightforward calculation $\hat{w} = \hat{\mu} \hat{v}$.

It follows that

$$\mathcal{Q}_{\mu D}^{1/\mu} \mathcal{P} \mathcal{Q}_{\mu D}^{1/\mu}(w') = \sum_{j=1}^m \hat{w}^T \hat{P} W^{-1}(j) \varphi_j \quad (5.47a)$$

$$= \sum_{j=1}^m (\hat{\mu} \hat{v})^T \hat{P} W^{-1}(j) \varphi_j \quad (5.47b)$$

$$= \sum_{j=1}^m \varphi_j \sum_{k=1}^m \hat{\mu}(k) \hat{v}(k) \hat{P} W^{-1}(k, j) \quad (5.47c)$$

$$= \sum_{j=1}^m \varphi_j \sum_{k=1}^m \hat{v}(k) \hat{\mu}(j) W^{-1} \hat{P}(j, k) \quad (5.47d)$$

$$= \sum_{j=1}^m W^{-1} \hat{P} \hat{v}(j) \hat{\mu}(j) \varphi_j \quad (5.47e)$$

$$= \sum_{j=1}^m W^{-1} \hat{P} \hat{v}(j) (\mu q_j) \quad (5.47f)$$

$$= \mu \sum_{j=1}^m W^{-1} \hat{P} \hat{v}(j) q_j \quad (5.47g)$$

$$= \mu \mathcal{Q}_D^\mu \mathcal{T} \mathcal{Q}_D^\mu(v'), \quad (5.47h)$$

where we used the fact that

$$\hat{\mu}(j) \hat{P} W^{-1}(j, i) = \hat{\mu}(i) W^{-1} \hat{P}(i, j), \quad (5.48)$$

since $(\hat{P}, \hat{\mu})$ are in detailed balance as is $(W^{-1}, \hat{\mu})$:

$$\hat{\mu}(i) W^{-1}(i, j) = \hat{\mu}(i) \hat{\mu}(j) S^{-1}(i, j) \quad (5.49a)$$

$$= \hat{\mu}(i) \hat{\mu}(j) S^{-1}(j, i) \quad (5.49b)$$

$$= \hat{\mu}(j) W^{-1}(j, i). \quad (5.49c)$$

We summarize the results in the following

Proposition 5.3. *Let $(X_k)_{k \in \mathbb{N}}$ be a reversible Markov chain with transition matrix P and stationary distribution μ . Let $D \subset \mathbb{R}^l$ with $\mathbf{1} \in D$ and let q_1, \dots, q_m be a basis of D . Let \hat{P} , W and $\hat{\mu}$ be defined according to Eqs. (5.38), (5.39), and (5.35), respectively. Moreover, let μD be the subspace of \mathbb{R}^l with basis $\{\varphi_1, \dots, \varphi_m\}$, where $\varphi_j = \mu q_j / \hat{\mu}(j)$. Then $\hat{P} W^{-1}$ is the matrix representation of the projected transfer operator $\mathcal{Q}_{\mu D}^{1/\mu} \mathcal{P} |_{\mu D}$ with respect to the multiplication of vectors to the matrix from the left:*

$$\begin{aligned} \mathcal{Q}_{\mu D}^{1/\mu} \mathcal{P} |_{\mu D} : \mu D &\rightarrow \mu D \\ w &\mapsto \sum_{j=1}^m (\hat{w}^T \hat{P} W^{-1})(j) \varphi_j \quad \text{where } w = \sum_{j=1}^m \hat{w}(j) \varphi_j. \end{aligned}$$

Similarly, $W^{-1} \hat{P}$ is the matrix representation of the projected transfer operator

$\mathcal{Q}_D^\mu \mathcal{T}|_D$ with respect to the multiplication of vectors to the matrix from the right:

$$\begin{aligned} \mathcal{Q}_D^\mu \mathcal{T}|_D : D &\rightarrow D \\ v &\mapsto \sum_{j=1}^m (W^{-1} \hat{P} \hat{v})(j) q_j \quad \text{where } v = \sum_{j=1}^m \hat{v}(j) q_j. \end{aligned}$$

Moreover, we have the relationship

$$\mathcal{Q}_{\mu D}^{1/\mu} \mathcal{P} \mathcal{Q}_{\mu D}^{1/\mu}(w') = \mu \mathcal{Q}_D^\mu \mathcal{T} \mathcal{Q}_D^\mu(v') \quad (5.50)$$

for all $v', w' \in \mathbb{R}^l$ such that $w' = \mu v'$.

5.3.2 General Markov State Models

On the basis of Proposition 5.3, we can define general Markov state models in such a way that there is again a correspondence between the matrix representation of the projected transfer operators and the transition matrix of the Markov chain. To do so, let D be the subspace of \mathbb{R}^l with basis q_1, \dots, q_m and let \hat{P}, W , and $\hat{\mu}$ be defined according to Eqs. (5.38), (5.39), and (5.35), respectively. Let $\hat{Z} = \{1, \dots, m\}$.

Now, let us define the general Markov state model as the Markov chain $(\hat{X}_k)_{k \in \mathbb{N}}$ with quasi-transition matrix $\hat{P}W^{-1}$. We write “quasi”-transition matrix since we are not sure that $\hat{P}W^{-1}$ is always a stochastic matrix. We only know that its rows sum up to one since this is the case for both \hat{P} and W , and thus also for W^{-1} as well as $\hat{P}W^{-1}$. In the examples, the entries of $\hat{P}W^{-1}$ are also non-negative, but we are not sure whether this holds in general.

Furthermore, notice that $\hat{\mu}$ is the stationary distribution of (\hat{X}_k) since $(\hat{P}, \hat{\mu})$ as well as $(W^{-1}, \hat{\mu})$ are in detailed balance, which implies that $\hat{\mu} \hat{P} = \hat{\mu}$ and $\hat{\mu} W^{-1} = \hat{\mu}$, respectively. Thus,

$$\hat{\mu} \hat{P} W^{-1} = \hat{\mu} W^{-1} = \hat{\mu}. \quad (5.51)$$

Therefore, the time-reversed matrix (see Section 5.1.1) with (i, j) th entry given by $\frac{\hat{\mu}(j)}{\hat{\mu}(i)} \hat{P} W^{-1}(j, i)$ is well-defined. By Eq. (5.48), the time-reversed matrix thus corresponds to $W^{-1} \hat{P}$. This implies that (\hat{X}_k) is reversible with respect to $\hat{\mu}$ if and only if \hat{P} and W^{-1} commute; that is,

$$\hat{P} W^{-1} = W^{-1} \hat{P}. \quad (5.52)$$

By Proposition 5.3, the definition of (\hat{X}_k) is made in such a way that, just as in the case of full partition Markov state models, both a) the transfer operator $\hat{P} : \mathbb{R}^m \rightarrow \mathbb{R}^m$ of (\hat{X}_k) has the same matrix representation $\hat{P}W^{-1}$ (with respect to the canonical basis of \mathbb{R}^m) as the projected transfer operator $\mathcal{Q}_{\mu D}^{1/\mu} \mathcal{P}$ (with respect to the φ_j 's), and b) the transfer operator $\hat{\mathcal{T}} : \mathbb{R}^l \rightarrow \mathbb{R}^l$ of (\hat{X}_k) has the same matrix representation $W^{-1} \hat{P}$ (with respect to the canonical basis) as the projected transfer operator $\mathcal{Q}_D^\mu \mathcal{T}$ (with respect to the q_j 's).

Lastly, notice that we can again consider the spaces $l^2(\hat{\mu})$ and $l^2(\frac{1}{\hat{\mu}})$, which correspond to \mathbb{R}^m equipped with the scalar products $\langle \cdot, \cdot \rangle_{\hat{\mu}}$ and $\langle \cdot, \cdot \rangle_{\frac{1}{\hat{\mu}}}$, respectively. We again get the relationship

$$\langle v_1, v_2 \rangle_{\hat{\mu}} = \langle \hat{\mu} v_1, \hat{\mu} v_2 \rangle_{\frac{1}{\hat{\mu}}}. \quad (5.53)$$

However, in general,

$$\langle \hat{v}_1, \hat{v}_2 \rangle_{\hat{\mu}} \neq \langle v_1, v_2 \rangle_{\mu}, \quad (5.54)$$

where \hat{v}_1, \hat{v}_2 are the vector representations of v_1, v_2 with respect to the basis $\{q_1, \dots, q_m\}$ of D , i.e., $v_1 = \sum_{j=1}^m \hat{v}_1(j)q_j$ and $v_2 = \sum_{j=1}^m \hat{v}_2(j)q_j$. Thus, in general,

$$\|\hat{v}\|_{\hat{\mu}} \neq \|v\|_{\mu}, \quad (5.55)$$

where \hat{v} is the vector representation of $v \in D$ with respect to the q_j 's. Equality holds in (5.54), e.g., if the basis vectors q_1, \dots, q_m are orthogonal with respect to $\langle \cdot, \cdot \rangle_{\mu}$ and

$$\|q_j\|_{\mu}^2 = \hat{\mu}(j). \quad (5.56)$$

This is, for instance, the case for full partition Markov state models. Similarly, in general

$$\langle \hat{w}_1, \hat{w}_2 \rangle_{\frac{1}{\hat{\mu}}} \neq \langle w_1, w_2 \rangle_{\frac{1}{\mu}}, \quad (5.57)$$

where \hat{w}_1, \hat{w}_2 are vector representations of $w_1, w_2 \in \mu D$ with respect to the basis $\{\varphi_1, \dots, \varphi_m\}$, as well as

$$\|\hat{w}\|_{\frac{1}{\hat{\mu}}} \neq \|w\|_{\frac{1}{\mu}}, \quad (5.58)$$

where \hat{w} is the vector representation of $w \in \mu D$ with respect to the φ_j 's. Equality holds in (5.57) under the same conditions as equality in (5.54), for instance, in the case of full partition Markov state models.

5.3.3 Core Set Markov State Models: Interpretation and Estimation

We now focus on special subspaces D . For this, let us consider pairwise disjoint sets $C_1, \dots, C_m \subset Z$, which we call *core sets*. Let $C = \cup_{i=1}^m C_i$. In contrast to the case of full partition Markov state models, we here do not assume that these sets C_1, \dots, C_m form a full partition of state space. In this case, the definition of the process \tilde{X} in Eq. (5.15) has to be adjusted since it is not well defined anymore. We set

$$\tilde{X}_k = i \Leftrightarrow X_{\sigma(k)} = i, \text{ where } \sigma(k) = \max\{t \leq k \mid X_t \in C\}. \quad (5.59)$$

We call the thus defined process (\tilde{X}_k) the *milestoning process*. Equation (5.59) means that the milestoning process remains in state i as long as the original Markov chain (X_k) last visited core set i .

Now, let q_1, \dots, q_m be the committor functions on the core sets C_1, \dots, C_m . Note that the q_i 's are linearly independent and

$$\sum_{i=1}^m q_i(x) = 1 \quad \text{for each } x \in Z, \quad (5.60)$$

thus $\sum_{i=1}^m q_i = \mathbf{1}$. We can thus consider the subspace $D = \text{span}\{q_1, \dots, q_m\}$ for the construction of Markov state models. Markov state models that result from projections onto subspaces D spanned by committor functions on core sets are called *core set Markov state models*.

Note that we can interpret a state $i \in \hat{Z}$ of the core set Markov state model as representing the *affiliation* with core set C_i . In the case of full partition Markov state models, which are a special case of core set Markov state models (see Remark 5.1), a state $i \in \hat{Z}$ is interpreted as “being in subset C_i ”. In the case that the C_i ’s do not constitute a full partition of state space, the committor functions on these set can, because of Eq. (5.60), still be interpreted as affiliation functions that associate to each state $x \in Z$ its affiliation with the respective core set. In this way, being in a state $i \in \hat{Z}$ can be interpreted as giving for each $z \in Z$ the degree q_i of belonging to subset C_i .

Example 5.3. Let $(X_t)_{t \in \mathbb{T}}$ be again the Markov chain on state space $Z = \{0, \frac{1}{n}, \dots, 1\}$ with transition matrix P in our running example (Example 3.6). Let us consider the core sets $C_1 = \{0\}$ and $C_2 = \{1\}$. The respective committor function q_1 is shown in Figure 5.2 and $q_2 = 1 - q_1$. The core set Markov state model is thus a Markov chain on the state space $\hat{Z} = 1, 2$ where state $i \in \hat{Z}$ refers to q_i , $i = 1, 2$. The matrices \hat{P} and W given by

$$\hat{P} = \begin{pmatrix} .9327 & .0673 \\ .0673 & .9327 \end{pmatrix}, \quad (5.61)$$

$$W = \begin{pmatrix} .9333 & .0667 \\ .0667 & .9333 \end{pmatrix}. \quad (5.62)$$

We can thus calculate

$$\hat{P}W^{-1} = \begin{pmatrix} .9993 & .0007 \\ .0007 & .9993 \end{pmatrix}, \quad (5.63)$$

$$W^{-1}\hat{P} = \begin{pmatrix} .9993 & .0007 \\ .0007 & .9993 \end{pmatrix}. \quad (5.64)$$

Thus, $\hat{P}W^{-1}$ is a stochastic matrix and $\hat{P}W^{-1} = W^{-1}\hat{P}$. The Markov chain $(\hat{X}_t)_{t \in \mathbb{T}}$ representing the core set Markov state model is thus reversible with stationary distribution $\hat{\mu} = (.5, .5)$. Table 5.2 shows the two largest eigenvalues in absolute value. It shows that, in terms of the approximation of dominant eigenvalues, the core set Markov state model is better than the full partition model considered in Example 5.2.

(X_t)	(\hat{X}_t)
$\lambda_1 = 1$	$\tilde{\lambda}_1 = 1$
$\lambda_2 = .99863$	$\tilde{\lambda}_2 = .99857$

Table 5.2: Dominant eigenvalues of (X_t) and (\hat{X}_t) , respectively.

Using the milestoning process (\tilde{X}_k) , we can give an interpretation of the matrices \hat{P} and W given in Eqs. (5.38) and (5.39), respectively, which determine the matrix representation of the core set Markov state model, see Proposition

5.3. In order to do so, we first note the following relationship between the committor functions and the milestoning process:

$$q_i(z) = \mathbb{P}(\tilde{X}_k = i \mid X_k = z). \quad (5.65)$$

In words, the probability of the milestoning process being in state i at time k conditional on the original Markov chain being in state z is given by $q_i(z)$. This implies

$$\mathbb{P}(\tilde{X}_k = i) = \sum_{z \in Z} \mathbb{P}(\tilde{X}_k = i \mid X_k = z) \quad (5.66a)$$

$$= \sum_{z \in Z} q_i(z) \mu(z) \quad (5.66b)$$

$$= \{ \text{Eq. (5.35)} \} \quad (5.66c)$$

$$\hat{\mu}(i). \quad (5.66d)$$

Moreover, using Bayes Theorem, we get

$$\mathbb{P}(X_k = x \mid \tilde{X}_k = i) = \frac{\mathbb{P}(\tilde{X}_k = i \mid X_k = x) \mathbb{P}(X_k = x)}{\mathbb{P}(\tilde{X}_k = i)} \quad (5.67a)$$

$$= \frac{q_i(x) \mu(x)}{\hat{\mu}(i)}, \quad (5.67b)$$

which implies

$$\mathbb{P}(X_k \in A \mid \tilde{X}_k = i) = \sum_{z \in A} \frac{q_i(z) \mu(z)}{\hat{\mu}(i)}. \quad (5.68)$$

With this we get

$$\begin{aligned} & \mathbb{P}(\tau_{C_j}^k < \tau_{C \setminus C_j}^k \mid \tilde{X}_k = i) \\ &= \sum_{z \in Z} \mathbb{P}(\tau_{C_j}^k < \tau_{C \setminus C_j}^k \mid X_k = z) \mathbb{P}(X_k = z \mid \tilde{X}_k = i) \quad (5.69a) \end{aligned}$$

$$= \sum_{z \in Z} \frac{q_i(z) \mu(z)}{\hat{\mu}(i)} \mathbb{P}(\tau_{C_j}^k < \tau_{C \setminus C_j}^k \mid X_k = z) \quad (5.69b)$$

$$= \sum_{z \in Z} \frac{q_i(z) q_j(z) \mu(z)}{\hat{\mu}(i)} \quad (5.69c)$$

$$= \frac{\langle q_i, q_j \rangle_\mu}{\hat{\mu}(i)} \quad (5.69d)$$

$$= W(i, j). \quad (5.69e)$$

Furthermore,

$$\mathbb{P}(\tau_{C_j}^{k+1} < \tau_{C \setminus C_j}^{k+1} \mid \tilde{X}_k = i) \quad (5.70a)$$

$$= \sum_{z \in Z} \mathbb{P}(\tau_{C_j}^{k+1} < \tau_{C \setminus C_j}^{k+1} \mid X_k = z) \mathbb{P}(X_k = z \mid \tilde{X}_k = i) \quad (5.70b)$$

$$= \sum_{z \in Z} \frac{q_i(z) \mu(z)}{\hat{\mu}(i)} \sum_{x \in Z} \mathbb{P}(\tau_{C_j}^{k+1} < \tau_{C \setminus C_j}^{k+1} \mid X_{k+1} = x) \mathbb{P}(X_{k+1} = x \mid X_k = z) \quad (5.70c)$$

$$= \sum_{z \in Z} \frac{q_i(z) \mu(z)}{\hat{\mu}(i)} \sum_{x \in Z} q_j(x) P(z, x) \quad (5.70d)$$

$$= \sum_{z \in Z} \frac{q_i(z) \mu(z)}{\hat{\mu}(i)} (Pq_j)(z) \quad (5.70e)$$

$$= \frac{\langle q_i, \mathcal{T}(q_j) \rangle_\mu}{\hat{\mu}(i)} \quad (5.70f)$$

$$= \frac{\langle \mathcal{T}(q_i), q_j \rangle_\mu}{\hat{\mu}(i)} \quad (5.70g)$$

$$= \hat{P}(i, j), \quad (5.70h)$$

where the second to last equality follows by the self-adjointness of \mathcal{T} in $l^2(\mu)$.

We have thus just shown the following relationship between the core set Markov state models and the milestoning process (\tilde{X}_k) :

Proposition 5.4. *Let (X_k) be a reversible Markov chain. Let C_1, \dots, C_m be some given core sets and let (\tilde{X}_k) denote the associated milestoning process. Let $D = \text{span}\{q_1, \dots, q_m\}$, where the q_i 's are the committor functions associated with the core sets. Moreover, let \hat{P} and W be defined according to Eqs. (5.38) and (5.39), respectively. Then,*

$$W(i, j) = \mathbb{P}(\tau_{C_j}^k < \tau_{C \setminus C_j}^k \mid \tilde{X}_k = i), \quad (5.71)$$

$$\hat{P}(i, j) = \mathbb{P}(\tau_{C_j}^{k+1} < \tau_{C \setminus C_j}^{k+1} \mid \tilde{X}_k = i). \quad (5.72)$$

Thus, in words, $W(i, j)$ for $j \neq i$ gives the probability that the Markov chain next hits C_j while being in a state in $Z \setminus C$ at some time k and last came from core set C_i , where $C = \cup_{j=1}^m C_j$. Similarly, $\hat{P}(i, j)$ gives the probability that the next core set hit is C_j conditional on having hit the core set C_i last at some time k . Note that

$$W(i, j) = \begin{cases} \mathbb{P}(\tau_{C_j}^k < \tau_{C \setminus C_j}^k, X_k \in Z \setminus C \mid \tilde{X}_k = i) & \text{if } j \neq i \\ 1 - \sum_{l=1, l \neq i}^m \mathbb{P}(\tau_{C_l}^k < \tau_{C \setminus C_l}^k, X_k \in Z \setminus C \mid \tilde{X}_k = i) & \text{if } j = i. \end{cases} \quad (5.73)$$

Remark 5.1. The full partition Markov state models present a special case of core set Markov state models. In this case, the committor functions on the core sets are again the indicator functions on the respective sets. Thus the matrix W reduces to the identity matrix Id . Moreover, $\hat{\mu}(i) = \sum_{z \in C_i} \mu(z) = \mu(C_i)$, and

$$\hat{P}(i, j) = \mathbb{P}_\mu(\tilde{X}_1 = j \mid \tilde{X}_0 = i), \quad (5.74)$$

just as the matrix \hat{P} was defined for the full partition Markov state models. Note that Eq. (5.74) does not generally hold for Markov state models with core sets that do not form a full partition of state space. Instead,

$$\mathbb{P}(\tilde{X}_{k+1} = j \mid \tilde{X}_k = i) = \begin{cases} \hat{P}(i, j) - W(i, j) & \text{if } j \neq i \\ \hat{P}(i, i) - W(i, i) + 1 & \text{if } j = i, \end{cases} \quad (5.75)$$

see Djurdjevac et al. (2010) for a proof.

Using Proposition 5.4, it follows that both matrices W and \hat{P} can be estimated from trajectory data in the following way: given a realization (x_0, \dots, x_K) of (X_k) of length K , we can estimate

$$W^{*,K}(i, j) = \begin{cases} \frac{R_{ij}^K}{r_i^K} & \text{if } j \neq i, \\ 1 - \sum_{j \neq i} W^{*,K}(i, j) & \text{otherwise,} \end{cases} \quad (5.76)$$

$$\hat{P}^{*,K}(i, j) = \frac{R_{ij}^{+,K}}{r_i^K}, \quad (5.77)$$

where R_{ij}^K denotes the number of times where the chain came from core set C_i , is in a state in $Z \setminus C$ and hits C_j next, r_i^K is the total number of time steps the trajectory was in i ; that is, $\tilde{X}_k = i$, and $R_{ij}^{+,K}$ denotes the number of times where the chain came from core set C_i and hit C_j next.

5.4 Approximation Quality

In this section, we use the results of the previous sections on Markov state models and orthogonal projections to consider how well Markov state models actually approximate the original Markov chain.

5.4.1 Stationary Distribution and Stochastic Stability

We have shown in Eq. 5.51 that $\hat{\mu}$ as defined according to Eq. (5.35) is a stationary distribution of both full partition and core set Markov state models. Given a model of evolution with noise $(X^\varepsilon(k))_{k \in \mathbb{N}}$, we can thus consider the sequence of stationary distributions $\hat{\mu}^\varepsilon$ and ask whether the construction of Markov state models preserves stochastic stability in the sense that if $x \in C_i$ for some core set $i = 1, \dots, m$ then the state i is stochastically stable for $(\hat{X}^\varepsilon(k))_{k \in \mathbb{N}}$. This is obvious in the case of a full partition. In the case of core set Markov state models, it seems reasonable to assume that every stochastically stable state $x \in Z$ is contained in some core set, say C_i . Then $q_i(x) = 1$ while $q_j(x) = 0$ for all $j \neq i$ and thus the state i is stochastically stable with respect to $\hat{\mu}^\varepsilon$ for $\varepsilon \rightarrow 0$.

Example 5.4. Let $(X_t^\varepsilon)_{t \in \mathbb{T}}$ be the Markov chain of our running example (Example 3.6) with population size 11 and revision protocol parameter ε . Let $(\hat{X}_t^{\varepsilon,f})_{t \in \mathbb{T}}$ denote the Markov state model with full partition (A_1, A_2) , where $A_1 = \{0, \dots, 5/11\}$, $A_2 = \{6/11, \dots, 1\}$. From Example 3.7, we know that the stochastically stable population states of (X_t^ε) are

- 0 and 1 if $a = b = 1$;
- 1 if $a > b$; and
- 0 if $b > a$.

This is reflected in the values of $\hat{\mu}^\varepsilon$, see Figure 5.4.




Figure 5.4: Value of $\hat{\mu}^\varepsilon(2)$ for the full partition Markov state model $(\hat{X}_t^{\varepsilon,f})_{t \in \mathbb{T}}$ for various values of the payoff parameter a and as a function of ε , where the state 2 of $(\hat{X}_t^{\varepsilon,f})_{t \in \mathbb{T}}$ refers to the set A_2 .

5.4.2 Propagation of Probability Distributions

In the case of a full partition (Section 5.2.1), we defined the process $(\tilde{X}_k)_{k \in \mathbb{N}}$ on the basis of which we defined the transition matrix \hat{P} . We noted, however, that the process (\tilde{X}_k) is, however, in general non-Markovian. The question thus arises what error we make when approximating (\tilde{X}_k) by the Markovian (\hat{X}_k) . In general, this question translates to considering the approximation error between $\mathcal{Q}_{\mu D}^{1/\mu} \mathcal{P}_{|\mu D}^k$ and $(\mathcal{Q}_{\mu D}^{1/\mu} \mathcal{P}_{|\mu D})^k$. The following theorem gives an upper bound on this approximation error:

Theorem 5.5. *Let $(X_k)_{k \in \mathbb{N}}$ be a reversible discrete-time Markov chain on a finite state space Z with transition matrix P and unique stationary distribution μ . Let $1 = \lambda_1 > |\lambda_2| \geq \dots \geq |\lambda_{l_d}|$ denote the $l_d \leq l$ dominant eigenvalues of P . That is, for every other eigenvalue λ we have $|\lambda_{l_d}| > r \geq |\lambda|$ such that r is an upper bound on the remaining spectrum. Let $\mathbf{1} = u_1, u_2, \dots, u_{l_d}$ be the corresponding orthonormal left eigenvectors, i.e.,*

$$u_j P = \lambda_j u_j, \quad \langle u_i, u_j \rangle_{\frac{1}{\mu}} = \delta_{ij}. \quad (5.78)$$

Furthermore, set $\eta = \frac{r}{|\lambda_2|}$ and define

$$\delta = \max_{i=2, \dots, l_d} \|\mathcal{Q}_{\mu D}^{1/\mu, \perp}(u_i)\|_{\frac{1}{\mu}}, \quad (5.79)$$

where $\mathcal{Q}_{\mu D}^{1/\mu, \perp} = Id - \mathcal{Q}_{\mu D}^{1/\mu}$ and $\mathcal{Q}_{\mu D}^{1/\mu}$ is the orthogonal projection onto some subspace $D \subset l^2(\mu)$ with $\mathbf{1} \in D$, i.e., $\mathcal{Q}_{\mu D}^{1/\mu, \perp}$ is the projection onto the orthogonal complement of μD in $l^2(\frac{1}{\mu})$ and δ is the maximal projection error of the eigenvectors onto the space μD . Then, we have for $k \in \mathbb{N}$,

$$E(k) = \|\mathcal{Q}_{\mu D}^{1/\mu} \mathcal{P}_{|\mu D}^k - (\mathcal{Q}_{\mu D}^{1/\mu} \mathcal{P}_{|\mu D})^k\|_{\frac{1}{\mu}} \leq |\lambda_2|^k \min\{2, C(\delta, \eta, k)\}, \quad (5.80)$$

where $\|\cdot\|_{\frac{1}{\mu}}$ refers to the operator norm corresponding to the $l^2(\frac{1}{\mu})$ -norm on vectors; that is,

$$\|B\|_{\frac{1}{\mu}} = \max_{\|w\|_{\frac{1}{\mu}}=1} \|B(w)\|_{\frac{1}{\mu}} \quad (5.81)$$

for operators $B : l^2(\frac{1}{\mu}) \rightarrow l^2(\frac{1}{\mu})$; and

$$C(\delta, \eta, k) = ((l_d - 1)\delta + \eta) \left(C_{space}(\delta, k) + C_{spec}(\eta, k) \right) \quad (5.82)$$

$$C_{space}(\delta, k) = (l_d - 1)^{1/2} (k - 1) \delta \quad (5.83)$$

$$C_{spec}(\eta, k) = \frac{\eta}{1 - \eta} (1 - \eta^{k-1}). \quad (5.84)$$

Proof. Note that by Proposition 5.3, we have $\mathcal{Q}_{\mu D}^{1/\mu} \mathcal{P}_{|\mu D} = \mu \mathcal{Q}_D^\mu \mathcal{T}_{|D}$ and thus

$$\|\mathcal{Q}_{\mu D}^{1/\mu} \mathcal{P}_{|\mu D}^k - (\mathcal{Q}_{\mu D}^{1/\mu} \mathcal{P}_{|\mu D})^k\|_{\frac{1}{\mu}} = \|\mathcal{Q}_D^\mu \mathcal{T}_{|D}^k - (\mathcal{Q}_D^\mu \mathcal{T}_{|D})^k\|_{\mu}, \quad (5.85)$$

where $\|\cdot\|_{\mu}$ refers to the operator norm associated with the vector norm on $l^2(\mu)$. Moreover,

$$\|\mathcal{Q}_{\mu D}^{1/\mu, \perp}(w)\|_{\frac{1}{\mu}} = \|\mathcal{Q}_D^{\mu, \perp}(w/\mu)\|_{\mu} \quad \text{for each } w \in \mathbb{R}^l, \quad (5.86)$$

and therefore

$$\|\mathcal{Q}_{\mu D}^{1/\mu, \perp}(u_i)\|_{\frac{1}{\mu}} = \|\mathcal{Q}_D^{\mu, \perp}(u_i/\mu)\|_{\mu}, \quad (5.87)$$

where u_i/μ is the normalized right eigenvector of P corresponding to λ_i .

The proof of the theorem now closely follows the proof of Theorem 5 of Sarich (2011) with the only difference being to consider the absolute value of eigenvalues when it comes to estimating the upper bound of the norm of involved quantities since in our case we have to consider the possibility of negative eigenvalues whereas in Sarich (2011) only non-negative eigenvalues are accounted for. \square

In case equality holds in (5.57), as is the case for full partition models, we can reformulate the error $E(k)$ in terms of the Markov state model $(\hat{X}_k)_{k \in \mathbb{N}}$ on state space $\hat{Z} = \{1, \dots, m\}$ whose quasi-transition matrix is $\hat{P}W^{-1}$, where

\hat{P} and W are defined according to Eqs. (5.38) and (5.39), respectively. The approximation error $E(k)$ in inequality (5.80) then translates to

$$E(k) = \|\widehat{\mathcal{P}}_k - \hat{\mathcal{P}}^k\|_{\frac{1}{\mu}} \quad (5.88)$$

$$= \|\widehat{P}_k W^{-1} - (\hat{P}W^{-1})^k\|_{\frac{1}{\mu}}, \quad (5.89)$$

where $\widehat{\mathcal{P}}_k$ is the transfer operator of $(\hat{X}_k)_{k \in \mathbb{N}}$ that transfers probability distributions k time steps ahead; that is,

$$\begin{aligned} \widehat{\mathcal{P}}_k : \mathbb{R}^m &\rightarrow \mathbb{R}^m \\ \hat{w} &\mapsto \hat{w}^T \widehat{P}_k W^{-1}, \end{aligned}$$

with $\widehat{P}_k W^{-1}$ being the matrix representation of the projected transfer operator $\mathcal{Q}_{\mu D}^{1/\mu} \mathcal{P}_{\mu D}^k$. According to Eq. (5.38), \widehat{P}_k thus has entries

$$\widehat{P}_k(i, j) = \hat{\mu}(j) \langle \mathcal{P}^k(\varphi_i), \varphi_j \rangle_{\frac{1}{\mu}}. \quad (5.90)$$

Theorem 5.5 shows that the projection error δ plays a crucial role in the upper bound of the approximation error $E(k)$; in order to make $E(k)$ as small as possible we should thus find a subspace D such that this projection error δ is as small as possible. For full partition models, this is the case if the right eigenvectors are as constant as possible on the sets of the partition. This relationship between the approximation qualities of full partition Markov state models and the projection error δ motivates approaches that partition state space by clustering algorithms (as has, for example, been done in the molecular dynamics context by Krivov and Karplus, 2004; Noé et al., 2007b; Rao and Caffisch, 2004). Similarly, in terms of core set Markov state models, finding core sets C_1, \dots, C_m so that the projection error δ is as small as possible can be interpreted as a fuzzy clustering problem (Djurdjevac, 2012; Sarich, 2011). Note, moreover, that we encountered the projection error δ already in Theorem 4.2 where it determines the lower bound of the joint metastability of the partition. The smaller δ , the higher the joint metastability. Thus, a small projection error δ in the case of a full partition Markov state model implies not only a small approximation error $E(k)$ but also a high joint metastability of the partition.

Moreover, the upper bound in (5.80) depends on $\eta = r/|\lambda_2|$ as well as on $|\lambda_2|$ itself. More specifically, decreasing η and decreasing $|\lambda_2|$ leads to an improved upper bound. There are now two ways to arrive at a smaller η : one point of departure is to notice that η decreases with an increasing number of eigenvalues since the upper bound on the remaining spectrum r decreases. We could thus make the approximation error $E(k)$ as small as possible by

1. making η small by choosing an appropriate number of eigenvalues, and then
2. choosing a partition of state space or the core sets in such a way that the resulting projection error on the associated eigenvectors is small.

However, if we fix the dimension of the subspace D we consider for projection, the projection error δ usually increases with the number of eigenvalues taken into consideration. Thus, in this case, increasing the number of eigenvalues does

not lead to an improved upper bound. Fixing the dimension of the subspace D is interesting since from Theorem 4.2 we know that we should not choose the number of sets in the partition higher than the number of dominant eigenvalues. Thus, for a full partition Markov state model with high joint metastability, we would fix the dimension of the subspace to the number of dominant eigenvalues.

Another approach to improving the upper bound is to introduce a *lag time* κ , where we require that $\kappa = \nu k$ with $\nu \in \mathbb{N}$, and to consider the snapshot dynamics (X_κ) with associated transition matrix P^κ . In this case, the second and third largest eigenvalues of P^κ are just λ_2^κ and λ_3^κ , respectively, and thus, we get an improved $\eta^\kappa = (|\lambda_3|/|\lambda_2|)^\kappa$. Moreover, introducing a lag time does not only decrease the ratio between the second and third eigenvalues, but also the second eigenvalue itself, and this also lowers the upper bound as given in (5.80). Notice, however, that this approach might be impractical for applications, i.e., when estimating the matrices \hat{P}_κ and W from trajectory data, because we might need a κ so large that it is not feasible to generate enough trajectory data to hold the resulting statistical error within acceptable bounds.

Example 5.5. Let us consider the Markov chain $(X_t)_{t \in \mathbb{T}}$ of our running example (Example 3.6) with parameters $a = 1, b = 1, \varepsilon = .3, n = 11$. For this case, we analyze the error $E(t)$ where $E(t)$ corresponds to the error $E(k)$ in (5.80) with $k = nt$, and we will focus on Markov state models with two states only. In this case, the error as well as the upper bound in (5.80) is smallest for $l_d = 2$, which is the reason we set $l_d = 2$ in what follows. Notice that all eigenvalues of the associated transition matrix are positive.

The state space of (X_k) is $Z = \{0, \frac{1}{11}, \dots, 1\}$ and we consider the core set Markov state models with core sets $C_1 = \{0, \dots, r/11\}$ and $C_2 = \{1 - r/11, \dots, 1\}$ for $r = 0, \dots, 6$. Notice that $r = 6$ corresponds to the full partition Markov state model considered in Example 5.2. In Figure 5.5, we display the projection error δ for the different core set models. It is smallest for $r = 2$, i.e., for core sets $C_1 = \{0, 1/11\}$ and $C_2 = \{10/11, 1\}$, and largest for the full partition model.

We next consider the upper bound in (5.80) for various values of η and for the projection errors δ^f and δ^c corresponding to the full partition Markov state model $(\hat{X}_t^f)_{t \in \mathbb{T}}$, i.e., $r = 6$, and to the core set Markov state model $(\hat{X}_t^c)_{t \in \mathbb{T}}$ with core sets $C_1 = \{0, 1/11\}$, $C_2 = \{10/11, 1\}$, respectively. Figure 5.6 displays these upper bounds as functions of time. Notice that for fixed δ , the upper bound is a monotonically increasing function of η . The figure shows that in the case of (\hat{X}_t^f) the upper bound is – in fact, even for small values of η – close to the worst case upper bound, which is given according to (5.80) by $2\lambda_2^{nt}$, whereas considering the core set model already improves on this a lot. However, since in our example $\eta = \lambda_3/\lambda_2 = .9079$, the upper bound given by (5.80) just coincides in both cases, except for very small times, with the worst case upper bound. As Figure 5.7 demonstrates, we can improve the upper bound considerably by considering the snapshot dynamics (X_t) for $t = \nu t$ with $\nu \in \mathbb{N}$.

In addition to the upper bound, we investigated the approximation error $E(t)$ for both (\hat{X}_t^f) and (\hat{X}_t^c) , see Figure 5.8. It shows that the upper bound considerably overestimates the error $E(t)$ for both Markov state models. In fact, the approximation quality is especially good in the core set model.

Lastly, notice that the projection error δ is itself a function of the popula-




Figure 5.5: Projection error δ for the core set Markov state models with core sets $C_1 = \{0, \dots, r/11\}$ and $C_2 = \{1 - r/11, \dots, 1\}$ for $r = 0, \dots, 6$.

tion size n and the parameter ϵ of the revision protocol, see Figure 5.9. The projection error δ is monotonically decreasing in population size and the difference in the projection errors between the full partition model and the core set model gets less pronounced for increasing n . In contrast, the projection error is a monotone increasing function of ϵ and the difference between the projection errors between the two models is less pronounced for small ϵ . Notice that this actually does not come as a surprise since the second eigenvector u_2 as well as the committor function q_1 approach a step function for larger n and small ϵ , see Figures 5.10 and 5.11. However, at the same time as the projection error decreases, i.e., for increasing population size n and decreasing ϵ , the value of η increases, see Figure 5.12. The impact of changing the population size or the revision protocol parameter ϵ thus has to be judged from case to case.





Figure 5.6: Upper bound of $E(t)$ as provided in (5.80) for (a) the full partition model (\hat{X}_t^f) and (b) the core set model (\hat{X}_t^c).




Figure 5.7: Upper bound of the $E(t)$, where $E(t)$ corresponds to $E(k)$ with $k = nvt$ as provided in (5.80).






Figure 5.8: Approximation error $E(t)$ for the Markov state models (\hat{X}_t^f) and (\hat{X}_t^c) .




(a)




(b)

Figure 5.9: The projection error δ as a function of (a) the population size n and (b) as a function of the revision protocol parameter ε for the Markov state models (\hat{X}_t^f) and (\hat{X}_t^c) based on the Markov chain of our running example (Example 3.6) with parameter values $a = 1, b = 1, \varepsilon = .3$ and core sets $C_1 = \{0, 1/n\}, C_2 = \{n - 1/n, 1\}$.




(a)




(b)

Figure 5.10: The second eigenvector u_2 normalized such that $u_2(0) = 1$ of our running example (Example 3.6, with parameter values $a = 1, b = 1, n = 11$) for various values of (a) the population size n and (b) of the revision protocol parameter ε .




(a)




(b)

Figure 5.11: Committor function q_1 of our running example (Example 3.6, with parameter values $a = 1, b = 1, n = 11$) for the core sets $C_1 = \{0, 1/n\}$, $C_2 = \{n - 1/n, 1\}$ for various values of (a) the population size n and (b) of the revision protocol parameter ε .



(a)



(b)

Figure 5.12: The second and third largest eigenvalue of \mathcal{P} as a function of (a) population size and (b) of the revision protocol parameter ε for our running example (Example 3.6) with parameters $a = 1, b = 1, \varepsilon = .3, n = 11$.

5.4.3 Eigenvalues

Next, we consider how well the constructed Markov state models approximate the eigenvalues and thus the timescales of the original Markov chain $(X_k)_{k \in \mathbb{N}}$:

Theorem 5.6. *Let $(X_k)_{k \in \mathbb{N}}$ be a reversible Markov chain on a finite state space Z with transition matrix P and unique stationary distribution μ . Let $1 = \lambda_1 > \lambda_2 > \dots > \lambda_{l_d}$ be the l_d dominating eigenvalues of the transfer operator \mathcal{P} , i.e., for every other eigenvalue λ of \mathcal{P} we have $\lambda < \lambda_{l_d}$. Let $\mathbb{1} = u_1, \dots, u_{l_d}$ be the corresponding orthonormal left eigenvectors. Moreover, let $D \subset l^2(\mu)$ be a subspace of dimension $m \geq l_d$ with $\mathbb{1} \in D$ and let $1 = \hat{\lambda}_1 > \hat{\lambda}_2 > \dots > \hat{\lambda}_{l_d}$ be the dominating eigenvalues of the projected transfer operator $\mathcal{Q}_{\mu D}^{1/\mu} \mathcal{P}|_D$. Then,*

$$\max_{i=2, \dots, l_d} |\lambda_i - \hat{\lambda}_i| \leq (\max\{0, \lambda_2\} - \min\{0, \lambda_{l_d}\}) (l_d - 1) \delta^2, \quad (5.91)$$

where

$$\delta = \max_{i=2, \dots, l_d} \|\mathcal{Q}_{\mu D}^{1/\mu, \perp}(u_i)\|_{\frac{1}{\mu}} \quad (5.92)$$

with $\mathcal{Q}_{\mu D}^{1/\mu, \perp} = Id - \mathcal{Q}_{\mu D}^{1/\mu}$.

In terms of Markov state models, Theorem 5.6 translates straightforwardly into the following

Corollary 5.7. *Let $(X_k)_{k \in \mathbb{N}}$ be a reversible Markov chain on a finite state space Z with transition matrix P and unique stationary distribution μ . Let $(\hat{X}_k)_{k \in \mathbb{N}}$ be an associated Markov state model on state space $\hat{Z} = \{1, \dots, m\}$ with quasi-transition matrix $\hat{P}W^{-1}$ and stationary distribution $\hat{\mu}$ defined according to Eqs. (5.38), (5.39) and (5.35), respectively. Let $1 = \lambda_1 > \lambda_2 > \dots > \lambda_{l_d}$ be the $l_d \leq l$ dominating eigenvalues of P and $1 = \hat{\lambda}_1 > \hat{\lambda}_2 > \dots > \hat{\lambda}_{l_d}$ be the $l_d \leq m$ dominating eigenvalues of $\hat{P}W^{-1}$. Then,*

$$\max_{i=2, \dots, l_d} |\lambda_i - \hat{\lambda}_i| \leq (\max\{0, \lambda_2\} - \min\{0, \lambda_{l_d}\}) (l_d - 1) \delta^2, \quad (5.93)$$

where

$$\delta = \max_{i=2, \dots, l_d} \|\mathcal{Q}_{\mu D}^{1/\mu, \perp}(u_i)\|_{\frac{1}{\mu}} \quad (5.94)$$

with $\mathcal{Q}_{\mu D}^{1/\mu, \perp} = Id - \mathcal{Q}_{\mu D}^{1/\mu}$.

Proof of Theorem 5.6. We have by Proposition 5.3 that λ is an eigenvalue of $\mathcal{Q}_{\mu D}^{1/\mu} \mathcal{P}(\mu v)$ if and only if λ is an eigenvalue of $\mathcal{Q}_D^\mu \mathcal{T}_D$:

$$\lambda v = \lambda(\mu v) / \mu \quad (5.95a)$$

$$= \mathcal{Q}_{\mu D}^{1/\mu} \mathcal{P}(\mu v) / \mu \quad (5.95b)$$

$$= \mathcal{Q}_D^\mu \mathcal{T}_D(v). \quad (5.95c)$$

Moreover,

$$\|\mathcal{Q}_{\mu D}^{1/\mu, \perp}(u_i)\|_{\frac{1}{\mu}} = \|\mathcal{Q}_D^{\mu, \perp}(u_i / \mu)\|_{\mu}, \quad (5.96)$$

for every u_i , and u_i/μ being the corresponding normalized right eigenvector of P . Using these relationships, the proof of the theorem follows the proof of Theorem 7 of Sarich (2011) closely. The only difference is that we have to take account of possibly negative eigenvalues when bounding the eigenvalue error using the Rayleigh-Ritz majorization error bounds from Knyazev and Argentati (2010). This requires us to change Equation (3.46) of Sarich (2011) to²

$$\max_{i=2,\dots,l_d} |\lambda_i - \hat{\lambda}_i| \leq (\max\{0, \lambda_2\} - \min\{0, \lambda_{l_d}\}) \max_j \sin^2(\theta_j(U, D)), \quad (5.97)$$


where $\Theta = \Theta(U, D) = (\theta_1, \dots, \theta_{l_d})$ is a vector of principal angles between the subspaces $U = \text{span}\{u_1, \dots, u_{l_d}\}$ and D . \square

Thus, the projection error δ again plays a crucial role in that a small projection error δ implies good approximation quality of our Markov state models in terms of the dominant eigenvalues.


Example 5.6. We again consider the Markov chain $(X_t)_{t \in \mathbb{T}}$ of our running example (Example 3.6) and the Markov state models (\hat{X}_t^f) and (\hat{X}_t^c) corresponding to the full partition $(\{0, 1/11, \dots, 5/11\}, \{6/11, \dots, 1\})$ of state space $Z = \{0, \frac{1}{n}, \dots, 1\}$ and to the core sets $C_1 = \{0, 1/11\}$, $C_2 = \{10/11, 1\}$, respectively. The projection error δ is shown in Figure 5.9 as a function of the population size n and the revision protocol parameter ε . Figure 5.13 shows the eigenvalue error $|\lambda_2 - \hat{\lambda}_2|$ as well as the upper bound of this error as provided by Equation 5.93. It shows that the upper bound is a good approximation for the eigenvalue error, especially in the case of larger population sizes and smaller values of ε . Moreover, it shows that the eigenvalue error as well as the upper bound are both considerably smaller for the core set model, which makes a difference in the case of smaller population sizes and larger values of ε . This is precisely the case where the other approximation approaches of stochastic evolutionary games (see Chapter 3.2) cannot be applied. This example thus demonstrates that core set Markov state models fill a gap and constitute an important complement to existing approximation approaches.

²There is a slight difference in notation between this thesis and Sarich (2011). Equation 5.97 above is stated in the notation used here. For ease of comparison with Sarich (2011) we also give the equation in the notation of Sarich (2011):

$$\max_{i=1,\dots,m-1} |\lambda_i - \hat{\lambda}_i| \leq (\max\{0, \lambda_1\} - \min\{0, \lambda_{m-1}\}) \max_j \sin^2(\theta_j(U, D)). \quad (5.98)$$



(a)



(b)

Figure 5.13: The eigenvalue error $|\lambda_2 - \hat{\lambda}_2|$ and the upper bound as provided by (5.93) for the Markov state models (\hat{X}_t^f) and (\hat{X}_t^c) as a function of (a) population size n , with parameters $a = 1, b = 1, \varepsilon = .3$, and (b) of the revision protocol parameter ε , with parameters $a = 1, b = 1, n = 11$.

Chapter 6

Identification of Metastable Sets for Agent-Based Evolutionary Models

In Chapter 3 we introduced stochastic evolutionary games and showed that there is a need for the characterization and identification of metastable dynamic behavior in these models. We subsequently discussed in Chapter 4 different characterizations of metastability and introduced in Chapter 5 an approach to the construction of Markov models on a reduced state space, which is appealing because the resulting Markov chains can be estimated from trajectory data and because we can give upper bounds for the resulting approximation errors. We learned that an important quantity in the upper bound for both the approximation error $E(k)$ as well as the eigenvalue error is the projection error δ defined via Equation (5.79) as the maximal projection error of the dominant eigenvectors onto the subspace μD considered for the construction of the Markov state model. Thus, if we want to construct core set Markov state models of reduced complexity that reproduce well the dominant timescales, and in this sense the metastability of our original Markov chain, we have to find core sets such that this projection error δ is minimized.

Against this background, we consider in this chapter the approach introduced by Sarich (2011) to the identification of core sets which result in a small projection error. This approach seems well suited since it is developed for Markov chains in which the overall metastability is controlled by a parameter $\varepsilon > 0$ in the sense that for smaller ε the metastability of the system under investigation should increase – and this is just the case for the models of evolution of noise which we introduced in Chapter 3.2.4. More specifically, in Section 6.1 we motivate the approach, building on which we introduce in Section 6.2 the algorithmic strategy to the identification of core sets and the construction of the resulting Markov state model. The algorithmic strategy is based only on trajectory data and seems especially interesting in the context of agent-based modeling since agent-based models usually lack model specifications and are, in addition, too complex to be investigated analytically, see Chapter 1. Moreover, we show that the identification algorithm to be presented preserves stochastic stability in the sense that, if a population state x is stochastically stable, it will

be identified as an element of the core set region C . One limitation of the approach, however, is that it depends on the initial Markov chain to be reversible. In Section 6.4 we therefore discuss this limitation and give an outlook for further research in this general direction.

6.1 Motivation

Starting point for the algorithmic strategy to the identification of core sets is the observation that in order to construct Markov state models which reproduces well the dominant timescales, we have to have a small maximal projection error of the dominant eigenvectors onto the subspace $\mu D \subset l^2(\frac{1}{\mu})$ considered for the construction of the Markov state model. That is, we would like a small projection error

$$\delta = \max_{i=2, \dots, l_d} \|\mathcal{Q}_{\mu D}^{1/\mu, \perp}(u_i)\|_{\frac{1}{\mu}}, \quad (6.1)$$

where u_2, \dots, u_{l_d} are the dominant, orthonormal left eigenvectors of the transition matrix P corresponding to the dominant eigenvalues $1 > \lambda_2 > \dots > \lambda_{l_d}$, and where $\mathcal{Q}_{\mu D}^{1/\mu, \perp} = Id - \mathcal{Q}_{\mu D}^{1/\mu}$ with $\mathcal{Q}_{\mu D}^{1/\mu}$ denoting the orthogonal projection onto μD , see Theorem 5.5 and 5.6.

Using this result as a starting point for an approach to the identification of core sets may seem awkward – since it assumes the subspace D from the outset. Moreover, in general, we cannot assume to know the dominant eigenvectors u_2, \dots, u_{l_d} . Fortunately, Sarich (2011) provides the following upper bound on $\|\mathcal{Q}_{\mu D}^{1/\mu, \perp}(u)\|_{\frac{1}{\mu}}$ for u being a normalized left eigenvector of P and under the condition that the subspace D is spanned by committor functions on core sets, which nicely translates into an algorithmic strategy for the identification of core sets knowing neither the core sets nor the dominant eigenvectors:

Proposition 6.1. *Let $(X_k)_{k \in \mathbb{N}}$ be an irreducible, reversible Markov chain on state space Z with transition matrix P and stationary distribution μ . Let λ be an eigenvalue of P and let u denote the corresponding normalized, left eigenvector. In addition, let D denote the subspace of \mathbb{R}^l , $l = |Z|$, spanned by the committor functions on the core sets C_1, \dots, C_m , and let $\mathcal{Q}_{\mu D}^{1/\mu, \perp} = Id - \mathcal{Q}_{\mu D}^{1/\mu}$ with $\mathcal{Q}_{\mu D}^{1/\mu}$ being the orthogonal projection onto μD . Then*

$$\|\mathcal{Q}_{\mu D}^{1/\mu, \perp}(u)\|_{\frac{1}{\mu}} \leq p(u) + 2\mu(T)p_{\max}(u) + r(T)(1 - \lambda) \left(\sum_{z \in T} u(z)^2 / \mu(z) \right)^{1/2}, \quad (6.2)$$

where $T = Z \setminus \cup_{j=1}^m C_j$, denotes the transition region, and with

$$r(T) = \sup_{\|w\|_{\frac{1}{\mu}}=1, w=0 \text{ on } C} \left(\frac{1}{\sum_{z \in T} (w(z) - (wP)(z))^2 / \mu(z)} \right)^{1/2} \quad (6.3)$$

$$p(u) = \|e\|_{\frac{1}{\mu}} \quad (6.4)$$

$$p_{\max}(u) = \|e/\mu\|_{\infty} = \max_{z \in Z} |e(z)/\mu(z)| \quad (6.5)$$

$$e(z) = \begin{cases} 0 & \text{if } z \in T, \\ u(z) - \frac{1}{\mu(C_j)} \sum_{x \in C_j} u(x) & \text{if } z \in C_j. \end{cases} \quad (6.6)$$

For the interpretation of this upper bound, notice first that we can decompose it into the terms $r(T)(1-\lambda)\left(\sum_{z \in T} u(z)^2/\mu(z)\right)^{1/2}$ and $p(u)+2\mu(T)p_{\max}(u)$. The first of these terms depends only on whether or not a state $z \in Z$ is an element of a core set (and not on the specific core set it might be in). Moreover, $r(T)$ measures how quickly the process leaves the transition region T . This interpretation of $r(T)$ can be seen by noticing that

$$r(T) = \frac{1}{\|S(Id - \mathcal{P})S\|_{\frac{1}{\mu}}}, \quad (6.7)$$

where S is the orthogonal projection of \mathbb{R}^l onto $Z \setminus \cup_{j=1}^m C_j$ with respect to the standard scalar product. Then, if λ_C is the largest eigenvalue of $\mathcal{P}_C = S\mathcal{P}S$, we get that

$$r(T) = \frac{1}{\|S(Id - \mathcal{P})S\|_{\frac{1}{\mu}}} \quad (6.8a)$$

$$\leq \frac{1}{1 - \lambda_C} \quad (6.8b)$$

$$\leq R_C \quad (6.8c)$$

where the last equality follows from Theorem 4.1 with

$$R_C = \sup_{z \in T} \mathbb{E}_z \tau_C \quad (6.9)$$

denoting the return time to $C = \cup_{j=1}^m C_j$. Now, in the upper bound, $r(T)$ is compared to $(1 - \lambda)$. Thus in order for $r(T)(1 - \lambda)$ to be small, we either need λ close to 1, or $r(T)$ to be small. Thus, the more eigenvalues we want to approximate, the faster the transition region T has to be left.

As for the term $p(u) + 2\mu(T)p_{\max}(u)$, notice that it depends only on the assignment of $x \in C$ to the specific core set C_j . Moreover, both $p(u)$ and $p_{\max}(u)$ measures how constant the eigenvectors are in each core set. Thus, for a small upper bound on the projection error, we want the dominant eigenvectors to be almost constant on the core sets C_1, \dots, C_m .

As we will see in the next section, this interpretation of Proposition 6.1 nicely motivates an algorithmic strategy whose first step consists of identifying the transition region T . Secondly, the remaining states $C = Z \setminus T$ are clustered into core sets.

Example 6.1. We consider again the Markov chain $(X_t^\varepsilon)_{t \in \mathbb{T}}$ of our running example (Example 3.6) with parameters $a = 1, b = 1, \varepsilon = .15, n = 11$. Let μ denote its stationary distribution and let μ^* be the stationary distribution of the evolutionary process with the same parameters except for $\varepsilon^* = .3$. Note that both stationary distributions have the same form in the sense of local minima and maxima, but the stationary distribution μ^* with increased noise intensity is less peaked. The distributions $\mathcal{P}^t(\mu^*)$ of μ^* under the transfer operator \mathcal{P} of $(X_t^{.15})$ converge to μ for $t \rightarrow \infty$ and the comparison of μ^* with $\mathcal{P}^t(\mu^*)$ for some $t \in \mathbb{T}$ will provide the information we need in order to identify the transition region T , see Figure 6.1.

More specifically, we set

$$T^\alpha = \{x \in \Delta_{11}^1 : \mu^*(x) > \mathcal{P}^\alpha(\mu^*)(x)\} \quad (6.10)$$




Figure 6.1: Weights of the stationary distribution μ^* for the stochastic evolutionary game of our running example (Example 3.6) with parameters $a = 1, b = 1, \varepsilon^* = .3, n = 11$ and its propagation $\mathcal{P}^t(\mu^*)$ under the stochastic evolutionary game with parameters $a = 1, b = 1, \varepsilon = .15, n = 11, t = 10/11$.

as this is the set of all population states that are less attractive in the stochastic evolutionary game with decreased noise parameter $\varepsilon = .15$. The other population states, i.e., the set

$$C^\alpha = \{x \in \Delta_{11}^1 : \mu^*(x) < \mathcal{P}^\alpha(\mu^*)(x)\}, \quad (6.11)$$

is more attractive and thus we consider it as the union of core sets.

In our example, C^α is the set $\{0, 1/11, 10/11, 1\}$ for all $\alpha \in \mathbb{T}$. This suggests the partition of C^α into the core sets $C_1 = \{0, 1/11\}$ and $C_2 = \{10/11, 1\}$. In Example 5.6 we discussed the eigenvalue error we make by considering this core set Markov state model; it turned out to be a good approximation and considerably better than the full partition model we considered in this example as well.

The resulting core set Markov state model $(\hat{X}_t^3)_{t \in \mathbb{T}}$ has the transition matrix

$$\hat{P}W^{-1} = \begin{pmatrix} .9993 & .0007 \\ .0007 & .9993 \end{pmatrix}. \quad (6.12)$$

Table 6.1 shows that the eigenvalues corresponding to the Markov chain (X_t^3) and the Markov state model (\hat{X}_t^3) .

For the approximation quality in terms of the propagation of probability distributions see Example 5.5.

(X_t)	(\hat{X}_t)
$\lambda_1 = 1$	$\hat{\lambda}_1 = 1$
$\lambda_2 = .99863$	$\hat{\lambda}_2 = .99861$

Table 6.1: Dominant eigenvalues of (X_t^3) and (\hat{X}_t^3) , respectively.

6.2 Algorithmic Strategy

In this section we consider the *simulation-based* approach to the identification of core sets and the construction of the resulting Markov state model based on Sarich (2011). That is, we will neither for the identification of core sets nor for the construction of the Markov state model assume that we have an expression for the dynamics of the evolutionary process. Instead, we take the setting of the models of evolution with noise introduced in Chapter 3.2.4 in which each Markov chain (X_k^ε) depends on a noise parameter ε that determines the metastability of the stochastic evolutionary game such that the metastability increases for decreasing ε , see Chapter 4.3.

The algorithm proceeds as suggested by the interpretation of Proposition 6.1: Given a realization (x_k) of the Markov chain (X_k^ε) for $k = 1, \dots, K$, we first identify the transition region T ; second, we cluster the remaining states $Z \setminus T$ into core sets; and third, we construct the Markov state model such that its transition matrix $\hat{P}W^{-1}$ corresponds to the matrix representation of the projected transfer operators as described in Chapter 5.3. As for the last of these steps, remember that we gave in Proposition 5.4 an interpretation of the matrices \hat{P} and W in terms of the milestone process $(\tilde{X}_k^\varepsilon)$, which allows their estimation from trajectory data. Moreover, as we want to use the given trajectory x_1, \dots, x_K for the estimation of \hat{P} and W , we only have to tell for these data points whether they belong to T or not – and not for the whole state space, which is advantageous in the case of large state spaces. In the last section, we showed that, in order for a state to be included into the core sets, it should be an element of the most metastable regions and should be relatively attractive if compared with the dominant timescales of interest. Now, given the set T , the second step, that is, the clustering of $Z \setminus T$ into core sets, is usually straightforward as constantness of the eigenvectors on the core sets implies that the core sets are dynamically well separated. In more complex cases, one can apply, for instance, spectral clustering methods like PCCA (see Deuffhard and Weber, 2005).

Algorithm 1: Simulation-Based Core Set Identification and Construction of Markov State Model

Data: Trajectory $(x_k), k = 1, \dots, K$ of (X_k^ε) ; noise parameter $\varepsilon < \varepsilon^*$;
time step $\alpha > 0$; $\nu \in \mathbb{N}$; radius $r > 0$

Result: Core sets C_1, \dots, C_m ; transition matrix $\hat{P}W^{-1}$

- 1 $y_\kappa \leftarrow x_{\nu\kappa}, \kappa = 1, \dots, \tilde{K}$ with $\tilde{K} \leftarrow \lfloor K/\nu \rfloor$;
- 2 $C \leftarrow \emptyset, T \leftarrow \emptyset$;
- 3 **for** $\kappa \leftarrow 1$ **to** \tilde{K} **do**
- 4 | Simulate $y_\kappa^\varepsilon \leftarrow X_\alpha^\varepsilon$ with $X_0^\varepsilon \leftarrow y_\kappa$
- 5 **end**
- 6 **for** $k \leftarrow 1$ **to** K **do**
- 7 | **if** $\#\{y_\kappa^\varepsilon \in B_r(x_k)\} > \#\{y_\kappa \in B_r(x_k)\}$ **then**
- 8 | | $C \leftarrow C \cup \{x_k\}$;
- 9 | **else**
- 10 | | $T \leftarrow T \cup \{x_k\}$;
- 11 | **end**
- 12 **end**
- 13 Cluster C ;
- 14 Estimate \hat{P}, W as explained in Chapter 5.3.3;
- 15 Calculate $\hat{P}W^{-1}$;

The interpretation of the algorithm is as follows. First, notice that since the Markov chain $(X_k^{\varepsilon^*})$ is irreducible and aperiodic, we can estimate $\mu^*(A)$ by

$$\mu^*(A) = \lim_{K \rightarrow \infty} \frac{\#\{y_\kappa \in A\}}{\lfloor K/\nu \rfloor} \quad (6.13)$$

for every measurable set $A \subset Z$. With this we get

$$\mathcal{P}^\alpha(\mu^*)(A) = \lim_{K \rightarrow \infty} \frac{\#\{y_\kappa^\varepsilon \in A\}}{\lfloor K/\nu \rfloor}, \quad (6.14)$$

where \mathcal{P} is the transfer operator of the Markov chain (X_k^ε) with decreased noise intensity $\varepsilon < \varepsilon^*$ and, thus, increased metastability. The appropriate noise parameter ε for the identification depends on the specific model. It should be chosen such that there is a noticeable difference in the metastable dynamic behavior between the Markov chain with parameter ε^* and $\varepsilon < \varepsilon^*$. In our running example, for instance, we might choose $\varepsilon = \varepsilon^*/2$.

According to Equations (6.13) and (6.14), we can thus approximately say that $\mathcal{P}^\alpha(\mu^*)(B_r(x_k)) > \mu^*(B_r(x_k))$ if the number of data points y_κ^ε in $B_r(x_k)$ is larger than the number of points y_κ in $B_r(x_k)$, where $B_r(x_k)$ denotes a ball of small radius $r > 0$ around the data point x_k . This results in assigning a point x_k to the core sets, i.e., $x_k \in C$, if

$$\#\{y_\kappa^\varepsilon \in B_r(x_k)\} > \#\{y_\kappa \in B_r(x_k)\}. \quad (6.15)$$

Having determined the set C , we cluster it into the core sets C_1, \dots, C_m , estimate \hat{P}, W and calculate $\hat{P}W^{-1}$ as discussed before. Notice that the parameter α is associated with the timescale we want to approximate, i.e., the larger α , the

more attractive the states need to be in order to be identified as an element of C . Furthermore, in the case of a small state space, as in our running example, we can simplify the above by considering r so small that the balls around the data points x_k consist only of x_k itself. Thus, in order to determine whether $x \in C$, we compare $\#\{y_{\kappa}^{\varepsilon} = x_k\}$ with $\#\{y_{\kappa} = x_k\}$. As should be clear, the computational effort of the algorithm increases with increasing α and increasing population size n as well as decreasing ν . However, for large state spaces, which ensue rather quickly in population games from increasing population size and increasing size of the strategy set¹, considering the balls around the data points is vital since the weight of a single state in the stationary distribution μ^* as well as in the propagated stationary distribution $\mathcal{P}^{\alpha}(\mu^*)$ might get so small that it is difficult to statistically estimate this difference from simulated trajectory data.

Remark 6.1. The identification of the core set region C by the algorithm preserves stochastic stability in the sense that if a state x is stochastically stable it will be identified as an element of C . The reason is that the weight of a stochastically stable state cannot decrease under \mathcal{P}^{α} compared to its weight in μ^* , where \mathcal{P} denotes the transfer operator of (X_k^{ε}) and μ^* denotes the stationary distribution of $(X_k^{\varepsilon^*})$, because a decreasing weight would mean that the states becomes less attractive with decreasing ε which contradicts the notion of stochastic stability. Note, however, that not every core set identified needs to contain a stochastically stable state as we will demonstrate in Example 6.3 below.

Example 6.2. Let us consider again the model of evolution with noise $(X_t^{\varepsilon^*})_{t \in \mathbb{T}}$ from our running example (Example 3.6). We simulated the Markov chain (X_t^{ε}) with $a = 1, b = 1, \varepsilon^* = .3, n = 11$ to get a trajectory of population states (x_t) for $t = 1/11, \dots, 50000$ and thus of length $5.5 \cdot 10^5$. We applied the algorithm for the identification of core sets and the construction of the Markov state model with parameters $\varepsilon = \varepsilon^*/2, \alpha = 10/11, \nu = 10$ to this trajectory of data points. It resulted in $C = \{0, 1/11, 10/11, 1\}$, which suggests straightforwardly the core sets $C_1 = \{0, 1/11\}$ and $C_2 = \{10/11, 1\}$. Remember that we singled out specifically these core sets in Examples 5.5 and 5.6 as the core sets with the least projection error and thus best approximation quality in terms of the propagation of probability distributions as well as in terms of the eigenvalue error. See Figure 6.2 for the data points of the trajectory with $t = 0, 1/11, \dots, 1000$ and the respective data points which are identified as elements of core sets.

Based on the core sets and the trajectory of data points, the estimated transition matrix $\hat{P}^*W^{*, -1}$ is given by

$$\hat{P}^*W^{*, -1} = \begin{pmatrix} .9993 & .0007 \\ .0007 & .9993 \end{pmatrix}, \quad (6.16)$$

¹Remember that the size of the set Δ_n^{m-1} of population states in a population game with population size n and size m of the strategy set is given by $\binom{n+m-1}{m-1} = \binom{n+m-1}{n}$ (Scheffé, 1958). This means that the state space Δ_n^{m-1} grows polynomially in m with exponent n for $m \rightarrow \infty$ and fixed n , and polynomially in n with exponent m for $n \rightarrow \infty$ and fixed m . More formally, $\Delta_n^{m-1} = \Theta(m^n)$ for $m \rightarrow \infty$ and fixed n ; and $\Delta_n^{m-1} = \Theta(n^m)$ for $n \rightarrow \infty$ and m fixed.

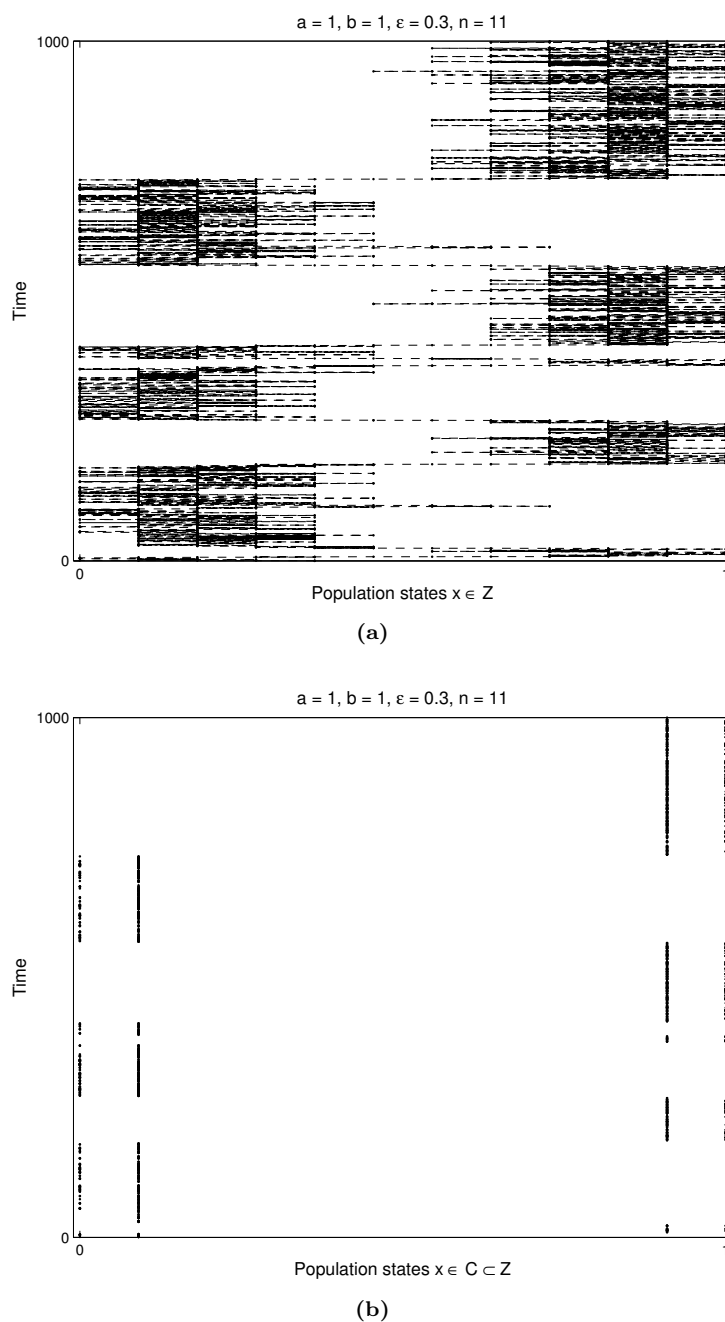


Figure 6.2: (a) Piece of the trajectory (x_t) and (b) corresponding data points which are identified as belonging to a core set by the identification algorithm.

with second eigenvalue $\hat{\lambda}_2^* = .99863$, which is thus a very good approximation of the transition matrix $\hat{P}W^{-1}$ and its second eigenvalue, compare Example 6.1.

Example 6.3. We consider again the model of evolution with noise (X_t^ε) of our running example (Example 3.6), this time with parameters $a = 1.2, b = 1, \varepsilon^* = .3, n = 11$. From Example 3.7 we know that in this case, the only stochastically stable state in the population state space $Z = \{0, \frac{1}{n}, \dots, 1\}$ is 1. We simulated this Markov chain to get a trajectory of population states $x_t, t = 1, \dots, 50000$ and applied the identification algorithm with parameters $\varepsilon = \varepsilon^*/2, \alpha = 10/11, \nu = 10$. It resulted in $C = \{0, 1/11, 10/11, 1\}$ which suggests again the clustering $C_1 = \{0, 1/11\}$ and $C_2 = \{10/11, 1\}$, see Figure 6.3. As the core set C_1 in this example demonstrates, the core sets identified by the algorithm need not contain stochastically stable states. Moreover, for small ε^* and large enough α , the algorithm recaptures the stochastically stable states as is to be expected. For instance, in our running example with parameters $a = 1.2, b = 1, n = 11$, this is the case for $\varepsilon^* = .1, \varepsilon = \varepsilon^*/2$ and $\alpha \approx 10^5$ where $C = \{1\}$ while for $\alpha = 10$ the algorithm results in $C = \{0, 1\}$ (based on a simulated trajectory of states $(x_t), t = 1/11, \dots, 10^5$ and $\nu = 1$).

6.3 Extended Example

So far, we have only considered the simple example of the currency game under the best response with mutations revision protocol. Another example of reversible stochastic evolutionary games are finite-population potential games with clever agents under a logit choice revision protocol. We start with a few definitions:

Definition 6.1. Let $(n, S, \Delta_n^{m-1}, F)$ be a finite-population game. It is a *potential game* if there is a function $f : \Delta_n^{m-1} \rightarrow \mathbb{R}$, called the *potential function*, such that

$$F_j(x + \frac{1}{n}(e_j - e_i)) - F_i(x) = f(x + \frac{1}{n}(e_j - e_i)) - f(x) \quad (6.17)$$

for all $i, j \in S$ and $x \in \Delta_n^{m-1}$ with $x_i > 0$.

Thus, in a potential game, the change in payoff for an agent that changes from strategy i to strategy j equals the resulting change in potential. Potential games are an important class of games in transportation science, economics and population genetics. Examples of finite-population potential games are:

- Two-strategy games: Let $F : \Delta_n^{m-1} \rightarrow \mathbb{R}^2$ be any payoff function for a finite population game with strategy set $S = \{0, 1\}$. F admits the following potential function $f : \Delta_n^{m-1} \rightarrow \mathbb{R}$:

$$f(x) = \sum_{j=1}^{nx_1} \left[F_1(x + \frac{j-1}{n}(e_0 - e_1)) - F_0(x + \frac{j}{n}(e_0 - e_1)) \right] \quad (6.18)$$

where $x = (x_0, x_1)$.

- Matching with self-matching in common interest games: Let $A \in \mathbb{R}^{m \times m}$ be the payoff matrix of a common interest two-player game in strategic form, i.e., $A = A^T$. Let $F : \Delta_n^{m-1} \rightarrow \mathbb{R}^m$ denote the associated payoff

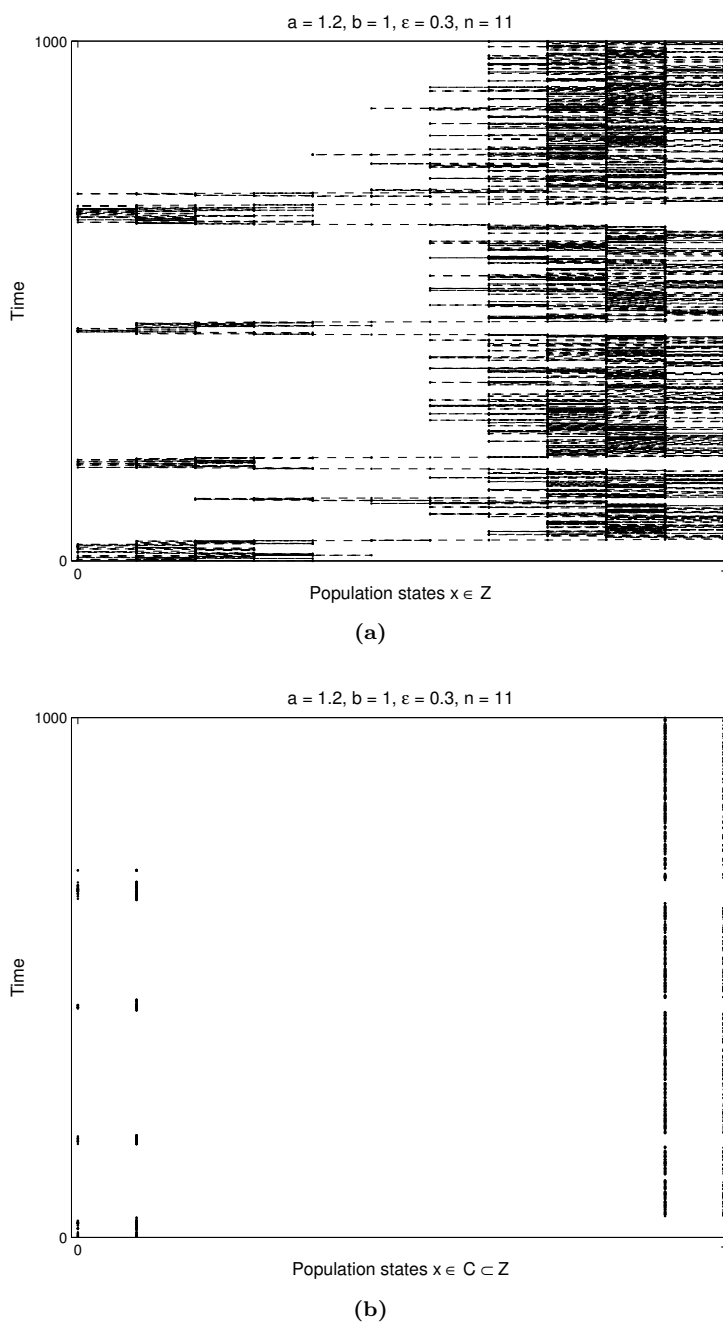


Figure 6.3: (a) Piece of the trajectory (x_t) and (b) corresponding data points which are identified as belonging to a core set by the identification algorithm.

function in the finite-population game under matching with self-matching, i.e., $F(x) = Ax$. Then F admits the following potential function f :

$\Delta_n^{m-1} \rightarrow \mathbb{R}$:

$$f(x) = \frac{1}{2} \left(nx^T Ax + \sum_{k \in S} a_{kk} x_k \right). \quad (6.19)$$

Example 6.4. In the examples that follow, we are going to consider the potential game that results from matching with self-matching in the strategic game with strategy set $S = \{1, 2, 3\}$ and payoff matrix

$$A = \begin{pmatrix} a & 0 & 0 \\ 0 & b & 0 \\ 0 & 0 & c \end{pmatrix}. \quad (6.20)$$

For $n = 5$, its potential function $f : \Delta_5^2 \rightarrow \mathbb{R}$ is given as in Eq. (6.19). In Figure 6.4 we depicted f as a function of the first two components (x_1, x_2) of the population state $(x_1, x_2, x_3) \in \Delta_5^2$; this is possible since the third component x_3 (and thus also the value of f at (x_1, x_2, x_3)) is uniquely determined by the condition $\sum_{j \in S} x_j = 1$.

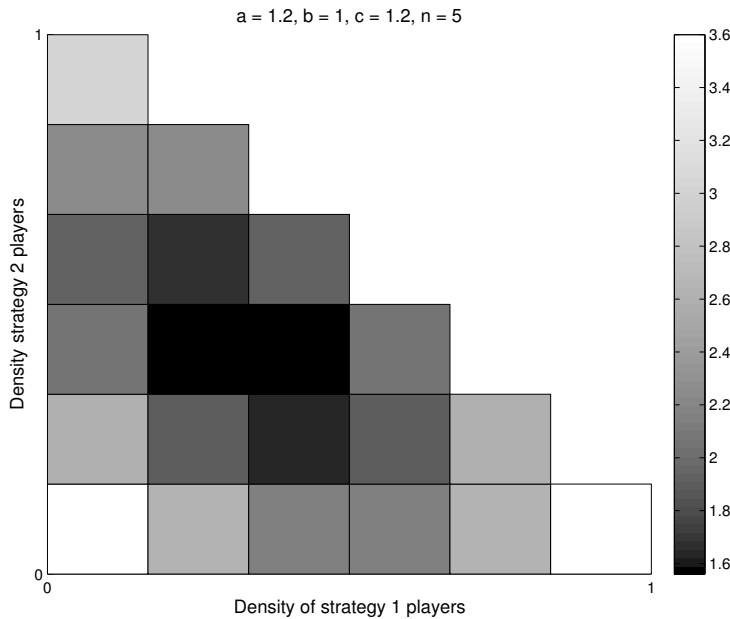


Figure 6.4: Potential function of the finite-population potential game that results from matching with self-matching in a common interest game with payoff matrix A and parameters $a = 1.2, b = 1, c = 1.2, n = 5$.

In general, the revision of an agent's strategy involves the comparison of the payoffs that the agent receives from using each strategy in the population game. In a finite-population, a switching agent also changes the population state and

thus the payoffs to be received. More specifically, if an i -player switches to strategy j , the population state changes from x and payoffs $F(x)$ to the state $x + \frac{1}{n}(e_j - e_i)$ and to the payoffs $F(x + \frac{1}{n}(e_j - e_i))$, where the e_j 's denote the standard unit vectors. Thus, it seems reasonable to assume that an i -agent compares his current payoffs $F_i(x)$ to the payoffs $F_j(x + \frac{1}{n}(e_j - e_i))$ rather than to $F_j(x)$ when considering switching from strategy i to strategy $j \in S$. Agents that revise in this way are called *clever*, otherwise *simple*.

Sandholm (2010, p. 430 f.) proves the following proposition:

Proposition 6.2. *Let $F : \Delta_n^{m-1} \rightarrow \mathbb{R}^m$ be the payoff function of a finite-population potential game with potential function $f : \Delta_n^{m-1} \rightarrow \mathbb{R}$. Suppose that agents are clever and use the logit choice revision protocol with noise level σ for strategy updating, i.e., $\rho : \Delta_n^{m-1} \rightarrow \mathbb{R}_{\geq 0}^{m \times m}$ is defined according to Eq. 3.11. The resulting discrete-time Markov chain $(X_t)_{t \in \mathbb{T}}$ is reversible with stationary distribution*

$$\mu(x) = \frac{1}{C} \frac{n!}{\prod_{k \in S} (nx_k)!} \exp(\sigma f(x)) \quad (6.21)$$

for $x \in \Delta_n^{m-1}$, where C is determined by the constraint $\sum_{x \in \Delta_n^{m-1}} \mu(x) = 1$.

Thus, the stationary distribution of the resulting stochastic evolutionary process at state x is proportional to the product of a multinomial coefficient and the exponential of the potential function weighted with the noise level. The multinomial coefficient counts the number of ways the players can be assigned to strategies such that the resulting proportion of strategies is x , i.e., the likelihood that state x would result from a random assignment of agents to strategies.

Example 6.5. Consider the potential game from Example 6.4. From Proposition 6.2, we know that, assuming clever agents and the logit choice revision protocol, the resulting stochastic evolutionary game $(X^\sigma(t))_{t \in \mathbb{T}}$ with transition matrix P is reversible with stationary distribution μ as given by Eq. 6.21, see Figure 6.5.

Note that for $\sigma \rightarrow \infty$, the stochastically stable states correspond to those population states that maximize the potential function f (see, e.g., Sandholm, 2010, Theorem 12.2.1; more generally, see Chapter 12 for how the finite-population potential games under the logit choice revision protocol fit into the context of models of evolution with noise). Thus, in our example, the stochastically stable population states correspond to those population states in which players uniformly play strategy $i = 1, 2$, or 3 , and for which $i \in \operatorname{argmax}_{j=1,2,3} A(j, j)$.

We applied the identification algorithm for two different values of the timescale parameter α and $\nu = 1$ to a simulated trajectory (x_t) of length $5 \cdot 10^6$ of the stochastic evolutionary game given in terms of the potential game that results from matching with self-matching in the strategic game with payoff matrix A , parameters $a = 1.2, b = 1, c = 1.2, n = 5$, and assuming clever agents that behave according to the logit choice revision protocol with noise intensity $\sigma^* = 3.5$. See Figure 6.6 for a graphical display of the initial piece of the trajectory; the noise intensity σ corresponding to a higher metastability was set to $\sigma = \sigma^* + 1 = 4.5$.

We obtained the following results:

- For $\alpha = 20$, the algorithm identified $C = \{(0, 0, 1), (1, 0, 0)\}$ as the core set region, which suggests the core sets $C_1 = \{(0, 0, 1)\}$ and $C_2 = \{(1, 0, 0)\}$.




Figure 6.5: Stationary distribution of the stochastic evolutionary game given in terms of the finite-population potential game that results from matching with self-matching in a common interest game with payoff matrix A and parameters $a = 1.2, b = 1, c = 1.2, n = 5$ and assuming clever agents and the logit choice revision protocol with noise intensity $\sigma = 3.5$.

The estimated transition matrix $\hat{P}^*W^{*, -1}$ of the associated core set Markov state model $(\hat{X}_t^{3.5})$ is given by

$$\hat{P}^*W^{*, -1} = \begin{pmatrix} .9931 & .0069 \\ .0067 & .9933 \end{pmatrix} \quad (6.22)$$

with eigenvalues $\hat{\lambda}_i^*$ compared to the eigenvalues $\lambda_i, i = 1, 2$ of the transition matrix P given by

$(X_t^{3.5})$	$(\hat{X}_t^{3.5})$
$\lambda_1 = 1$	$\hat{\lambda}_1^* = 1$
$\lambda_2 = .98630$	$\hat{\lambda}_2^* = .98634$

See Figure 6.7 for the data points of the simulated trajectory (x_t) identified as core sets.

- For $\alpha = 2$. the core set region identified by the algorithm is $C = \{(0, 0, 1), (0, 1, 0), (1, 0, 0)\}$, which suggests the core sets $C_1 = \{(0, 0, 1)\}$, $C_2 = \{(0, 1, 0)\}$, and $C_3 = \{(1, 0, 0)\}$; that is, each core set corresponds to a population state with uniform strategy choice among the players. The estimated transition matrix of the corresponding core set Markov state




Figure 6.6: Initial piece of the simulated trajectory $(x_t)_{t=0,\dots,1000}$ of the stochastic evolutionary game given in terms of the finite-population potential game that results from matching with self-matching in a common interest game with payoff matrix A and parameters $a = 1.2, b = 1, c = 1.2, n = 5$ and assuming clever agents and the logit choice revision protocol with noise intensity $\sigma = 3.5$.

model $(\hat{X}_t^{3.5})$ is given by

$$\hat{P}^* W^{*, -1} = \begin{pmatrix} .9915 & .0030 & .0055 \\ .0156 & .9686 & .0158 \\ .0053 & .0030 & .9916 \end{pmatrix}. \quad (6.23)$$

Its eigenvalues $\hat{\lambda}_i^*$ compared to the eigenvalues λ_i of P are

(X_t)	(\hat{X}_t)
$\lambda_1 = 1$	$\hat{\lambda}_1^* = 1$
$\lambda_2 = .98630$	$\hat{\lambda}_2^* = .986160$
$\lambda_3 = .966355$	$\hat{\lambda}_3^* = .965564$

See Figure 6.8 for the data points of the simulated trajectory (x_t) identified as core sets.

Note that in both cases, i.e., for $\alpha = 10$ as well as $\alpha = 100$, the results of the identification algorithm correspond to the results obtained by applying the analytical approach outlined in Example 6.1. The example nicely demonstrates the

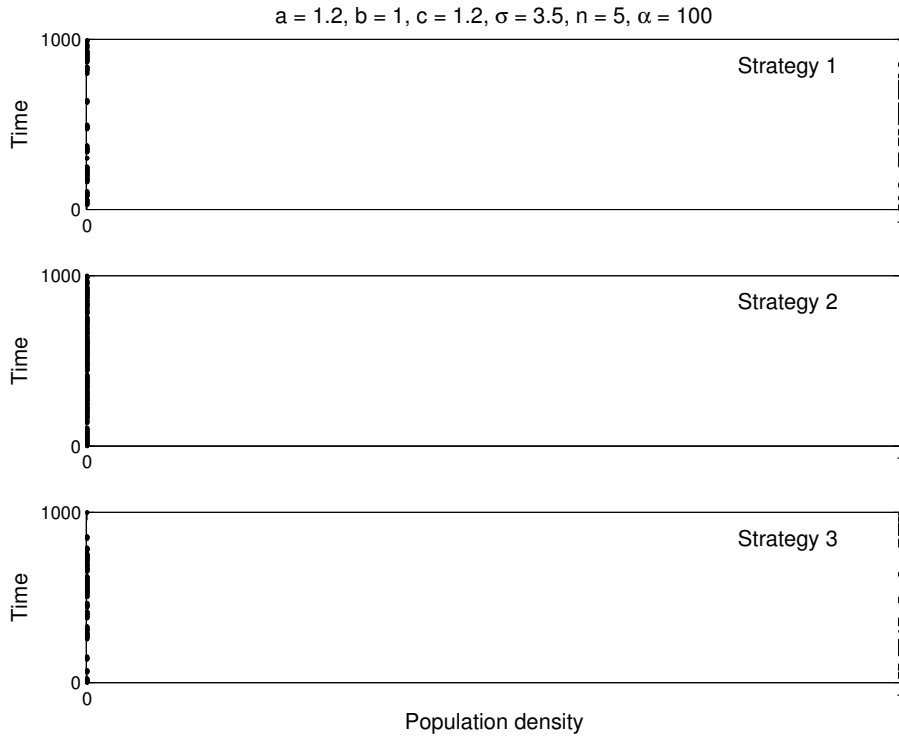


Figure 6.7: Data points corresponding to the identified core set region C for $\alpha = 20$ of the initial piece of the simulated trajectory $(x_t)_{t=0, \dots, 1000}$ of the stochastic evolutionary game given in terms of the finite-population potential game that results from matching with self-matching in a common interest game with payoff matrix A and parameters $a = 1.2, b = 1, c = 1.2, n = 5$ and assuming clever agents and the logit choice revision protocol with noise intensity $\sigma^* = 3.5$.

role of the timescale parameter α . For the smaller α , the algorithm identified another core set which corresponds to a set of population states that became attractive on this shorter timescale.

6.4 Discussion and Outlook

The identification algorithm presented in Section 6.2 is appealing with regard to agent-based modeling because it does not rely on analytical expressions for the dynamics as it is a simulation-based approach. However, the algorithm and, more generally, the approach presented in the second part of this thesis relied heavily on the assumption that the considered Markov chain is reversible.

We gave two examples of classes of stochastic evolutionary games with reversible dynamics: the class of population games with two strategies under full support revision protocols (see Section 3.2.3) and the class of finite-population potential games with clever agents under a logit choice revision protocol. Sand-

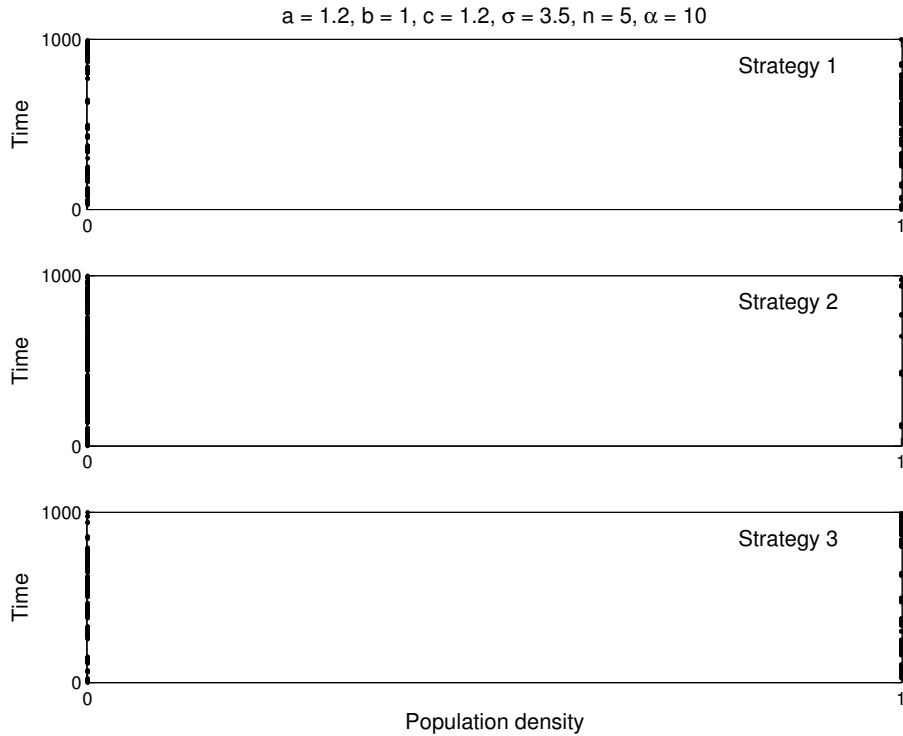


Figure 6.8: Data points corresponding to the identified core set region C for $\alpha = 2$ of the initial piece of the simulated trajectory $(x_t)_{t=0, \dots, 1000}$ of the stochastic evolutionary game given in terms of the finite-population potential game that results from matching with self-matching in a common interest game with payoff matrix A and parameters $a = 1.2, b = 1, c = 1.2, n = 5$ and assuming clever agents and the logit choice revision protocol with noise intensity $\sigma^* = 3.5$.

holm (2010, Chapter 11.5.3) provides more general conditions on revision protocols under which finite-population potential games with clever agents result in reversible dynamics.

In general, it will be difficult to say whether it is reasonable to assume that an agent-based model results in a reversible Markov chain. One reason for this difficulty is that, if we estimate the transition matrix from simulated trajectory data, it does not need to fulfill the detailed balance equation, even if the underlying Markov chain is reversible (Noé, 2008; Prinz et al., 2011). In the context of molecular dynamics, however, it was possible to derive approximative models that can be proven to be reversible although the original model is not. An example is the diffusion model, which represents an approximation to the Langevin model in the limit of high friction (see, e.g., Schütte and Sarich, 2013, Chapter 2 and references therein). As a future research question, it seems worthwhile to explore whether similar results can be obtained for agent-based models; that is, whether there are approximations of certain agent-based models that can be shown to be reversible.

Beyond that, we would like an approach that applies also to non-reversible

Markov dynamics. Notice that it is not difficult to derive a construction of a matrix representation of the core set Markov state models for given core sets in the case of non-reversible Markov chains (see, e.g., Djurdjevac, 2012; Djurdjevac et al., 2010). However, we neither have results with respect to their approximation quality nor an approach to the identification of metastable subsets for non-reversible Markov chains. One fundamental problem is that the eigenvalues and eigenvectors of the transfer operators corresponding to non-reversible Markov chains need not be real anymore. In this case, the interpretation of the spectral information for the identification of metastable sets is unclear. Up to now, there are few approaches that apply also to non-reversible Markov chains (Eckhoff, 2002; Gaudillière and Landim, 2011; Horenko, 2011; Sarich and Schütte, 2014).

The identification of Markov state models for general agent-based models is therefore an open problem and will be a topic of future research. The examples of this chapter nevertheless constitute a proof of concept for the usefulness and applicability of the Markov state modeling approach for general agent-based models.

Summary

This thesis is concerned with the formal description (Part I) and metastability analysis (Part II) of agent-based evolutionary models. Part I discussed the need for specifications as an intermediate layer between implementations and narrative descriptions of computer-based models – an intermediate layer which is necessary but often not provided in the context of agent-based modeling. We presented the basic structure of a functional framework for the specification of agent-based models of exchange. Introducing such a framework as a basis for model description, analysis and reimplementations is novel for computational economics. The focus of the presentation was to expose the relationship of agent-based models of exchange with existing economic theory, in particular, with evolutionary game theory and general equilibrium theory. We explicitly represented these models as discrete-time Markov processes and showed that they differ from stochastic evolutionary games only in that their fitness function is in general not deterministic. We applied the framework to demonstrate how to formulate the research question behind agent-based models of exchange precisely, which establishes a direct relationship with general equilibrium theory. Moreover, we discussed how it constitutes a starting point for formal model analysis and further numerical investigations by allowing us to deduce specific formal model properties from the specifications. This part of the thesis made clear that although specifications might seem obvious and at the same time tedious, they are an indispensable basis for a rigorous understanding and analysis of agent-based models. In this way, specifications are fundamental for agent-based modeling to get out of its infancy.

Part II presented a novel approach to the analysis of stochastic evolutionary games, which are simple agent-based models. We motivated this approach by the observation that stochastic evolutionary games often exhibit metastable dynamics but the existing approaches to their analysis, such as Nash equilibrium solution concepts, deterministic approximation, and stochastic stability analysis, are not able to capture this behavior. We derived the aggregated strategy updating process for general stochastic evolutionary games as a discrete-time Markov chain on a finite state space. We presented two characterizations of metastability, one in terms of hitting times and one in terms of transition probabilities. Both characterizations build on the representation as a Markov chain and relate to spectral properties of the transfer operator associated with the Markov chain. We investigated both of these characterizations for stochastic evolutionary games. In particular, we showed that every decomposition of state space into limit sets of the unperturbed Markov chain of the stochastic evolu-

tionary game is metastable according to both characterizations.

We furthermore considered the core set Markov state modeling approach to the construction of Markov models on a reduced state space for stochastic evolutionary games. Given some disjoint subsets of the set of population states, which are called core sets, the underlying idea of the construction is the orthogonal projection, and thus best approximation, of the transfer operators of the original Markov chain onto the subspace spanned by the committor functions on core sets or onto this subspace weighted with the stationary distribution of the original Markov chain, depending on the considered transfer operator. We derived the matrix representation of both projected transfer operators and elaborated on the relationship between the projected transfer operators as well as on the relationship of the original Markov chain with the core set Markov state model, which is defined as the Markov chain on the state space $\{1, \dots, m\}$ with transition matrix given by the matrix representation of the transfer operators, where m is the number of core sets considered. This construction of core set Markov state models is appealing, because the matrix representation of the projected transfer operators can be estimated from simulated trajectory data. Moreover, the construction preserves stochastic stability in the sense that, if a core set contains a stochastically stable state of the original Markov chain, the corresponding state in the core set Markov state model is also stochastically stable. We showed how to use the results on the approximation error we make with respect to the propagation of probability distributions as well as with respect to the approximation of eigenvalues for an algorithmic strategy to the simulation-based identification of core sets in such a way that the approximation error is small and the resulting core set Markov state model reproduces well the dominant timescales and in this sense the metastability of the original Markov chain. Lastly, we discussed the fact that the presented simulation-based approach to the identification of core sets and the construction of core set Markov state models is appealing in the context of agent-based modeling because it does not rely on analytical expressions for the dynamics. However, the approach is built on the assumption that the original Markov chain is reversible, an assumption which is usually not fulfilled for general agent-based models. We gave two important examples of classes of stochastic evolutionary games with reversible dynamics and gave an overview of the state-of-the-art of Markov models on reduced state space for Markov chains with non-reversible dynamics.

Taken as a whole, the second part of this thesis presents not only a novel approach to the analysis of stochastic evolutionary games but constitutes a proof of concept for the applicability of the core set Markov state modeling approach to agent-based models. In this respect, it is a foundation for future research into the analysis of agent-based models.

Translation between Notations

Table 2 provides a translation of selective expressions from functional notation to the vector-based notation as used in mathematical economics (see, e.g., Varian, 1992).

Table 2: Translation table. The page number in the first column refers to the page of the first occurrence of the expression.

Expression		Functional notation	Vector-Based Notation
Name	Page		
Stocks	24	$q : G \rightarrow \mathbb{R}_{\geq 0}$	$q \in \mathbb{R}_{\geq 0}^n$, where $n = G $.
Allocations	24	$x : A \rightarrow Q$	$x = (q_1, \dots, q_k) \in \mathbb{R}_{\geq 0}^{n \times k}$, where $k = A $.
Representation of functions of more than one variable	24	Example $x : A \rightarrow G \rightarrow \mathbb{R}_{\geq 0}$	$x : A \times G \rightarrow \mathbb{R}_{\geq 0}$.
Function application	24	Example $x a \in Q$ for x function of type $A \rightarrow Q$ Example $x a g \in \mathbb{R}_{\geq 0}$ for function x of type $x : A \rightarrow G \rightarrow \mathbb{R}_{\geq 0}$	$x(a) \in Q$. $x(a)(g)$ or $x(a, g)$ depending on whether x is represented in curried form $x : A \rightarrow G \rightarrow \mathbb{R}_{\geq 0}$ or uncurried form $x : A \times G \rightarrow \mathbb{R}_{\geq 0}$, respectively.
Multiplication	24	$r_1 * r_2$	$r_1 r_2$
Stocks of an agent in an allocation	24	$x a \in Q$	$x_i \in \mathbb{R}_{\geq 0}^n$, where a is the i th agent.
Amount of a good of an agent in an allocation	24	$x a g \in \mathbb{R}_{\geq 0}$	x_{ij} where agent a is the i th agent and good g is the j th good.
Lists of type X	25	<i>List</i> X	Set of all sequences (x_0, \dots, x_l) in X of finite length
Empty list	25	\square	$()$

Table 2: Translation table (continued).

Expression		Functional notation	Vector-Based Notation
Name	Page		
k th element of a list	25	$xs!!k$ for list xs	x_k for list (x_1, \dots, x_k)
List with a certain first element	25	$(x : xs)$	(x, x_1, \dots, x_l)
Membership function elem	25	Example elem $x xs$	There is no standard notation for such a membership function in the vector-based notation. We could define one similarly to the definition given on page 25 in the functional notation.
fold	25	Example $x' = \text{fold } f \ x \ ys$, where $f : X \rightarrow Y \rightarrow X, x \in X, ys \in \text{List } Y$	There is no standard function for fold in the vector-based notation. What is often done, instead, is to define for a function $f : X \times Y \rightarrow X, x \in X$ and a finite sequence (y_0, \dots, y_l) of elements in y the resulting x' recursively as follows: $x' = x_{l+1} = f(x_l, y_l)$, where $x_0 = x$ and $x_i = f(x_{i-1}, y_{i-1})$ for $i = 1, \dots, l$.
Realizations of a stochastic process $(X_t)_{t \in \mathbb{N}}$ with sample space Θ and state space Z	33	$z \theta : \mathbb{N} \rightarrow Z$ defined by $z \theta t = X_t \theta$ for $t \in \mathbb{N}$	$z(\theta) \in Z^{\mathbb{N}}$ defined by $z(\theta)(i) = X_i(\theta)$.
Expected utility	39	$\mathbb{E} \ u.rvC$	$\mathbb{E}(u \circ rvC)$.

Table 2: Translation table (continued).

Expression		Functional notation	Vector-Based Notation
Name	Page		
Outcome of a trading round	39	$x' = \text{fold } ebt \ x \ ts$, where $ebt : (A \rightarrow Q) \rightarrow A \times A \rightarrow (A \rightarrow Q)$, $x : A \rightarrow Q$, and $ts \in List \ (A \times A)$	Let $K = \{1, \dots, k\}$. Define x' for the function $ebt : \mathbb{R}_{\geq 0}^{n \times k} \times (K \times K) \rightarrow \mathbb{R}_{\geq 0}^{n \times k}$, $x \in \mathbb{R}_{\geq 0}^{n \times k}$ and a finite sequence of pairs $((i_0, j_0), (i_1, j_1), \dots, (i_l, j_l))$ in K , which denotes a finite sequence of agent pairs, recursively by $x' = x_{l+1} = ebt(x_l, (i_l, j_l))$, where $x_0 = x$ and $x_m = ebt(x_{m-1}, (i_{m-1}, j_{m-1}))$ for $m = 1, \dots, l$.
Average utility of n samples	40	$\frac{1}{n} \sum_{j=1}^n (u_{ext} \ a) . rvC_j$	$\frac{1}{n} \sum_{j=1}^n u_{ext, i} \circ rvC_j$

Bibliography

- Aho, A. V., Ullman, J. D., and Hopcroft, J. E. (1982). *Data Structures and Algorithms*. Addison-Wesley, Reading, MA.
- Arrow, K. J. and Debreu, G. (1954). Existence of an equilibrium for a competitive economy. *Econometrica*, 22:265–90.
- Arrow, K. J. and Hahn, F. H. (1971). *General Competitive Analysis*. North-Holland.
- Aubin, J.-P. and Cellina, A. (1984). *Differential Inclusions*. Springer, Berlin.
- Bala, V. and Majumdar, M. (1992). Chaotic tatonnement. *Economic Theory*, 2:437–445.
- Ben-Shoham, A., Serrano, R., and Volij, O. (2004). The evolution of exchange. *Journal of Economic Theory*, 114(2):310 – 328.
- Benaïm, M. and Weibull, J. (2003). Deterministic approximation of stochastic evolution in games. *Econometrica*, 71:873–903.
- Berman, A. and Plemmons, R. (1979). *Nonnegative Matrices in the Mathematical Sciences*. Academic Press, New York. Reprinted by SIAM, Philadelphia, 1994.
- Bernergård, A. and Wärneryd, K. (2011). Finite-population “mass-action” and evolutionary stability. CESifo Working Paper Series No. 3378. Available at SSRN: <http://ssrn.com/abstract=1788828>.
- Binmore, K. (1995). Foreword. pp. ix-xi. In Weibull, J. W. *Evolutionary Game Theory*. Cambridge, US: The MIT Press.
- Binmore, K. and Samuelson, L. (1997). Muddling through: Noisy equilibrium selection. *Journal of Economic Theory*, 74:235–265.
- Binmore, K., Samuelson, L., and Young, H. P. (2003). Equilibrium selection in bargaining models. *Games and Economic Behavior*, 45:296–328.
- Bird, R. (1998). *Introduction to Functional Programming using Haskell*. International Series in Computer Science, Prentice Hall.
- Blume, L. E. (1993). The statistical mechanics of strategic interaction. *Games and Economic Behavior*, 5:387–424.

- Botta, N., Mandel, A., Hofmann, M., Schupp, S., and Ionescu, C. (2013). Mathematical specification of an agent-based model of exchange. In Lange, C., Rowat, C., and Kerber, M., editors, *Do-Form: Enabling Domain Experts to use Formalised Reasoning (Proceedings to the Do-Form symposium at the AISB Annual Convention in April 2013)*, Exeter, UK. Society for the Study of Artificial Intelligence and Simulation of Behaviour (AISB). ISBN 978-1-908187-32-1.
- Botta, N., Mandel, A., Ionescu, C., Hofmann, M., Lincke, D., Schupp, S., and Jaeger, C. (2011). A functional framework for agent-based models of exchange. *Applied Mathematics and Computation*, 218(8):4025–4040.
- Bovier, A. (2006). Metastability: A potential theoretic approach. In *Proceedings of the International Congress of Mathematicians, Madrid, Spain, 2006*. European Mathematical Society.
- Bovier, A. (2009). Metastability. In Kotecký, R., editor, *Methods of Contemporary Mathematical Statistical Physics*, Lecture Notes in Mathematics 1970. Springer-Verlag Berlin Heidelberg.
- Bovier, A., Eckhoff, M., Gayrard, V., and Klein, M. (2001). Metastability in stochastic dynamics of disordered mean field models. *Probability Theory and Related Fields*, 119:99–161.
- Bovier, A., Eckhoff, M., Gayrard, V., and Klein, M. (2002). Metastability and low lying spectra in reversible Markov chains. *Communications in Mathematical Physics*, 228:219–255.
- Brémaud, P. (1999). *Markov Chains: Gibbs Fields, Monte Carlo Simulation, and Queues*. Springer.
- Brenner, T. (2006). Agent learning representation: Advice on modelling economic learning. In Tesfatsion, L. and Judd, K. L., editors, *Handbook of Computational Economics: Agent-Based Computational Economics*, volume 2, chapter 18, pages 895–947. North-Holland.
- Davies, E. B. (1982a). Metastable states of symmetric Markov semigroups i. *Proceedings of the London Mathematical Society*, 45(3):133–150.
- Davies, E. B. (1982b). Metastable states of symmetric Markov semigroups ii. *Journal of the London Mathematical Society*, 26(2):541–556.
- Davies, E. B. (1983). Spectral properties of metastable Markov semigroups. *Journal of Functional Analysis*, 52:315–329.
- Dawid, H. and Fagiolo, G. (2008). Agent-based models for economic policy design: Introduction to the special issue. *Journal of Economic Behavior & Organization*, 67:351–354.
- Debreu, G. (1974). Excess demand functions. *Journal of Mathematical Economics*, 1:15–23.
- Deffuant, G. and Gilbert, N., editors (2011). *Viability and Resilience of Complex Systems: Concepts, Methods and Case Studies from Ecology and Society (Understanding Complex Systems)*. Springer.

- Deissenberg, C., Hoog, S. V. D., and Dawid, H. (2008). Eurace: A massively parallel agent-based model of the European economy. Working Paper Nr. halshs-00339756_v1, HAL.
- Dellnitz, M. and Junge, O. (1999). On the approximation of complicated dynamical behavior. *SIAM Journal on Numerical Analysis*, 36(2):491515.
- Deuffhard, P., Huisinga, W., Fischer, A., and Schütte, C. (2000). Identification of almost invariant aggregates in reversible nearly uncoupled Markov chains. *Linear Algebra and its Applications*, 315:39–59.
- Deuffhard, P. and Weber, M. (2005). Robust perron cluster analysis in conformation dynamics. *Linear Algebra and its Applications*, 398:161184.
- Diaconis, P. and Saloff-Coste, L. (1996). Logarithmic Sobolev inequalities for finite Markov chains. *The Annals of Applied Probability*, 6(3):695–750.
- Djurdjevac, N. (2012). *Methods for analyzing complex networks using random walker approaches*. PhD thesis, Fachbereich Mathematikund Informatik, FU Berlin.
- Djurdjevac, N., Sarich, M., and Schütte, C. (2010). On Markov state models for metastable processes. In *Proceedings of the International Congress of Mathematicians, Hyderabad, India*.
- Eckhoff, M. (2002). The low lying spectrum of irreversible, infinite state Markov chains in the metastable regime. Preprint.
- Edenhofer, O. (2010). The economics of climate change – concepts, methods, solutions. Lecture at Technical University Berlin.
- Ellison, G. (1993). Learning, local interaction, and coordination. *Econometrica*, 61:1047–1071.
- Ellison, G. (2000). Basins of attraction, long-run stochastic stability, and the speed of step-by-step evolution. *The Review of Economic Studies*, 67(1):17–45.
- Epstein, J. M. and Axtell, R. (1996). *Growing Artificial Societies*. MIT Press, Cambridge, MA.
- Evensen, P. and Märdin, M. (2009). An extensible and scalable agent-based simulation of barter economies. Master’s thesis, Chalmers Technical University and University of Gothenburg, Sweden.
- Farmer, D. and Foley, D. (2009). The economy needs agent-based modelling. *Nature*, 460:685–86.
- Fisher, F. M. (1983). *Disequilibrium Foundations of Equilibrium Economics*. Cambridge University Press.
- Foster, D. and Young, P. (1990). Stochastic evolutionary game dynamics. *Theoretical Population Biology*, 38:219–232.
- Freidlin, M. and Wentzell, A. (1984). *Random Perturbations of Dynamical Systems*. New York: Springer Verlag.

- Fudenberg, D. and Levine, D. (1998). *The Theory of Learning in Games*. MIT Press, Cambridge, MA.
- Gale, D. (1986a). Bargaining and competition, Part i: Characterization. *Econometrica*, 54(4):785–806.
- Gale, D. (1986b). Bargaining and competition, Part ii: Existence. *Econometrica*, 54(4):807–818.
- Gale, D. (2000). *Strategic Foundations of General Equilibrium: Dynamic Matching and Bargaining Games*. Cambridge University Press, Cambridge, UK.
- Gallegati, M. and Richiardi, M. (2009). Agent-based models in economics and complexity. In Meyers, R. A., editor, *Encyclopedia of Complexity and Systems Science*, pages 200–223. Springer.
- Gaudillière, A. and Landim, C. (2011). A Dirichlet principle for non reversible Markov chains and some recurrence theorems. *Probability Theory and Related Fields*, 158:55–89.
- Gaveau, B. and Schulman, L. S. (1998). Theory of nonequilibrium first-order phase transitions for stochastic dynamics. *Journal of Mathematical Physics*, 39(3):1517–1533.
- Gilboa, I. and Matsui, A. (1991). Social stability and equilibrium. *Econometrica*, 59:859–867.
- Gintis, H. (2006). The emergence of a price system from decentralized bilateral exchange. *Contributions to Theoretical Economics*, 6(1):Article 13.
- Grimm, V., Berger, U., Bastiansen, F., Eliassen, S., Ginot, V., Giske, J., Goss-Custard, J., Grand, T., Heinz, S. K., Huse, G., Huth, A., Jepsen, J. U., Jørgensen, C., Mooij, W. M., Müller, B., Pe'er, G., Piou, C., Railsback, S. F., Robbins, A. M., Robbins, M. M., Rossmanith, E., Rüger, N., Strand, E., Souissi, S., Stillman, R. A., Vab, R., Visser, U., and DeAngelis, D. L. (2006). A standard protocol for describing individual-based and agent-based models. *Ecological Modeling*, 198:115–126.
- Grimm, V., Berger, U., DeAngelis, D. L., Polhill, J. G., Giske, J., and Railsback, S. F. (2010). The odd protocol: A review and first update. *Ecological Modelling*, 221:27602768.
- Hildenbrand, W. and Kirman, A. P. (1988). *Equilibrium analysis: Variations on themes by Edgeworth and Walras*, volume 28 of *Advanced Textbooks in Economics*. North-Holland.
- Hofbauer, J. and Sigmund, K. (2003). Evolutionary game dynamics. *Bulletin (New Series) of the American Mathematical Society*, 40(4):479–519.
- Hofmann, M. E. (2007). Resilience – a formal approach to an ambiguous concept. Diplomarbeit am Institut für Mathematik der Freien Universität Berlin und FAVAIA Working Paper 5, Potsdam Institute for Climate Impact Research, Potsdam, Germany.

- Horenko, I. (2011). Parameter identification in nonstationary Markov chains with external impact and its application to computational sociology. *SIAM Multiscale Modeling & Simulation*, 9(4):1700–1726.
- Huisinga, W. (2001). *Metastability of Markovian Systems: A transfer operator based approach in application to molecular dynamics*. PhD thesis, Fachbereich Mathematik und Informatik, FU Berlin.
- Huisinga, W., Meyn, S., and Schütte, C. (2004). Phase transitions & metastability in Markovian and molecular systems. *Annals of Applied Probability*, 14:419458.
- Huisinga, W. and Schmidt, B. (2006). Metastability and dominant eigenvalues of transfer operators. In Leimkuhler, B., Chipot, C., Elber, R., Laaksonen, A., Mark, A., Schlick, T., Schütte, C., and Skeel, R., editors, *New Algorithms for Macromolecular Simulation*, volume 49 of *Lecture Notes in Computational Science and Engineering*, pages 167–182. Springer.
- Ionescu, C. (2008). *Vulnerability Modeling and Monadic Dynamical Systems*. PhD thesis, Freie Universität Berlin.
- Ionescu, C., Klein, R. J. T., Hinkel, J., Kumar, K. S. K., and Klein, R. (2009). Towards a formal framework of vulnerability to climate change. *Environmental Modelling and Assessment*, 14:1–16.
- Izquierdo, L. R. (2008). *Advancing Learning and Evolutionary Game Theory with an Application to Social Dilemmas*. PhD thesis, Manchester Metropolitan University and The Macaulay Institute, University of Burgos.
- Jacod, J. and Protter, P. (2004). *Probability Essentials*. Berlin, Heidelberg: Springer.
- Jaeger, C. (2008). Mathematics and the social sciences: The challenge ahead. Introductory contribution to the Dahlem Conference: "Is There a Mathematics of Social Entities"?, December 2008, Berlin, Germany. Unpublished manuscript.
- Jaeger, C. C. (2012). Scarcity and coordination in the global commons. In Jaeger, C. C., Hasselmann, K., Leipold, G., Mangalagiu, D., and Tábara, J. D., editors, *Reframing the Problem of Climate Change: From Zero Sum Game to Win-Win Solutions*, pages 85–101. Routledge.
- Kandori, M., Mailath, G. J., and Rob, R. (1993). Learning, mutation, and long run equilibria in games. *Econometrica*, 61(1):29–56.
- Kandori, M. and Rob, R. (1995). Evolution of equilibria in the long run: A general theory and applications. *Journal of Economic Theory*, 65:383–414.
- Kandori, M. and Rob, R. (1998). Bandwagon effects and long run technology choice. *Games and Economic Behavior*, 22(1):30–60.
- Kehoe, T. J. (1991). Computation and multiplicity of equilibria. In Hildenbrand, W. and Sonnenschein, H., editors, *Handbook of Mathematical Economics*, volume IV, chapter 38, pages 2049–2144. Elsevier Science Publishers B.V.

- Kirman, A. P. (1992). Whom or what does the representative individual represent? *Journal of Economic Perspectives*, 6(2):117–136.
- Kirman, A. P. (2008). Economy as a complex system. In Durlauf, S. N. and Blume, L. E., editors, *The New Palgrave Dictionary of Economics*, volume 2, pages 723–729. PALGRAVE MACMILLIAN.
- Knyazev, A. V. and Argentati, M. E. (2010). Rayleigh-Ritz majorization error bounds with applications to fem. *SIAM Journal on Matrix Analysis and Applications*, 31(3):15211537.
- Kreps, D. (1988). *Notes on the Theory of Choice*. Westview Press.
- Kreps, D. M. (1990). *A course in microeconomic theory*. Princeton University Press, Princeton, New Jersey.
- Krivov, S. V. and Karplus, M. (2004). Hidden complexity of free energy surfaces for peptide (protein) folding. *Proceedings of the National Academy of Sciences of the United States of America*, 101:1476614770.
- Kurtz, T. (1970). Solutions of ordinary differential equations as limits of pure jump Markov processes. *Journal of Applied Probability*, 7:49–58.
- Lawrence, F. and Rosebrugh, R. (2003). *Sets for Mathematics*. Cambridge University Press.
- Levin, D. A., Peres, Y., and Wilmer, E. L. (2008). *Markov Chains and Mixing Times*. American Mathematical Society.
- Lewis, D. (1969). *Convention*. Cambridge: Harvard University Press.
- Littman, M. L. (1994). Markov games as a framework for multi-agent reinforcement learning. In *Proceedings of the Eleventh International Conference on Machine Learning*, page 157163. Morgan Kaufman Publishers.
- Mandel, A. and Botta', N. (2009). A note on Herbert Gintis' "Emergence of a price system from decentralized bilateral exchange". *The B.E. Journal of Theoretical Economics*, 9(1). Article 44.
- Mantel, R. (1976). Homothetic preferences and community excess demand functions. *Journal of Economic Theory*, 12:197–201.
- McFadden, D. (2001). Economic choices. *American Economic Review*, 91:351–378.
- Meintrup, D. and Schäffler, S. (2005). *Stochastik: Theorie und Anwendungen*. Berlin, Heidelberg: Springer.
- Metzner, P., Noé, F., and Schütte, C. (2009a). Estimating the sampling error: Distribution of transition matrices and functions of transition matrices for given trajectory data. *Phys. Rev. E*, 80:021106.
- Metzner, P., Schütte, C., and Vanden-Eijnden, E. (2009b). Transition path theory for Markov jump processes. *Multiscale Modeling and Simulation*, 7(3):11921219.

- Meyn, S. and Tweedie, R. L. (2009). *Markov chains and stochastic stability*. Cambridge, UK: Cambridge University Press, 2 edition.
- Nash, J. F. (1950). *Non-cooperative Games*. PhD thesis, Princeton University.
- Noé, F. (2008). Probability distributions of molecular observables computed from Markov models. *The Journal of Chemical Physics*, 128:244103.
- Noé, F., Horenko, I., Schütte, C., and Smith, J. (2007a). Hierarchical analysis of conformational dynamics in biomolecules: Transition networks of metastable states. *Journal of Chemical Physics*, 126:155102.
- Noé, F., Horenko, I., Schütte, C., and Smith, J. C. (2007b). Hierarchical analysis of conformational dynamics in biomolecules: Transition networks of metastable states. *The Journal of Chemical Physics*, 125:155102.
- Noé, F., Schütte, C., Vanden-Eijnden, E., Reich, L., and Weikl, T. (2009). Constructing the equilibrium ensemble of folding pathways from short off-equilibrium trajectories. *PNAS*, 106(45):19011–19016.
- Nöldeke, G. and Samuelson, L. (1993). An evolutionary analysis of backward and forward induction. *Games and Economic Behavior*, 5:424–454.
- Nöldeke, G. and Samuelson, L. (1997). A dynamic model of equilibrium selection in signaling markets. *Journal of Economic Theory*, 73:118–156.
- Norris, J. R. (1998). *Markov Chains (Cambridge Series in Statistical and Probabilistic Mathematics)*. Cambridge University Press.
- Osborne, M. J. and Rubinstein, A. (1990). *Bargaining and Markets*. Academic Press.
- Osborne, M. J. and Rubinstein, A. (1994). *A Course in Game Theory*. The MIT Press, Cambridge, Massachusetts and London, England.
- Prinz, J.-H., Wu, H., Sarich, M., Keller, B., Senne, M., Held, M., Chodera, J. D., Schütte, C., and Noé, F. (2011). Markov models of molecular kinetics: Generation and validation. *The Journal of Chemical Physics*, 134(174105).
- Rao, F. and Caffisch, A. (2004). The protein folding network. *Journal of Molecular Biology*, 342(1):299306.
- Röblitz, S. (2008). *Statistical Error Estimation and Grid-free Hierarchical Refinement in Conformation Dynamics*. PhD thesis, Fachbereich Mathematik und Informatik, FU Berlin.
- Robson, A. and Vega-Redondo, F. (1996). Efficient equilibrium selection in evolutionary games with random matching. *Journal of Economic Theory*, 70:65–92.
- Roth, G. and Sandholm, W. H. (2013). Stochastic approximations with constant step size and differential inclusions. *SIAM Journal on Control and Optimization*, 51:525–555.
- Rubinstein, A. and Wolinsky, A. (1985). Equilibrium in a market with sequential bargaining. *Econometrica*, 53:1133–50.

- Saari, D. (1985). Iterative price mechanisms. *Econometrica*, 53:1117–1131.
- Sandholm, W. H. (2010). *Population Games and Evolutionary Dynamics*. MIT Press, Cambridge.
- Sandholm, W. H. (2011). Stochastic evolutionary game dynamics: foundations, deterministic approximation, and equilibrium selection. Sigmund, Karl (ed.), *Evolutionary game dynamics*. American Mathematical Society short course, January 4–5, 2011, New Orleans, Louisiana, USA. Providence, RI: American Mathematical Society (AMS). *Proceedings of Symposia in Applied Mathematics* 69, 111-141 (2011).
- Sarich, M. (2011). *Projected Transfer Operators*. PhD thesis, Fachbereich Mathematik und Informatik, FU Berlin.
- Sarich, M., Banisch, R., Hartmann, C., and Schütte, C. (2014). Markov state models for rare events in molecular dynamics. *Entropy*, 16(1):258–286.
- Sarich, M., Noé, F., and Schütte, C. (2010). On the approximation quality of Markov state models. *Multiscale Modeling and Simulation*, 8(4):1154–1177.
- Sarich, M. and Schütte, C. (2014). Utilizing hitting times for finding metastable sets in non-reversible Markov chains. Technical Report 14-32, Konrad-Zuse-Zentrum für Informationstechnik Berlin.
- Scarf, H. (1960). Some examples of global instability of competitive equilibrium. *International Economic Review*, 1(3):157–172.
- Scheffé, H. (1958). Experiments with mixtures. *Journal of the Royal Statistical Society. Series B (Methodological)*, 20(2):344–360.
- Schütte, C. (1998). *Conformational Dynamics: Modelling, Theory, Algorithm, and Applications to Biomolecules*. Habilitation thesis, Fachbereich Mathematik und Informatik, FU Berlin.
- Schütte, C., Fischer, A., Huisinga, W., and Deuffhard, P. (1999). A direct approach to conformational dynamics based on hybrid monte carlo. *Journal of Computational Physics*, 151:146–168. Special Issue on Computational Biophysics.
- Schütte, C. and Huisinga, W. (2003). Biomolecular conformations can be identified as metastable sets of molecular dynamics. In Bris, C. L., editor, *Special Volume, Computational Chemistry*, volume 10 of *Handbook of Numerical Analysis*, pages 699 – 744. Elsevier.
- Schütte, C., Noé, F., Lu, J., Sarich, M., and Vanden-Eijnden, E. (2011). Markov state models based on milestoning. *The Journal of Chemical Physics*, 134.
- Schütte, C. and Sarich, M. (2013). *Metastability and Markov State Models in Molecular Dynamics: Modeling, Analysis, Algorithmic Approaches*, volume 24 of *Courant Lecture Notes*. American Mathematical Society.
- Seneta, E. (1981). *Non-negative Matrices and Markov Chains*. Series in Statistics. Springer, 2 edition.

- Serrano, R. and Volij, O. (2008). Mistakes in cooperation: The stochastic stability of Edgeworth's recontracting. *The Economic Journal*, 118(532):1719–1741.
- Simon, H. A. (1957). *Models of Man*. Wiley, New York.
- Singhal, N. and Pande, V. S. (2005). Error analysis and efficient sampling in Markovian state models for molecular dynamics. *The Journal of Chemical Physics*, 123:204909.
- Sonnenschein, H. (1972). Market excess demand functions. *Econometrica*, 40(3):549–63.
- Stern, N. (2006). *The Economics of Climate Change: The Stern Review*. Cambridge University Press, Cambridge, UK.
- Tesfatsion, L. (2006). Agent-based computational economics: A constructive approach to economic theory. In Tesfatsion, L. and Judd, K. L., editors, *Handbook of Computational Economics: Agent-Based Computational Economics*, volume 2, chapter 16, pages 831–880. North-Holland.
- Tesfatsion, L. (2014). Agent-based computational economics: Growing economies from the bottom up. <http://www2.econ.iastate.edu/tesfatsi/ace.htm>.
- Tesfatsion, L. and Judd, K. L., editors (2006). *Handbook of Computational Economics – Agent-Based Computational Economics*. North-Holland.
- Thurstone, L. (1927). A law of comparative judgement. *Psychological Review*, 34:273–286.
- Train, K. (2003). *Discrete Choice Methods with Simulations*. Cambridge, Cambridge University Press.
- Varian, H. R. (1992). *Microeconomic analysis*. W. W. Norton & Company, Inc., 3 edition.
- Weibull, J. W. (1995). *Evolutionary Game Dynamics*. MIT Press, Cambridge.
- Wolf, S. (2009). *From Vulnerability Formalization to Finitely Additive Probability Monads*. PhD thesis, Freie Universität Berlin.
- Wolf, S., Hinkel, J., Hallier, M. E., Bisaro, A., Ionescu, C., Lincke, D., and Klein, R. (2013). Clarifying vulnerability definitions and assessments using formalisation. *To appear in: International Journal of Climate Change Strategies and Management*, 5(1).
- Young, H. P. (1993a). The evolution of conventions. *Econometrica*, 61(1):57–84.
- Young, H. P. (1993b). An evolutionary model of bargaining. *Journal of Economic Theory*, 59:145–168.
- Young, H. P. (1998). *Individual Strategy and Social Structure*. Princeton University Press.
- Young, H. P. (2006). Social dynamics: Theory and applications. In Tesfatsion, L. and Judd, K. L., editors, *Handbook of Computational Economics: Agent-Based Computational Economics*, volume 2, chapter 22, pages 1081–1108. North-Holland.

Zusammenfassung

Inhalt dieser Arbeit ist die Formalisierung (Teil I) und Metastabilitätsanalyse (Teil II) agenten-basierter, evolutionärer Modelle. Teil I diskutiert die Notwendigkeit von Spezifikationen als komplementärer Bestandteil der Beschreibung agenten-basierter Modelle neben der Implementierung und einer Beschreibung in natürlicher Sprache. Es wurde die Grundstruktur eines funktionalen Frameworks für die Spezifizierung agenten-basierter Austauschmodelle präsentiert. Der Hauptfokus dabei lag auf der Herausarbeitung des Zusammenhangs agenten-basierter Austauschmodelle mit traditioneller ökonomischer Theorie. Es zeigte sich, dass die untersuchten Modelle sich von evolutionären Spielen nur darin unterscheiden, dass die Fitnessfunktion nicht deterministisch ist. Das Framework wurde weiterhin benutzt um die Fragestellung der Modelle präzise herauszuarbeiten und damit einen direkten Bezug zu allgemeinen Gleichgewichtsmodellen herzustellen. Darüberhinaus wurde aufgezeigt, wie das Framework unerlässliche Grundlage für die formale und numerische Analyse ist.

Teil II beinhaltet die Metastabilitätsanalyse evolutionärer Spiele in endlichen Populationen, welche einfache agenten-basierte Modelle sind. Die Motivation basiert auf der Beobachtung, dass evolutionäre Spiele häufig metastabile Dynamik aufweisen, bisherige Analyseansätze diese Dynamik jedoch nicht beschreiben können, wie z.B. die Analyse von Nash-Gleichgewichten, die deterministische Approximation oder die Analyse stochastischer Stabilität. Die evolutionären Spiele wurden als zeit-diskrete Markovketten repräsentiert. Darauf aufbauend wurden zwei formale Charakterisierungen von Metastabilität vorgestellt und in den Kontext evolutionärer Spiele gebettet. Es zeigte sich, dass jede Partition des Zustandsraumes metastabil hinsichtlich beider Charakterisierungen ist.

Weiterhin betrachteten wir für die evolutionären Spiele den Ansatz der Core-Set Markov-State-Modellierung zur Konstruktion von zeit-diskreten Markovketten, welche die essentielle, metastabile Dynamik erfassen, dabei jedoch einen wesentlich kleineren Zustandsraum und dadurch eine reduzierte Komplexität aufweisen. Grundlegende Idee ist die Bestapproximation und damit orthogonale Projektion eines Transfer-Operators des evolutionären Spiels auf den Unterraum, welcher aufgespannt wird durch die Committor-Funktionen bzgl. gegebener, disjunkter Teilmengen des Zustandsraumes, der sogenannten Core-Sets. Die darstellende Matrix dieser Projektion ist die Übergangsmatrix der reduzierten Markovkette. Zusammenhänge zwischen der ursprünglichen und der reduzierten Markovkette wurden detailliert herausgearbeitet. Es wurde gezeigt, dass die Konstruktion stochastische Stabilität erhält. Aufbauend auf der Analyse des Approximationsfehlers wird ein Algorithmus zur Identifikation von Core-Sets präsentiert, sodass das resultierende Core-Set Modell die dominanten Zeitskalen und in diesem Sinne die Metastabilität gut wiedergibt. Sowohl die Identifikation von Core-Sets als auch die Konstruktion des entsprechenden Core-Set Markov-State-Modells können simulationsbasiert umgesetzt werden; d.h. sowohl die Core-Sets als auch die Übergangsmatrix des reduzierten Modells können aus simulierten Trajektionsdaten geschätzt werden und sind somit auch im Rahmen der agenten-basierten Modellierung reizvoll. Obwohl der Anwendungsbereich durch die Annahme reversibler Dynamik für allgemeine agenten-basierte Modelle eingeschränkt ist, legt diese Arbeit einen Grundstein für zukünftige Arbeiten in diesem Bereich. Wir geben zwei Beispiele für Klassen reversibler evolutionärer Spiele und einen Überblick über Ansätze für nicht-reversible Markovprozesse.

Erklärung

Ich versichere, dass ich die vorliegende Arbeit selbstständig und nur unter Verwendung der angegebenen Literatur und Hilfsmittel angefertigt habe.

Berlin, den 10. Dezember 2014

Mareen Hallier