Quasi-Stationary Distributions in Stochastic Reaction Networks

PhD thesis

Mads Christian Hansen

Department of Mathematical Sciences
University of Copenhagen

This thesis has been submitted to the PhD School of the Faculty of Science,
University of Copenhagen, October 2018.

Mads Christian Hansen
Email: mads@math.ku.dk

Department of Mathematical Sciences
University of Copenhagen
Universitetsparken 5
2100 Copenhagen
Denmark

Academic Advisor:	Prof. Carsten Wiuf University of Copenhagen, Denmark
Assessment committee:	Prof. Jeffrey F. Collamore (chair) University of Copenhagen, Denmark
	Sen. Lec. Simon L. Cotter University of Manchester, England
	Asst. Prof. Enrico Bibbona Politecnico di Torino, Italy

This work is part of the Dynamical Systems Interdisciplinary Network, University of Copenhagen, and supported by the Danish Research Council and Lundbeckfonden.

ISBN 978-87-7078-881-6

Acknowledgements

This work was carried out as part of my studies at the Department of Mathematical Sciences, University of Copenhagen, from October 2015 to September 2018.

The transition from a background in differential topology to applied mathematics was initiated at the University of Oxford, and the past three years have been a tour de force in probability theory and stochastic dynamical systems. I wish to express my deepest gratitude to my advisor Prof. Carsten Wiuf, for guiding me through (most of) the wilderness, and for his patience when I got lost.

This thesis draws on a great deal of various subjects I have been fortunate enough to have picked up various places in the world during my mathematical upbringing. The inspiration from my stays at MIT, Oxford and most notably here in Copenhagen have shaped the path into and across the mathematical landscape I have undertaken. In particular, the teachings of Viviane Baladi on dynamical systems, of Henrik Schlichtkrull on differential geometry and ODEs, of Richard Hepworth on topology, of Nathalie Wahl on Morse theory and of Steffen Lauritzen on probabilistic graphical models have all found their way into the present work. Furthermore, I wish to give a warm thanks to Prof. Sebastian Schreiber for welcoming me to the University of California, Davis, and for taking me on a journey merging the fields of quasi-stationarity and dynamical systems.

I would like to thank my research group "Mathematics of Reactions Networks" for a welcoming atmosphere, many great discussions and enjoyable trips to various conferences from the snowy French alps to the sunny Deep South. In particular, I thank Daniele Cappelletti and Abhishek Pal Majumder for the many interesting reading sessions and discussions about the strange world of stochastic reaction networks that we have conducted during the years, and for their openness to include me into the international research community of reaction networks.

Finally, I would like to thank my closest family for all their love and support during the years and their indulgence with my at times very absent presence as my thoughts (presumably) embarked on yet another impassable stroll into mathmagic land.

Mads Christian Hansen

Abstract

Stochastic reaction networks compose a broad class of applicable continuous-time Markov processes with a particularly rich structure defined through a corresponding graph. As such, they pose a general and natural framework for representing non-linear stochastic dynamical systems where the interactions among types of entities are themselves of transformational form. Many such systems, in particular when they model real world phenomena, are certain to go "extinct" eventually, yet appear to be stationary over any reasonable time scale. This phenomenon is termed quasi-stationarity. A stationary measure for the stochastic process conditioned on non-extinction, called a quasi-stationary distribution, assigns mass to states in a way that mirrors this observed quasi-stationarity.

In the paper (Hansen and Wiuf, 2018a), we are concerned with providing sufficient conditions for the existence and uniqueness of quasi-stationary distributions in reaction networks. Specifically, for any reaction network we introduce the inferred notion of an absorbing set, and prove through the use of Foster-Lyapunov theory, sufficient conditions for the associated Markov process to have a globally attracting quasi-stationary distribution in the space of all probability measures on the complement of the absorbing set.

The manuscript (Hansen and Schreiber, 2018) deals with connections to the corresponding deterministic system, where qualitative information about the dynamics is often much easier to obtain. Through the use of Morse-decompositions, we show that under the classical scaling, the limit of quasi-stationary measures converges weakly to a probability measure whose support is contained in the attractors of the deterministic system lying entirely within the strictly positive orthant.

Having shown that a specific network at hand has a unique quasi-stationary distribution, the manuscript (Hansen and Wiuf, 2018b) provides an inductive procedure to analytically determine this. Exploiting a center manifold structure, we show that, when one considers the full system as a linear perturbation of a particular sub-network and the coupling to the absorbing set is sufficiently weak, the quasi-stationary distribution may be written as a formal sum with the stationary distribution of the sub-network as a first approximation. We furthermore characterize such stationary distributions for one-species networks.

Resumé

Stokastiske reaktionsnetværk udgør en alsidig klasse af anvendelige kontinuert-tids Markov-processer, som besidder en særlig rig struktur defineret gennem en tilhørende graf. Disse danner således en generel og intuitiv ramme for modelleringen af ikke-lineære stokastiske dynamiske systemer, hvori interaktionerne mellem forskellige typer af entiteter i sig selv er af foranderlig karakter. Omend mange sådanne systemer, særligt når de beskriver processer fra den virkelige verden, med sikkerhed vil "uddø" til sidst, synes disse ofte at opføre sig stationært over enhver rimelig tidshorisont. Dette fænomen kaldes quasi-stationaritet. Et stationært mål for den stokastiske process betinget på ikke-absorption, kaldet en quasi-stationær fordeling, tildeler sandsynlighedsmasse til tilstande på en måde, der afspejler denne observerede quasi-stationaritet.

I artiklen (Hansen and Wiuf, 2018a) beskæftiger vi os med at opnå tilstrækkelige betingelser for eksistensen og entydigheden af quasi-stationære distributioner i reaktions-netværk. Mere nøjagtigt introducerer vi for ethvert reaktionsnetværk den afledte definition af en absorberende mængde, og fastlægger, ved brug af Foster-Lyapunov-teori, tilstrækkelige betingelser for at den tilhørende Markov-proces har en globalt tiltrækkende quasi-stationær fordeling i rummet af sandsynlighedsmål på komplementet af den absorberende mængde.

Manuskriptet (Hansen and Schreiber, 2018) omhandler forbindelser til det tilsvarende deterministiske system, hvor kvalitative oplysninger om dynamikken ofte er lettere at opnå. Gennem brugen af Morse-dekompositioner viser vi, at under den klassiske skalering, konvergerer den tilhørende følge af quasi-stationære fordelinger svagt til et sandsynlighedsmål, hvis støtte er indeholdt i foreningen af de af det deterministiske systems attraktorer, der er helt indeholdt i den strengt positive kvadrant.

Efter at have vist, at et givet netværk har en unik quasi-stationær fordeling, fremføres i manuskriptet (Hansen and Wiuf, 2018b) en induktiv procedure til analytisk at bestemme denne. Ved udnyttelse af en center-mangfoldighedsstruktur, viser vi, at når det fulde system betragtes som en lineær perturbation af et bestemt del-netværk, og koblingen til den absorberende mængde er tilstrækkelig svag, kan den quasi-stationære fordeling skrives som en formel sum, hvori den stationære distribution for del-netværket optræder som en første approksimation. Vi karakteriserer desuden sådanne stationære distributioner for en-dimensionelle netværk.

Contents

1 Introduction

This work falls within the interdisciplinary field of probability theory, graph theory and dynamical systems commonly referred to as *stochastic reaction network theory* (Anderson and Kurtz, 2015; Erdi and Toth, 1989). We may think of reaction networks in generality as a natural framework for representing systems where the interactions among entities are of transformational form (Veloz and Razeto-Barry, 2017). The set of entities (species) may in principle be of any nature, and specifying not just which ones interact (reactions) but also quantifying how fast they interact (kinetics), we obtain the dynamical system of a reaction network. The aim of reaction network theory is in a nutshell to exploit the graphical properties of an underlying network to infer some qualitative properties of the system's associated dynamics. One may argue that the foundation of reaction networks, as it is known today, was laid in the seminal papers (Horn and Jackson, 1972; Horn, 1972; Feinberg, 1972). Examples abound in biochemistry, where the language originated, however the true power of this approach is the ability to model a range of diverse real world processes as dynamical systems on networks (Newman, 2010). To wit, examples are found in biological (Turing, 1952; Arkin et al., 1998; Barkai and Leibler, 2000), medical (Anderson et al., 2006), political (Dittrich and Winter, 2008), sociological (Weidlich and Haag, 2012), computational (Cook et al., 2009), economical (Veloz et al., 2014), ecological (Shakil et al., 2015) or epidemiological (Nåsell, 2011; Kermack and McKendrick, 1927) contexts.

The complexity of systems modeled by reaction networks is illustrated by the fact that changes or actions of some of the entities may have implications spreading across the graph, effecting other entities and *vice versa*, thus causing feedback effects and thereby non-linearity. One therefore needs to consider the system as a whole. This line of thought has broadly been dubbed *systems theory* (Emery, 1969) and more specifically *synergetics* when one considers a unifying mathematical framework within which quantitative descriptions of complex, *self-organizing* systems can be made (Haken, 1983). Indeed, we may think of reaction networks as the prototype of nonlinear systems, capable of exhibiting bifurcations, fluctuations and chaos (Erdi and Toth, 1989).

Many real world systems are certain to "die out" eventually, yet appear to be stationary over any reasonable time scale. This phenomenon is termed *quasi-stationarity* (van Doorn and Pollett, 2013; Collet et al., 2013). If the system that we are studying has been running for a long time, and if the only information available to us is that it has not reached extinction, then the so called *quasi-stationary distribution* (QSD), if it exists, is the likely distribution of the state variable (Nåsell, 2011; Méléard and Villemonais, 2012; Renshaw, 2011; Collet et al., 2013). Thus, while the limiting distribution necessarily assigns all of its probability mass to the extinction states, the quasi-stationary distribution assigns mass to states in a way that mirrors the quasi-stationarity observed (van Doorn and Pollett, 2013). As such, when the

time to extinction is large, this effectively ultimate distribution, rather than the degenerate limit distribution, is the natural object of study. The term quasi-stationary distribution was coined by Bartlett (1960), although thoughts roaming around the same general idea can be found decades before in various scientific fields, bearing testimony to the universal relevance of quasi-stationarity (van Doorn and Pollett, 2013).

During the last couple of decades, the interest in reaction networks and quasi-stationary distributions have independently acquired much attention (For general references see Angeli, 2009; Pollett, 2015). However, beyond a few special cases, not much work has been done in trying to understand quasi-stationarity in reaction networks (Anderson et al., 2014). The papers and manuscripts constituting this thesis is an attempt to fill in some of this gap.

1.1 Motivation

Whether the universe is inherently deterministic or stochastic in nature, the lack of complete information in complex systems inevitably introduces some degree of stochasticity. Thus, a stochastic description is not an alternative to the deterministic approach, but a more complete one (Qian, 2011; Anishchenko et al., 2007). From a modeling standpoint, when the population of interacting entities are small and reaction rates are slow, it is important to recognize that the individual reaction steps occur discretely and are separated by time intervals of random length (Arkin et al., 1998). This is for example the case at the cellular level (Elowitz et al., 2002), where stochastic effects resulting from these small numbers may be physiologically significant (Cook et al., 2009), or at the level of populations describing the spread of disease or opinions. Such models are often referred to as individual- or agent-based models. In this inherently discrete domain of individuals (the microscale), stochastic models are ubiquitous. In contrast, large populations are traditionally modeled by ordinary differential equations, when the spatial structure, the age-structure, the fluctuations of the environment, etc, are ignored. These population-level models are supposed to account for the deterministic trends of large populations (the macroscale), and are inherently incapable of describing extinction phenomena (Chazottes et al., 2017; Johnston et al., 2017). Indeed, given a reaction network of interest, the associated deterministic model solution is an approximation of the corresponding solution for the stochastic model, improving with the system size, and in general only remaining valid on finite time intervals (Kurtz, 1970, 1972, 1978). Thus, the long-term behavior of a given reaction network may depend crucially on whether it is modeled deterministically or stochastically (Gupta et al., 2014).

One of the indisputable facts about life, as we know it, is that it has an end. In this light, the notion of extinction would appear to be a very natural and widely occurring phenomenon. Generally, the finite nature of the resources available prevents the system from growing without limit. Thus, provided we wait sufficiently long, a downward fluctuation in population size that is sufficiently violent to drive the species extinct is bound to occur (Renshaw, 2011). Therefore, given that the model under consideration is sufficiently accurate, the stationary measure (if it exists) would not provide us with any new information. As a consequence, the counterpart to a stable stationary solution in the associated deterministically modeled system is not generally a stationary distribution of the corresponding stochastic model. Instead, when it exists and the time to extinction is sufficiently long, the quasi-stationary distribu-

tion, which is the stationary measure when we condition on the process not going extinct, has shown to be the natural object of study (van Doorn and Pollett, 2013; Collet et al., 2013).

Luckily, extinction is often a rare event in the sense that the expected time until its occurrence is exceedingly large compared to the timeframe of interest. When this is the case, one may view the dynamics on different timescales. In the short term, the stochastic dynamical system undergoes a quick transformation from its initial state to a seemingly stationary behavior. In the long term, the process keeps wandering around the state space according to this seemingly stationary behavior; this is the fluctuation captured by the QSD. Finally, in the very long term, $t \to \infty$, the system reaches extinction. Note that this quasi-stationary behavior might well reflect a more complex dynamics in the deterministic reaction network such as limit cycles or even chaos. For example a multimodal QSD generally occurs when there are multiple stationary states in the deterministic system, a property referred to as *multistationarity* and well studied in the reaction network literature (Joshi and Shiu, 2015; Wiuf and Feliu, 2013). In the corresponding stochastic system, fluctuations allow *switching* from one stable stationary state to another, which again is an inherently stochastic phenomenon. These *noise-induced* effects are often paramount in understanding nonlinear dynamics, and may give rise to stochastic resonance, sustained oscillations or, contrary to intuition, a more ordered temporal behavior (Anishchenko et al., 2007).

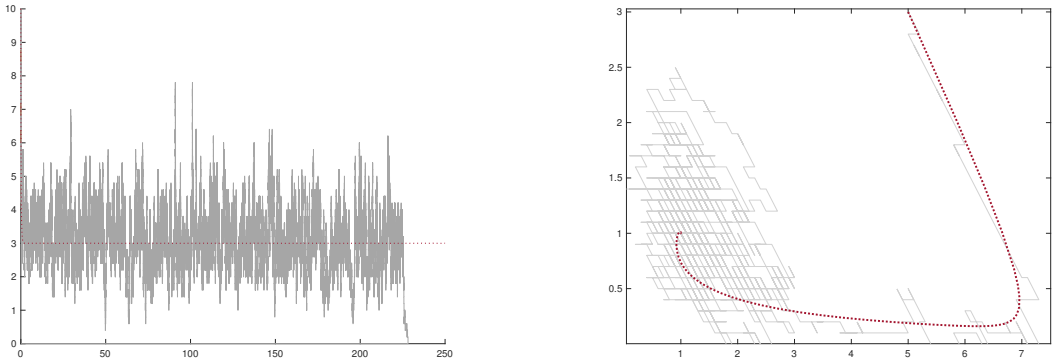

Figure 1.1: Two sample paths of stochastic reaction networks exhibiting quasi-stationary behavior (grey). The deterministic dynamics is included for comparison (red).

Results on quasi-stationary distributions are not only of theoretical interest. Real world processes are often described through complex reaction networks with a large number of species and reactions. In order to gain information of such systems, it is of great importance to find suitable model reductions. This is in itself an active area of research (Cappelletti and Wiuf, 2016). Often the reactions involved in biochemical processes occur at very different time-scales (Kang and Kurtz, 2013), which poses several problems. In particular, numerical simulations of the systems may become infeasible (Cotter, 2016). However, quasi-stationary distributions are exactly the objects required in order to perform the multi-scale reduction methods for stochastic reaction networks currently being explored in the probability litera-

ture (Kang and Kurtz, 2013; Anderson et al., 2014). For example, in order to determine the behavior of species operating on a timescale that is slower than other species, it is necessary to first understand the long-term dynamics of the species on the fast timescale, which may be described by a suitable quasi-stationary distribution, in order to perform either the necessary stochastic averaging, or to recognize that some species will have gone extinct (Anderson et al., 2014).

As with all types of mathematical modeling, the aim is to make predictions about the behavior of real world phenomena, and to design ways to interact with them. To give some examples, suppose one considers a reaction network modeling an endemic disease in a population, say ebola (Nieddu et al., 2017). Then the quasi-stationary distribution answers the question of how likely it is to see n people infected. Further, one may use the QSD to estimate parameters in the model based on observations, determining how infectious the disease is. Knowing the QSD in terms of the model parameters would in turn allow governments to determine the degree of action needed. This example underlines the value of having analytic mathematical results rather than just computational simulations, although the latter may in itself greatly reduce the time and resources needed for experiments. Likewise, suppose one considers a stochastic reaction network describing the political opinions of a population, to the degree of determining which political party an individual votes for. Provided the model is sufficiently accurate, the only possible long term outcome corresponds to extinction of some or all parties (species). However, if one wishes to predict the behavior and variability of the system until the next election, the quasi-stationary distribution is of interest. Similar considerations could be made for market crashes in economics or the extinction of species in ecology (Schreiber, 2016). Although such events are rare, it is a general trait of complex systems that they have a fat-tailed distribution (Holland, 2014).

Stochastic reaction networks seem especially well suited for modeling systems in biochemistry where the language originated, and where quasi-stationary behavior naturally occurs. A better description of their behavior can thus aid the understanding of the functioning of the cell, including gene regulatory systems, signaling systems and metabolic systems which in itself may have far reaching consequences in the long run. It has been suggested that reaction networks provide the language for constructing new technology such as bio-computers which need no electricity to run and can function in a hydrous environment such as our bodies (Cappelletti, 2015). Indeed, stochastic reaction networks are essentially *Turing universal*, meaning they can compute any computable function with arbitrarily small error (Cook et al., 2009). Furthermore, with the rise of synthetic biology these networks are being implemented through DNA computing (Shin et al., 2017). Reaction network theory may thus provide the theoretical framework for designing controlled interventions which for example could cure or ease a disease by prescribing targeted drugs (Cappelletti, 2015). In order to aid this development, solid mathematical insights are needed. This thesis can be seen as a small theoretical contribution towards such an understanding.

1.2 Contribution

The papers and manuscripts contained in this thesis provide my contribution to the field of stochastic reaction networks, developed at the University of Copenhagen, from October 2015 to September 2018. Generally speaking, the papers are concerned with examining quasi-stationarity in stochastic reaction networks from three different angles. Existence and uniqueness (Hansen and Wiuf, 2018a), analytic computability (Hansen and Wiuf, 2018b) and the relationship between deterministic and stochastic models (Hansen and Schreiber, 2018). In the following, we will describe in more detail the content of this work as well as its relation to the international state-of-the-art. The reader may advantageously consult Chapter 2 where the relevant mathematical theory is introduced.

When considering quasi-stationarity, one is interested in the long-term behavior before entry to a particular *absorbing set* A. Thus, given a reaction network, the first natural step is to give a characterization of this set. In the paper (Hansen and Wiuf, 2018a), we introduce the inferred notion of an *endorsed set*, E, and provide an automated way of decomposing the state space, D, of the Markov process corresponding to a reaction network into the disjoint union $D = A \sqcup E$.

The questions of sufficient and necessary conditions for the existence of QSDs are fundamental and have been extensively studied (Ferrari et al., 1995; van Doorn, 1991; van Doorn and Pollett, 2013). Unfortunately, the current results are rather intractable, preventing any swift application but in a few special situations. Recently it was shown that a couple of abstract conditions was equivalent to the existence and uniqueness of a globally attracting QSD in the space of probability distributions on E (Champagnat and Villemonais, 2016, Theorem 2.1). Using Foster-Lyapunov theory (Meyn and Tweedie, 1993, 2009), a series of assumptions on the process $(X_t : t \geq 0)$ has been shown to be sufficient for these conditions to hold (Champagnat and Villemonais, 2017). This approach has been applied to a particular case of multidimensional birth-death processes, giving sufficient conditions, in terms of the parameters of the process, for the existence and uniqueness of a QSD. In the paper (Hansen and Wiuf, 2018a), we extend this result, not just to a larger set of parameter values in the birth-death process case, but to the *much* broader class of stochastic processes associated to stochastic reaction networks. In proving the main theorem, we also provide a sufficient condition for a reaction network to go extinct almost surely, as this has attracted much attention on its own (Johnston et al., 2017). In the process, we gain more intuition of the importance of "coming down from infinity in finite time", which has been studied in the QSD literature (Martinez et al., 2014).

Acknowledging that an exact analytical description of the QSD remains a utopian dream, the aim of the manuscript (Hansen and Schreiber, 2018) is to gain qualitative information about the QSD by analyzing the corresponding deterministic system, which is often a much easier task. The seminal work of Kurtz (Kurtz, 1972, 1978; Ethier and Kurtz, 1986) allows one to obtain a tight connection between the two descriptions through the so called *classical scaling*. When increasing a system parameter, ϵ, usually related to the system size, the stochastic process converges in probability to the deterministic solution on *compact* time

intervals. With this at our disposal, we may consider the stochastic dynamical system as a *random perturbation* of the corresponding deterministic system, a notion originally due to Kifer (1988). Under the assumption of an attractor contained strictly inside the positive orthant, we show that the expected time to extinction scales exponentially with system size, ϵ. This in particular implies that for any rate constants, if the system is sufficiently large, quasi-stationarity is to be expected. Further, for a decreasing series of ϵ, we show that the corresponding series of QSDs converge to an invariant measure of the deterministic flow. This work builds directly on Faure and Schreiber (2014) and Marmet (2013) which studied quasi-stationary distributions of randomly perturbed dynamical system in discrete time and continuous time on compact sets respectively. These papers furthermore required a series of assumptions which may in general be difficult to check. In particular, one needs a *large deviation principle*, which we show is always satisfied for sufficiently large reaction networks.

We exploit the notion of *Morse decompositions* to show, following the ideas of Marmet (2013) extended to the general case of reaction networks, that the weak limit of the QSDs has support contained in the union of those attractors for the deterministic system that are contained strictly inside the positive orthant, whenever a minimal such Morse decomposition exists. We illustrate this result through a series of examples with varying deterministic dynamics.

Once we know that the stochastic process associated to a particular reaction network has a unique QSD, the natural question is how to determine it, or at least its form, analytically. This question is of course even harder than just guaranteeing existence, and the positive cases where the form of the QSD can be given may be counted on a single hand. Another fact bearing testimony to the delicate nature of the question is that in reaction network theory, the results pertaining to ordinary stationary distributions exploit the notion of complex balanced - a property which eventual absorption precludes. For these reasons, one is usually forced to make some kind of approximation. In the manuscript (Hansen and Wiuf, 2018b) we try to circumvent this and at the same time exploit some of the theory available to us in reaction network theory, by means of separation of time-scales and a *center manifold* (Childs and Keener, 2012; Pollett and Roberts, 1990).

We prove, under certain conditions, that the usual stationary distribution of a reduced reaction network may be used as a first approximation of the QSD, and provide an explicit iterative expression for the exact QSD through certain rate matrices, using results from Pollett and Roberts (1990). For one-species reaction networks, we provide another iterative procedure determining the stationary distribution. Together, these in principle determine analytically the QSD of a substantial class of reaction networks. However, the formulas are so involved that an implementation on a computer is necessary. In particular, one has to truncate the matrices whereby only an approximate solution may be found. We examine the performance of the method by comparing it to other means of calculating the QSD approximately and find excellent agreement.

1.3 Thesis Structure

The overall structure of the thesis is as follows. In Chapter 2 the necessary mathematical background on the fields of reaction networks and quasi-stationary distributions are given. Some smaller new results are also included. The purpose is to familiarize the reader with the notation and context of the subjects and provide an overview of the most essential results needed for the understanding of the subsequent papers. Chapters 3-5 each contain a research paper constituting the main contributions of the thesis. More specifically, Chapter 3 contains the foundational paper on existence and uniqueness of quasi-stationary distributions in reaction networks. This provides a natural starting point for further investigation of the properties of such a guaranteed quasi-stationary distribution. In Chapter 4 we examine in a manuscript, by using random perturbations, the question of what can be inferred about the quasi-stationary distribution from the dynamical properties of the corresponding deterministic system. The last and least developed manuscript is contained in Chapter 5. It is concerned with the question of how one may describe a quasi-stationary distribution through a slow manifold reduction when there exists a separation of time-scales. Finally, Chapter 6 contains some interesting perspectives for further research as well as an overall outlook of future directions within the field.

2 Dynamics of Reaction Networks

In this chapter the necessary definitions and results of classical reaction network theory and the field of quasi-stationary distributions are introduced (for general expositions, see Anderson and Kurtz, 2015; Feinberg, 1979; Erdi and Toth, 1989; Gunawardena, 2003; Collet et al., 2013; Méléard and Villemonais, 2012; van Doorn and Pollett, 2013; Champagnat and Villemonais, 2018). A few new results are also presented. All used notation and definitions are also employed in the papers and manuscripts collected in this work. This chapter may therefore be used as a consistent reference for the remaining parts of the thesis.

2.1 Notation and Prerequisites

Let $\mathbb{Z}$, $\mathbb{N}$, $\mathbb{R}$ denote the integers, natural numbers and real numbers respectively. Let

$$\mathbb{N}_0 = \{0, 1, 2, \dots\}, \qquad \mathbb{R}_0 = [0, \infty).$$

For two real numbers $c_1, c_2 \in \mathbb{R}$, we let $c_1 \wedge c_2$ and $c_1 \vee c_2$ denote $\min\{c_1, c_2\}$ and $\max\{c_1, c_2\}$ respectively. The transpose of a matrix Q is denoted Q'. Let $d \in \mathbb{N}$. For any vector $v \in \mathbb{R}^d$, we denote by v_i the ith component of v and by $\|v\|_1$ its L^1-norm. Further, by $[v]$ we denote the vector of the floor functions of the entries of v, that is $[v]_i = \lfloor v_i \rfloor$. For two vectors $v, w \in \mathbb{R}^d$, we write $v < w$ and $v \leq w$ if the inequality holds component wise. Furthermore, $\langle v, w \rangle$ denotes the usual inner product. For $v, w \in \mathbb{R}^d$ define

$$v^w = \prod_{i=1}^{d} v_i^{w_i}, \qquad \text{and} \qquad v! = \prod_{i=1}^{d} v_i! \tag{2.1.1}$$

with the conventions that $0! = 1$ and $0^0 = 1$. Finally, for any set B we denote its cardinality by $|B|$ and by $\mathbb{1}_B : D \to \{0, 1\}$ the indicator function of a subset $B \subseteq D$.

We shall tacitly assume that the reader is familiar with the standard results of continuous time stochastic processes on countable spaces (Norris, 2009; Ethier and Kurtz, 1986; Rogers and Williams, 2000). Of particular use will be the Poisson process, birth-death processes and stationary distributions. When a sequence of probability measures $(\mu_n)_{n\in\mathbb{N}}$ converges weakly to μ, we will use the notation $\mu_n \Rightarrow \mu$. Further, we will assume familiarity with basic notions from graph theory which may be found in (van Lint and Wilson, 2006), as well as some continuous time dynamical systems theory (Brin and Stuck, 2002; Katok and Hasselblatt, 1995). In particular we assume familiarity with attractors, vector fields, Lyapunov functions and the classical results of existence and uniqueness of ordinary differential equations (See Taylor, 2011).

2.2 Reaction Network Theory

A reaction network is a triple $\mathcal{N} = (\mathcal{S}, \mathcal{C}, \mathcal{R})$, where $\mathcal{S}$ is a finite ordered set of *species*, $\mathcal{C}$ is a finite ordered set of *complexes*, consisting of linear combinations over $\mathbb{N}_0$ of the species and $\mathcal{R}$ is a non-empty irreflexive relation on $\mathcal{C}$, referred to as the set of *reactions* (Anderson and Kurtz, 2015; Feinberg, 1979; Gunawardena, 2003).

We define the dimension of the reaction network, $d = |\mathcal{S}|$. Any species $S_i \in \mathcal{S}$ can be identified with the unit vector $e_i \in \mathbb{N}_0^d$, thus any complex $y \in \mathcal{C}$ will be identified with a vector in $\mathbb{N}_0^d$. One may then regard the zero vector as a complex, and we denote it by $\emptyset$. It is customary to denote an element $(y_k, y_k') \in \mathcal{R}$ by $y_k \to y_k' \in \mathcal{R}$ in which case we refer to y_k as the *source complex* and to y_k' as the *product complex* of reaction k. We may thus write $\mathcal{R} = \{y_k \to y_k' : y_k, y_k' \in \mathcal{C}, y_k \neq y_k', k = 1, \ldots, r\}$. Employing a standard, although slight abuse of, notation, we identify $\mathcal{S} = \{S_1, \ldots, S_d\}$ with the set $\{1, \ldots, d\}$, $\mathcal{C}$ with the set $\{1, \ldots, m\}$ and $\mathcal{R}$ with $\{1, \ldots, r\}$. In particular, $r = |\mathcal{R}|$. We write the kth reaction with the notation

$$\sum_{i \in \mathcal{S}} y_{ki} S_i \to \sum_{i \in \mathcal{S}} y_{ki}' S_i,$$

where $y_{ki} = (y_k)_i$ and $y_{ki}' = (y_k')_i$ are the stoichiometric coefficients associated with the source and product complexes of reaction k, respectively. One may completely describe a reaction network $\mathcal{N}$ (without kinetics) in terms of its reaction graph $\mathcal{G}$, whose nodes is the set of complexes, $\mathcal{C}$, and whose directed edges are given by the set of reactions, $\mathcal{R}$. This concise description will be employed in the rest of this thesis.

The connected components of $\mathcal{G}$ are called the *linkage classes* of the reaction network. We say that a reaction network is *weakly reversible* if and only if the linkage classes are strongly connected. Such strongly connected linkage classes are simply called *strong linkage classes*. A strong linkage class is called *terminal* if no path out of the linkage class exists, that is any path starting in a terminal strong linkage class is confined to that strong linkage class.

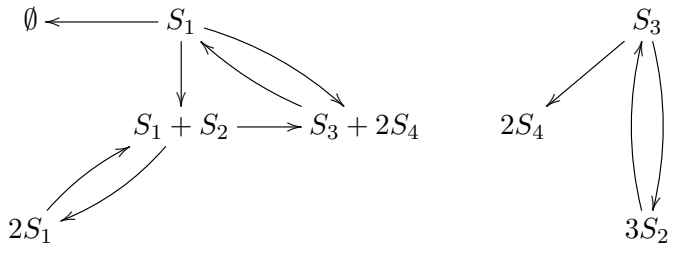

Figure 2.1: The reaction graph of a reaction network with species $\mathcal{S} = \{S_1, S_2, S_3, S_4\}$ and complexes $\mathcal{C} = \{\emptyset, S_1, S_1 + S_2, 2S_1, S_3 + 2S_4, S_3, 2S_4, 3S_2\}$. There are two linkage classes and four strong linkage classes. $\{2S_4\}$ and $\{\emptyset\}$ are terminal.

Define the *reaction vectors* $\xi_k = y'_k - y_k$ corresponding to the net gain of species obtained by the occurrence of the reaction $y_k \to y'_k$. The *stoichiometric subspace* of a reaction network, $\mathcal{N}$, is the linear subspace of $\mathbb{R}^d$ given by

$$\Psi = \{y'_k - y_k \,|\, y_k \to y'_k \in \mathcal{R}\} = \{\xi_k \,|\, k \in \mathcal{R}\}.$$

For $x \in \mathbb{R}^d$, the sets $(x + \Psi) \cap \mathbb{R}^d_0$ are called the *stoichiometric compatibility classes* of $\mathcal{N}$. If there exists a positive vector orthogonal to Ψ, then the network is called *conservative*. Equivalently, a reaction network is conservative if every stoichiometric compatibility class is bounded. Defining the *stoichiometric matrix* as

$$\Xi = \begin{pmatrix} \xi_1 & \dots & \xi_r \end{pmatrix} \in \mathbb{N}_0^{d \times r}.$$

Thus, a sufficient condition for a network to be non-conservative is for Ξ to have full rank. The *order* of reaction k is the sum of the stoichiometric coefficients of the source complex, $\sum_{i \in \mathcal{S}} y_{ki}$. Finally, we define the maximum of a vector over the set $\mathcal{R}$, $x = \max_{k \in \mathcal{R}} y_k$, as the entry-wise maximum, $x_i = \max_{k \in \mathcal{R}} y_{ki}$.

One of the most well studied graphical properties is the *deficiency* of a reaction network $\mathcal{N}$ (Horn and Jackson, 1972; Feinberg, 1979). It is the non-negative integer given by

$$\delta = |\mathcal{C}| - \ell - s,$$

where $|\mathcal{C}|$ is the number of complexes, ℓ is the number of linkage classes, and s is the dimension of the stoichiometric subspace Ψ (Anderson and Kurtz, 2015). One may view it as a specialized measure of the linear independence of the reaction vectors, or more informatively and surprisingly, as being closely related to the dimension of a linear subspace, upon which the associated dynamical system is in equilibrium (Gunawardena, 2003). This way of inferring information on the dynamics based solely on graphical properties lies, as we shall see, at the heart of modern reaction network theory.

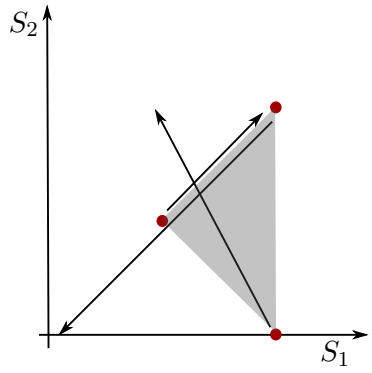

Figure 2.2: $S_1 + S_2 \to 2S_1 + 2S_2 \to \emptyset$, $2S_1 \to S_1 + 2S_2$, endotactic but not strongly endotactic.

Another graphical property, which has attracted much interest in both the deterministic and the stochastic descriptions is *endotacticity*. As there exist endotactic networks which admit extinction events, these might pose a sufficiently well behaved class of systems for

which quasi-stationarity may be explored. As the definition of endotacticity is rather technical and as we will only touch upon this class in passing, we shall only give the geometric interpretation and refer the reader to Craciun et al. (2013); Johnston et al. (2016) for details. Suffice it to say that a reaction network is endotactic if, roughly speaking, every reaction vector points "inward" relative to the polytope in $\mathbb{R}^d_0$ consisting of the convex hull of the source complexes. Adding the requirement that each proper face of the polytope contains at least one reactant in a reaction that points out of that face, one obtain the class of *strongly endotactic* networks (Gopalkrishnan et al., 2014).

2.2.1 Associated Deterministic and Stochastic Dynamical Systems

Given a reaction network $\mathcal{N}$, we may consider either its corresponding continuous time deterministic dynamical system on the continuous state space $\mathbb{R}^d_0$ or its continuous time stochastic dynamical system on the discrete state space $\mathbb{N}^d_0$. To obtain either of the dynamical systems, one has to endow the graphical structure with a kinetics.

Deterministic Reaction Systems. When all species are present in large abundances, the concentration of each species is considered, instead of the exact counts. The state space is $\mathbb{R}^d_0$, and each point herein is a state. A unique solution $\varphi_t(x)$ to an ordinary differential equation with initial condition $\varphi_0(x) = x$ defines a flow or continuous time dynamical system. We denote by $x_t := \varphi_t(x)$ the non-negative real vector of $\mathbb{R}^d_0$ whose entries are the concentrations of interest at time t, when the initial concentration is $x_0 \in \mathbb{R}^d$. The evolution of x_t is modeled as the solution to the ordinary differential equation (ODE)

$$\dot{x}_t = \sum_{k\in\mathcal{R}} (y'_k - y_k)\lambda_k(x_t) = \sum_{k\in\mathcal{R}} \lambda_k(x_t)\xi_k =: F(x_t) \tag{2.2.1}$$

for some so called *rate functions* $\lambda_k : \mathbb{R}^d_0 \to \mathbb{R}_0$ and an initial condition $x_0 \in \mathbb{R}^d_0$ (Gunawardena, 2003). We require that the functions λ_k are continuously differentiable and that $\lambda_k(x) > 0$ if and only if $\operatorname{supp} y \subseteq \operatorname{supp} x$. When integrating, we may write this as

$$x_t = x_0 + \sum_{k\in\mathcal{R}} \left(\int_0^t \lambda_k(x_s)\, ds \right) \xi_k. \tag{2.2.2}$$

The rate functions determine a deterministic kinetics $\lambda = (\lambda_1, \ldots, \lambda_r)$ for $\mathcal{N}$, and the pair $(\mathcal{N}, \lambda)$ is called a deterministic reaction system (Cappelletti, 2015). The rate functions may in general be highly non-linear, yielding interesting and complex dynamical behavior and thus preventing a straight forward solution to the ODE system. In the most common case, the rate functions take the form

$$\lambda_k(x) = \alpha_k x^{y_k}, \qquad k \in \mathcal{R},$$

where $\alpha_k > 0$ are called *rate constants*. The model is then referred to as deterministic mass-action kinetics, and the pair $(\mathcal{N}, \alpha)$ as a mass-action system. With this choice of kinetics, one obtains a polynomial dynamical system. As such, reaction network theory

is closely connected to a classical branch of dynamical systems. Indeed, Hilbert's 16th problem concerns the determination of an upper bound for the number of limit cycles in two-dimensional polynomial vector fields of degree n and their relative positions. This is still unresolved for any $n > 1$ (Ilyashenko, 2002).

What aids reaction network theory is a hidden linearity stemming from the stoichiometry. Let $Y \in \mathbb{N}_0^{d\times m}$ be the matrix of stoichiometric coefficients, that is, $Y_{i,j}$ is the ith entry in the vector representing complex j. Note that some refer to Y as the stoichiometric matrix, whereas we use that terminology for Ξ. The Laplacian matrix $L_\alpha \in \mathbb{R}^{m\times m}$ is given by

$$(L_\alpha)_{i,j} = \begin{cases} -\sum_{l=1}^m \alpha(i,l), & i = j \\ \alpha(j,i), & \text{otherwise} \end{cases}$$

where $\alpha(i,j) > 0$ is the rate constant associated to the reaction from complex i to complex j. Finally, $\psi(x) \in \mathbb{R}_0^m$ is the vector with $\psi_j(x) = \prod_{i=1}^d x_i^{y_{ji}}$ where y_j is the vector corresponding to the jth complex. Then the following diagram commutes

$$\begin{array}{ccc} \mathbb{R} & \xrightarrow{L_\alpha} & \mathbb{R} \\ {\scriptstyle\psi}\uparrow & & \downarrow{\scriptstyle Y} \\ \mathbb{R}^d & \xrightarrow[F]{} & \mathbb{R}^d \end{array}$$

that is, one may represent the vector field as

$$\dot{x} = Y L_\alpha \psi(x).$$

Note that the non-linearity only comes from the term $\psi(x)$, which is just a bookkeeping function. What carries the dynamics is L_α which depends both on the network structure and the kinetics Gunawardena (2003). The fact that this is a linear function is key to the insights of deterministic reaction network theory.

Stochastic Reaction Systems. When the copy-numbers of interacting species are low, one needs to track the exact species counts. This microscopic view has an inherent discreetness of the state space, which is $\mathbb{N}_0^d$. We specify the stochastic process $(X_t \colon t \geq 0)$ associated to the reaction network as follows. Let X_t be the vector in $\mathbb{N}_0^d$ whose entries are the species counts at time t. If reaction $y_k \to y_k'$ occurs at time t, then the new state is $X_t = X_{t-} + y_k' - y_k = X_{t-} + \xi_k$, where X_{t-} denotes the previous state. The stochastic process then follows

$$X_t = X_0 + \sum_{k\in\mathcal{R}} Y_k \left(\int_0^t \lambda_k(X_s)\, ds \right) \xi_k, \tag{2.2.3}$$

where Y_k are independent and identically distributed unit-rate Poisson processes and for each reaction $k \in \mathcal{R}$, we specify an intensity function $\lambda_k \colon \mathbb{N}_0^d \to [0,\infty)$(Anderson and Kurtz, 2015; Ethier and Kurtz, 1986; Norris, 2009). This is sometimes referred to as a *random time change representation*. Here, the random variable $Y_k(\int_0^t \lambda_k(X_s)\, ds)$ counts the number of

times the reaction $k \in \mathcal{R}$ has occurred up to time t. Note the similarity to the deterministic case of (2.2.2). We require that the rate functions satisfy the *stoichiometric admissibility condition*,

$$\lambda_k(x) > 0 \quad \Leftrightarrow \quad x \geq y_k,$$

where we use the usual vector inequality notation; $x \geq y$ if $x_i \geq y_i$ for all $i \in \mathcal{S}$. Thus, reactions are only allowed to take place whenever the copy-numbers of each species in the current state is at least as great as those of the corresponding source complex. A widely used example is *stochastic mass-action kinetics* given by

$$\lambda_k(x) = \alpha_k \frac{x!}{(x-y_k)!} \mathbb{1}_{\mathbb{N}_0^d}(x - y_k) = \alpha_k \prod_{i=1}^{d} y_{ki}! \binom{x_i}{y_{ki}} \mathbb{1}_{\mathbb{N}_0}(x_i - y_{ki}),$$

for some reaction rate constants $\alpha_k > 0$ (Anderson and Kurtz, 2015). The idea is that the rate is proportional to the number of distinct subsets of the molecules present that can form the inputs for the reaction. Again, it reflects the assumption that the system is well-stirred (Anderson and Kurtz, 2015), but note the slight difference from the deterministic case, due to the combinatorics of the discrete nature of the copy-numbers. Other examples include power law kinetics or generalized mass action kinetics (Horn and Jackson, 1972; Müller and Regensburger, 2012). A particular choice of such rate functions constitute a stochastic kinetics $\lambda = (\lambda_1, \ldots, \lambda_r)$ for the reaction network $\mathcal{N}$, and the pair $(\mathcal{N}, \lambda)$ is again referred to as a stochastic reaction network, or simply a network.

The generator, L, for the Markov process $(X_t \colon t \geq 0)$ associated to a reaction network is given by

$$Lf(x) = \sum_{k \in \mathcal{R}} \lambda_k(x) \left(f(x + \xi_k) - f(x) \right)$$

for functions $f \colon \mathbb{N}_0^d \to \mathbb{R}$ (Anderson and Kurtz, 2015). Assuming that the solution of (2.2.3) exists for all times, that is $(X_t \colon t \geq 0)$ jumps only finitely many times in a finite time interval, then

$$f(X_t) - f(X_0) - \int_0^t Lf(X_s)\, ds \tag{2.2.4}$$

is at least a local martingale for all functions $f \colon \mathbb{N}_0^d \to \mathbb{R}$ (Kang and Kurtz, 2013). In fact, for the reaction networks we will consider, it is a true martingale (Anderson and Kurtz, 2015, Cond. 1.20). In particular its expectation is zero hence letting $p_x(t) = \mathbb{P}(X_t = x)$ we obtain

$$\sum_x f(x) p_x(t) = \sum_x f(x) p_x(0) + \int_0^t \sum_x Lf(x) p_x(s)\, ds.$$

Choosing $f(x) = \mathbb{1}_{\{y\}}(x)$ it follows that

$$p_y(t) = p_y(0) + \int_0^t \sum_{k\in\mathcal{R}} (\lambda_k(y-\xi_k)p_{y-\xi_k}(s) - \lambda_k(y)p_y(s))\, ds,$$

which it the integral form of the *Kolmogorov forward equation*, also known as the *master equation*, describing the rate of change of the probability of being in state y as the probability of coming into state y minus the probability of jumping out of that state

$$\dot{p}_y(t) = \sum_{k\in\mathcal{R}} \lambda_k(y-\xi_k)p_{y-\xi_k}(t) - \sum_{k\in\mathcal{R}} \lambda_k(y)p_y(t). \tag{2.2.5}$$

The random time-change formula (2.2.3), the martingale description (2.2.4) and the master equation (2.2.5) provide three different ways of specifying the same model (Kang and Kurtz, 2013; Anderson and Kurtz, 2015).

Results from Reaction Network Theory. In either description of the reaction network, one is interested in determining the qualitative properties of the associated dynamics. It follows from (2.2.2) and (2.2.3) that x_t and X_t are confined to the stoichiometric compatibility classes

$$x_t \in (x_0 + \Psi) \cap \mathbb{R}^d_0, \qquad X_t \in (X_0 + \Psi) \cap \mathbb{N}^d_0, \qquad t \geq 0 \tag{2.2.6}$$

respectively. Therefore, questions of existence and uniqueness of steady states or multistationarity as well as (quasi)stationary distributions are considered within each stoichiometric compatibility class. Similarly, one is interested in the stability of a given equilibrium relative to the stoichiometric compatibility class (Feinberg, 2015). Note that in the case of conservative reaction networks, the solution x_t is bounded, as it is confined within the compact stoichiometric compatibility class. In the stochastic case, the state space becomes finite, which, as we shall see, implies that questions of long-term behavior has been resolved. For this reason, we shall mostly be concerned with non-conservative systems.

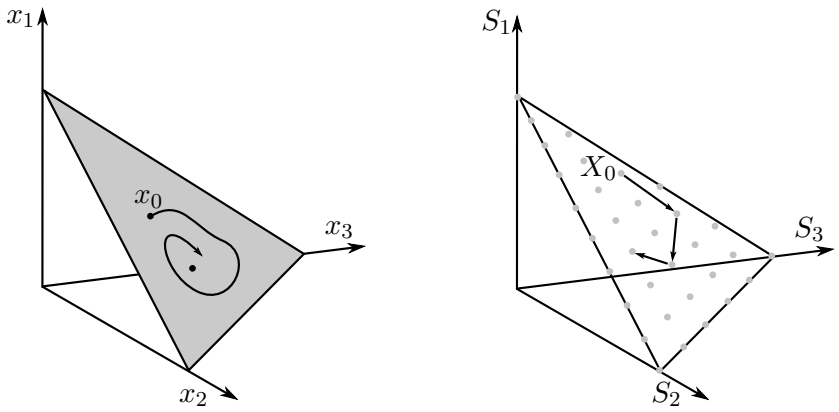

Figure 2.3: A stoichiometric compatibility class in the deterministic (left) and stochastic (right) case. The questions of interest

Central to the analysis of deterministic reaction networks are the following properties of the long-term dynamics, which may be seen as various degrees of stability and thus required for the corresponding stochastic system to display quasi-stationarity (Craciun et al., 2013; Schreiber, 2000).

Definition 2.2.7. A d-dimensional dynamical system is *persistent* if for any forward trajectory $\varphi_t(x) = x(t)$ with positive initial condition $x \in \mathbb{R}^d_0$,

$$\liminf_{t\to\infty} x_i(t) > 0, \qquad \forall 1 \leq i \leq d.$$

It is called *permanent* on $\mathbb{R}^d_0$ if there exists $\epsilon > 0$ such that for any forward trajectory $\varphi_t(x) = x(t)$ with positive initial condition $x \in \mathbb{R}^d_0$,

$$\epsilon < \liminf_{t\to\infty} x_i(t), \qquad \limsup_{t\to\infty} x_i(t) < 1/\epsilon, \qquad \forall 1 \leq i \leq d$$

Thus, persistence means that no forward trajectory with positive initial condition approaches the coordinate axes arbitrarily close. Permanence is a stronger property and means that there exists a compact set $K \subset \mathbb{R}^d_0$ not intersecting $\partial\mathbb{R}^d_0$ such that any forward trajectory with positive initial condition eventually enters K. It is natural to suspect that endotactic mass-action systems are indeed both persistent and permanent, although a formal proof only exists for $d \leq 2$ (Craciun et al., 2013). On the other hand, strongly endotactic networks are known to be both persistent and permanent in general (Gopalkrishnan et al., 2014). Although the strongly endotactic deterministic reaction networks are well behaved, this does not carry over to the corresponding stochastic reaction networks. Indeed, these are not necessarily positive recurrent, and both transitive and explosive examples are known (Anderson et al., 2018a).

To get a flavor for the type of results available in the theory of reaction networks, we first introduce the following central property. An equilibrium point $c \in \mathbb{R}^d_0$ is said to be *complex balanced* if and only if for each complex $z \in \mathcal{C}$, the total inflow equals the total outflow

$$\sum_{k\in\mathcal{R}:\, y'=z} \alpha_k c^{y_k} = \sum_{k\in\mathcal{R}:\, y=z} = \alpha_k c^{y_k}. \tag{2.2.8}$$

This condition may be traced back to Boltzmann. Note that a necessary condition for having complex balance is weak reversibility of the network. A mass-action system with complex balancing states is also called a *toric dynamical system*. We may now state the following classical theorem (Horn and Jackson, 1972; Feinberg, 1979).

Theorem 2.2.9 (Deficiency Zero Theorem)**.** Consider a mass action system $(\mathcal{N}, \alpha)$ for which the underlying reaction network has deficiency zero.

(i) If the network is weakly reversible, then for all choices of rate constants, the system has exactly one equilibrium concentration in each positive stoichiometric compatibility class and that equilibrium concentration is locally asymptotically stable.

(ii) If the network is not weakly reversible then no equilibrium nor limit cycles intersecting the positive orthant exist.

Thus, from the reaction graph alone, one is capable of inferring strong implications for the dynamics. In fact, more is believed to hold. The *Global Attractor Conjecture* states that the unique positive steady state in every stoichiometric compatibility class of a toric dynamical system is the *global* attractor for its stoichiometric compatibility class (Craciun, 2015). The proof of the deficiency zero theorem is based on the Lyapunov function

$$V(x) = \sum_{i \in \mathcal{S}} x_i(\ln(x_i) - \ln(c_i) - 1) + c_i,$$

with $c \in \mathbb{R}^d_0$ being a complex balanced equilibrium. This function originates from classical thermodynamics, where it is known as the Helmholtz or Gibbs free energy function (Anderson et al., 2015). Although V is strictly convex in the interior of the first orthant, it is bounded along $\partial\mathbb{R}^d_0$, which is what causes problems in concluding global attraction. Naturally, one would like to obtain other Lyapunov functions, getting rid of the restriction of complex balanced. This is an active area of research (Ke et al., 2017; Al-Radhawi and Angeli, 2016). Results in the same vein exist for higher deficiency systems. We include the following (Feinberg, 1995).

Theorem 2.2.10 (Deficiency One Theorem)**.** Consider a mass-action system, $(\mathcal{N}, \lambda)$, with deficiency δ and ℓ linkage classes. Assume that it satisfies

(i) $\delta_j \in \{0, 1\}$, $1 \leq j \leq \ell$

(ii) $\delta = \sum_{j=1}^{\ell} \delta_j$

(iii) each linkage class contains a single terminal strong linkage class.

Then, if there exists a positive equilibrium, each positive stoichiometric compatibility class contains exactly one equilibrium. Furthermore, if the network is weakly reversible, it has a positive equilibrium.

There are few general results available about the form of even the stationary distribution in a general reaction network. Exceptions are the class of monomolecular reaction networks (Jahnke and Huisinga, 2007) and more remarkably the following (Anderson et al., 2010).

Theorem 2.2.11. Let $(\mathcal{N}, \lambda)$ be a reaction system endowed with mass-action kinetics. Suppose that, modeled deterministically, the system is complex balanced with a complex balanced equilibrium $c \in \mathbb{R}^d_0$. Then, the stochastically modeled system has a stationary distribution π_I on each irreducible component $I \subseteq \mathbb{N}^d_0$ given by

$$\pi_I(x) = \frac{1}{Z_I} \prod_{i=1}^{d} \frac{c_i^{x_i}}{x_i!} e^{-c_i}, \qquad x \in I, \tag{2.2.12}$$

and $\pi_I(x) = 0$ otherwise, where Z_I is a positive normalizing constant.

As we shall see, in the setting of quasi-stationarity, a reaction network can never be complex balanced due to eventual absorption. Needless to say, even fewer results for quasi-stationary distributions exist. However, one particular class of systems which has been of interest in the intersection of quasi-stationarity and reaction network theory is the following (Anderson et al., 2014).

Definition 2.2.13. Consider a deterministic mass-action system $(\mathcal{N}, \lambda)$. The system is *Absolute Concentration Robust* (ACR) if there exists an index $1 \leq i \leq d$ and a real number $u \in (0, \infty)$ such that all $x > 0$ with $F(x) = 0$ satisfy $x_i = u$. In this case, the ith species is called an ACR species.

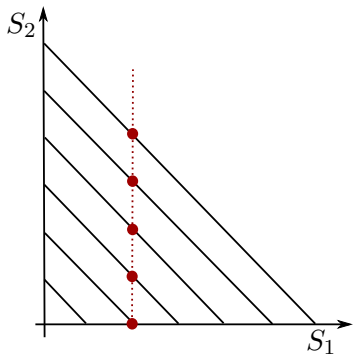

Figure 2.4: The reaction network $S_1 + S_2 \to 2S_2$, $S_2 \to S_1$ exhibiting ACR in species S_1.

Absolute concentration robustness is interesting from a biological perspective, as it states that whenever the system is at a positive equilibrium (and there could be many such positive equilibria), some special species are always expressed at the same level, and are therefore "robust" to environmental changes (Anderson and Cappelletti, 2018). One may determine ACR purely from the network structure. The following theorem is due to Shinar and Feinberg (2010).

Theorem 2.2.14. Let a deterministic mass action system $(\mathcal{N}, \lambda)$ with deficiency $\delta = 1$ be given and suppose there exists at least one positive equilibrium and that there exists two non-terminal complexes $y \neq y'$ such that only the ith entry of $y' - y$ is different from 0. Then the ith species is ACR.

In the stochastic description, this property is connected to extinction events. More precisely we have the following theorem (Anderson et al., 2014; Anderson and Cappelletti, 2018).

Theorem 2.2.15. Let a stochastic mass action system $(\mathcal{N}, \lambda)$ be given and suppose there exists at least one positive equilibrium, the deficiency $\delta = 1$ and there exists two non-terminal complexes $y \neq y'$ such that only the ith entry of $y' - y$ is different from 0. Furthermore, assume that the network is conservative. Then the associated stochastic process $(X_t \colon t \geq 0)$ goes extinct almost surely. In particular, with probability 1, the process enters a closed set $A \subseteq \mathbb{N}_0^d$ such that $\lambda_k(x) = 0$ for all $x \in A$ and all $y_k \to y_k' \in \mathcal{R}$ with y_k being a non-terminal complex.

It has recently been proven (Anderson and Cappelletti, 2018) that one cannot remove any of the assumptions made in Theorem 2.2.15. In particular, this guarantee of extinction is confined to the case of *conservative* systems, which, as we shall see, is less interesting for quasi-stationary distributions. Only few other results connecting graphical conditions to extinction events in reaction networks exist (Johnston et al., 2017).

2.2.2 Connection Through Classical Scaling

There is a strong link between the stochastic and deterministic models, as would be expected since they both describe the same underlying phenomenon. This relationship was first described in the seminal papers by Kurtz (1970, 1972, 1978).

As this construction is essential to the development in (Hansen and Schreiber, 2018), we shall provide the details. Let ϵ denote some scaling parameter, say the inverse of the volume in which the system evolves. Since the concentration is given by the copy-number of the species divided by the volume of the system, we may track the concentrations of the system by considering the scaled process $X_t^\epsilon = \epsilon X_t$. Consequently, using (2.2.3) we arrive at the stochastic equation for the concentrations

$$X_t^\epsilon = X_0^\epsilon + \sum_{k\in\mathcal{R}} \epsilon Y_k \left(\epsilon^{-1}\int_0^t \lambda_k(X_s^\epsilon)\,ds\right)\xi_k.$$

Recall from the deterministic setting the vector field $F(x) = \sum_{k\in\mathcal{R}} \lambda_k(x)\xi_k$ and let $\tilde{Y}_k(u) = Y_k(u) - u$ be the centered process. Then it follows that

$$\begin{aligned} X_t^\epsilon &= X_0^\epsilon + \sum_{k\in\mathcal{R}} \epsilon\tilde{Y}_k \left(\epsilon^{-1}\int_0^t \lambda_k(X_s^\epsilon)\,ds\right)\xi_k + \int_0^t \sum_{k\in\mathcal{R}} \lambda_k(X_s^\epsilon)\xi_k\,ds \\ &= X_0^\epsilon + M_t^\epsilon + \int_0^t F(X_s^\epsilon)\,ds, \end{aligned} \tag{2.2.16}$$

where M_t^ϵ defined as the middle term is a martingale, cf. Watanabe's characterization of Poisson processes (Watanabe, 1964). Using $\dot{x}(t) = F(x(t))$ we find upon integration that

$$x(t) = x(0) + \int_0^t F(x(s))\,ds. \tag{2.2.17}$$

Here we have implicitly assumed global existence of this solution. Now, suppose we have the local Lipshitz condition

$$|F(x) - F(y)| \le K_c|x-y|, \qquad |x|,|y| \le c \tag{2.2.18}$$

for some constant $c > 0$ and define the first exit time

$$\tau_c^\epsilon = \inf\{t \ge 0 : |X_t^\epsilon| \vee |x(t)| \ge c\}$$

Using in turn (2.2.17), the triangle inequality and the Lipschitz condition (2.2.18), we find

by Gronwall's inequality (Ethier and Kurtz, 1986, p.498),

$$\begin{aligned}
|X^\epsilon_{t\wedge\tau^\epsilon_c} - x(t\wedge\gamma^\epsilon_a)| &= \left| X^\epsilon_0 + M^\epsilon_{t\wedge\tau^\epsilon_c} + \int_0^{t\wedge\gamma^\epsilon_a} F(X^\epsilon_s)\,ds - x(0) - \int_0^{t\wedge\tau^\epsilon_c} F(x(s))\,ds \right| \\
&\le |X^\epsilon_0 - x(0)| + |M^\epsilon_{t\wedge\tau^\epsilon_c}| + \int_0^{t\wedge\tau^\epsilon_c} |F(X^\epsilon_s) - F(x(s))|\,ds \\
&\le |X^\epsilon(0) - x(0)| + |M^\epsilon_{t\wedge\tau^\epsilon_c}| + \int_0^{t\wedge\tau^\epsilon_c} K_c |X^\epsilon_s - x(s)|\,ds \\
&\le \left(|X^\epsilon_0 - x(0)| + \sup_{s\le t\wedge\tau^\epsilon_c} |M^\epsilon_s| \right) e^{K_c t}.
\end{aligned}$$

Our aim is to have the terms in the parenthesis tend to 0 as $\epsilon \to 0$. The first is implied by imposing the assumption, $X^\epsilon_0 \to x(0)$ for $\epsilon \to 0$. To show that $\sup_{s\le t\wedge\tau^\epsilon_c} |M^\epsilon_s| \to 0$ almost surely, as $\epsilon \to 0$, we note that M^ϵ_t is a martingale where, by the Poisson property,

$$\mathbb{E}(|M^\epsilon_{t\wedge\tau^\epsilon_c}|^2) = \epsilon \mathbb{E}\left(\int_0^{t\wedge\tau^\epsilon_c} \sum_{k\in\mathcal{R}} \lambda_k(X^\epsilon_s) |\xi_k|^2 \right). \tag{2.2.19}$$

By Doob's inequality (Ethier and Kurtz, 1986, p. 64) we then find

$$\mathbb{E}\left(\sup_{s\le t} |M^\epsilon_{s\wedge\tau^\epsilon_c}|^2 \right) \le 4\mathbb{E}((M^\epsilon_{t\wedge\tau^\epsilon_c})^2) \le 4\epsilon t \sup_{|x|\le c} \sum_{k\in\mathcal{R}} \lambda_k(x)|\xi_k|^2 \to 0 \tag{2.2.20}$$

for $\epsilon \to 0$. Thus we have convergence in mean square and in particular convergence in probability. We arrive at the following essential theorem (Kurtz, 1972, 1978; Ethier and Kurtz, 1986).

Theorem 2.2.21 (Kurtz)**.** Suppose that for each $c > 0$, the Lipshitz condition

$$|F(x) - F(y)| \le K_c |x - y|, \qquad |x|, |y| \le c$$

holds, and that the solution of the deterministic equation holds for all time. If $X^\epsilon_0 \to x(0)$ as $\epsilon \to 0$ then for each $\delta > 0$ and each $T > 0$,

$$\lim_{\epsilon\to 0} \mathbb{P}\left(\sup_{t\le T} |X^\epsilon_t - x(t)| \ge \delta \right) = 0.$$

One may view this result as a Law of Large Numbers. When it holds, one often refers to $x(t)$ as the *fluid limit* or *thermodynamic limit* of X^ϵ_t. The following figure illustrates the convergence for a three species reaction network.

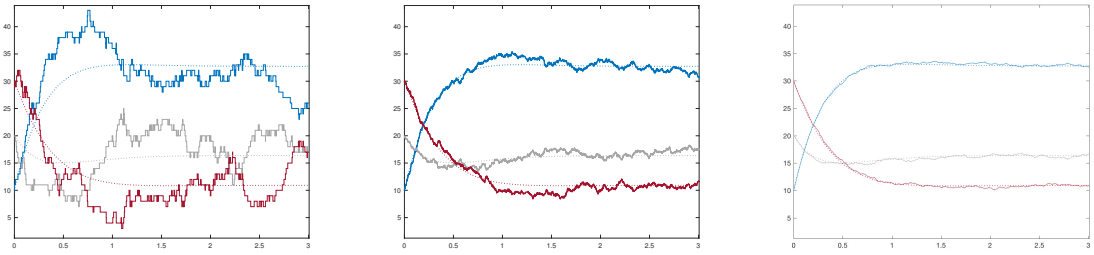

Figure 2.5: Stochastic simulation of the reaction network $S_1 \xrightarrow{1} S_2 \xrightarrow{2} S_3 \xrightarrow{3} S_1$ with $X_0 = (10, 20, 30)$ and $\epsilon = 1, 0.1, 0.01$ (solid). The solution to the ODE included for comparison (dotted).

Returning to (2.3.23), and using that by the functional central limit theorem, $W_k^\epsilon(u) = \tilde{Y}_k(\epsilon^{-1}u)\sqrt{\epsilon}$ converges in distribution to the standard *Brownian motion* W_k as $\epsilon \to 0$ (Øksendal, 2010), we obtain

$$\begin{aligned} X_t^\epsilon &= X_0^\epsilon + \sum_{k\in\mathcal{R}} \epsilon \tilde{Y}_k \left(\epsilon^{-1} \int_0^t \lambda_k(X_s^\epsilon)\, ds \right) \xi_k + \int_0^t \sum_{k\in\mathcal{R}} \lambda_k(X_s^\epsilon)\xi_k\, ds \\ &\approx X_0^\epsilon + \int_0^t \sum_{k\in\mathcal{R}} \lambda_k(X_s^\epsilon)\xi_k\, ds + \sum_{k\in\mathcal{R}} \sqrt{\epsilon} W_k \left(\int_0^t \lambda_k(X_s^\epsilon)\, ds \right) \xi_k, \end{aligned}$$

where W_k are independent Brownian motions. Using that the Itô integral $\int_0^t g(s,\omega)dW(s)$ is distributional equivalent to $W(\int_0^t g(s,\omega)^2\, ,ds)$ we conclude that a good approximation to X_t^ϵ is in differential form given by

$$dZ_t^\epsilon = F(Z_t^\epsilon)\, dt + \sum_{k\in\mathcal{R}} \xi_k \sqrt{\epsilon\lambda_k(Z_t^\epsilon)}\, dW_k$$

This is the *diffusion approximation* or *Langevin equation*. In the same vein as Theorem 2.2.21 one may prove (Ethier and Kurtz, 1986), under further regularity assumptions, that for fixed $T > 0$, there exists a random variable Γ_ϵ^T such that

$$\sup_{t\in[0,T\wedge\tau_\epsilon]} |X_t^\epsilon - Z_t^\epsilon| \le -\Gamma_\epsilon^T \epsilon \log \epsilon,$$

where $\tau_\epsilon = \inf\{t \ge t\colon Z_t^\epsilon \in \partial\mathbb{R}_0^d\}$. Thus, the diffusion approximation is only valid away from the boundary of the orthant. Recently, several suggestions to overcome this issue have been proposed (Bibbona and Sirovich, 2017; Anderson et al., 2018b)

2.2.3 Approximations to the Markov Process

Having the classic scaling, one could be tempted to use the deterministic approach, at least as an approximation, to the behavior of the corresponding stochastic system. However, this naïve approach may fail in many non-trivial ways. Even for very simple reaction networks endowed with mass-action kinetics, the deterministic ordinary differential equations and the

stochastic formulation given by the master equation may predict qualitatively different long term outcomes (Keizer, 1987; Anderson and Cappelletti, 2018). This phenomenon is referred to as *Keizer's paradox*, as it seemingly contradicts Theorem 2.2.21.

Surely, when the copy-numbers of the interacting species are low, the scaling parameter ϵ is large, implying that the system is far from thermodynamic equilibrium. In this non-equilibrium regime, stochastic effects, as illustrated by the leftmost diagram of Figure 2.5, are clearly significant. Surprisingly, even in large systems close to equilibrium, in the long time limit, stochastic fluctuations can still play an important role, and deterministic equilibria may be perturbed in significant ways (Childs and Keener, 2012). In particular, one may find reaction networks having extinction events for any value of $\epsilon > 0$. As ϵ decreases, these become increasingly rare events. Formally, there is no paradox however. Note that Theorem 2.2.21 is a statement about the behavior of the stochastic system over a *compact* time interval. As such, it stands silent on the (very) long time limit, $t \to \infty$.

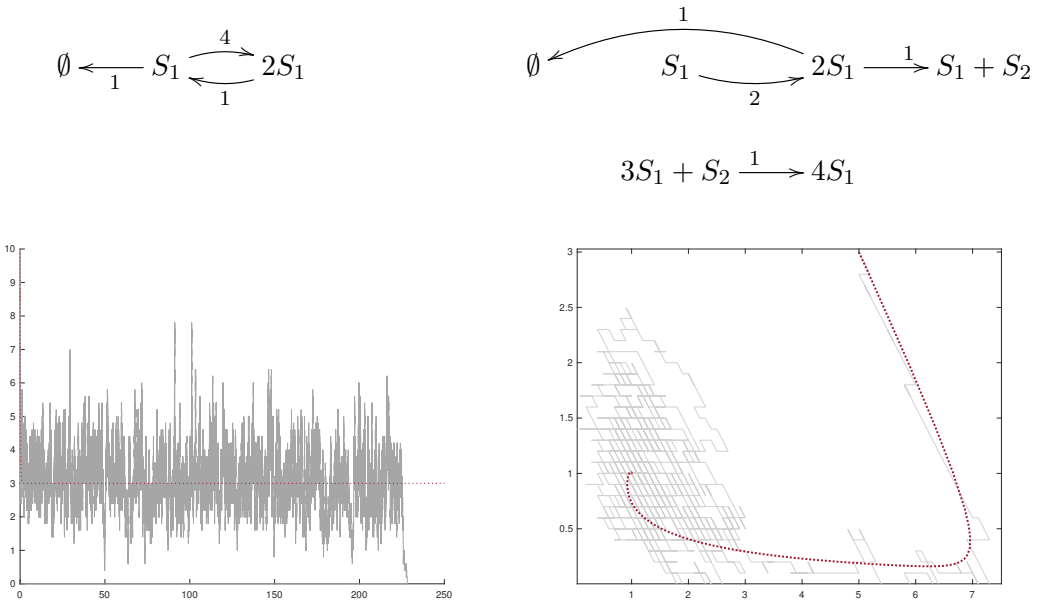

Figure 2.6: Two reaction networks with corresponding sample paths computed by Gillespie's algorithm below.

The diffusion approximation often shows good agreement with the Markov process in the bulk of the positive orthant, even for moderate values of ϵ (Beccuti et al., 2018). However, it may also fail to capture the exact stochastic dynamics of the reaction network, although for more subtle reasons. Indeed, it was shown in (Pakdaman et al., 2010) that one may choose parameters in the network $\emptyset \leftarrow S \rightarrow 2S$ such that the escape points from a domain depends on whether one studies X_t^ϵ or its diffusion approximation Z_t^ϵ. Further, it is known that the escape times from a domain differs asymptotically (Pollett, 2001). As we shall be concerned with such rare events, we will in the study of quasi-stationary distributions only examine the exact stochastic equation.

2.3 Quasi-Stationary Distributions

Consider a right-continuous time-homogenous Markov process $(X_t \colon t \geq 0)$ (Rogers and Williams, 2000), that evolves in a domain D, where there is a set of absorbing states, a "trap", $A \subset D$. The process is absorbed, also referred to as being killed, when it hits the absorbing set, meaning $X_t \in A$ for all $t \geq \tau_A$, where $\tau_A = \inf\{t \geq 0 : X_t \in A\}$ is the hitting time of A. We refer to the complement,

$$E := D \backslash A,$$

as the set of endorsed states. For any probability distribution on E, we let $\mathbb{P}_\mu$ and $\mathbb{E}_\mu$ be the probability and expectation respectively, associated with the process $(X_t \colon t \geq 0)$, initially distributed with respect to μ. For any $x \in E$, we let $\mathbb{P}_x = \mathbb{P}_{\delta_x}$ and $\mathbb{E}_x = \mathbb{E}_{\delta_x}$. Assuming the process hits the absorbing set a.s., that is $\mathbb{P}_x(\tau_A < \infty) = 1$ for all $x \in E$, we investigate the behavior of the process before being absorbed (Collet et al., 2013).

Definition 2.3.1. A probability measure ν on E is called a *quasi-stationary distribution* (QSD) for the process $(X_t \colon t \geq 0)$ absorbed at A, if for every measurable set $B \subseteq E$

$$\mathbb{P}_\nu(X_t \in B \,|\, t < \tau_A) = \nu(B), \qquad t \geq 0,$$

Note that if the absorbing set A is empty, this reduces to the standard definition of a stationary measure. However, when A is non-empty, the possibility for the network of being weakly reversible is ruled out. In turn, the reaction networks of interest to us will initially not be complex balanced, thus much of the known machinery from reaction network theory is not readily available to us.

Definition 2.3.2. A probability measure, ν, on E is a *quasi-limiting distribution* (QLD) for $(X_t \colon t \geq 0)$ if there exists a probability measure μ on E such that for any measurable set $B \subseteq E$,

$$\lim_{t\to\infty} \mathbb{P}_\mu(X_t \in B \,|\, t < \tau_A) = \nu(B)$$

The QLD is sometimes referred to as a limiting conditional distribution (van Doorn and Pollett, 2013). We may advantageously consider the QSD as the long time behavior of the conditioned distribution. Indeed it is a classical result of (Vere-Jones, 1969) that the two terms are equivalent for countable state spaces. A more recent proof, which we include for completeness, extends the equivalence to the general case (Méléard and Villemonais, 2012).

Proposition 2.3.3. *A probability measure ν on E is a QLD if and only if it is a QSD.*

Proof. If ν is a QSD for $(X_t \colon t \geq 0)$ then in particular, it is a QLD with initial distribution ν. To prove the reverse direction, assume ν is a QLD for $(X_t \colon t \geq 0)$ with initial distribution μ. For any measurable and bounded function f on E,

$$\int_E f(x)\nu(dx) = \lim_{t\to\infty} \mathbb{E}_\mu(f(X_t) \,|\, t < \tau_A) = \lim_{t\to\infty} \frac{\mathbb{E}_\mu(f(X_t), t < \tau_A)}{\mathbb{P}_\mu(t < \tau_A)} \tag{2.3.4}$$

Letting $f(x) = \mathbb{P}_x(s < \tau_A)$, it follows from (2.3.4) and the Markov property that

$$\mathbb{P}_\nu(s < \tau_A) = \lim_{t\to\infty} \frac{\mathbb{E}_\mu(\mathbb{P}_{X_t}(s < \tau_A), t < \tau_A)}{\mathbb{P}_\mu(t < \tau_A)} = \lim_{t\to\infty} \frac{\mathbb{P}_\mu(t + s < \tau_A)}{\mathbb{P}_\mu(t < \tau_A)}. \tag{2.3.5}$$

Now, letting $f(x) = \mathbb{P}_x(X_s \in B, s < \tau_A)$ it follows from (2.3.4) and the Markov property again that

$$\begin{aligned}
\mathbb{P}_\nu(X_s \in B, s < \tau_A) &= \lim_{t\to\infty} \frac{\mathbb{P}_\mu(X_{t+s} \in B, t + s < \tau_A)}{\mathbb{P}_\mu(t < \tau_A)} \\
&= \lim_{t\to\infty} \frac{\mathbb{P}_\mu(X_{t+s} \in B, t + s < \tau_A)}{\mathbb{P}_\mu(t + s < \tau_A)} \frac{\mathbb{P}_\mu(t + s < \tau_A)}{\mathbb{P}_\mu(t < \tau_A)} \\
&= \nu(B)\mathbb{P}_\nu(s < \tau_A),
\end{aligned}$$

where we in the last equality have used the property of ν being a QLD and (2.3.5) for the two factors respectively. Dividing by $\mathbb{P}_\nu(s < \tau_A)$ we obtain

$$\mathbb{P}_\nu(X_s \in B \,|\, s < \tau_A) = \nu(B),$$

thus ν is a QSD as desired. □

Definition 2.3.6. The process $(X_t \colon t \geq 0)$ has a *Yaglom limit* if there exists a probability measure ν on E such that for any $x \in E$ and any measurable set $B \subseteq E$,

$$\lim_{t\to\infty} \mathbb{P}_x(X_t \in B \,|\, t < \tau_A) = \nu(B). \tag{2.3.7}$$

Note that a Yaglom limit is in particular a QLD, and thus a QSD. The reverse is not generally the case. Indeed, starting from different initial distributions δ_{x_1} and δ_{x_2} may lead to different quasi-limiting distributions. This raises the question of the *domain of attraction* for a quasi-stationary distribution, ν, which we define as follows

$$\mathcal{D}(\nu) = \left\{ \mu \in \mathcal{P}(E) \colon \lim_{t\to\infty} \mathbb{P}_\mu(X_t \in \cdot \,|\, t < \tau_A) = \nu(\cdot) \right\}. \tag{2.3.8}$$

Thus, the Yaglom limit, when it exists, is the unique quasi-stationary distribution whose domain of attraction contains $\{\delta_x \colon x \in E\}$ (Villemonais, 2015). In general, the rate of convergence to a QSD, ν, may depend greatly upon its domain of attraction $\mathcal{D}(\nu)$.

2.3.1 Countably Infinite State Space

The stochastic processes associated to reaction networks will, unless they are conservative, take place on countably infinite state spaces. These will therefore be our focus. Furthermore the questions of existence and uniqueness of a QSD on a finite state space is well known (Collet et al., 2013, Chapter 3), and a complete descriptions of the QSD is given by the normalized left Perron-Frobenius eigenvector of the transition rates matrix restricted to E. This characterization will be of importance to us only for computational purposes where a truncation of the full state space is required. For the infinite dimensional case, most work has

been carried out for birth-death processes in one dimension (van Doorn and Pollett, 2013), where classification results yielding information about the set of QSDs exist (van Doorn, 1991). Indeed, as we shall see, a birth-death process can have no quasi-stationary distribution, one unique quasi-stationary distribution or a whole continuum of quasi-stationary distributions.

To set notation, we consider a continuous time Markov process $(X_t \colon t \geq 0)$ on a state space $D = A \sqcup E$ where A is absorbing and E is a countably infinite set of transient states. Without loss of generality, we may make the identification $D = \{0\} \sqcup \mathbb{N}$ by collapsing A to a single state $\{0\}$. The jump-rate matrix is then given by

$$\tilde{Q} = \begin{pmatrix} 0 & 0' \\ a & Q \end{pmatrix} \tag{2.3.9}$$

where we assume that $a = -Q \cdot 1 \geq 0$, $a \neq 0$, such that absorption is possible. We have the following fundamental theorem in the QSD literature (Méléard and Villemonais, 2012; van Doorn and Pollett, 2013; Collet et al., 2013).

Theorem 2.3.10. A probability measure ν is a QSD if and only if

$$\nu' Q = -\theta(\nu)\nu', \qquad \theta(\nu) = \sum_{x \in E} \nu(x) a(x) > 0. \tag{2.3.11}$$

Thus, to find the QSD, one has to solve this non-linear eigenvalue problem. Unfortunately, for countably infinite state spaces, this is in general a very hard problem. In fact, for stochastic reaction networks, the shape of a corresponding QSD is only known for a small handful of the most trivial cases.

The Decay Parameter and Minimal QSD. By means of the following proposition, one may interpret $\theta(\nu)$ as the exponential rate of survival (Méléard and Villemonais, 2012).

Theorem 2.3.12. If ν is a QSD then

$$\mathbb{P}_\nu(\tau_A > t) = e^{-\theta(\nu)t}, \qquad \forall t \geq 0. \tag{2.3.13}$$

Hence, starting from ν, the time to absorption τ_A is exponentially distributed with parameter $\theta(\nu) \in (0.\infty)$.

It follows from the form of the moment generating function of the exponential distribution that if ν is a QSD then for any $0 < \theta < \theta(\nu)$,

$$\mathbb{E}_\nu(e^{\theta \tau_A}) = \frac{\theta(\nu)}{\theta(\nu) - \theta} < \infty. \tag{2.3.14}$$

In particular, ν-a.s. in x, $\mathbb{E}_x(e^{\theta \tau_A}) < \infty$. This suggests that if the population can escape extinction for too long times with positive probability, then the process has no QSD.

Assume for simplicity that E is irreducible. As the reaction vectors in a reaction network are finite, the set of states from which one may jump to the absorbing set is finite as well. Thus, the exponential rate of survival equals the exponential rate of transition probabilities

(Collet et al., 2013), and we shall refer to this quantity simply as the *decay parameter* of $(X_t : t \geq 0)$. It characterizes the time-scale on which absorption takes place and is given by

$$\theta^* = \sup\left\{\theta \geq 0 : \mathbb{E}_x(e^{\theta\tau_A}) < \infty\right\} = \liminf_{t\to\infty} -\frac{1}{t}\log \mathbb{P}_x(\tau_A > t) = \lim_{t\to\infty} -\frac{1}{t}\log \mathbb{P}_x(X_t = y),$$

independent of $x, y \in E$. A necessary condition for the existence of a QSD is thus that $\theta^* > 0$, and any QSD, ν satisfies $0 < \theta(\nu) \leq \theta^*$. A QSD satisfying $\theta(\nu) = \theta^*$ is called an *extremal* or *minimal* QSD.

Few sufficient conditions for the existence of a QSD exist, and these are rather intangible for general reaction networks. In the case where absorption happens almost surely, and the process satisfies the condition of *asymptotic remoteness*, that is

$$\lim_{x\to\infty} \mathbb{P}_x(\tau_A \leq t) = 0, \qquad \forall t > 0 \tag{2.3.15}$$

then $\theta^* > 0$ is a necessary and sufficient condition for the existence of a QSD (Ferrari et al., 1995). In general, as we shall discover in the next section for birth-death processes, obtaining information about θ^* is a delicate matter. However, finding sufficient conditions in terms of the rate matrix Q is still largely an open problem. One exception is when Q is *uniformizable*, that is the rates satisfy $-q_{ii} = a_i + \sum_{j\in E, j\neq i} q_{ij} \leq C$, for all $i \in E$ where C is some constant and furthermore *skip-free to the left*, that is $q_{ij} = 0$ for $j < i - 1$. Then the existence of a QSD is guaranteed (van Doorn and Pollett, 2013). Unfortunately, uniformizability does not occur for reaction networks, and being skip-free to the left greatly narrows the class of such networks. Curiously, there exists more tangible conditions ensuring both existence *and* uniqueness (Martinez et al., 2014; Champagnat and Villemonais, 2016, 2017). In particular, we have the following theorem (Champagnat and Villemonais, 2016).

Theorem 2.3.16. There exists a probability measure ν on E such that

(A1) there exists $t_0, c_1 > 0$ such that for all $x \in E$,

$$\mathbb{P}_x(X_{t_0} \in \cdot \,|\, t_0 < \tau_A) \geq c_1\nu(\cdot),$$

(A2) there exists $c_2 > 0$ such that for all $x \in E$ and $t \geq 0$,

$$\mathbb{P}_\nu(t < \tau_A) \geq c_2\mathbb{P}_x(t < \tau_A),$$

if and only if there exists a probability measure ν on E and two constants $C, \gamma > 0$ such that, for all initial distributions μ on E,

$$\|\mathbb{P}_\mu(X_t \in \cdot \,|\, t < \tau_A) - \nu(\cdot)\|_{TV} \leq Ce^{-\gamma t}, \qquad \forall t \geq 0.$$

Condition (A1) says, in the case where $\nu = \mathbb{1}_K$ for some compact set $K \subset E$, that the process must come down from infinity and go away from 0 conditionally on non-absorption. (A2) is a Harnack type inequality stating that no matter where the process starts, it will behave essentially as starting from ν. We shall exploit this results in the paper Hansen and Wiuf (2018a).

2.3.2 Birth-Death Processes

When the associated stochastic process of a reaction network can be characterized as a birth-death process, that is, each reaction vector ξ_k, $k \in \mathcal{R}$ is a one-step jump in the sense that $\xi_k = \pm e_i$ for some $i = 1, \dots, d$, the QSD literature is much richer (Cavender, 1978; van Doorn, 1991; Champagnat and Villemonais, 2015; Villemonais, 2015). We shall use this as a guiding example of the type of behavior one may expect for general reaction networks.

In one dimension, the description of QSDs is best understood, and this shall serve as a guideline for the type of behavior one may observe. Furthermore, we hope to illustrate how the angle of attack hinges fundamentally on the birth-death structure, making a generalization to larger classes of networks a difficult task. Thus, suppose the associated stochastic process $(X_t : t \geq 0)$ is a one-dimensional birth-death process taking values in $D = \mathbb{N}_0$. Without loss of generality, we may take the absorbing set to be $A = \{0\}$ and $E = \mathbb{N}$. Its jump rate matrix Q may then be characterized by $q(i, i-1) = \mu_i$ and $q(i, i+1) = \lambda_i$, $q(i,i) = -(\lambda_i + \mu_i)$ and all other coefficients vanish.

$$\emptyset \xleftarrow{\alpha_1} S \xrightarrow{\alpha_2} 2S \underset{\alpha_4}{\overset{\alpha_3}{\rightleftarrows}} 3S \xleftarrow{\alpha_5} 4S$$

The study of birth-death processes relevant to our presentation goes back to the work of Karlin and McGregor (1957). Define the sequence $\sigma = (\sigma_x)_{x \in \mathbb{N}}$ where

$$\sigma_1 = 1, \qquad \sigma_x = \prod_{i=1}^{x-1} \frac{\lambda_i}{\mu_{i+1}}, \qquad x > 1. \tag{2.3.17}$$

One may completely describe when the birth-death process goes extinct almost surely through consideration of a tangible series (Anderson, 1991; Méléard and Villemonais, 2012, Prop. 12)

Lemma 2.3.18. *The birth-death process is absorbed almost surely, that is $\mathbb{P}_x(\tau_A < \infty) = 1$ for all $x \in \mathbb{N}$, if and only if*

$$\sum_{x \geq 1} \frac{1}{\lambda_x \sigma_x} = \sum_{x=1}^{\infty} \frac{\mu_1 \dots \mu_x}{\lambda_1 \dots \lambda_x} = \infty. \tag{2.3.19}$$

From the characterization in Theorem 2.3.10, it follows that $\theta(\nu) = \mu_1 \nu_1$, thus, in the setting of birth-death processes, one may obtain the following lemma.

Proposition 2.3.20. *The sequence $(\nu_x)_{x \in \mathbb{N}}$ is a QSD if and only if*

1. $\nu_x \geq 0$, $\forall x \geq 1$ *and* $\sum_{x \geq 1} \nu_x = 1$
2. $\forall x \geq 1$,

$$\begin{aligned} \lambda_{x-1}\nu_{x-1} - (\lambda_x + \mu_x)\nu_x + \mu_{x+1}\nu_{x+1} &= -\mu_1 \nu_1 \nu_i \\ -(\lambda_1 + \mu_1)\nu_1 + \mu_2 \nu_2 &= -\mu_1 \nu_1^2 \end{aligned} \tag{2.3.21}$$

Polynomials. We now study the set of sequences which fulfills the above criteria. Let the polynomials $(\psi_i(\theta) : i \geq 0)$ be defined by $\psi_0 \equiv 0$, $\psi_1 \equiv 1$ and the following recurrence relation

$$\lambda_i \psi_{i+1}(\theta) - (\lambda_i + \mu_i - \theta)\psi_i(\theta) + \mu_i \psi_{i-1}(\theta) = 0, \qquad i \in \mathbb{N}. \tag{2.3.22}$$

By definition the polynomials verify $Q\psi(\theta) = -\theta\psi(\theta)$. These polynomials have been well studied and provide a spectral representation of the transition probabilities (Karlin and McGregor, 1957; van Doorn, 1991). Letting $\theta_{n,x}$ denote the nth root of the polynomial $\psi_x(\theta)$, we may write

$$\zeta_1 = \lim_{x\to\infty} \theta_{1,x} = \sup\{\theta \geq 0 \colon \psi_x(\theta) > 0 \, \forall x \geq 1\}. \tag{2.3.23}$$

It has been shown that, in fact, $\theta^* = \zeta_1$ (Collet et al., 2013). Looking at the spectral properties of the semigroup (P_t) and the polynomials ψ_x yields a tractable necessary and sufficient condition for the existence of a QSD based on the birth and death rates only. Indeed, define the quantity

$$D := \sup_{x\geq 1} \mathbb{E}_x(\tau_A) = \sum_{x=1}^{\infty} \frac{1}{\lambda_x \sigma_x} \sum_{i=x+1}^{\infty} \sigma_i.$$

where the last equality is only possible as we are concerned with birth-death processes. When D is finite, this implies that the process comes down from infinity in finite time, which is closely related to the existence and uniqueness of a QSD (Martinez et al., 2014). We then have the following classification theorem (van Doorn, 1991, Theorems 3.2, 4.1).

Theorem 2.3.24. We have the convergence to the Yaglom limit

$$\lim_{t\to\infty} \mathbb{P}_i(X_t = j \,|\, \tau_A > t) = \frac{1}{\mu_1} \sigma_j \theta^* \psi_j(\theta^*), \tag{2.3.25}$$

and the classification

(i) If $\theta^* = 0$, there is no QSD.

(ii) If D converges, then $\theta^* > 0$ and the Yaglom limit is the unique QSD.

(iii) If D diverges and $\theta^* > 0$ then there is a continuum of QSDs, given by the one parameter family $(\hat{\nu}_x(\theta))_{0<\theta\leq\theta^*}$

$$\hat{\nu}_x(\theta) = \frac{1}{\mu_1} \sigma_x \theta \psi_x(\theta)$$

Note that in the last case, the Yaglom limit is the extremal distribution, explaining this terminology for θ^*. One may even compare the QSDs in the case where there are infinitely many. Given two probability measures, μ_1, μ_2 on $\mathbb{N}$ we say that μ_1 is stochastically smaller than μ_2, written $\mu_1 \preceq \mu_2$, if for all $j \in \mathbb{N}$, $\sum_{i=1}^{j} \mu_1(i) \geq \sum_{i=1}^{j} \mu_2(i)$. We refer to (Cavender, 1978; Collet et al., 2013, Proposition 5.16) for a proof of the following proposition.

Proposition 2.3.26. *Let $\theta_1, \theta_2 \in (0, \theta^*]$ with $\theta_1 < \theta_2$. Then $\nu_{\theta_2} \preceq \nu_{\theta_1}$.*

In this sense, the extremal QSD $\nu(\theta^*)$ is indeed minimal, explaining the common notion for Yaglom limits as the minimal quasi-stationary distribution. The classification of Theorem 2.3.24 further provides a way to compute the QSD. Indeed, if we can find the polynomials $(\psi_x)_{x\in\mathbb{N}}$ and then from these $\theta^* = \zeta_1$, we may provide a calculation of the QSD. If fact, we may write the polynomials explicitly, rather than inductively, in the following way.

Proposition 2.3.27. *Define the polynomials*

$$h_1(\theta) = -\mu_2 + (\lambda_2 + \mu_2 - \theta)\left(1 + \frac{\mu_1}{\lambda_1} - \frac{\theta}{\lambda_1}\right),$$
$$h_2(\theta) = -\mu_3\left(1 + \frac{\mu_1}{\lambda_1} - \frac{\theta}{\lambda_1}\right),$$

and the matrix

$$M(\theta) = \begin{pmatrix} h_1(\theta) & \lambda_2 & 0 & 0 & 0 & \dots \\ h_2(\theta) & -(\lambda_3+\mu_3-\theta) & \lambda_3 & 0 & 0 & \dots \\ 0 & \mu_4 & -(\lambda_4+\mu_4-\theta) & \lambda_4 & 0 & \dots \\ 0 & 0 & \mu_5 & -(\lambda_5+\mu_5-\theta) & \lambda_5 & \\ & & \ddots & \ddots & \ddots & \ddots \end{pmatrix}.$$

Then we have the following explicit expression for the polynomials

$$\psi_{2+i}(\theta) = \frac{(-1)^{i-1}}{\prod_{k=1}^{i}\lambda_{k+1}} \det M_i(\theta), \qquad i \geq 1, \tag{2.3.28}$$

where $M_i(\theta)$ is the matrix $M(\theta)$ restricted to the first i rows and i columns.

Proof. We apply (Kittappa, 1993, Theorem 2). In this notation, we find from comparing (2.3.22) with (Kittappa, 1993, (2a)) that $n = 2$ and

$$q(i, i+2) = \lambda_{i+1}, \qquad q(i, i+1) = -(\lambda_{i+1} + \mu_{i+1} - \theta), \qquad q(i,i) = \mu_{i+1},$$

and the initial values $c_1 = \psi_1(\theta) = 1$, $c_2 = \psi_2(\theta) = 1 + \frac{\mu_1}{\lambda_1} - \frac{\theta}{\lambda_1}$. Further, $g(k) = 0$ for all $k \in \mathbb{N}$. Then as

$$h(k) = g(k) - \sum_{i=0}^{n-k} q(k, i+k)c_{i+k},$$

we obtain the expressions

$$h(2) = -\sum_{i=0}^{0} q(2, i+2)c_{i+2} = -q(2,2)c_2 = -\mu_3\left(1 + \frac{\mu_1}{\lambda_1} - \frac{\theta}{\lambda_1}\right),$$

$$h(1) = -\sum_{i=0}^{1} q(1, i+1)c_{i+1} = -q(1,1)c_1 - q(1,2)c_2 = -\mu_2 + (\lambda_2 + \mu_2 - \theta)\left(1 + \frac{\mu_1}{\lambda_1} - \frac{\theta}{\lambda_1}\right).$$

The result now follows directly from (Kittappa, 1993, Theorem 2). □

We may now implement a suitable numerical algorithm to find the first root of the polynomials $(\psi_i)_i$ and see whether it converges, thus giving us θ^*. If indeed $\theta^* > 0$ we may obtain the Yaglom limit by Theorem 2.3.24.

2.3.3 A Dynamical Systems Perspective

Given that a QSD exists, one may not necessarily observe any quasi-stationary behavior. From a dynamical viewpoint, this behavior is characterized by having a fast convergence to the quasi-limiting distribution followed by a very slow convergence to extinction. In other words, if given an initial distribution μ on $\mathbb{N}_0$, we let $p' = (p_0, p_1, \dots)$ denote the total probabilities,

$$p_j(t) = \mathbb{P}_\mu(X_t = j), \qquad j \geq 0, t \geq 0,$$

then one needs to have a separation of the time-scale on which the conditional probabilities

$$\nu_j(t) := \mathbb{P}_\mu(X_t = j \,|\, t < \tau_A),$$

converge to the QLD, ν_j, and the time-scale on with absorption occurs. The last is characterized through θ^* while the former may be connected to a generalization of the notion of the spectral gap, $\sigma(Q)$, with the aid of operator theory. Unfortunately, not much is known about this object for general absorbing Markov processes. However, other methods considering a truncation of Q to a finite space and then taking the limit exist (Childs and Keener, 2012).

Note that $p' = (p_0, p_1, \dots)$ satisfies the Kolmogorov forward equation

$$\frac{dp'}{dt} = p'\tilde{Q}, \qquad p(0) = \mu.$$

Further, we have the following representation which is often found in earlier or more applied work on quasi-stationary distributions

$$\nu_j(t) = \frac{\mathbb{P}_\mu(X_t = j, t < \tau_A)}{\mathbb{P}_\mu(t < \tau_A)} = \frac{p_j(t)}{1 - p_0(t)}, \quad j > 0, t \geq 0.$$

It then follows that $\nu(t)$ converges to a QLD, (and thus a QSD), ν, for $t \to \infty$, where $\nu' = (\nu_1, \nu_2, \dots)$. Following Pollett and Roberts (1990), substituting this into the forward equation we find

$$\frac{dp_0(t)}{dt} = \nu'(t)a(1 - p_0(t)),$$

and, letting $p_E = (p_1, p_2, \dots)$ one obtains

$$\frac{d\nu'}{dt}(1-p_0(t)) - \nu'(t)\frac{dp_0}{dt} = \frac{dp'_E}{dt} = p'_E Q = \nu'(t)Q(1-p_0(t)).$$

Thus we find the non-linear Riccati system of differential equations for $\nu(t)$,

$$\frac{d\nu'}{dt} = \nu' Q + (\nu' a)\nu'. \tag{2.3.29}$$

One should therefore expect that potential QLDs are stationary points for (2.3.29). As QLDs are QSDs, we rediscover the condition for being a QSD,

$$\nu' Q = -(\nu' a)\nu' = -\theta(\nu)\nu'.$$

We shall exploit this dynamical view on the QSD in the manuscript Hansen and Wiuf (2018b), by considering a slow manifold.

2.3.4 The Inverse Problem

Suppose we have a given QSD, ν, and wish to determine whether there exists a reaction network whose associated stochastic process admits ν as a QSD. Following (O'neill, 2007), when $(X_t \colon t \geq 0)$ is a birth-death process, then (2.3.21) is satisfied for some birth and death rates, which we wish to determine.

Lemma 2.3.30. *For $x \geq 2$*

$$\mu_x \nu_x = \lambda_{x-1}\nu_{x-1} + \mu_1 \nu_1 \left(1 - \sum_{i=1}^{x-1} \nu_i\right) \tag{2.3.31}$$

To illustrate how Lemma 2.3.30 may be applied, suppose we wish to examine whether there exists a reaction network whose QSD is the geometric distribution, $\nu_x = (1-p)p^{x-1}$ with parameter $p \in (0,1)$. Inserting in (2.3.31) we obtain

$$\mu_x(1-p)p^{x-1} = \lambda_{x-1}(1-p)p^{x-2} + \mu_1(1-p)p^{x-1}.$$

Dividing through, we arrive at

$$\mu_x = \lambda_{x-1}p^{-1} + \mu_1$$

As the geometric distribution is supported on $\mathbb{N}$, there must be a reaction with S as source complex. Letting $\lambda_x = \alpha_2 x$ and using that all rates must be a function of x, we arrive at $p = \alpha_2/\mu_1$, and thus in turn $\mu_x = \mu_1 x$. Letting $\alpha_1 = \mu_1$ we find the linear reaction network

$$\emptyset \overset{\alpha_1}{\leftarrow} S \overset{\alpha_2}{\rightarrow} 2S$$

and may conclude that it has the quasi-stationary distribution given by

$$\nu_x = \left(1 - \frac{\alpha_2}{\alpha_1}\right)\left(\frac{\alpha_2}{\alpha_1}\right)^{x-1}, \qquad x \in \mathbb{N},$$

whenever $\alpha_1 > \alpha_2$. This was also obtained in (O'neill, 2007) and in (Méléard and Villemonais, 2012) by other means. Had we imposed a higher order on the birth rates, say $\lambda_x = \alpha_2 x + \alpha_3 x(x-1)$ one would arrive at the following network,

$$\emptyset \xleftarrow{\alpha_1} S_1 \underset{\alpha_4}{\overset{\alpha_2}{\rightleftarrows}} 2S_1 \xrightarrow{\alpha_3} 3S_1$$

However, now the rates are no longer free. Choosing $\alpha_1, \alpha_2, \alpha_3$ arbitrarily forces $\alpha_4 = \frac{\alpha_1\alpha_3}{\alpha_2 - 2\alpha_3}$ under the further constraint $\alpha_2 - 2\alpha_3 < \alpha_1$. Thus, one may argue that the degree of freedom in the geometric distribution is only sufficient to capture the simple linear network. The argument does suggest, though, that one should look for a generalization of the geometric distribution to provide further examples of QSDs for reaction networks. What makes the geometric distribution especially well suited for examination is the simple form of the cumulative density function, yielding the approach of Lemma 2.3.30 tractable.

2.3.5 Approximations and Computational Aspects

In order to numerically compute a QSD, there are various approaches (see Blanchet et al., 2016). Whenever E is finite, we may calculate the QSD as the normalized left eigenvector ν of the rate matrix Q corresponding to the eigenvalue of maximal real part $-\theta(\nu)$. Ways to implement a numerical algorithm for this is well known (van Doorn and Pollett, 2013). One of the only existing overlaps between reaction networks and quasi-stationary distributions is for the class of conservative ACR networks. For a specific network, which could be reduced to a 1-dimensional birth-death process, it was shown (Anderson et al., 2014) that when the conserved quantity $M \to \infty$ the quasi-stationary distribution converged to a poisson distribution with mean equal to the deterministic ACR value. This was later generalized by means of a *variable freezing* method to a whole class of conservative systems (Enciso, 2016).

However, for non-conservative reaction networks, the state space is countably infinite, which poses several problems. The standard approach in this case is to employ a truncation procedure, whereby one approximates the full rate matrix Q with a sequence of square $n \times n$ matrices $\{Q^{(n)}\}$. When n is chosen sufficiently large, one would anticipate that the corresponding left eigenvector, $\nu^{(n)}$, is close to a true quasi-stationary distribution ν. While this is generally the case, and certainly when there is a unique globally attracting QSD, it has been proven that the method does not work in generality (Breyer and Hart, 2000). Further, as the complexity of this method is $\mathcal{O}(n^3)$ the approach quickly becomes infeasible, especially for high-dimensional systems.

For the special case of one-dimensional birth-death processes, we may apply the result of Proposition 2.3.27. As an example, consider the logistic network $\emptyset \overset{\alpha_1}{\leftarrow} S \overset{\alpha_2}{\to} 2S, 2S \overset{\alpha_3}{\to} S$, which has been well studied in the literature, which we shall return to several times in the sequel. Let $\alpha_1 = 1, \alpha_2 = 5, \alpha_3 = 1$. By means of Proposition 2.3.27, we may find the

polynomials $(\psi_x)_x$, and by exploiting a simple Newton algorithm to find the first root $\theta_{1,x}$ for each x, we see that the sequence $(\theta_{1,x})$ converges rapidly. By (2.3.23) we conclude that $\theta^* \approx 0.1081 > 0$. It follows from Proposition 2.3.24 that the Yaglom limit is given by (2.3.25). This is illustrated in Figure 2.7. We shall later see that the Yaglom limit is indeed unique for this reaction network without calculating the series D.

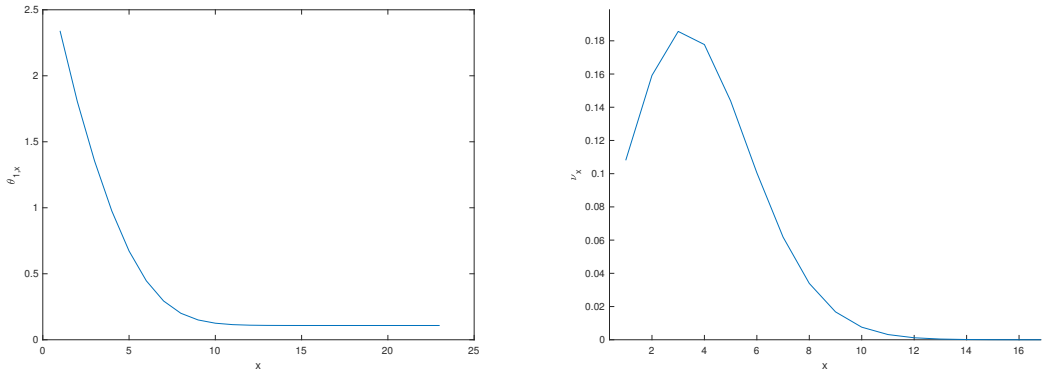

Figure 2.7: Left: The sequence $\theta_{1,x}$ converging to a positive number θ^*. Right: The QSD or Yaglom limit.

A different approach would be to simulate the stochastic system using the Doob-Gillespie algorithm (Gillespie, 1977; Erban et al., 2007; Wilkinson, 2012). However, unlike for stationary distributions, it is challenging to simulate from QSDs due to the fact that the absorption event becomes increasingly likely through time but, to be representative of the QSD, any sample paths must survive long enough to have forgotten their starting state (Griffin et al., 2017). For this reason several variants of a renewal process have been considered. One well studied approach is the Fleming-Viot method. Here, N particles evolves according to the defining stochastic equation, $(X_t^1, \ldots, X_t^N) \in \mathbb{N}_0^d$. When a (unique) particle reaches the absorbing set, it jumps instantaneously to the position of a particle chosen uniformly among the remaining $N-1$ ones. As both t and N tends to infinity, the empirical measure

$$\nu_t^N = \frac{1}{N} \sum_{i=1}^{N} \delta_{X_t^i}$$

converges almost surely to the underlying quasi-stationary distribution (Villemonais, 2015; Blanchet et al., 2016). In another approach (Groisman and Jonckheere, 2012) particles are redrawn from a given distribution over E once the process hits the absorbing set. The stationary distribution of this resurrected process, sometimes referred to as a pseudo-transient distribution, may then be connected with the QSD of the original system (Barbour and Pollett, 2010).

3

Existence of a Unique Quasi-stationary Distribution for Stochastic Reaction Networks

Mads Christian Hansen and Carsten Wiuf

University of Copenhagen

Abstract. In the setting of stochastic dynamical systems that eventually go extinct, the quasi-stationary distributions are often useful in understanding the long-term behavior before evanescence. We obtain sufficient conditions for the existence and uniqueness of a globally attracting quasi-stationary distribution for a broad class of applicable continuous-time Markov processes on countably infinite state spaces, called reaction networks. Furthermore, under these conditions, the convergence from any initial distribution to the quasi-stationary distribution is exponential in the total variation norm.

Keywords: *Quasi-statonary distribution, Lyapunov theory, Stochastic reaction netoworks*

Publication details: Submitted (2018) arXiv:1808.06833

3.1 Introduction

We may think of reaction networks in generality as a natural framework for representing systems of transformational interactions of entities (Veloz and Razeto-Barry, 2017). The set of entities (species) may in principle be of any nature, and specifying not just which ones interact (stoichiometry and reactions) but also quantifying how frequent they interact (kinetics), we obtain the dynamical system of a reaction network. Examples abound in biochemistry, where the language originated, however the true power of this approach is the ability to model diverse processes such as found in biological (Arkin et al., 1998; Barkai and Leibler, 2000), medical (Anderson et al., 2006), social (Dittrich and Winter, 2008), computational (Cook et al., 2009), economical (Veloz et al., 2014), ecological (Shakil et al., 2015) or epidemiological (Nåsell, 2011) contexts.

Whether the universe is inherently deterministic or stochastic in nature, the lack of complete information in complex systems inevitably introduces some degree of stochasticity. Thus, a stochastic description is not an alternative to the deterministic approach, but a more complete one (Qian, 2011). Indeed, the deterministic model solution is an approximation of the solution for the stochastic model, improving with the system size, and in general only remaining valid on finite time intervals (Kurtz, 1970). Thus, the long-term behavior of a given reaction network may depend crucially on whether it is modeled deterministically or stochastically (Gupta et al., 2014). In particular, the possibility of extinction, which is a widely occurring phenomenon in nature, may sometimes only be captured by the latter (Johnston et al., 2017). As a consequence, the counterpart to a stable stationary solution in the deterministically modeled system is not generally a stationary distribution of the corresponding stochastic model. Instead, a so-called quasi-stationary distribution, which is a stationary measure when conditioned on the process not going extinct, has shown to be the natural object of study. A concise overview of the history and current state of this field can be found in van Doorn and Pollett (2013), while Pollett (2015) contains a comprehensive bibliography on quasi-stationary distributions and related work.

From a modeling standpoint, when the copy-numbers of interacting entities are low and reaction rates are slow, it is important to recognize that the individual reaction steps occur discretely and are separated by time intervals of random length (Arkin et al., 1998). This is for example the case at the cellular level (Elowitz et al., 2002), where stochastic effects resulting from these small numbers may be physiologically significant (Cook et al., 2009). Furthermore, stochastic variations inherent to the system may in general be beneficial for identifying system parameters (Munsky et al., 2009). The quasi-stationary distribution possesses several desirable properties in this domain. Most importantly, if the system under study has been running for a long time, and if the only available knowledge about the system is that it has not reached extinction, then we can conclude that the quasi-stationary distribution, if it exists and is unique, is the likely distribution of the state variable (Nåsell, 2011).

Consider a right-continuous time-homogenous Markov process $(X_t \colon t \geq 0)$ (Rogers and Williams, 2000), that evolves in a domain $D \subseteq \mathbb{R}^d$, wherein there is a set of absorbing states, a "trap", $A \subset D$. The process is absorbed, also referred to as being killed, when it hits the

set of absorbing states, implying $X_t \in A$ for all $t \geq \tau_A$, where $\tau_A = \inf\{t \geq 0 : X_t \in A\}$ is the hitting time of A. As we are interested in the process before reaching A, there is no loss of generality in assuming $X_t = X_{t\wedge\tau_A}$. We refer to the complement,

$$E := D\backslash A,$$

as the set of endorsed states. For any probability distribution, μ, on E, we let $\mathbb{P}_\mu$ and $\mathbb{E}_\mu$ be the probability and expectation respectively, associated with the process $(X_t \colon t \geq 0)$, initially distributed with respect to μ. For any $x \in E$, we let $\mathbb{P}_x = \mathbb{P}_{\delta_x}$ and $\mathbb{E}_x = \mathbb{E}_{\delta_x}$. Under suitable conditions, the process hits the absorbing set almost surely (a.s.), that is $\mathbb{P}_x(\tau_A < \infty) = 1$ for all $x \in E$, and we investigate the behavior of the process before being absorbed (Collet et al., 2013).

Definition 3.1.1. A probability measure ν on E is called a quasi-stationary distribution (QSD) for the process $(X_t \colon t \geq 0)$ absorbed at A, if for every measurable set $B \subseteq E$

$$\mathbb{P}_\nu(X_t \in B \,|\, t < \tau_A) = \nu(B), \qquad t \geq 0,$$

or equivalently, if there exists a probability measure μ on E such that

$$\lim_{t\to\infty} \mathbb{P}_\mu(X_t \in B \,|\, t < \tau_A) = \nu(B),$$

in which case we also say that ν is a quasi-limiting distribution.

We refer to Méléard and Villemonais (2012) for a proof of the equivalence of quasi-limiting and quasi-stationary distributions. Existence and uniqueness of a QSD on a finite state space is well known (Collet et al., 2013, Chapter 3), and it is given by the normalized left Perron-Frobenius eigenvector of the transition rates matrix restricted to E. For the infinite dimensional case, most work has been carried out for birth-death processes in one dimension (van Doorn and Pollett, 2013), where classification results yielding information about the set of QSDs exist (van Doorn, 1991).

In the present paper, we will focus on a special case of multidimensional processes on countable infinite state spaces which can be viewed as reaction networks. We will prove as the main result in Theorem 4.5.5 and Corollary 3.5.2 sufficient conditions for the existence of a unique globally attracting QSD in the space of probability distributions on E, equipped with the total variation norm, $\|\cdot\|_{TV}$. Recall that this norm may be defined as (Pollard, 2005)

$$\|\mu\|_{TV} = 2 \sup_{B\subseteq E} |\mu(B)|.$$

Thus, informally, the metric associated to this norm is the largest possible difference between the probabilities that two probability distributions can assign to the same event. Our result is based on the following recent result (Champagnat and Villemonais, 2016, Theorem 2.1).

Theorem 3.1.2. The following are equivalent

- There exists a probability measure ν on E and two constants $C, \gamma > 0$ such that, for all initial distributions μ on E,

$$\|\mathbb{P}_\mu(X_t \in \cdot \,|\, t < \tau_A) - \nu(\cdot)\|_{TV} \leq Ce^{-\gamma t}, \qquad \forall t \geq 0.$$

- There exists a probability measure ν on E such that

  (A1) there exists $t_0, c_1 > 0$ such that for all $x \in E$,

$$\mathbb{P}_x(X_{t_0} \in \cdot \,|\, t_0 < \tau_A) \geq c_1 \nu(\cdot),$$

  (A2) there exists $c_2 > 0$ such that for all $x \in E$ and $t \geq 0$,

$$\mathbb{P}_\nu(t < \tau_A) \geq c_2 \mathbb{P}_x(t < \tau_A).$$

Now, using Foster-Lyapunov theory (Meyn and Tweedie, 1993, 2009), a series of assumptions on the process $(X_t \colon t \geq 0)$ has been shown to be sufficient for (A1) and (A2) to hold (Champagnat and Villemonais, 2017). This approach has been applied to a particular case of multidimensional birth-death processes, giving sufficient conditions, in terms of the parameters of the process, for the existence and uniqueness of a QSD. Here, we extend this result, not just to a larger set of parameter values in the birth-death process case, but to the *much* broader class of stochastic processes known as stochastic reaction networks.

The outset of the paper is as follows. In section 2, we introduce the setup and notation of reaction network theory, and define the central inferred notions of endorsed and absorbing states for this class of processes. Section 3 contains the terminology and main assumptions that we shall use throughout the paper. We then move on in section 4, to prove that the processes associated with stochastic reaction networks do indeed satisfy all the required assumptions made by (Champagnat and Villemonais, 2017, Corollary 2.8). Section 5 contains the main result, Theorem 4.5.5. Finally, we give some examples in section 6, illustrating the applicability of the results.

3.2 Reaction Network Setup

Denote the real numbers by $\mathbb{R}$, the integers by $\mathbb{Z}$, the natural numbers by $\mathbb{N} = \{1, 2, \dots\}$ and the nonnegative integers by $\mathbb{N}_0 = \{0, 1, 2, \dots\}$. Further, for any set, B, let $|B|$ denote its cardinality and denote by $\mathbb{1}_B \colon D \to \{0, 1\}$ the indicator function of a subset $B \subseteq D$.

A reaction network is a triple $\mathcal{N} = (\mathcal{S}, \mathcal{C}, \mathcal{R})$, where $\mathcal{S}$ is a finite ordered set of species[1], $\mathcal{C}$ is a finite set of complexes, consisting of linear combinations over $\mathbb{N}_0$ of the species, and $\mathcal{R} \subset \mathcal{C} \times \mathcal{C}$ is an irreflexive relation on $\mathcal{C}$, referred to as the set of reactions (Anderson and Kurtz, 2015; Feinberg, 1979; Gunawardena, 2003). Furthermore, $\mathcal{R}$ is assumed to be ordered.

[1]The terminology "species" is standard, although one may equally think of them as general entities or agents.

We define the dimension of the reaction network, $d = |\mathcal{S}|$. Any species $S_i \in \mathcal{S}$ can be identified with the unit vector $e_i \in \mathbb{N}_0^d$, thus any complex $y \in \mathcal{C}$ can be identified with a vector in $\mathbb{N}_0^d$. It is customary to denote an element $(y_k, y_k') \in \mathcal{R}$ by $y_k \to y_k' \in \mathcal{R}$ in which case we refer to y_k as the source complex and to y_k' as the product complex of reaction k. We may thus write $\mathcal{R} = \{y_k \to y_k' \colon k = 1, \dots, r\}$. Employing a standard, although slight abuse of, notation, we identify $\mathcal{S} = \{S_1, \dots, S_d\}$ with the set $\{1, \dots, d\}$ and $\mathcal{R}$ with $\{1, \dots, r\}$. We write the k'th reaction with the notation

$$\sum_{i \in \mathcal{S}} y_{ki} S_i \to \sum_{i \in \mathcal{S}} y_{ki}' S_i,$$

where $y_{ki} = (y_k)_i$ and $y_{ki}' = (y_k')_i$ are the stoichiometric coefficients associated with the source and product complexes of reaction k, respectively. Define the reaction vectors $\xi_k = y_k' - y_k$ and the stoichiometric matrix

$$\Xi = (\xi_1 \, \xi_2 \, \dots \, \xi_r) \in \mathbb{N}_0^{d \times r}.$$

The order of reaction k is the sum of the stoichiometric coefficients of the source complex, $\sum_{i \in \mathcal{S}} y_{ki}$. Finally, we define the maximum of a vector over the set $\mathcal{R}$, $x = \max_{k \in \mathcal{R}} y_k$, as the entry-wise maximum, $x_i = \max_{k \in \mathcal{R}} y_{ki}$.

A set of reactions $\mathcal{R}$ induces a set of complexes and a set of species, namely the complexes and species that appear in the reactions. We will assume that a reaction network is always given in this way by $\mathcal{R}$, and one may then completely describe a reaction network in terms of its reaction graph, whose nodes are the complexes and whose directed edges are the reactions. This concise description will be employed in the rest of the paper. To avoid trivialities, we assume $\mathcal{R} \neq \emptyset$.

For each reaction we specify an intensity function $\lambda_k \colon \mathbb{N}_0^d \to [0, \infty)$, $k \in \mathcal{R}$, which satisfies the stoichiometric admissibility condition:

$$\lambda_k(x) > 0 \quad \Leftrightarrow \quad x \geq y_k,$$

where we use the usual vector inequality notation; $x \geq y$ if $x_i \geq y_i$ for all $i \in \mathcal{S}$. Thus, reactions are only allowed to take place whenever the copy-numbers of each species in the current state is at least as great as those of the corresponding source complex. A widely used example is stochastic mass action kinetics given by

$$\lambda_k(x) = \alpha_k \prod_{i=1}^{d} y_{ki}! \binom{x}{y_k} = \alpha_k \prod_{i=1}^{d} \frac{x_i!}{(x_i - y_{ki})!},$$

for some reaction rate constants $\alpha_k > 0$ (Anderson and Kurtz, 2015). The idea is that the rate is proportional to the number of distinct subsets of the molecules present that can form the input of the reaction. It reflects the assumption that the system is well-stirred (Anderson and Kurtz, 2015). Other examples include power law kinetics or generalized mass action kinetics (Anderson, 2008; Horn and Jackson, 1972; Müller and Regensburger, 2012). A particular choice of such rate functions constitute a stochastic kinetics $\lambda = (\lambda_1, \dots, \lambda_r)$ for

the reaction network $\mathcal{N}$, and the pair $(\mathcal{N}, \lambda)$ is referred to as a stochastic reaction system, or simply a reaction network with kinetics λ.

We may then specify the stochastic process $(X_t \colon t \geq 0)$ on the state space $D := \mathbb{N}_0^d$ related to the reaction system $(\mathcal{N}, \lambda)$. Let X_t be the vector in $\mathbb{N}_0^d$ whose entries are the species counts at time t. If reaction $y_k \to y_k'$ occurs at time t, then the new state is $X_t = X_{t-} + y_k' - y_k = X_{t-} + \xi_k$, where X_{t-} denotes the previous state. The stochastic process then follows,

$$X_t = X_0 + \sum_{k \in \mathcal{R}} Y_k \left(\int_0^t \lambda_k(X_s)\, ds \right) \xi_k, \tag{3.2.1}$$

where Y_k are independent and identically distributed unit-rate Poisson processes (Anderson and Kurtz, 2015; Ethier and Kurtz, 1986; Norris, 2009). This stochastic equation is referred to as a random time change representation. We assume throughout the paper that the process is non-explosive, so that the process is well defined. Assumption 2, though, will imply non-explosiveness.

3.2.1 The State Space

To define the set of endorsed states and absorbing states in the setting of stochastic reaction networks, we recall some terminology from stochastic processes. We say that there is a path from x to y, denoted $x \mapsto y$, if there exists $t \geq 0$ such that $\mathbb{P}_x(X_t = y) > 0$. We extend this notion to sets as follows; $B_1 \mapsto B_2$ if there exists $x \in B_1$ and $y \in B_2$ such that $x \mapsto y$. Finally, we introduce the region of large copy numbers, where all reactions may take place, defined as

$$R = \{x \in \mathbb{N}_0^d \,|\, \lambda_k(x) > 0 \,\forall\, k \in \mathcal{R}\}.$$

Any network satisfies $R \neq \emptyset$. Indeed, by the stoichiometric compatibility condition, $\{x \in D \,|\, x \geq M\} \subseteq R$ where $M = \max_{k \in \mathcal{R}} y_k \in \mathbb{N}_0^d$. Letting $D_E = \{x \in D \,|\, x \mapsto R\}$, we may decompose the state space into a disjoint union

$$D = D_E \sqcup D_A.$$

A state space D is irreducible if for all $x, y \in D$ we have $\mathbb{P}_x(X_{t_1} = y) > 0$ and $\mathbb{P}_y(X_{t_2} = x) > 0$ for some $t_1, t_2 > 0$ (Gupta et al., 2014). Thus, D is irreducible if for all $x, y \in D$ there exists a path $x \mapsto y$. Irreducibility induces a class structure on the state space (Norris, 2009), and we denote the classes by $\mathcal{I}_1, \mathcal{I}_2, \ldots$ (potentially infinitely many). Let $\mathscr{I}$ denote the set of irreducible classes. Obviously, either $\mathcal{I}_i \subseteq D_E$ or $\mathcal{I}_i \subseteq D_A$, $i \geq 1$.

Lemma 3.2.2. *The pair $(\mathscr{I}, \preceq)$, where $\preceq$ is given by*

$$\mathcal{I}_j \preceq \mathcal{I}_i \Leftrightarrow \mathcal{I}_i \mapsto \mathcal{I}_j, \qquad i, j \geq 1,$$

is a well defined poset. The irreflexive kernel $(\mathscr{I}, \prec)$ gives a well defined strict poset.

Proof. Since all elements $\mathcal{I} \in \mathscr{I}$ are irreducible, there exists a path between any two points in $\mathcal{I}$ hence $\mathcal{I} \preceq \mathcal{I}$ yielding the relation reflexive.

Suppose $\mathcal{I}_i \preceq \mathcal{I}_j$ and $\mathcal{I}_j \preceq \mathcal{I}_i$ for some $i, j \geq 1$. Let $x \in \mathcal{I}_i$ and $y \in \mathcal{I}_j$ be given. By assumption, we may find a path from x to some $z_2 \in \mathcal{I}_j$, and by irreducibility of $\mathcal{I}_j$ there is a path from z_2 to y. Similarly, we may by assumption find a path from y to some $z_1 \in \mathcal{I}_i$ and by irreducibility of $\mathcal{I}_i$ a path from z_1 to x. As x, y were arbitrary, we conclude that there exists a path between any two points in $\mathcal{I}_i \cup \mathcal{I}_j$ hence $\mathcal{I}_i = \mathcal{I}_j$, yielding the relation antisymmetric.

Finally, suppose $\mathcal{I}_k \preceq \mathcal{I}_j$ and $\mathcal{I}_j \preceq \mathcal{I}_i$ for some $i, j, k \geq 1$. Then there exists a path from some $x \in \mathcal{I}_i$ to some $z_1 \in \mathcal{I}_j$ and a path from some $z_2 \in \mathcal{I}_j$ to some $y \in \mathcal{I}_k$. By irreducibility of $\mathcal{I}_j$ there is a path from z_1 to z_2, and concatenation of the three paths yield one from x to y. We conclude that $\mathcal{I}_k \preceq \mathcal{I}_i$, hence the relation is transitive. □

A similar ordering has been considered in van Doorn and Pollett (2009). However, their further analysis rests on the setting of discrete time, rendering the approach insufficient for stochastic reaction networks. To exploit the graphical structure induced by $\preceq$, define the marked directed acyclic graph $\mathcal{D} = (\mathscr{I}, \mathscr{E})$ as follows. The set of directed edges is

$$\mathscr{E} = \{(\mathcal{I}_i, \mathcal{I}_j) \in \mathscr{I}^2 \,|\, \mathcal{I}_j \prec \mathcal{I}_i, i, j \geq 1\},$$

while the marking $\mathscr{I} = \mathscr{I}_A \sqcup \mathscr{I}_E$ is given by

$$\mathscr{I}_E = \{\mathcal{I} \in \mathscr{I} \,|\, \mathcal{I} \subseteq D_E\}, \qquad \mathscr{I}_A = \{\mathcal{I} \in \mathscr{I} \,|\, \mathcal{I} \subseteq D_A\}.$$

Let $\mathscr{V}_1, \mathscr{V}_2, \ldots$ denote the vertex set of the respective connected components of the induced subgraph $\mathcal{D}[\mathscr{I}_E]$, the graph with vertex set $\mathscr{I}_E$ and edges from $\mathscr{E}$ with start and end nodes in $\mathscr{I}_E$.

Definition 3.2.3. The endorsed sets and absorbing sets are defined, respectively, by

$$E_n = \bigcup_{\mathcal{I} \in \mathscr{V}_n} \mathcal{I}, \qquad A_n = \bigcup_{\mathcal{I} \in \mathscr{I}_A \colon E_n \mapsto \mathcal{I}} \mathcal{I}, \qquad n \geq 1.$$

The corresponding state space is defined by $D_n = E_n \sqcup A_n$.

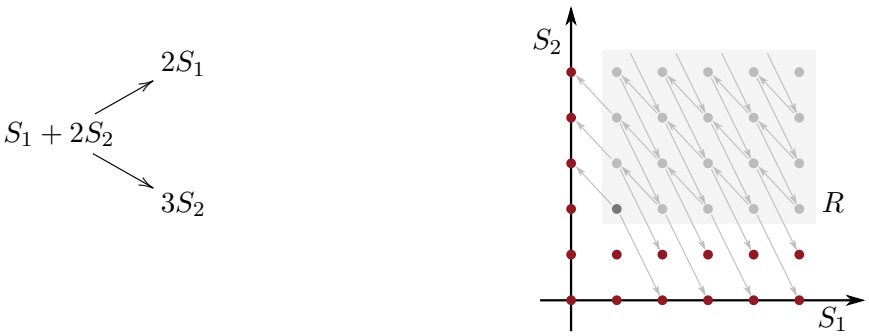

Figure 3.1: Left: The reaction graph of a stochastic reaction network. Right: The state space with the region R in shaded grey. Points in D_A are marked red.

As D_E is non-empty, the existence of at least one endorsed set is guaranteed. By construction, the endorsed sets are disjoint, their union is D_E and their number, N_E, may in general be countable infinite. Furthermore, any absorbing set is confined to a subset of $\{x \in \mathbb{N}_0^d \,|\, x \not\geq M\}$, lying "close" to the boundary of $\mathbb{N}_0^d$. If extinction is possible from an endorsed class E_n then this absorption will take place in A_n. Note, however, that the set $\mathscr{I}_A$ may in general be empty, in which case no absorbing set exist. To further illuminate the structure of the endorsed sets, we provide the following classification result.

Proposition 3.2.4. *For $M \in \mathbb{N}^d$ sufficiently large, the set $\{x \in D_E \,|\, x \geq M\}$ intersects*

(i) finitely many endorsed sets if and only if rank $\Xi = d$.

(ii) a single endorsed set if and only if $\operatorname{span}_{\mathbb{Z}} \Xi = \mathbb{Z}^d$.

Proof. Suppose first that rank $\Xi < d$. Let $x \in D_E$. Then $x \in E_1$, say and there exists a $y \in D_E$ such that $y \notin (x + \operatorname{span}_{\mathbb{R}} \Xi)$. In particular, $y \in E_2$ where $E_1 \neq E_2$. This procedure can be repeated indefinitely, yielding infinitely many endorsed sets.

Now, suppose rank $\Xi = d$. Then one may choose a linear combination of the reaction vectors yielding a strictly positive point,

$$\sum_{k \in \mathcal{R}} a_k \xi_k > 0, \qquad a_k \in \mathbb{Z}.$$

Let $a = \sum_{k \in \mathcal{R}} |a_k|$ and $a_0 = 0$. Define the sequence $(w_\ell)_{\ell=1,\ldots,a}$ by

$$w_\ell = \xi_k, \qquad \ell = 1 + \sum_{j=0}^{k-1} |a_j|, \ldots, \sum_{j=1}^{k} |a_j|, \qquad k \in \mathcal{R}.$$

As the reaction vectors are finite, the partial sums $P_j = \sum_{\ell=1}^{j} w_\ell$ are finite for each $j \leq a$. Let

$$m = \min_{i \in \mathcal{S}, j \leq a} (P_j)_i.$$

Choosing each coordinate $M_i > |m| + \max_{k \in \mathcal{R}} y_k$ for each $i \in \mathcal{S}$, it follows that any point $x \in D_E$ with $x \geq M$ satisfies $x + P_j \in R$ for any $j \leq a$. We say that a sequence of states $(x_1, \ldots, x_n)$ is an undirected walk from x_1 to x_n if for all $1 \leq i \leq n-1$ there exists $k(i) \in \mathcal{R}$ such that $x_{i+1} = x_i \pm \xi_{k(i)}$. As all reactions may occur in R, we conclude that x has an undirected walk to the point

$$x' := x + P_a = x + \sum_{k \in \mathcal{R}} a_k \xi_k > x.$$

Thus, by definition, x and x' belong to the same endorsed set, say E_1. To determine the number of endorsed sets, let $B = (b_i)$ denote the basis matrix of the free $\mathbb{Z}$-module generated

by the stoichiometric matrix Ξ, and define the lattice generated by Ξ to be

$$\mathcal{L}(\Xi) = \left\{ \sum_{i=1}^{\operatorname{rank}\Xi} z_i b_i \colon z_i \in \mathbb{Z} \right\}.$$

By the construction in the previous paragraph, it follows that all points in the region $\{y \in E_1 \colon y \geq M\}$ belong to the same translated lattice $x + \mathcal{L}(\Xi)$. By assumption, $\operatorname{span}_{\mathbb{R}} \Xi = \mathbb{R}^d$ hence the lattice $\mathcal{L}(\Xi)$ has rank d. The number of ways one may translate a rank d lattice in $\mathbb{R}^d$ to an integer lattice point without any points intersecting is given by considering the number of integer lattice points inside the fundamental parallelotope,

$$\mathcal{P}(\Xi) = \left\{ \textstyle\sum_{i=1}^{d} \theta_i b_i \,|\, \theta_i \in [0,1), b_i \in \operatorname{span}_{\mathbb{Z}} \Xi, \det(B) \neq 0 \right\}.$$

Indeed, as $\mathcal{P}(\Xi)$ tiles $\mathbb{R}^d$, that is for any point $z \in \mathbb{R}^d$ there exists a unique $z' \in \mathcal{L}(\Xi)$ such that $z \in z' + \mathcal{P}(\Xi)$, the problem is reduced to a single fundamental parallelotope, which by definition contains exactly one point from each translated lattice (Dadush, 2013). Further, the number of integer lattice points inside $\mathcal{P}(\Xi)$ is exactly equal to the volume of the parallelotope (Cohen et al., 1993, p. 97), hence, by finiteness of the reaction vectors,

$$N_E(M) = |\det(B)| < \infty,$$

for M sufficiently large, where $N_E(M)$ is the number of endorsed sets intersecting $\{x \in D_E \,|\, x \geq M\}$. This proves (i) of the proposition.

Finally, if $\operatorname{span}_{\mathbb{Z}} \Xi = \mathbb{Z}^d$ then $\mathcal{L}(\Xi) = \mathbb{Z}^d$ and the unit vectors $e_1, \ldots, e_d \in \mathcal{L}(\Xi)$. As $\mathcal{P}(I_d) \cap \mathcal{L}(\Xi) = \{0\}$ we conclude from Dadush (2013) that $e_1, \ldots, e_d$ is a basis for $\mathcal{L}(\Xi)$ hence $N_E(M) = 1$ as desired. □

Note that, in particular, a reaction network whose associated stochastic process is a birth-death process, that is, a process where for each $i = 1, \ldots, d$ either $e_i \in \mathcal{R}$ or $-e_i \in \mathcal{R}$, have a single endorsed set for x sufficiently large. In practice, one may find the endorsed sets by picking $x \in R$ and adding states by a backtracking algorithm (Paulevé et al., 2014). Verification of $\operatorname{span}_{\mathbb{Z}} \Xi = \mathbb{Z}^d$ can be done by calculation of the Hermite normal form (Paulevé et al., 2014).

One may suspect that Proposition 3.2.4 could be strengthened to hold on the entire set D_E. This is only partially true. Consider as an example the three-dimensional reaction network given by the reaction graph

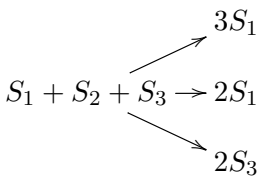

It follows that $\operatorname{rank}\Xi = 3$. However, $D_E = \mathbb{N}^3$, $D_A = \mathbb{N}_0^3 \setminus \mathbb{N}^3$, and each singleton $\{(1, m, 1)\}$,

$m \geq 1$, constitutes its own endorsed set. Thus, in the generic picture, close to the absorbing set, there may be infinitely many endorsed sets. We do, however, have the following corollary.

Corollary 3.2.5. If $d \leq 2$, there are finitely many endorsed sets if and only if $\operatorname{rank} \Xi = d$.

Proof. We only need to prove that if $\operatorname{rank} \Xi = d$ then there are finitely many endorsed sets. For this, it suffices to prove that at most finitely many $x \in D_E$ do not have an undirected walk (as introduced in the proof of Proposition 3.2.4) to a point $z \geq M$. Indeed, by the proof of Proposition 3.2.4, if such a path exists, then by definition x belongs to one of finitely many endorsed sets. With the remaining set being finite, the total number of endorsed sets is therefore finite.

Now, let $x \in D_E$ be given. By the definition of endorsed sets, there exists a path $x \mapsto y$ with $\lambda_k(y) > 0$ for all $k \in \mathcal{R}$. As $\operatorname{rank} \Xi = d$, for each $1 \leq j \leq d$, there exists $k(j) \in \mathcal{R}$ such that $\langle e_j, \xi_{k(j)} \rangle \neq 0$. Keeping the jth coordinate fixed and increasing the possible other if necessary, thus arriving at a point y' with $y_j = y'_j$ and $y_i \leq y'_i$ for $i \neq j$, by the stoichiometric compatibility condition we may by repeated use of reaction $k(j)$ find a path $y' \mapsto z$ or $z \mapsto y'$ with $z \geq M$. Repeating the argument for the possible remaining coordinate, we conclude that there exists an $M' \in \mathbb{N}^d$ such that if $x \not< M'$ then x has an undirected walk to a point $z \geq M$. As the set $\{x \in D_E \,|\, x < M'\}$ is finite, this concludes the proof. □

One may easily verify, that the second part of Proposition 3.2.4 can also be extended for $d = 1$. Indeed, for any point $x \in D_E$ there exists a reaction $k \in \mathcal{R}$ such that $\lambda_k(x) > 0$ and either $x + \xi_k > x$ or $x + \xi_k < x$. Otherwise $x \in D_A$. Consequently, there is a point $z \geq M$ for any $M \in D_E$ such that either $x \mapsto z$ or $z \mapsto x$. However, note that the network in Figure 3.1 shows that this result does not hold in the case $d = 2$.

An endorsed set, E_n, $n \geq 1$, is only irreducible if it consists of a single irreducible class. If there is more than one irreducible class in E_n, then we need that there is a smallest one to ensure uniqueness of a QSD.

Assumption 1. For a given endorsed class E_n, $n \geq 1$, we assume:

(i) E_n contains a unique minimal irreducible class, $\mathcal{I}^n_{\min}$.

(ii) if $A_n \neq \emptyset$ then $\mathcal{I}^n_{\min} \mapsto A_n$.

We shall see that Assumption 1(i) is equivalent to a more technical property of the state space, which is necessary for our results to hold. Thus no generality is lost in having Assumption 1(i).

Networks without any minimal class exists, for example $\emptyset \to S_1$, which does not have an absorbing set either. Furthermore, networks with more than one minimal class also exist, for example, $S_1 + S_2 \to \emptyset$, $S_2 \to \emptyset$. Thus Assumption 1(i) is indeed not superfluous. We believe Assumption 1(ii) is always met if Assumption 1(i) is. It ensures that one may always reach the absorbing set, if it is non-empty.

Definition 3.2.3 accommodates the general case where uniqueness of a QSD does not necessarily hold, in which case the support of the QSD may stretch the entire endorsed set rather than, as we shall see, the unique minimal irreducible class. We remark that rather

than investigating an entire endorsed set, one may be interested in a particular irreducible component, say $\mathcal{I}$. Letting the state space be $D = E \sqcup A$ where

$$E = \mathcal{I}, \qquad A = \bigcup_{\mathcal{J} \in \mathscr{I}_A \colon E \mapsto \mathcal{J}} \mathcal{J},$$

the theory to be developed in this paper applies to this case as well.

As an illuminating example consider the generalized death process $mS_1 \to \emptyset$ with $m \in \mathbb{N}$, where each point in the state space $D = \mathbb{N}_0$ constitutes its own irreducible class. Here, the endorsed and absorbing sets are $E_n = \{n + pm - 1 \,|\, p \in \mathbb{N}\}$ and $A_n = \{n - 1\}$ respectively, for $n = 1, \dots, m$, and Assumption 1 is satisfied for all n. Thus, $D_E = \{m, m+1, \dots\}$ and $D_A = \{0, \dots, m-1\}$. It is known that in the simple death case, $m = 1$, uniqueness does not hold on $D_E = \mathbb{N}$. Indeed, there is a continuum of QSDs with support larger than $\{1\}$, the unique minimal class (Griffin, 2016).

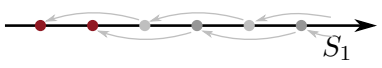

Figure 3.2: State space of the reaction network $2S_1 \to \emptyset$. There are two endorsed sets.

In the setting of birth-death processes in one dimension, which has an infinite state space, it is known that there will be either none, a unique or a continuum of QSDs (van Doorn, 1991). Consider the two reaction networks

$$\emptyset \overset{\alpha_1}{\leftarrow} S_1 \overset{\alpha_2}{\to} 2S_1, \qquad \emptyset \overset{\alpha_1}{\leftarrow} S_1 \underset{\alpha_3}{\overset{\alpha_2}{\rightleftharpoons}} 2S_1, \tag{3.2.6}$$

endowed with mass action kinetics. In both cases we conclude, according to Definition 3.2.3, that the set of absorbing states and the set of endorsed states are

$$D_A = A_1 = \{0\}, \qquad D_E = E_1 = \{1, 2, \dots\},$$

respectively. For the network on the left in (3.2.6), assuming $\alpha_1 > \alpha_2$, there is a continuum of QSDs on E_1, while for the network on the right, there is a unique QSD on E_1, for all parameter values (Méléard and Villemonais, 2012). This fits well with our result – the necessity of having reactions of order higher than one to ensure uniqueness permeates to higher dimensions.

3.3 Extension of Arguments

In this and the following sections, we shall simply use the notation E to refer to a single endorsed set, with corresponding non-empty absorbing set A, when there is no ambiguity. Further, as existence and uniqueness is known on finite state spaces, we shall assume without loss of generality that E is countably infinite. We make the following definitions inspired by Champagnat and Villemonais (2017) and Meyn and Tweedie (1993).

Definition 3.3.1. For any vector $v \in \mathbb{N}^d$, we define a corresponding function $\langle v, \cdot \rangle \colon \mathbb{Z}^d \to \mathbb{Z}$ given by the standard inner product

$$\langle v, x \rangle = x \cdot v.$$

This function may in general take negative values, however, when restricting $\langle v, \cdot \rangle$ to E, one obtains a norm-like function, (Meyn and Tweedie, 1993). Choosing $v = (1, \ldots, 1)$ we recover the function used in (Champagnat and Villemonais, 2017). In general, we shall choose $v \in \mathbb{N}^d$ based on the particular reaction network at hand, and will in the following consider it fixed. For $n \in \mathbb{N}$, define the sets

$$O_n = \{x \in E \colon \langle v, x \rangle \leq n\}.$$

which, irrespectively of v, are compact subsets of E, satisfying $O_n \subseteq O_{n+1}$ and $E = \bigcup_{n \in \mathbb{N}} O_n$. We denote the first hitting time of A, the first hitting time of O_n and the first exit time of O_n by

$$\tau_A = \inf\{t \geq 0 \colon X_t \in A\}, \qquad \tau_n = \inf\{t \geq 0 \colon X_t \in O_n\}, \qquad T_n = \inf\{t \geq 0 \colon X_t \notin O_n\},$$

respectively. Note that all of these are stopping times and might be infinite. As we will be concerned with the application of unbounded functions serving the purpose of a Lyapunov function, we introduce the weakened generator, L, for the Markov process (Meyn and Tweedie, 1993; Champagnat and Villemonais, 2016).

Definition 3.3.2. A measurable function $W \colon D \to \mathbb{R}$ belongs to the domain $\mathcal{D}(L)$ of the weakened generator L of $(X_t \colon t \geq 0)$ if there exists a measurable function $U \colon E \to \mathbb{R}$ such that, for all $n \in \mathbb{N}$, $t \geq 0$ and $x \in E$

$$\mathbb{E}_x W(X_{t \wedge T_n}) = W(x) + \mathbb{E}_x \left(\int_0^{t \wedge T_n} U(X_s) \, ds \right),$$

and

$$\mathbb{E}_x \left| \int_0^{t \wedge T_n} U(X_s) \, ds \right| < \infty, \tag{3.3.3}$$

and we define $LW = U$ on E and $LW \equiv 0$ on A.

As the state space of interest is always countable, all functions $f \colon D \to \mathbb{R}$ are measurable. Moreover, as the state space is discrete and O_n is finite for all $n \in \mathbb{N}$, all functions $f \colon D \to \mathbb{R}$ are in the domain of the weakened generator, $\mathcal{D}(L)$ (Meyn and Tweedie, 1993). In particular, $\mathbb{E}_x W(X_{t \wedge T_n})$ is well-defined and finite. Generally, the weakened and the infinitesimal generator need not agree, and the infinitesimal generator may not exist (Meyn and Tweedie, 1993).

However, if f is bounded, then they do agree. In particular, $E_x|f(X_t)| < \infty$ and it follows that as $t \to 0$,

$$\begin{aligned}\mathbb{E}_x f(X_t) &= \left(\sum_{k\in\mathcal{R}} f(x+\xi_k)\mathbb{P}_x(X_t = x+\xi_k)\right) + f(x)\mathbb{P}_x(X_t = x) + o(t)\\ &= \left(\sum_{k\in\mathcal{R}} f(x+\xi_k)\lambda_k(x)t + o(t)\right) + f(x)\left(1 - \sum_{k\in\mathcal{R}}\lambda_k(x)t + o(t)\right) + o(t)\\ &= \sum_{k\in\mathcal{R}} \lambda_k(x)(f(x+\xi_k) - f(x))t + f(x) + o(t).\end{aligned}$$

Hence $\mathbb{E}_x f(X_t)$ is differentiable and from the fundamental theorem of calculus we conclude that the weakened generator coincides with the (weak) infinitesimal generator (Meyn and Tweedie, 1993),

$$\widehat{L}f(x) = \lim_{t\to 0}\frac{\mathbb{E}_x f(X_t) - f(x)}{t} = \sum_{k\in\mathcal{R}}\lambda_k(x)(f(x+\xi_k) - f(x)),$$

for $x \in E$.

Moreover, setting $W(x) = \langle v, x\rangle$ as in Definition 3.3.1, it follows from the Poisson characterization of the process (5.2.1), that

$$\begin{aligned}\mathbb{E}_x W(X_{t\wedge T_n}) &= \mathbb{E}_x\langle v, X_{t\wedge T_n}\rangle = \mathbb{E}_x\langle v, X_0\rangle + \mathbb{E}_x\langle v, \sum_{k\in\mathcal{R}} Y_k\left(\int_0^{t\wedge T_n}\lambda_k(X_s)\,ds\right)\xi_k\rangle\\ &= W(x) + \mathbb{E}_x\left(\int_0^{t\wedge T_n}\sum_{k\in\mathcal{R}}\lambda_k(X_s)\langle v,\,\xi_k\rangle\,ds\right),\end{aligned}$$

such that

$$LW(x) = \sum_{k\in\mathcal{R}}\lambda_k(x)\langle v, \xi_k\rangle, \qquad x \in E. \tag{3.3.4}$$

Note that (3.3.3) is fulfilled as O_n is finite.

Definition 3.3.5. Define the functions $d_v, d^v\colon \mathbb{N} \to \mathbb{R}$ by

$$\begin{aligned}d_v(n) &= -\max_{x\in E,\langle v,x\rangle = n}\sum_{k=1}^{r}\lambda_k(x)\langle v,\xi_k\rangle\mathbb{1}_E(x+\xi_k),\\ d^v(n) &= \max_{x\in E,\langle v,x\rangle = n} n\sum_{k=1}^{r}\lambda_k(x)\mathbb{1}_A(x+\xi_k).\end{aligned}$$

All networks, for which extinction is possible, have the property that there exists a $v \in \mathbb{N}^d$ such that $\langle v, \xi_k\rangle \le 0$ for some reaction $k \in \mathcal{R}$. Indeed, suppose that $\langle v, \xi_k\rangle > 0$ for all $v \in \mathbb{N}^d$ and $k \in \mathcal{R}$. Then $\xi_k \in \mathbb{N}_0^d$ for all $k \in \mathcal{R}$, hence $x + \xi_k \ge x$ for any $x \in D$. In particular, if $x \in E$ then $x + \xi_k \in E$ and we conclude, from the observation that any absorbing set is

confined to a subset of $\{x \in \mathbb{N}_0^d \,|\, x \not\geq M\}$ for some M sufficiently large, that the process is not absorbed. By contraposition the desired claim holds. Note that for fixed $n \in \mathbb{N}$, the set $\{x \in E, \langle v, x\rangle = n\}$ might be empty, thus we define

$$\mathscr{N} = \{n \in \mathbb{N} \,|\, \exists\, x \in E\colon \langle v, x\rangle = n\},$$

and make the following central assumption.

Assumption 2. There exists $v \in \mathbb{N}^d$ and $\eta > 0$, $N \in \mathbb{N}$, such that, for $n \geq N$,

$$d_v(n) \geq \eta\, d^v(n),$$

and, with the limit being taken over $\mathscr{N} \subseteq \mathbb{N}$,

$$\lim_{n\to\infty} \frac{d_v(n)}{n^{1+\eta}} = \infty.$$

We note that, as $d^v(n)$ is always non-negative, this assumption assures that $d_v(n)$ is non-negative for n sufficiently large. We shall see that this assumption further ensures the ability to "come down from infinity" in finite time. In the case where there is no absorbing set, the empty sum in Definition 3.3.5 yields $d^v(n) = 0$, and we may reformulate a result of Gupta et al. (2014) (see below). In this paper, we will extend the result to the case where E is not necessarily irreducible, but satisfies Assumption 1(i), and where there may exist a non-empty absorbing set of states (see Theorem 3.5.3).

Theorem 3.3.6. For a reaction network satisfying Assumption 2, with $A = \emptyset$ and E irreducible, the associated stochastic process $(X_t : t \geq 0)$ is exponentially ergodic and thus admits a unique stationary distribution π. Further, there exist constants $C, \gamma > 0$ such that, for all probability measures μ on E,

$$\|\mathbb{P}_\mu(X_t \in \cdot) - \pi(\cdot)\|_{TV} \leq Ce^{-\gamma t}, \qquad t \geq 0.$$

Proof. For $x \in D$, let $n = \langle v, x\rangle$. Thus, by Assumption 2(ii), for any constant $c > 0$,

$$-\sum_{k=1}^{r} \lambda_k(x)\langle v, \xi_k\rangle \geq -\max_{x' \in E, \langle v, x'\rangle = n} \sum_{k=1}^{r} \lambda_k(x')\langle v, \xi_k\rangle = d_v(n) \geq cn^{1+\eta} \geq c\langle v, x\rangle,$$

for n larger than some $N \in \mathbb{N}$. The set of $x' \in D$ such that $\langle v, x'\rangle = n \leq N$ is compact, hence there is $c_1 > 0$ such that

$$d_v(n) \geq cn - c_1 = c\langle v, x\rangle - c_1.$$

We conclude that for all $x \in D$,

$$\sum_{k=1}^{r} \lambda_k(x)\langle v, \xi_k\rangle \leq c_1 - c\langle v, x\rangle,$$

hence from (Gupta et al., 2014, Proposition 4) it follows, due to irreducibility of E, that

there exist constants $C, \gamma > 0$ such that for all $x_0 \in E$,

$$\|\mathbb{P}_{x_0}(X_t \in \cdot) - \pi(\cdot)\|_{TV} \leq Ce^{-\gamma t}$$

for all $t \geq 0$. Finally, if we consider the random starting point $X_0 \sim \mu$, we find

$$\begin{aligned}\|\mathbb{P}_{\mu}(X_t \in \cdot) - \pi(\cdot)\|_{TV} &= \left\| \sum_{x_0 \in E} \mu(x_0)(\mathbb{P}_{x_0}(X_t \in \cdot) - \pi(\cdot)) \right\|_{TV} \\ &\leq \sum_{x_0 \in E} \mu(x_0)\|\mathbb{P}_{x_0}(X_t \in \cdot) - \pi(\cdot)\|_{TV} \leq \sum_{x_0 \in E} \mu(x_0) Ce^{-\gamma t} = Ce^{-\gamma t},\end{aligned}$$

as required. □

It is sufficient to have $\eta = 0$ for Theorem 3.3.6 to hold. However, as we shall see, if $A \neq \emptyset$ then $\eta > 0$ is required. The intuitive meaning is that the quasi-stationary distribution exists on the long-time, but not infinite time horizon, where the process will be absorbed. Thus if the process does not "come down from infinity in finite time", that is if $\eta = 0$, starting close to A will almost surely result in absorption while starting at "infinity" will not, contradicting uniqueness of the QSD. When no absorbing set exist, however, the quasi-stationary distribution reduces to the stationary distribution which exists on the infinite time horizon.

3.4 Verifying Assumptions

We start by introducing some notation and definitions from Champagnat and Villemonais (2017) for ease of reference.

Definition 3.4.1. A couple (V, φ) of measurable functions V and φ from $D = E \cup A$ to $\mathbb{R}$ is an admissible couple of functions if

(i) V and φ are bounded and nonnegative on D, positive on E, satisfy $V(x) = \varphi(x) = 0$ for all $x \in A$, and further

$$\inf_{x \in E} \frac{V(x)}{\varphi(x)} > 0.$$

(ii) For all sequences $(x_p)_{p \geq 1}$ in E such that $\{p \in \mathbb{N} \colon x_p \in O_n\}$ is finite for all $n \geq 1$,

$$\lim_{p \to \infty} \frac{V(x_p)}{\varphi(x_p)} = \infty, \qquad \text{and} \qquad \lim_{n \to \infty} V(X_{T_n}) = 0 \quad \mathbb{P}_x\text{-a.s. for all } x \in E.$$

(iii) LV is bounded from above and $L\varphi$ is bounded from below.

The definition of a couple of admissible functions in Champagnat and Villemonais (2017) further requires that V and φ belong to the domain of the weakened infinitesimal generator

of $(X_t\colon t \geq 0)$. However, since any function $f\colon \mathbb{N}_0^d \to \mathbb{R}$ is in this domain for discrete state spaces (Meyn and Tweedie, 1993), the requirement is automatically satisfied. Furthermore, as V, φ are bounded, the infinitesimal generator $\widehat{L}$ is defined hereon and agrees with the weakened generator L.

The question of extinction has recently attracted much attention on its own (Johnston et al., 2017). Therefore, we provide the following proposition which renders an explicit criterion for when the stochastic process associated to a stochastic reaction network goes extinct almost surely. The assumption in the proposition is weaker than Assumption 2.

Proposition 3.4.2. *Under Assumption 1, with $A \neq \emptyset$, the process $(X_t\colon t \geq 0)$ is absorbed $\mathbb{P}_x$-a.s. for all $x \in E$ if $d_v(n) > \zeta \frac{d^v(n)}{n}$ for n sufficiently large, where*

$$\zeta = \max_{k \in \mathcal{R}_A} \langle v, \xi_k \rangle, \qquad \mathcal{R}_A = \{k \in \mathcal{R} : (E + \xi_k) \cap A \neq \emptyset\}.$$

Proof. Define the norm-like function $W(x) = \langle v, x \rangle$ on $\mathbb{N}_0^d$ as in Definition 3.3.1, and let L be the weakened infinitesimal generator of $(X_t\colon t \geq 0)$. It follows from (3.3.4) and the assumption that $d_v(n) > C\frac{d^v(n)}{n}$ for n sufficiently large that for each $x \in E$ with $\langle v, x \rangle = n$,

$$\begin{aligned}
LW(x) &= \sum_{k=1}^{r} \lambda_k(x) \langle v, \xi_k \rangle \\
&= \sum_{k=1}^{r} \lambda_k(x) \langle v, \xi_k \rangle \mathbb{1}_E(x + \xi_k) + \sum_{k=1}^{r} \lambda_k(x) \langle v, \xi_k \rangle \mathbb{1}_A(x + \xi_k) \\
&\leq \max_{x' \in E, \langle v, x' \rangle = n} \sum_{k=1}^{r} \lambda_k(x') \langle v, \xi_k \rangle \mathbb{1}_E(x' + \xi_k) + \max_{x' \in E, \langle v, x' \rangle = n} \sum_{k=1}^{r} \lambda_k(x') \langle v, \xi_k \rangle \mathbb{1}_A(x' + \xi_k) \\
&\leq -d_v(n) + \zeta \max_{x' \in E, \langle v, x' \rangle = n} \sum_{k=1}^{r} \lambda_k(x') \mathbb{1}_A(x' + \xi_k) = -d_v(n) + \frac{\zeta}{n} d^v(n) < 0, \qquad (3.4.3)
\end{aligned}$$

for n sufficiently large. In particular, there exists an $N \in \mathbb{N}$ such that for $n = \langle v, x \rangle \geq N$, we have $LW(x) < 0$, hence, setting $M = \max_{x \in E\colon 1 \leq \langle v, x \rangle \leq N} \{0, LW(x)\}$, yields

$$LW(x) \leq M \cdot \mathbb{1}_{O_N}(x), \qquad x \in E.$$

Since O_N is compact, we may apply (Meyn and Tweedie, 1993, Theorem 3.1) to conclude that the process $(X_t\colon t \geq 0)$ is non-evanescent, that is,

$$\mathbb{P}_x \left(\langle v, X_t \rangle \xrightarrow{t \to \infty} \infty \right) = 0, \qquad x \in E. \qquad (3.4.4)$$

Define the discrete time jump chain $(Y_n\colon n \in \mathbb{N}_0)$ by $Y_n = X_{J_n}$, where $J_0, J_1, \ldots$ denote the jump times of $(X_t\colon t \geq 0)$ given by

$$J_0 = 0, \qquad J_{n+1} = \inf\{t \geq J_n : X_t \neq X_{J_n}\}.$$

Let $B_m = \{Y_n \in O_m \text{ i.o.}\}$ and $F = \{Y_n \in A \text{ i.o.}\} = \{Y_n \in A \text{ for some } n\}$, where the last

equality follows from A being an absorbing set. By Assumption 1, all states in $O_m \subseteq E$ have a shortest path to A (there is a path to the minimal irreducible class, and then to A, for any $x \in O_m$), which has some positive probability. For each state $x \in O_m$, let b_x be the probability of this shortest path, and define $\beta_m = \min_{x \in O_m} b_x$. As O_m is compact, $\beta_m > 0$. It follows that for each $n \in \mathbb{N}_0$ the conditioned process fulfils

$$\mathbb{P}\left(\bigcup_{k=1}^{\infty}(Y_{n+k} \in A)\Big| Y_n \in O_m\right) \geq \beta_m > 0.$$

By (Durrett, 1996, Theorem 2.3) we get

$$\mathbb{P}_y\left(B_m \backslash F\right) = 0, \tag{3.4.5}$$

with $y = Y_0 = X_0 = x$, for any $m \in \mathbb{N}$. Now, the complement of the event $\bigcup_{m=0}^{\infty} B_m$ is the event $G \cup F$, where $G = \left\{\langle v, Y_n\rangle \xrightarrow{n\to\infty} \infty\right\}$. As B_m is an increasing sequence of events in m, we obtain by monotone convergence and (3.4.5) that

$$1 = \mathbb{P}_y\left(G \cup F \cup \bigcup_m B_m\right) = \lim_{m\to\infty} \mathbb{P}_y(G \cup F \cup B_m) = \lim_{m\to\infty} \mathbb{P}_y(G \cup F) = \mathbb{P}_y(G \cup F).$$

Thus, $(Y_n \colon n \in \mathbb{N}_0)$ either tends to infinity or is eventually absorbed in A. The same holds for the full process $(X_t \colon t \geq 0)$, and by (3.4.4) we conclude that $\mathbb{P}_x(G) = 0$ hence $\mathbb{P}_x(\tau_A < \infty) = \mathbb{P}_x(X_t \in A \text{ for some } t) = 1$. In particular, we also have that

$$\lim_{n\to\infty} T_n = \tau_A,$$

thus the process is regularly absorbed, by definition. □

Note that ζ in Proposition 3.4.2 may be negative thus Assumption 2(i) is stronger and immediately provides the same conclusion of almost sure absorption of the process. Further, as we shall see in the next proposition, Assumption 2(ii) assures that the expected magnitude of X_t, in the form of $\langle v, X_t\rangle$ given $X_0 = x$, is uniformly bounded in $x \in E$ for any $t > 0$. This, in turn, implies that the time of "coming down from infinity" is finite, which is closely related to the uniqueness of QSDs. This is where $\eta > 0$ is required.

Proposition 3.4.6. *Under Assumptions 1-2 with $A \neq \emptyset$, the process $(X_t, t \geq 0)$ satisfies*

$$\tau_A = \lim_{n\to\infty} T_n < \infty \qquad \mathbb{P}_x\text{-a.s. for all } x \in E,$$

in particular, the process is absorbed $\mathbb{P}_x$-a.s. Further, $\sup_{x\in E} \mathbb{E}_x\langle v, X_t\rangle < \infty$ for any $t > 0$.

Proof. By Assumption 2(i), it follows, applying the same notation as in Proposition 3.4.2, that

$$d_v(n) \geq \eta d^v(n) > \zeta \frac{d^v(n)}{n},$$

for n sufficiently large and ζ as in the proposition. Thus by Proposition 3.4.2, the process $(X_t\colon t \geq 0)$ satisfies

$$\tau_A = \lim_{n\to\infty} T_n < \infty,$$

$\mathbb{P}_x$-a.s. for all $x \in E$. Hence the process is regularly absorbed.

The second claim is apparently a 'classical result' (Champagnat and Villemonais, 2016) but we are not aware of a proof in the literature, hence we provide one here. Let $W(x) = \langle v, x\rangle$ on $\mathbb{N}_0^d$ as in Definition 3.3.1. It follows from (3.4.3) of Proposition 3.4.2 and Assumption 2(i) that with $\langle v, x\rangle = n$,

$$\begin{aligned} LW(x) &\leq -d_v(n) + \zeta \max_{x'\in E, \langle v,x'\rangle = n} \sum_{k=1}^{r} \lambda_k(x')\mathbb{1}_A(x'+\xi_k) \\ &\leq -d_v(n) + \zeta\frac{1}{n}d^v(n) \leq -d_v(n) + \frac{\zeta}{n\eta}d_v(n) = -\left(1 - \frac{\zeta}{n\eta}\right) d_v(n). \end{aligned}$$

It follows, under Assumption 2(ii), that

$$\frac{LW(x)}{W(x)^{1+\eta}} \leq -\left(1 - \frac{\zeta}{n\eta}\right)\frac{d_v(n)}{n^{1+\eta}} \to -\infty,$$

as $n = \langle v, x\rangle \to \infty$ in $\mathscr{N}$, in which case $1 - \zeta/(n\eta)$ becomes positive. Hence, there exist constants $D_1, D_2 > 0$ such that

$$LW(x) \leq D_2 - D_1 W(x)^{1+\eta}, \qquad \text{for all } x \in E. \tag{3.4.7}$$

Since we have $\sum_{k\in\mathcal{R}} \lambda_k(x) < \infty$ for each $x \in D$, it follows from (Anderson and Kurtz, 2015, p. 12) and the equivalence of the weakened and infinitesimal generators on W that

$$W(X_t) - W(0) - \int_0^t LW(X_s)\, ds$$

is a martingale. Thus, by the martingale property we find

$$\mathbb{E}_x W(X_t) = W(x) + \int_0^t \mathbb{E}_x(LW(X_s))\, ds, \tag{3.4.8}$$

which is a form of Dynkin's formula (Kallenberg, 2001). Using the bound (3.4.7) combined with Jensen's inequality we obtain, upon differentiation of (3.4.8),

$$\frac{d}{dt}\mathbb{E}_x W(X_t) = \mathbb{E}_x(LW(X_t)) \leq \mathbb{E}_x(D_2 - D_1 W(X_t)^{1+\eta}) \leq D_2 - D_1(\mathbb{E}_x W(X_t))^{1+\eta}.$$

Define $f_x(t) = \mathbb{E}_x W(X_t)$ and choose $D_3 > D_2$. Consider the associated differential equations

$$g'_{x,\epsilon}(t) = D_3 - D_1 g_{x,\epsilon}(t)^{1+\eta}, \qquad g_{x,\epsilon}(0) = W(x) + \epsilon \tag{3.4.9}$$

$$h'_x(t) = -D_1 h_x(t)^{1+\eta}, \qquad h_x(0) = W(x). \tag{3.4.10}$$

Define $F(t,z) = D_3 - D_1 z^{1+\eta}$. It then follows that

$$\begin{aligned} f_x'(t) &< F(t, f_x(t)), & f_x(0) &= W(x), \\ h_x'(t) &< F(t, h_x(t)), & h_x(0) &= W(x), \\ g_{x,\epsilon}'(t) &= F(t, g_{x,\epsilon}(t)), & g_{x,\epsilon}(0) &= W(x) + \epsilon. \end{aligned}$$

By Petrovitsch' theorem (Mitrinovic et al., 1991, p. 316) applied twice and the fact that solutions to the ordinary differential equations are continuous in the initial value, we conclude that

$$f_x(t) \leq g_x(t), \qquad h_x(t) \leq g_x(t) \qquad t \in [0,T], \tag{3.4.11}$$

where $g_x(t) = \lim_{\epsilon \to 0} g_{x,\epsilon}(t)$. The solution to the initial value problem (3.4.9) cannot be given in explicit form, however, the associated simpler differential equation (3.4.10) does have an explicit solution for $\eta > 0$, given by

$$h_x(t) = \frac{1}{(D_1 \eta t + W(x)^{-\eta})^{1/\eta}}, \qquad t \geq 0.$$

In order to use this to bound $f_x(t)$, define the function

$$k_t(x) = D_3 t + h_x(t).$$

Then, using (3.4.11), we have

$$g_x'(t) = D_3 - D_1 g_x(t)^{1+\eta} \leq D_3 - D_1 h_x(t)^{1+\eta} = D_3 + h_x'(t) = k_x'(t).$$

Since $g_x(0) = W(x) = k_x(0)$ it follows that $g_x(t) \leq k_x(t)$ and we infer that

$$\sup_{x \in E} \mathbb{E}_x W(X_t) = \sup_{x \in E} f_x(t) \leq \sup_{x \in E} g_x(t) \leq \sup_{x \in E} k_x(t) < \infty,$$

for all fixed $t > 0$ as desired. □

Definition 3.4.12. Let $v \in \mathbb{N}^d$ and $\alpha, \beta > 1$. Define $V \colon \mathbb{N}_0^d \to \mathbb{R}$ and $\varphi \colon \mathbb{N}_0^d \to \mathbb{R}$ by

$$V(x) = \mathbb{1}_E(x) \sum_{j=1}^{\langle v,x \rangle} \frac{1}{j^\alpha}, \qquad \varphi(x) = \mathbb{1}_E(x) \sum_{j=\langle v,x \rangle + 1}^{\infty} \frac{1}{j^\beta}.$$

Lemma 3.4.13. *Under Assumption 2, for suitable choices of $\alpha, \beta > 1$, the pair (V, φ) satisfies*

(a) V, φ are bounded.

(b) There exists an integer n and a constant $C \geq 0$, such that

$$-L\varphi \leq C \mathbb{1}_{O_n}.$$

(c) There exists constants $\epsilon, C' > 0$ and $C'' \geq 0$, such that

$$LV + C' \frac{V^{1+\epsilon}}{\varphi^{\epsilon}} \leq C''\varphi.$$

Proof. As $v \in \mathbb{N}^d$, it follows that $\langle v, x\rangle \in \mathbb{N}_0$ for all $x \in E$. In particular, $V(x), \varphi(x) \leq \sum_{j=1}^{\infty} \frac{1}{j}^p$ with $p > 1$ which is a convergent hyperharmonic series. This proves (a). Thinking in terms of Riemann sums, using that $j \mapsto j^{-\beta}$ is decreasing in j for $\beta > 1$, we obtain the bound

$$\sum_{j=\langle v,x\rangle+1}^{\infty} \frac{1}{j^{\beta}} \leq \int_{\langle v,x\rangle}^{\infty} \frac{1}{y^{\beta}}\, dy = \frac{\langle v, x\rangle^{1-\beta}}{\beta - 1}.$$

Exploiting the linearity of $\langle v, \cdot\rangle$, we consider, for $x \in E$ with $n = \langle v, x\rangle$, and L the weakened generator which coincides with the infinitesimal generator since φ is bounded,

$$\begin{aligned}
L\varphi(x) =& \sum_{k=1}^{r} \lambda_k(x) \left(\varphi(x + \xi_k) - \varphi(x)\right) \\
=& \sum_{k=1}^{r} \lambda_k(x) \mathbb{1}_E(x + \xi_k) \left(\varphi(x + \xi_k) - \varphi(x)\right) + \sum_{k=1}^{r} \lambda_k(x) \mathbb{1}_A(x + \xi_k) \left(\varphi(x + \xi_k) - \varphi(x)\right) \\
=& \sum_{k=1}^{r} \lambda_k(x) \mathbb{1}_E(x + \xi_k) \left(\sum_{j=\langle v,x+\xi_k\rangle+1}^{\infty} \frac{1}{j^{\beta}} - \sum_{j=\langle v,x\rangle+1}^{\infty} \frac{1}{j^{\beta}} \right) \\
& - \sum_{k=1}^{r} \lambda_k(x) \mathbb{1}_A(x + \xi_k) \sum_{j=\langle v,x\rangle+1}^{\infty} \frac{1}{j^{\beta}} \\
=& - \sum_{k=1}^{r} \lambda_k(x) \mathbb{1}_E(x + \xi_k) \mathbb{1}_{(0,\infty)}(\langle v, \xi_k\rangle) \sum_{i=1}^{\langle v,\xi_k\rangle} \frac{1}{(\langle v, x\rangle + i)^{\beta}} \\
& + \sum_{k=1}^{r} \lambda_k(x) \mathbb{1}_E(x + \xi_k) \mathbb{1}_{(-\infty,0)}(\langle v, \xi_k\rangle) \sum_{i=\langle v,\xi_k\rangle}^{-1} \frac{1}{(\langle v, x\rangle + i + 1)^{\beta}} \\
& - \sum_{k=1}^{r} \lambda_k(x) \mathbb{1}_A(x + \xi_k) \sum_{j=\langle v,x\rangle+1}^{\infty} \frac{1}{j^{\beta}} \\
\geq& \frac{1}{\langle v, x\rangle^{\beta}} \left(- \sum_{k=1}^{r} \lambda_k(x) \mathbb{1}_E(x + \xi_k) \mathbb{1}_{(0,\infty)}(\langle v, \xi_k\rangle) \langle v, \xi_k\rangle \right. \\
& \left. - \sum_{k=1}^{r} \lambda_k(x) \mathbb{1}_E(x + \xi_k) \mathbb{1}_{(-\infty,0)}(\langle v, \xi_k\rangle) \langle v, \xi_k\rangle \right) - \sum_{k=1}^{r} \lambda_k(x) \mathbb{1}_A(x + \xi_k) \frac{\langle v, x\rangle^{1-\beta}}{\beta - 1}
\end{aligned}$$

$$
\begin{aligned}
=&\frac{1}{\langle v,x\rangle^{\beta}}\left(-\sum_{k=1}^{r}\lambda_k(x)\mathbb{1}_E(x+\xi_k)\langle v,\xi_k\rangle-\sum_{k=1}^{r}\lambda_k(x)\mathbb{1}_A(x+\xi_k)\frac{\langle v,x\rangle}{\beta-1}\right)\\
\geq&\frac{1}{\langle v,x\rangle^{\beta}}\left(-\max_{x'\in E,\langle v,x'\rangle=n}\sum_{k=1}^{r}\lambda_k(x')\langle v,\xi_k\rangle\mathbb{1}_E(x'+\xi_k)\right.\\
&\left.-\max_{x'\in E,\langle v,x'\rangle=n}\sum_{k=1}^{r}\lambda_k(x')\mathbb{1}_A(x'+\xi_k)\frac{\langle v,x'\rangle}{\beta-1}\right)\\
=&\frac{1}{n^{\beta}}\left(d_v(n)-\frac{d^v(n)}{\beta-1}\right).
\end{aligned}
$$

Using Assumption 2(i), we can choose $\beta>1$ large enough such that $L\varphi(x)\geq 0$ for all $n=\langle v,x\rangle$ sufficiently large. In particular, condition (b) is satisfied.

Similarly, as V is bounded, we find for $x\in E$ with $\langle x,v\rangle=n$,

$$
\begin{aligned}
LV(x)=&\sum_{k=1}^{r}\lambda_k(x)\left(V(x+\xi_k)-V(x)\right)\\
=&\sum_{k=1}^{r}\lambda_k(x)\mathbb{1}_E(x+\xi_k)\left(\sum_{j=1}^{\langle v,x+\xi_k\rangle}\frac{1}{j^{\alpha}}-\sum_{j=1}^{\langle v,x\rangle}\frac{1}{j^{\alpha}}\right)-\sum_{k=1}^{r}\lambda_k(x)\mathbb{1}_A(x+\xi_k)\sum_{j=1}^{\langle v,x\rangle}\frac{1}{j^{\alpha}}\\
\leq&\sum_{k=1}^{r}\lambda_k(x)\mathbb{1}_E(x+\xi_k)\mathbb{1}_{(0,\infty)}(\langle v,\xi_k\rangle)\sum_{j=\langle v,x\rangle+1}^{\langle v,x+\xi_k\rangle}\frac{1}{j^{\alpha}}\\
&-\sum_{k=1}^{r}\lambda_k(x)\mathbb{1}_E(x+\xi_k)\mathbb{1}_{(-\infty,0)}(\langle v,\xi_k\rangle)\sum_{j=\langle v,x+\xi_k\rangle+1}^{\langle v,x\rangle}\frac{1}{j^{\alpha}}\\
=&\sum_{k=1}^{r}\lambda_k(x)\mathbb{1}_E(x+\xi_k)\mathbb{1}_{(0,\infty)}(\langle v,\xi_k\rangle)\sum_{j=1}^{\langle v,\xi_k\rangle}\frac{1}{(j+\langle v,x\rangle)^{\alpha}}\\
&-\sum_{k=1}^{r}\lambda_k(x)\mathbb{1}_E(x+\xi_k)\mathbb{1}_{(-\infty,0)}(\langle v,\xi_k\rangle)\sum_{j=\langle v,\xi_k\rangle}^{-1}\frac{1}{(j+\langle v,x\rangle+1)^{\alpha}}\\
\leq&\sum_{k=1}^{r}\lambda_k(x)\mathbb{1}_E(x+\xi_k)\mathbb{1}_{(0,\infty)}(\langle v,\xi_k\rangle)\frac{\langle v,\xi_k\rangle}{\langle v,x\rangle^{\alpha}}\\
&+\sum_{k=1}^{r}\lambda_k(x)\mathbb{1}_E(x+\xi_k)\mathbb{1}_{(-\infty,0)}(\langle v,\xi_k\rangle)\frac{\langle v,\xi_k\rangle}{\langle v,x\rangle^{\alpha}}\\
\leq&\frac{1}{\langle v,x\rangle^{\alpha}}\sum_{k=1}^{r}\lambda_k(x)\langle v,\xi_k\rangle\mathbb{1}_E(x+\xi_k)\\
\leq&\frac{1}{\langle v,x\rangle^{\alpha}}\max_{x'\in E,\langle v,x'\rangle=n}\sum_{k=1}^{r}\lambda_k(x')\langle v,\xi_k\rangle\mathbb{1}_E(x'+\xi_k)=-\frac{d_v(n)}{n^{\alpha}}.
\end{aligned}
$$

Note that by treating $V(x)$ as a lower Riemann sum, for $x \in E$,

$$V(x) \leq \sum_{j=1}^{\infty} \frac{1}{j^\alpha} = 1 + \sum_{j=2}^{\infty} \frac{1}{j^\alpha} \leq 1 + \int_1^\infty \frac{1}{x^\alpha}\, dx = \frac{\alpha}{\alpha - 1},$$

and similarly, treating $\varphi(x)$ as an upper Riemann sum,

$$\varphi(x) = \sum_{j=\langle v,x\rangle+1}^{\infty} \frac{1}{j^\beta} \geq \int_{\langle v,x\rangle+1}^{\infty} \frac{1}{x^\beta}\, dx = \frac{(1+\langle v,x\rangle)^{1-\beta}}{\beta-1} \geq \frac{\langle v,x\rangle^{1-\beta}}{2(\beta-1)},$$

with the last inequality holding for $\langle v, x\rangle$ sufficiently large, using that for $\beta > 1$,

$$\frac{(1+\langle v,x\rangle)^{1-\beta}}{\langle v,x\rangle^{1-\beta}} = \left(\frac{1}{\langle v,x\rangle} + 1\right)^{1-\beta} \to 1 > \frac{1}{2}, \qquad \text{for} \quad n = \langle v,x\rangle \to \infty.$$

We infer that

$$LV(x) + \frac{V^{1+\epsilon}(x)}{\varphi^\epsilon(x)} \leq -\frac{d_v(n)}{n^\alpha} + Cn^{\epsilon(\beta-1)},$$

where $C = [\alpha/(\alpha-1)]^{1+\epsilon}[2(\beta-1)]^\epsilon$. Note that by definition $\varphi(x) > 0$ for $x \in E$ hence, choosing $\alpha = 1 + \eta/2$ and $\epsilon = \eta/[2(\beta-1)]$, we get

$$LV(x) + \frac{V^{1+\epsilon}(x)}{\varphi^\epsilon(x)} \leq -\frac{d_v(n)}{n^{1+\eta/2}} + Cn^{\eta/2} = \left(C - \frac{d_v(n)}{n^{1+\eta}}\right) n^{\eta/2}.$$

Using Assumption 2(ii), the first term becomes negative for $\langle v, x\rangle$ sufficiently large. In other words $LV(x) + \frac{V^{1+\epsilon}(x)}{\varphi^\epsilon(x)} \leq 0$ for $x \notin O_n$ with n sufficiently large. Since by definition $\varphi(x) \geq 0$, condition (c) holds. □

Lemma 3.4.14. *Under Assumpstions 1-2, the pair (V, φ) is an admissible couple of functions.*

Proof. Choose $\alpha, \beta \in \mathbb{R}$ such that the conclusions of Lemma 4.3.5 hold. In particular, $\alpha, \beta > 1$ hence the functions V and φ are bounded, non-negative on $E \cup A$ and positive on E, and by definition $V(x) = \varphi(x) = 0$ for $x \in A$. Furthermore, $\inf_{x\in E} V(x) > 0$, hence by non-negativity of φ,

$$\inf_{x\in E} \frac{V(x)}{\varphi(x)} > 0,$$

and Definition 3.4.1(1) is fulfilled. Let $(x_p)_{p\geq 1}$ be any sequence in E such that the set $\{p \in \mathbb{N} : x_p \in O_n\}$ is finite for all $n \geq 1$. Then $\langle v, x_p\rangle \to \infty$ as $p \to \infty$, hence

$$\lim_{p\to\infty} \frac{V(x_p)}{\varphi(x_p)} = \lim_{p\to\infty} \frac{\sum_{j=1}^{\langle v,x_p\rangle} \frac{1}{j^\alpha}}{\sum_{j=\langle v,x_p\rangle+1}^{\infty} \frac{1}{j^\beta}} = \infty.$$

Furthermore, since the process is regularly absorbed by Proposition 3.4.6, we have

$$\lim_{n\to\infty} V(X_{T_n}) = V(X_{\tau_A}) = 0 \qquad \mathbb{P}_x\text{-}a.s,$$

hence Definition 3.4.1(2) is fulfilled. Finally, from Lemma 4.3.5 and the fact that V and φ are both bounded functions, it follows that LV is bounded from above and $L\varphi$ is bounded from below. This concludes the proof. □

3.4.1 Lemmas

Lemma 3.4.15. *Assumption 1(i) is equivalent to the following: There exists $n_0 \in \mathbb{N}$, $\theta_0, \theta_1, a_1 > 0$ and a probability measure ν on E such that, for all $x \in O_{n_0}$ and all $s \in [\theta_0, \theta_0 + \theta_1]$,*

$$\mathbb{P}_x(X_s \in \cdot) \geq a_1 \nu, \tag{3.4.16}$$

and in addition, for all $n \geq n_0$, there exists $s_n \geq 0$ such that

$$\inf_{x\in O_n} \mathbb{P}_x(X_{s_n} \in O_{n_0}) > 0. \tag{3.4.17}$$

Proof. We first prove that Assumption 1 implies the existence of such constants and probability measure. Let $\mathcal{I}_{\min}$ be the unique minimal irreducible class contained in E. Set

$$n_0 = \inf\{n \in \mathbb{N} \colon O_n \cap \mathcal{I}_{\min} \neq \emptyset\} < \infty.$$

Pick $z \in \mathcal{I}_{\min} \cap O_{n_0}$ arbitrarily and let $\nu = \delta_z$. Pick arbitrary $\theta_0, \theta_1 > 0$ and let

$$a_1 = \inf_{s\in[\theta_0,\theta_0+\theta_1], x\in O_{n_0}} \mathbb{P}_x(X_s = z).$$

Note that for all $n \in \mathbb{N}$ the set O_n is finite and any $x \in O_n \subset E$ has a path to $\mathcal{I}_{\min}$ and thus, by irreducibility, to z. By continuity of $\mathbb{P}_x(X_\cdot = z)$, we conclude that $a_1 > 0$ and choosing $s_n = \theta_0 > 0$ for all n, the desired holds.

For the reverse direction, we first prove uniqueness of the minimal irreducible class in the endorsed set E. Suppose for contradiction that $\mathcal{I}_i \neq \mathcal{I}_j$ are irreducible minimal classes in E. Let $x_1 \in \mathcal{I}_i$ and $x_2 \in \mathcal{I}_j$. Then there is a path $x_1 \mapsto y_1$ and a path $x_2 \mapsto y_2$ with $y_1, y_2 \in O_{n_0}$. Indeed, if $\mathcal{I}_\ell \cap O_{n_0} \neq \emptyset$ for $l = i, j$ we may simply choose $y_\ell = x_\ell$, $\ell = 1, 2$. Otherwise, x_1, x_2 are in some O_{n_1} and O_{n_2} respectively with $n_1, n_2 > n_0$, hence by (3.4.17), there exist paths as described. By (3.4.16), there exist $\theta_0, \theta_1, a_1 > 0$ and a probability measure ν on E such that, for all $y \in O_{n_0}$ and all $s \in [\theta_1, \theta_1 + \theta_0]$,

$$\mathbb{P}_y(X_s \in \cdot) \geq a_1 \nu.$$

As ν is a probability measure on a countable space, there exists some $z \in E$ such that $\nu(\{z\}) > 0$. But then $\mathbb{P}_y(X_s = z) > 0$ for all $y \in O_{n_0}$. We conclude that there exist paths $y_1 \mapsto z$ and $y_2 \mapsto z$. By minimality of $\mathcal{I}_i$ and $\mathcal{I}_j$, we conclude that $z \in \mathcal{I}_i \cap \mathcal{I}_j$ which is a

contradiction. This proves the uniqueness in Assumption 1(i).

We now prove existence. Suppose for contradiction that no minimal class exists. Then, for all $\mathcal{I} \subseteq E$ there exists $\mathcal{J} \subseteq E$ such that $\mathcal{I} \to \mathcal{J}$. Repeating the argument results in an infinite path

$$\mathcal{J}_1 \to \mathcal{J}_2 \to \mathcal{J}_3 \to \dots \tag{3.4.18}$$

In the case where there are only finitely many irreducible classes in E, there must exist an $i \in \mathbb{N}$ such that $\mathcal{J}_1 \to \mathcal{J}_i = \mathcal{J}_1$ which contradicts the lack of cycles in $\mathcal{D}$.

Suppose therefore that there are infinitely many classes. Given n_0, the set O_{n_0} is finite. Thus it intersects at most finitely many irreducible classes. We conclude that in the infinite path (3.4.18), there are infinitely many $i \in \mathbb{N}$ such that

$$\mathcal{I}_i \cap O_{n_0} = \emptyset, \qquad \mathcal{I}_i \cap O_{n_i} \neq \emptyset,$$

for some $n_i > n_0$. However, by (3.4.17) each of these classes have a path to O_{n_0}. This implies the existence of at least one irreducible class intersecting O_{n_0} which appears more than once in the infinite path (3.4.18). This creates a cycle in $\mathcal{D}$ which is a contradiction. □

Lemma 3.4.19. *Under Assumpstions 1-2, for all $\lambda > 0$, there exists $n \geq 1$ such that*

$$\sup_{x \in E} \mathbb{E}_x \left(e^{\lambda(\tau_n \wedge \tau_A)} \right) < \infty. \tag{3.4.20}$$

Proof. It follows from Proposition 3.4.6 that $\sup_{x \in E} \mathbb{E}_x \langle v, X_t \rangle < M$ for all $t > 0$ and some constant $M > 0$. Let $\tau_{n,A} = \tau_n \wedge \tau_A$. Then, for $\epsilon > 0$

$$\mathbb{E}_x \langle v, X_\epsilon \rangle = \sum_{n=0}^{\infty} \mathbb{P}_x(\langle v, X_\epsilon \rangle > n) \geq \sum_{n=0}^{\infty} \mathbb{P}_x(\tau_{n,A} > \epsilon).$$

Suppose for contradiction that $\mathbb{P}_x(\tau_{n,A} > \epsilon)$ does not converge uniformly in x to 0 as $n \to \infty$. Then for any $\delta > 0$, $\sup_{x \in E} \mathbb{P}_x(\tau_{n,A} > \epsilon) > \delta$ for infinitely many $n \in \mathbb{N}$. Thus, choosing $\delta = \epsilon$, there exists a sequence $(x_i, n_i)_{i \geq 1}$ with $\langle v, x_i \rangle > n_i$ such that $\mathbb{P}_{x_i}(\tau_{n_i,A} > \epsilon) > \epsilon$ and $n_i > n_{i-1}$. But then, noting that $\mathbb{P}_x(\tau_{n-1,A} > \epsilon) > \mathbb{P}_x(\tau_{n,A} > \epsilon)$ for all $n \in \mathbb{N}$, we obtain

$$\mathbb{E}_{x_i} \langle v, X_\epsilon \rangle \geq \sum_{n=0}^{\infty} \mathbb{P}_{x_i}(\tau_{n,A} > \epsilon) \geq \sum_{n=0}^{n_i} \mathbb{P}_{x_i}(\tau_{n,A} > \epsilon) \geq n_i \mathbb{P}_{x_i}(\tau_{n_i,A} > \epsilon) > n_i \epsilon.$$

Letting $i \to \infty$ we would have $M > \lim_{i \to \infty} \mathbb{E}_{x_i} \langle v, X_\epsilon \rangle > \lim_{i \to \infty} n_i \epsilon = \infty$, which is a contradiction. We conclude that $\sup_{x \in E} \mathbb{P}_x(\tau_{n,A} > \epsilon) \leq \epsilon$ for $n \geq N(\epsilon)$.

Now, given $t > 0$ we can choose $0 < \epsilon < 1$ such that $t = p\epsilon$ for some $p \in \mathbb{N}$. Applying the Markov property we get, for $n \geq N(\epsilon)$,

$$\begin{aligned}
&\mathbb{P}_x(\tau_{n,A} \geq t) = \mathbb{P}_x(\tau_{n,A} \geq p\epsilon) \\
&= \sum_{y \in D} \mathbb{P}_x(\tau_{n,A} > p\epsilon \,|\, \tau_{n,A} > (p-1)\epsilon, X_{(p-1)\epsilon} = y)\mathbb{P}_x(\tau_{n,A} > (p-1)\epsilon, X_{(p-1)\epsilon} = y) \\
&= \sum_{y \in D} \mathbb{P}_y(\tau_{n,A} > \epsilon)\mathbb{P}_x(\tau_{n,A} > (p-1)\epsilon, X_{(p-1)\epsilon} = y) \leq \epsilon\, \mathbb{P}_x(\tau_{n,A} > (p-1)\epsilon) \leq \epsilon^p,
\end{aligned}$$

hence we conclude, as $t < p$, that

$$\sup_{x \in E} \mathbb{P}_x(\tau_{n,A} \geq t) \leq \epsilon^t.$$

Since $e^{\lambda(\tau_n \wedge \tau_A)}$ is non-negative, by choosing $\epsilon < e^{-\lambda} < 1$ we get

$$\begin{aligned}
\sup_{x \in E} \mathbb{E}_x \left(e^{\lambda(\tau_n \wedge \tau_A)} \right) &= \sup_{x \in E} \int_0^\infty \mathbb{P}_x \left(e^{\lambda(\tau_n \wedge \tau_A)} \geq t \right) dt \leq 1 + \sup_{x \in E} \int_1^\infty \mathbb{P}_x \left(\tau_{n,A} \geq \frac{\ln t}{\lambda} \right) dt \\
&\leq 1 + \int_1^\infty \epsilon^{\ln t/\lambda}\, dt = 1 + \int_0^\infty (e\epsilon^{1/\lambda})^u\, du = 1 + \frac{1}{1 + \ln(\epsilon)/\lambda} < \infty,
\end{aligned}$$

thus (3.4.20) holds as desired. □

Lemma 3.4.21. *For all $n \geq 0$, there exists a constant C_n such that, for all $t \geq 0$,*

$$\sup_{x \in O_n} \mathbb{P}_x(t < \tau_A) \leq C_n \inf_{x \in O_n} \mathbb{P}_x(t < \tau_A).$$

Further, with V from Definition 3.4.12, there exist constants $r_0, p_0 > 0$ such that for n sufficiently large,

$$\mathbb{P}_x(r_0 < \tau_A) \leq p_0 V(x), \qquad \text{for all } x \in E \backslash O_n.$$

Proof. Let $n \geq 0$ be given. If O_n is empty, then the statement is vacuously true for any C_n. If O_n is non-empty, then as $\sup_{x \in O_n} \mathbb{P}_x(t < \tau_A) \leq 1$ and $\inf_{x \in O_n} \mathbb{P}_x(t < \tau_A) > 0$ since O_n is finite, we may simply choose

$$C_n = \frac{\sup_{x \in O_n} \mathbb{P}_x(t < \tau_A)}{\inf_{x \in O_n} \mathbb{P}_x(t < \tau_A)} < \infty.$$

To see the second claim, note that letting

$$p_0 = \frac{1}{\inf_{x \in E} V(x)} > 0,$$

we have for any $r_0 > 0$,

$$\mathbb{P}_x(r_0 < \tau_A) \leq 1 = p_0 \inf_{x \in E} V(x) \leq p_0 V(x),$$

for all $x \in E$, and the desired holds. □

3.5 The Main Result

We are now ready to state and prove the main result of the paper. In the case where only a single endorsed set, E, is considered, we find that the unique QSD hereon is in fact globally attracting in the space of probability measures on E.

Theorem 3.5.1. A reaction network $(\mathcal{N}, \lambda)$ with associated stochastic process $(X_t \colon t \geq 0)$ on $D = E \sqcup A$, with $A \neq \emptyset$ and satisfying Assumption 1-2, admits a unique quasi-stationary distribution ν. Further, there exist constants $C, \gamma > 0$ such that, for all probability measures μ on E,

$$\|\mathbb{P}_\mu(X_t \in \cdot \,|\, t < \tau_A) - \nu\|_{TV} \leq Ce^{-\gamma t}, \qquad t \geq 0.$$

Proof. By Proposition 3.4.6, the process is regularly absorbed, and by Lemma 3.4.14 the pair (V, φ) given in Definition 3.4.12 is admissible, satisfying conditions (a) and (b) from Lemma 4.3.5. The result now follows from Lemma 5.3.14-3.4.21 together with (Champagnat and Villemonais, 2017, Cor. 2.8). □

Corollary 3.5.2. Let ν be the unique quasi-stationary distribution on E and $\mathcal{I}_{\min}$ the unique minimal class of E. Then supp $\nu = \mathcal{I}_{\min}$.

Proof. Suppose for contradiction that supp $\nu \not\subseteq \mathcal{I}_{\min}$. Then there exists a point $y \in \mathcal{I} \neq \mathcal{I}_{\min}$ for which $\nu(\{y\}) > 0$, where $\mathcal{I}$ is an irreducible class of E. As ν is globally attracting in $\mathcal{P}(E)$, the space of probability distributions on E, it follows by definition that

$$\lim_{t\to\infty} \mathbb{P}_\mu(X_t \in B \,|\, t < \tau_A) = \nu(B),$$

for any $\mu \in \mathcal{P}(E)$ and any measurable set $B \subseteq E$. In particular, letting $\mu = \delta_z$ with $z \in \mathcal{I}_{\min}$ yields, by minimality

$$0 = \lim_{t\to\infty} \mathbb{P}_{\delta_z}(X_t = y \,|\, t < \tau_A) = \nu(\{y\}) > 0,$$

which is a contradiction. We conclude that supp $\nu \subseteq \mathcal{I}_{\min}$. In particular, there exists $x' \in \mathcal{I}_{\min}$ such that $\nu(\{x'\}) > 0$. Further, since $\mathcal{I}_{\min}$ is irreducible, $\mathbb{P}_x(X_t = y) > 0$ for all $x, y \in \mathcal{I}_{\min}$. As ν is a QSD, it follows from (Collet et al., 2013, p. 48) that

$$e^{-\theta t}\nu(y) = \sum_{x\in E} \nu(x)\mathbb{P}_x(X_t = y) = \sum_{x\in\mathcal{I}_{\min}} \nu(x)\mathbb{P}_x(X_t = y) > \nu(x')\mathbb{P}_{x'}(X_t = y) > 0.$$

for some $\theta > 0$ and all $t \geq 0$, $y \in \mathcal{I}_{\min}$. Consequently, supp $\nu \supseteq \mathcal{I}_{\min}$ which in turn implies supp $\nu = \mathcal{I}_{\min}$ as desired. □

We now examine the general holistic setting with state space $D = D_E \cup D_A$, where D_E consists of possibly several endorsed sets, each with or without a corresponding non-empty absorbing set. As one might in practice not have complete information about the starting-point of the process, one may in general not know exactly which endorsed set the process

evolves in. However, one may have a qualitative guess in the form of an initial distribution, μ on D_E.

The following theorem shows that we may consider the problem of finding a unique QSD on each endorsed set independently and then piecing these together to form a unique limiting measure up to a choice of the initial distribution, μ, on D_E.

Theorem 3.5.3. Let $(\mathcal{N}, \lambda)$ be a reaction network and $U = E_1 \cup \cdots \cup E_m$ a finite union of endorsed sets. If Assumption 1-2 are satisfied, then the associated stochastic process $(X_t \colon t \geq 0)$ admits a unique quasi-stationary distribution, ν_n, on each endorsed set $E_n \subseteq U$, $n = 1, \ldots, m$. Furthermore, given an initial distribution μ on U, the measure ν_μ defined by

$$\nu_\mu(B) = \sum_{n=1}^{m} \mu(E_n)\nu_n(B \cap E_n),$$

is well defined and there exist constants $C, \gamma > 0$ such that,

$$\|\mathbb{P}_\mu(X_t \in \cdot \,|\, t < \tau_A) - \nu_\mu\|_{TV} \leq Ce^{-\gamma t}, \qquad t \geq 0.$$

Proof. That $(X_t \colon t \geq 0)$ admits a unique quasi-stationary distribution, ν_n, on each endorsed set $E_n \subseteq U$ with $A_n \neq \emptyset$ follows from Theorem 4.5.5. Now, suppose $A_n = \emptyset$ for some n. Then the definitions of quasi-stationary distribution and stationary distribution on E_n are equal, $\tau_{A_n} = \infty$ and all the previous proofs go through without changes. In particular, it follows from Corollary 3.5.2 that any stationary distribution, π_n, is supported by the unique minimal irreducible class as well. Furthermore, by Theorem 4.5.5, for any $\mu \in \mathcal{P}(E_n)$ there exists $C_1, \gamma_1 > 0$ such that for any $B \subseteq E_n$

$$\|\mathbb{P}_\mu(X_t \in B) - \pi_n(B)\|_{TV} = \|\mathbb{P}_\mu(X_t \in B \,|\, t < \tau_A) - \nu_n(B)\|_{TV} \leq C_1 e^{-\gamma_1 t}.$$

Clearly, ν_μ is a well defined probability measure. Finally, we may let

$$C = \max_{n \in \{1,\ldots,m\}} C_n > 0, \qquad \gamma = \min_{n \in \{1,\ldots,m\}} \gamma_n > 0,$$

where $C_n, \gamma_n > 0$ for $n \in \{1, \ldots, m\}$ are given through Theorem 4.5.5 and the above argument, for each class E_n. This proves the desired. □

Corollary 3.5.4. A one-species reaction network $(\mathcal{N}, \lambda)$ has $gcd(\xi_1, \ldots, \xi_r) < \infty$ endorsed sets. If in addition the kinetics λ is stochastic mass-action, up to a choice of the initial distribution μ on D_E, there is a unique QSD on D_E if the highest order of a reaction is at least 2 and

$$\sum_{k \in \mathcal{R}^*} \lambda_k(x)\xi_k < 0, \qquad \text{for } \langle v, x \rangle \text{ sufficiently large,}$$

where $\mathcal{R}^*$ is the set of reactions of highest order.

Proof. Note first that by Corollary 3.2.5, any one-species reaction network has a finite number of endorsed sets, as rank $\Xi = 1$. In fact, the number is given by $\gcd(\xi_1, \ldots, \xi_r)$. Indeed,

as $\gcd(\xi_1, \dots, \xi_r) \leq |\xi_k|$ it follows that $\mathcal{P}(\gcd(\xi_1, \dots, \xi_r)) \cap \mathcal{L}(\Xi) = \emptyset$ where $\mathcal{P}(B)$ is the fundamental parallelepiped generated by B and $\mathcal{L}(\Xi)$ is the lattice generated by Ξ as introduced in the proof of Proposition 3.2.4. Therefore, $b_1 = \gcd(\xi_1, \dots, \xi_r)$ is a basis for $\mathcal{L}(\Xi)$ and we conclude that the number of endorsed sets is $|\det(b_1)| = b_1$ (Dadush, 2013).

As there is only one species, we may without loss of generality take $v = 1$. It follows that $d_v(n) = -\sum_{k \in \mathcal{R}} \lambda_k(n)\xi_k > 0$ for n sufficiently large and $d_v(n) = \mathcal{O}(n^a)$ with $a \geq 2$ by assumption. Furthermore, $d^v(n) = 0$ for n sufficiently large. Thus Assumption 2 is satisfied. We infer the desired by Theorem 3.5.3. □

3.6 Examples

In this section, the main theorems and their applicability are illustrated through a series of examples. In particular, we show explicitly how the results of Champagnat and Villemonais (2017) are extended.

Example 3.6.1.

$$mS_1 \xrightarrow{\alpha_1} \emptyset$$

As discussed in Section 2, the endorsed sets and corresponding absorbing sets are

$$E_i = \{i + pm - 1 \,|\, p \in \mathbb{N}\}, \qquad A_i = \{i - 1\}, \qquad \text{for } i = 1, \dots, m,$$

respectively. These endorsed sets are evidently not irreducible. However, $\{i + m - 1\}$ is the unique minimal irreducible class in E_i from which one may jump directly to A_i, thus Assumption 1 is satisfied. Further, assuming mass-action kinetics,

$$\begin{aligned} d_v(n) &= \alpha_1 n(n-1)\dots(n-m+1)m\mathbb{1}_E(n-m) = \mathcal{O}(n^m), \\ d^v(n) &= \mathcal{O}(1), \end{aligned}$$

for n sufficiently large, hence we conclude by Theorem 3.5.3 that there exists a unique QSD on each E_i if $m \geq 2$. Further, in this case, for any initial distribution, μ, on the full set of endorsed states, $D_E = \{m, m+1, \dots\}$, the measure $\mathbb{P}_\mu(X_t \in \cdot \,|\, t < \tau_A)$ tends to ν_μ exponentially fast for $t \to \infty$. △

Example 3.6.2 (Lotka-Volterra)**.**

$$S_1 \xrightarrow{\alpha_1} 2S_1$$
$$S_1 + S_2 \xrightarrow{\alpha_2} 2S_2$$
$$S_2 \xrightarrow{\alpha_3} \emptyset$$

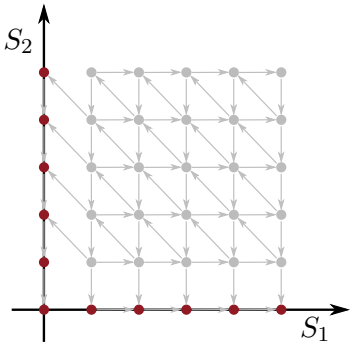

The Lotka-Volterra system describing competitive and predator-prey interactions has been of interest for approximately a century (Lotka, 1925; Volterra, 1926). In the stochastic description of the model, we find $R = \mathbb{N}^2$ hence it follows that the state space can be divided into the endorsed and absorbing sets given by

$$D_E = E = \mathbb{N}^2, \qquad D_A = A = \mathbb{N}_0^2 \backslash \mathbb{N}^2,$$

respectively. Using mass-action kinetics, it follows that for $v = (v_1, v_2) \in \mathbb{N}^2$,

$$\begin{aligned} d_v(n) &= -\max_{x \in E, \langle v, x \rangle = n} \sum_{k=1}^{r} \lambda_k(x) \langle v, \xi_k \rangle \mathbb{1}_E(x + \xi_k) \\ &= -\max_{x \in E, \langle v, x \rangle = n} (\alpha_1 v_1 x_1 + (v_2 - v_1)\alpha_2 x_1 x_2 \mathbb{1}_{\{2,3,\dots\}}(x_1) - v_2 \alpha_3 x_2 \mathbb{1}_{\{2,3,\dots\}}(x_2)). \end{aligned}$$

Letting $(v_1, v_2) = (\ell, 1)$ yields $d_v(n) = -\ell\alpha_1(n-1) + (\ell - 1)\alpha_2(n-1)$. Thus, choosing ℓ sufficiently large, $d_v(n) = \mathcal{O}(n)$ and $d_v(n) > 0$ for n sufficiently large, provided that $\alpha_2 > \alpha_1$. Note also that $\max_{k \in \mathcal{R}_A} \langle v, \xi_k \rangle < 0$ hence by Proposition 3.4.2 the process is $\mathbb{P}_x$-a.s. absorbed for all $x \in E$ if $\alpha_2 > \alpha_1$. However, $d^v(n) = \mathcal{O}(n^2)$, hence Assumption 2(i) is not satisfied, and one can not apply Theorem 4.5.5.

Using generalized mass-action (Müller and Regensburger, 2012) for the same standard Lotka-Volterra network, we may obtain a different result. Suppose for example that

$$\lambda_1(x) = \alpha_1 x_1 x_2, \quad \lambda_2(x) = \alpha_2 x_1^4 x_2^2, \quad \lambda_3(x) = \alpha_3 x_2^3.$$

Choosing $v = (2, 1)$, say, we find

$$d_v(n) = -\max_{x \in E, \langle v, x \rangle = n} \left(2\alpha_1 x_1 x_2 - \alpha_2 x_1^4 x_2^2 \mathbb{1}_{\{2,3,\dots\}}(x_2) - \alpha_3 x_2^3 \mathbb{1}_{\{2,3,\dots\}}(x_2)\right) = \mathcal{O}(n^3),$$

and likewise

$$d^v(n) = \max_{x \in E, \langle v, x \rangle = n} n \left(\alpha_2 x_1^4 x_2^2 \mathbb{1}_{\{1\}}(x_1) + \alpha_3 x_2^3 \mathbb{1}_{\{1\}}\right)(x_2) = \mathcal{O}(n^3).$$

Thus Assumption 2 is satisfied. As E is irreducible, Assumption 1 is also satisfied and we

conclude by Theorem 4.5.5 that there is a unique QSD on E.

$$\begin{aligned} S_1 &\underset{\alpha_2}{\overset{\alpha_1}{\rightleftharpoons}} 2S_1 \\ S_1 + S_2 &\underset{\alpha_4}{\overset{\alpha_3}{\rightleftharpoons}} 2S_2 \\ S_2 &\overset{\alpha_5}{\to} \emptyset \end{aligned} \tag{3.6.3}$$

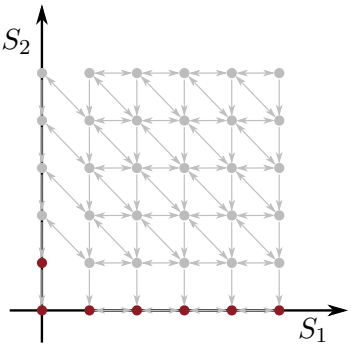

Let us now consider the slightly altered version of the original Lotka-Volterra system using mass-action kinetics, obtained by addition of the reactions $2S_1 \to S_1$ and $2S_2 \to S_1 + S_2$. In this case, there is still one endorsed set with corresponding absorbing set given by

$$E = \mathbb{N}_0 \times \mathbb{N} \backslash \{(0,1)\}, \qquad A = \mathbb{N}_0^2 \backslash E,$$

respectively.

For a general $v = (v_1, v_2) \in \mathbb{N}^2$ it follows that $d^v(n) = \mathcal{O}(n)$ and

$$\begin{aligned} d_v(n) &= -\max_{x \in E, \langle v, x \rangle = n} \sum_{k=1}^{r} \lambda_k(x) \langle v, \xi_k \rangle \mathbb{1}_E(x + \xi_k) \\ &= \min_{x \in E, \langle v, x \rangle = n} (-\alpha_1 x_1 v_1 + v_1 \alpha_2 x_1 (x_1 - 1) - (v_2 - v_1)\alpha_3 x_1 x_2 \\ &\qquad + (v_2 - v_1)\alpha_4 x_2 (x_2 - 1) + v_2 \alpha_5 x_2 \mathbb{1}_{\{2,3,\dots\}}(x_2)). \end{aligned}$$

We need $v_2 > v_1$ for the coefficient of the 4th reaction to be positive. Further, by the second derivative test, $d_v(n) = \mathcal{O}(n^2)$ exactly when

$$4 v_1 \alpha_2 (v_2 - v_1)\alpha_4 > (v_2 - v_1)^2 \alpha_3^2.$$

The set of possible v-vectors, $\mathcal{V}$, is therefore

$$\mathcal{V} = \left\{ v \in \mathbb{N}^2 \,\middle|\, v_1 < v_2 < v_1 \left(1 + \frac{4\alpha_2 \alpha_4}{\alpha_3^2}\right) \right\},$$

which is non-empty for any positive reaction rates. A particular choice would be $v = (\ell, \ell+1)$, for ℓ sufficiently large. As E is irreducible, Assumption 1 is satisfied and we conclude by Theorem 4.5.5 that the modified Lotka-Volterra system has a unique QSD on E for any reaction rates. △

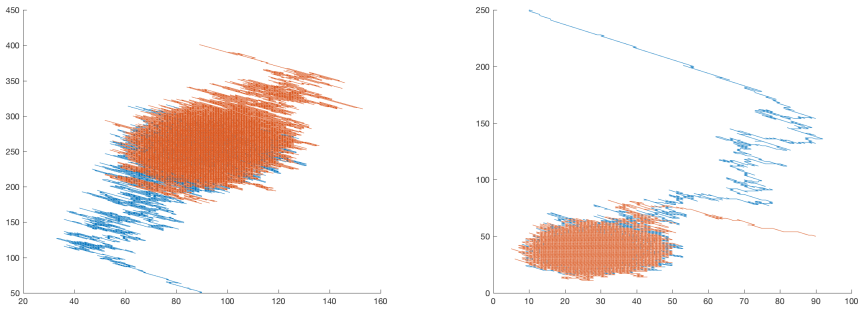

Figure 3.3: Two realizations of the process $(X_t : t \geq 0)$, associated with the reaction network (3.6.3), with differing X_0. Left: $\alpha_1 = 30, \alpha_2 = 0.3, \alpha_3 = 2, \alpha_4 = 0.7, \alpha_5 = 1$. Right: $\alpha_1 = 20, \alpha_2 = 0.7, \alpha_3 = 1, \alpha_4 = 0.7, \alpha_5 = 1$.

Example 3.6.4.

$$S_1 \xrightarrow{\alpha_1} 3S_1$$
$$2S_2 \xrightarrow{\alpha_2} \emptyset$$
$$2S_1 + S_2 \xrightarrow{\alpha_3} 2S_2$$

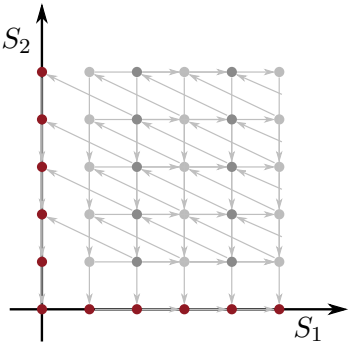

There are two endorsed sets given by

$$E_1 = \{x \in \mathbb{N}^2 : x_1 = 1 \mod 2\},$$
$$E_2 = \{x \in \mathbb{N}^2 : x_1 = 0 \mod 2\}.$$

The corresponding set of absorbing sets are given by

$$A_1 = \{x \in \mathbb{N}_0^d : x_2 = 0, x_1 = 1 \mod 2\}$$
$$A_2 = \{x \in \mathbb{N}_0^d : x_2 = 0, x_1 = 0 \mod 2\} \cup \{x \in \mathbb{N}_0^d : x_1 = 0\}.$$

Assuming mass action kinetics, it follows that for a general $v = (v_1, v_2) \in \mathbb{N}^2$ we obtain for E_1 and E_2 respectively,

$$\begin{aligned} d_v(n) = - \max_{x \in E, \langle v, x \rangle = n} \big(& 2v_1\alpha_1 x_1 - 2v_2\alpha_2 x_2(x_2 - 1)\mathbb{1}_{\{3,4,\dots\}}(x_2) \\ & + (-2v_1 + v_2)\alpha_3 x_1(x_1 - 1)x_2 \big), \\ d_v(n) = - \max_{x \in E, \langle v, x \rangle = n} \big(& 2v_1\alpha_1 x_1 - 2v_2\alpha_2 x_2(x_2 - 1)\mathbb{1}_{\{3,4,\dots\}}(x_2) \\ & + (-2v_1 + v_2)\alpha_3 x_1(x_1 - 1)x_2 \mathbb{1}_{\{3,4,\dots\}}(x_1) \big), \end{aligned}$$

which are both $\mathcal{O}(n^2)$ exactly if $v_2 < 2v_1$. A particular choice would be $v = (1, 1)$. Further,

$$d^v(n) = \max_{x \in E, \langle v, x \rangle = n} n\big(\alpha_2 x_2(x_2 - 1)\mathbb{1}_{\{2\}}(x_2) + \alpha_3 x_1(x_1 - 1)x_2\mathbb{1}_{\{2\}}(x_1)\big) = 2\alpha_3 n^2 = \mathcal{O}(n^2).$$

We conclude that there exists an $0 < \eta < 1$ such that Assumption 2 holds. Since both E_1 and E_2 are irreducible, Assumption 1 is satisfied hence Theorem 3.5.3 applies regardless of the rate constants – there exists a unique QSD, ν_n, on each E_n, and given an initial distribution, μ, on D_E, the measure $\mathbb{P}_x(X_t \in \cdot \,|\, t < \tau_A)$ approaches

$$\nu_\mu = \mu(E_1)\nu_1(\cdot \cap E_1) + \mu(E_2)\nu_1(\cdot \cap E_2)$$

exponentially fast in t.

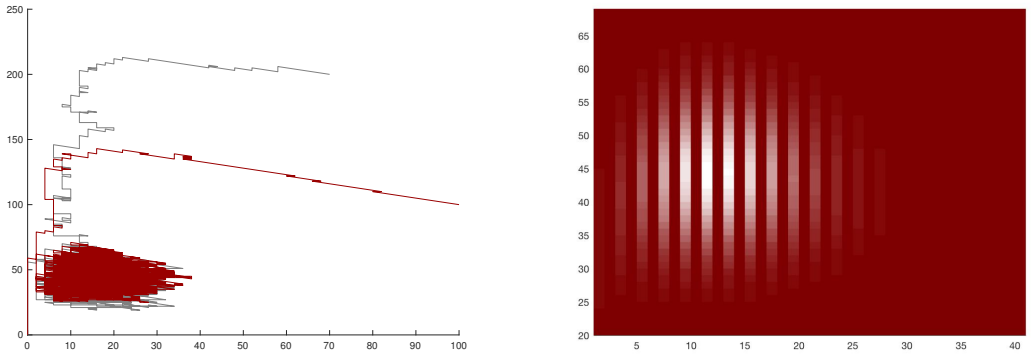

Figure 3.4: Left: Two realizations of $(X_t \colon t \geq 0)$ with $\alpha_1 = 300, \alpha_2 = 1, \alpha_3 = 0.5$. $X_0 = (70, 200)$ in grey and $X_0 = (100, 100)$ in red. Right: Approximate density of the QSD on one of the two endorsed sets.

Note that if we had modeled the network using deterministic mass action (Gunawardena, 2003), we would have found the attracting[2] fixed point

$$(x, y) = \left(\frac{(2\alpha_1\alpha_2)^{1/3}}{\alpha_3^{2/3}}, \frac{\alpha_1^{2/3}}{(2\alpha_2\alpha_3)^{1/3}} \right),$$

which seems to lie near the peak of the quasi-stationary distribution for the parameter values used in Figure 4.4. Indeed, we find the fixed point $(13.39, 44.81)$ which is a stable spiral. $\triangle$

Example 3.6.5 (Birth-death)**.**

$$S_2 \overset{\alpha_1}{\leftarrow} S_1 + S_2 \overset{\alpha_2}{\rightarrow} 2S_1 + S_2 \qquad S_1 \overset{\alpha_3}{\leftarrow} S_1 + S_2 \overset{\alpha_4}{\rightarrow} S_1 + 2S_2$$
$$S_1 \overset{\alpha_5}{\leftarrow} 2S_1 \overset{\alpha_6}{\rightarrow} 3S_1 \qquad S_2 \overset{\alpha_7}{\leftarrow} 2S_2 \overset{\alpha_8}{\rightarrow} 3S_2$$

[2]One can calculate the trace of the Jacobin matrix to be $-2\alpha_1 - (2\alpha_1\alpha_2)^{2/3}\alpha_3^{-1/3} < 0$, making the fixed point attracting for all parameter values.

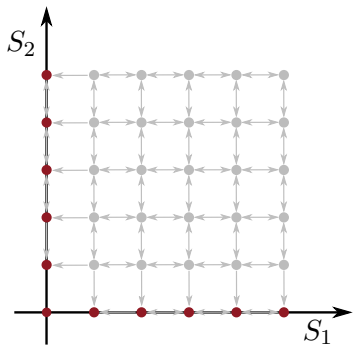

The full set of endorsed states is found to be $D_E = \mathbb{N}^2$ which is irreducible, and the full set of absorbing states is $D_A = \mathbb{N}_0^2 \backslash \mathbb{N}^2$. Thus Assumption 1 is satisfied. For a general $v \in \mathbb{N}^2$ we have, with $\alpha_5 - \alpha_6 = \rho_1, \alpha_7 - \alpha_8 = \rho_2, \alpha_1 - \alpha_2 = \rho_3$ and $\alpha_3 - \alpha_4 = \rho_4$,

$$d_v(n) = \min_{x \in E, \langle v, x \rangle = n} (v_1\rho_3 + v_2\rho_4)x_1x_2 - v_1\rho_1 x_1(x_1 - 1) - v_2\rho_2 x_2(x_2 - 1).$$

The second derivative test gives the sufficient and necessary criteria for $d_v(n)$ to be $\mathcal{O}(n^2)$,

$$4v_1v_2\rho_1\rho_2 - (v_1\rho_3 + v_2\rho_4)^2 > 0.$$

Solving for v_2 we find the set of possible v-vectors,

$$\mathcal{V} = \left\{ v \in \mathbb{N}^2 \,\middle|\, \frac{2\rho_1\rho_2 - \rho_3\rho_4 - 2\sqrt{\rho_1^2\rho_2^2 - \rho_1\rho_2\rho_3\rho_4}}{\rho_4^2} v_1 < v_2 < \frac{2\rho_1\rho_2 - \rho_3\rho_4 + 2\sqrt{\rho_1^2\rho_2^2 - \rho_1\rho_2\rho_3\rho_4}}{\rho_4^2} v_1 \right\},$$

which is non-empty exactly when $\rho_1\rho_2 > \rho_3\rho_4$. Choosing parameter values as, say,

$$X_2 \overset{1}{\leftarrow} X_1 + X_2 \overset{3}{\rightarrow} 2X_1 + X_2 \qquad X_1 \overset{0.75}{\leftarrow} X_1 + X_2 \overset{1}{\rightarrow} X_1 + 2X_2$$
$$X_1 \overset{2}{\leftarrow} 2X_1 \overset{1}{\rightarrow} 3X_1 \qquad X_2 \overset{2}{\leftarrow} 2X_2 \overset{1}{\rightarrow} 3X_2$$

we find

$$\mathcal{V} \approx \{ v \in \mathbb{N}^2 \,|\, 1.373v_1 < v_2 < 46.627v_1 \}.$$

Note that $v = (1, 1) \notin \mathcal{V}$, and Theorem 4.5.5 would not be applicable with this choice. We have therefore extended the result of Champagnat and Villemonais (2017). △

Example 3.6.6.

$$S_1 \overset{\alpha_1}{\rightarrow} 3S_1$$
$$S_1 + 3S_2 \overset{\alpha_2}{\rightarrow} 2S_2$$
$$3S_1 \overset{\alpha_3}{\rightarrow} 2S_1 + 3S_2$$

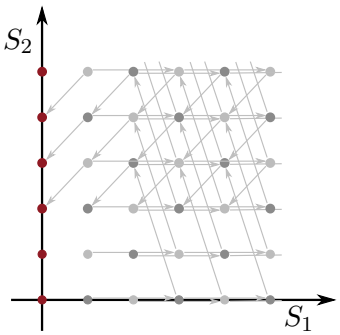

The stoichiometric matrix is in this example given by

$$\Xi = \begin{pmatrix} 2 & -1 & -1 \\ 0 & -1 & 3 \end{pmatrix}.$$

It follows that $\operatorname{rank}\Xi = 2 = d$ hence by Proposition 3.2.4, there is a finite number endorsed sets for x sufficiently large. Indeed, upon inspection, we find two endorsed sets given by

$$E_1 = \{x \in \mathbb{N}^2 : x_2 \geq 0, x_1 + x_2 = 0 \mod 2\},$$
$$E_2 = \{x \in \mathbb{N}^2 : x_2 \geq 0, x_1 + x_2 = 1 \mod 2\},$$

and these each have a unique minimal irreducible class hence Assumption 1 is satisfied. Note, however, that the sets E_1 and E_2 are not irreducible. Here $D_E = \{x \in \mathbb{N}_0^2 : x_1 \geq 1\}$ and $D_A = \{x \in \mathbb{N}_0^2 : x_1 = 0\}$. Taking $v = (4, 1)$, it follows that for both classes

$$d_v(n) = -\max_{x \in E, \langle v, x\rangle = n} (8\alpha_1 x_1 - 5\alpha_2 x_1 x_2 (x_2 - 1)(x_2 - 2)\mathbb{1}_{\{2,3,\dots\}}(x_1) - \alpha_3 x_1 (x_1 - 1)(x_1 - 2)),$$

which is $\mathcal{O}(1)$. Furthermore,

$$d^v(n) = \max_{x \in E, \langle v, x\rangle = n} n x_1 x_2 (x_2 - 1)(x_2 - 2)\mathbb{1}_{\{1\}}(x_1) = \mathcal{O}(n^4).$$

Thus, Assumption 2 is not satisfied and we can not apply Theorem 3.5.3. △

Example 3.6.7.

$$S_1 + S_2 \xrightarrow{\alpha_1} \emptyset$$
$$S_1 \xrightarrow{\alpha_2} 2S_1 + S_2$$
$$2S_1 + 2S_2 \xrightarrow{\alpha_3} S_1 + S_2$$

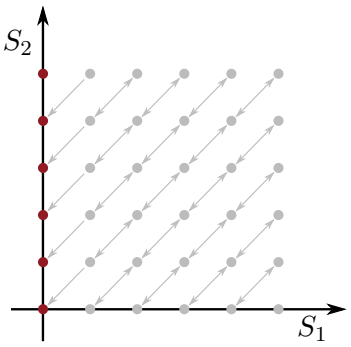

As $\operatorname{rank}\Xi = 1 < d$, there are infinitely many endorsed sets. These are each irreducible. Indeed, upon inspection, we find the endorsed sets

$$E_i = \{x \in \mathbb{N}_0^2 : x_1 \geq 1, x_2 = x_1 + (-1)^i \lfloor i/2 \rfloor\}, \qquad i \geq 1,$$

and the full set of absorbing states $D_A = \{x \in \mathbb{N}_0^2 : x_1 = 0\}$. Furthermore, some sets (E_i with $i \equiv 1 \mod 2$) have no corresponding absorbing set while all others do. For all classes however, we find

$$d_v(n) = \mathcal{O}(n^2), \qquad d^v(n) = \mathcal{O}(n),$$

hence, by Theorem 3.5.3, we conclude that for any initial distribution μ with support contained in a finite subset of the endorsed sets, the measure $\mathbb{P}_\mu(X_t \in \cdot \,|\, t < \tau_A)$ converges

exponentially fast in t to ν_μ. Had we instead looked at the slightly altered network

$$
\begin{aligned}
S_1 + S_2 &\xrightarrow{\alpha_1} \emptyset \\
S_1 &\xrightarrow{\alpha_2} 2S_1 + S_2 \\
2S_1 + 2S_2 &\xrightarrow{\alpha_3} S_1 + S_2 \\
3S_1 + S_2 &\xrightarrow{\alpha_4} 3S1
\end{aligned}
$$

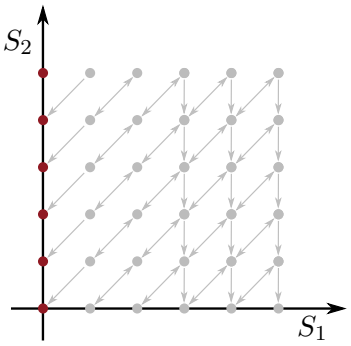

then Assumption 1 would no longer be satisfied as no minimal irreducible class exists. △

3.7 Conclusions

In this paper, we have provided sufficient conditions for the existence and uniqueness of a quasi-stationary distribution, within each endorsed set, for general stochastic reaction networks. In particular, we have provided sufficient conditions for the existence and uniqueness of a globally attracting quasi-stationary distribution in the space of probability distributions.

The requirement that for mass-action reaction systems there exists a stoichiometric coefficient strictly greater than one for each species is strong, however, it seems to be intrinsic to the problem of guaranteeing uniqueness of the QSD. Indeed, this can be seen already in 1-dimensional systems, and one can imagine that getting stuck near one of the axes would approximately reduce the multi-dimensional system to such a 1-dimensional case. It is an important question to determine sufficient conditions for just the existence of a QSD, in which case one would expect much weaker conditions to be satisfied. This is still an open problem.

The application of our results depends on the existence of a vector $v \in \mathbb{N}^d$, such that Assumption 2 holds. One would like to have easy graphical ways of guaranteeing this existence or an explicit algorithmic way of constructing such a vector. This may not be possible in general, since the problem is equivalent to determining the sign of a multivariate polynomial in a certain region of the positive orthant. However, the specific network at hand is often prone to analytical ad hoc methods, and even in lack of this, one may rely on numerical methods to find a suitable candidate for v. Another caveat is the exact calculation of the endorsed sets. Methods for finding these numerically, in the case where the intensity functions are positive in the positive orthant, exist (Gupta and Khammash, 2013; Paulevé et al., 2014). However, the problem becomes more complicated in the general case (Gupta et al., 2014).

Knowing the existence of a unique QSD, one would of course like to know the explicit analytical expression for this distribution on each endorsed set. However, this seems to be a very hard problem to solve in general, and even for 1-dimensional systems it has not yet been fully resolved. Thus, so far, one is forced to apply numerical methods or rely on analytical approximations, see Groisman and Jonckheere (2012); Villemonais (2015).

Another problem which is important for applications, is the question of observability. Indeed, the relative sizes of the time to extinction, τ_A, versus the time to reach the mean of the QSD, τ_ν, determines whether we are likely to observe the QSD or not. Only if $\tau_A \gg \tau_\nu$ would we expect the process to behave according to the QSD (Vellela and Qian, 2007). Few recent results on this matter exist for limited scenarios (Gupta et al., 2014), while methods exploiting the WKB approximation (Assaf and Meerson, 2017) are more general although not yet fully rigorous (Chazottes et al., 2017).

Finally, numerical evidence seems to suggest a strong connection between the deterministic and stochastic models of the same underlying reaction network. Indeed, for systems close to thermodynamic equilibrium, also referred to as the fluid limit (Kurtz, 1970), the modes of the QSD appear to be located near the deterministic steady states. However, far from equilibrium, the picture may be radically different. Our result can be seen as a stochastic analogue to the deterministic case of having an equilibrium point within each stoichiometric compatibility class. In this light, the QSD bridges the gap between the knowledge of extinction in the stochastic description and the existence of a stationary steady state in the deterministic setting. Future work lies in analyzing what can be inferred about the QSD from the corresponding deterministic dynamical system.

4

Quasi-Stationary Distributions for Randomly Perturbed Reaction Networks

Mads Christian Hansen and Sebastian Schreiber
University of Copenhagen
University of California, Davis

Abstract. Under the classical scaling of reaction networks, with ϵ being the system size variable, we may consider the family of associated Markov processes $\{X_t^\epsilon\}_{\epsilon>0}$ as a random perturbation of the corresponding deterministic system, described by the semi-flow $\varphi_t : D \to D$. Consistent with many reaction networks, we consider the case where one may decompose the state space into the disjoint union of invariant sets $D = D_0 \sqcup D_1$ with D_0 being closed and absorbing, and analyze the associated quasi-stationary distributions $\{\nu_\epsilon\}_{\epsilon>0}$. We show that the expected time to extinction scales exponentially with ϵ. Furthermore, if there exists a minimal Morse decomposition of the deterministic dynamics, with at least one Morse set being an attractor in D_1, then the weak* limit points of ν_ϵ are supported by the union of attractors in D_1. To illustrate the implications of the results, we provide numerical calculations for a series of reaction networks, exhibiting a range of qualitatively different dynamical properties.

Keywords: *Quasi-statonary distribution, Morse-decomposition, random perturbations, coexistence.*

Publication details: Unpublished

4.1 Introduction

The long-term behavior of a given reaction network (section 4.2.1) may depend crucially on whether it is modeled deterministically or stochastically (Gupta et al., 2014). This is unfortunate as the analysis is often much easier for the deterministic case where graphical properties of the network may yield valuable information of the dynamics (Craciun et al., 2013), while our true interest lies in the stochastic description which is more complete (Qian, 2011). To illustrate, consider the following logistic network,

$$\emptyset \xleftarrow{\alpha_1} S \underset{\alpha_3}{\overset{\alpha_2}{\rightleftarrows}} 2S$$

endowed with mass-action. The corresponding deterministic rate equation has, when $\alpha_2 > \alpha_1$ a unique positive stable equilibrium point. However, modeled stochastically, the associated Markov process goes extinct almost surely. This discrepancy is known as Keizer's paradox (Keizer, 1987; Vellela and Qian, 2007). The resolution comes partially from looking at the time scales. Although extinction is inevitable, the expected time to extinction may be extremely large compared to the time it takes the process to reach the (deterministically) stable equilibrium point. Thus, up until extinction the stochastic process will wander about the stable equilibrium point. It is therefore natural to introduce the quasi-stationary distribution, which is exactly the stationary distribution when we condition the process not to have gone extinct. In other words, the quasi-stationary distribution is supported in a neighborhood of the positive equilibrium point, and we have restored the link between the two descriptions of the underlying network, at least on a not too long time horizon.

To obtain a mathematically sound foundation for the above intuition, yielding a complete resolution to the paradox, we apply the classical results of Kurtz (1972, 1978) allowing us to view the stochastic model as a *random perturbation* of the deterministic system with absorbing states. In this setting, we show that for sufficiently large system size, the quasi-stationary distributions concentrate on the *positive attractors* for the unperturbed system. This idea was first developed in Faure and Schreiber (2014) for discrete time homogeneous Markov chains and in Marmet (2013) for stochastic approximation algorithms of continuous time dynamical systems on compact state spaces. Here, we review some of this theory and extend it to the class of continuous-time Markov processes on countable state spaces which can be associated to a reaction network, and show that the analog of the series of assumptions prescribed in Faure and Schreiber (2014) and Marmet (2013) are automatically satisfied. Furthermore, we rephrase the main theorem in terms of *Morse-decompositions*, providing a more tractable approach. For this reason, our results are very applicable, which we underline through a series of examples where the underlying deterministic system displays a variety of different dynamics.

From a modeling point of view, the concept of quasi-stationarity is very natural. Indeed, the metastable behavior captured hereby could correspond to persistence of an endemic disease in epidemiology, coexistence of interacting species in ecology or of opinions in political science or maintenance of a genetic polymorphism (Faure and Schreiber, 2014). Reaction networks may be seen as a natural framework for modeling these kinds of dynamical systems

where the interactions among species are themselves of transformational form. However, as the complex nature of most reaction networks prevents an explicit analytical description of a corresponding quasi-stationary distribution, it is of great interest to infer qualitative properties through the corresponding deterministic system (Anderson and Kurtz, 2015).

The structure of the paper is as follows. In Section 4.2 we introduce the central notion of random perturbations with absorbing states and define the deterministic and stochastic dynamical systems associated to a reaction network. We then introduce quasi-stationary distributions and show how one may reconcile the different views through the so called "classical scaling". In Section 4.3 we derive some preliminary results under the assumption of a positive attractor. In particular, we show that the expected time to extinction grows exponentially with system size. Section 4.4 is devoted to showing the main result; the support of the limiting measure concentrates on the union of positive attractors. The argument goes through a sample path large deviations result as well as so called absorption preserving chain recurrence, which for completeness are introduced in the appendix. A more applicable version of the main result is provided through Morse decompositions introduced in subsection 4.5.2. Examples and numerical calculations illustrating the main theorem is given in Section 4.6. Finally, in the last section contains a discussion of possible extensions to the field of Freidlin-Wentzell theory.

4.2 Random Perturbations

A reaction network may be considered as a deterministic dynamical system on the state space $D \subseteq \mathbb{R}_0^d$, in which the evolution of species concentrations $x = x(t) \in D$ at time t is described by the, possibly highly non-linear, first order system of ODEs,

$$\frac{dx}{dt} = F(x). \tag{4.2.1}$$

When a solution exists, it determines a *semi-flow* $\varphi_t(x) = x(t)$ for $t \geq 0$. However, mathematical models are derived from experimental data and thus only describe real phenomena approximately. In particular, as physical processes are being affected by a huge amount of small external fluctuations, by naturally considering these as random, the stability of dynamical systems under random perturbations is to be considered. For each $\epsilon > 0$, let

$$p^\epsilon(t, x, B) = \mathbb{P}(X_t^\epsilon \in B \,|\, X_0^\epsilon = x) = \mathbb{P}_x(X_t^\epsilon \in B), \qquad t \geq 0, x \in D, B \in \mathcal{B}(D),$$

denote the transition kernels of a homogeneous Markov process $(X_t^\epsilon \colon t \geq 0)$. We shall interpret ϵ as the inverse of the system size. The models we consider allow the possibility of a set of absorbing states $D_0 \subseteq \mathbb{R}_0^d$, corresponding to the loss of one or more species. We thus suppose that we can write the state space as the disjoint union

$$D = D_0 \sqcup D_1,$$

where D_0, D_1 are positively φ-invariant and D_0 is a closed and absorbing subset of D,

$$p^\epsilon(t, x, D_1) = 0 \qquad \forall \epsilon > 0, t > 0, x \in D_0.$$

Thus, once the Markov process hits the set D_0 it stays there forever. Let $N^\delta(B) = \{x \in D\colon \inf_{y\in B} \|x - y\| < \delta\}$ denote a δ-neighborhood of the set B.

Definition 4.2.2. A *random perturbation* of a semi-flow φ_t is a family of homogeneous Markov processes $\{(X_t^\epsilon\colon t \geq 0)\}_{\epsilon>0}$ on $D \subseteq \mathbb{R}_0^d$ whose transition kernels, $p^\epsilon(t, x, B)$, satisfy that for any $\delta > 0, T > 0$ and $K \subset D_1$ compact,

$$\beta_{\delta,K}(\epsilon) := \sup_{t\in[0,T]} \sup_{x\in K} p^\epsilon \left(t, x, D\backslash N^\delta(\varphi_t(x))\right) \to 0 \qquad \text{for } \epsilon \to 0.$$

Note that this in particular implies that $p^\epsilon(t, x, \cdot)$ converges uniformly to the Dirac measure $\delta_{\varphi_t(x)}$ at $\varphi_t(x)$ for any continuous function $g : M \to \mathbb{R}$ with compact support as $\epsilon \to 0$,

$$\lim_{\epsilon\to 0} \sup_{x\in K} \left| \int_M g(y) p^\epsilon(t, x, dy) - g(\varphi_t(x)) \right| = 0.$$

The meaning is that due to a random error, the particle starting at x misses the point $\varphi_t(x)$ and ends up at a random point whose distribution is close to the Dirac δ-function at $\varphi_t(x)$. Thus, for small $\epsilon > 0$, the asymptotic behavior of the Markov process $(X_t^\epsilon\colon t \geq 0)$ should be closely related to the dynamics of φ_t (Kifer, 1988).

Due to the seminal work of Kurtz (1970, 1972, 1978), the usual stochastic description of a reaction network can, as we shall make explicit, under the so called *"classical scaling"*, be seen as a random perturbation of the deterministic system. This has several interesting consequences. Indeed, many real world systems will exhibit eventual extinction, which is only possible to capture by the stochastic models (Chazottes et al., 2017). Looking at the quasi-stationary distributions, we are able to bridge the gap between the eventual outcomes of the two descriptions. We find that – under suitable assumptions – in this *fluid* or *thermodynamic limit*, the weak* limit of quasi-stationary distributions are supported by a subset of the attractors prescribed by the deterministic system.

4.2.1 Reaction Network Setup

A reaction network, stochastic or deterministic, is a triple $\mathcal{N} = (\mathcal{S}, \mathcal{C}, \mathcal{R})$, where $\mathcal{S}$ is a finite ordered set of species, $\mathcal{C}$ is a finite ordered set of complexes, consisting of linear combinations over $\mathbb{N}_0$ of the species and $\mathcal{R}$ is a non-empty irreflexive relation on $\mathcal{C}$, referred to as the set of reactions (Anderson and Kurtz, 2015; Feinberg, 1979; Gunawardena, 2003). Furthermore, $\mathcal{R}$ is ordered.

We define the dimension of the reaction network, $d = |\mathcal{S}|$. Any species $S_i \in \mathcal{S}$ can be identified with the unit vector $e_i \in \mathbb{N}_0^d$, thus any complex $y \in \mathcal{C}$ will be identified with a vector in $\mathbb{N}_0^d$. It is customary to denote an element $(y_k, y_k') \in \mathcal{R}$ by $y_k \to y_k' \in \mathcal{R}$ in which case we refer to y_k as the source complex and to y_k' as the product complex of reaction k. We may thus write $\mathcal{R} = \{y_k \to y_k' : y_k, y_k' \in \mathcal{C}, y_k \neq y_k', k = 1, \ldots, r\}$. Employing a standard,

although slight abuse of, notation, we identify $\mathcal{S} = \{S_1, \ldots, S_d\}$ with the set $\{1, \ldots, d\}$ and $\mathcal{R}$ with $\{1, \ldots, r\}$. In particular, $r = |\mathcal{R}|$. We write the k'th reaction with the notation

$$\sum_{i \in \mathcal{S}} y_{ki} S_i \to \sum_{i \in \mathcal{S}} y'_{ki} S_i,$$

where $y_{ki} = (y_k)_i$ and $y'_{ki} = (y'_k)_i$ are the stoichiometric coefficients associated with the source and product complexes of reaction k, respectively. Define the reaction vectors $\xi_k = y'_k - y_k$. One may completely describe a reaction network (without kinetics) in terms of its reaction graph, whose nodes is the set of complexes, $\mathcal{C}$, and whose directed edges are given by the set of reactions, $\mathcal{R}$. This concise description will be employed in the rest of the paper.

When all species are present in large abundances, it is natural to consider the concentration of each species, instead of the exact counts. This leads to the state space $D \subseteq \mathbb{R}_0^d$. We denote by x_t the non-negative real vector in $\mathbb{R}_0^d$ whose entries are the concentrations of the species at time t. The evolution of x_t is modeled as the solution to the ordinary differential equation (4.2.1), where the vector field is given by

$$F(x) = \sum_{k \in \mathcal{R}} \lambda_k(x) \xi_k,$$

for some functions $\lambda_k : [0, \infty)^d \to [0, \infty)$ called rate functions (Gunawardena, 2003). Here, we shall use the most common case of deterministic *mass-action kinetics*, which is given by

$$\lambda_k(x) = \alpha_k \prod_{i=1}^{d} x_i^{y_{ki}} = \alpha_k x^{y_k},$$

where $\alpha_k > 0$ are referred to as rate constants. This represents the idea that the system is well stirred. A unique solution $\varphi_t(x) = x_t$ with initial condition x_0 of (4.2.1) defines a continuous time dynamical system. When integrating, we may write this as

$$\varphi_t(x) = \varphi_0(x) + \sum_{k \in \mathcal{R}} \left(\int_0^t \lambda_k(\varphi_s(x))\, ds \right) \xi_k. \tag{4.2.3}$$

The rate functions determine a deterministic kinetics $\lambda = (\lambda_1, \ldots, \lambda_r)$ for $\mathcal{N}$, and the pair $(\mathcal{N}, \lambda)$ is called a deterministic reaction system (Cappelletti, 2015).

In the stochastic description of the same underlying reaction network, $\mathcal{N}$, we may specify the related continuous-time Markov process $(X_t \colon t \geq 0)$. Here, X_t is the vector in $\mathbb{N}_0^d$ whose entries are the species counts at time t, with the process taking place in a space of unit volume, $\epsilon = 1$. We are interested in the connection to the deterministic description, which follows the densities of each species, thus letting $1/\epsilon$ be the volume of the system, we are led to consider the scaled process

$$X_t^\epsilon := \epsilon X_t, \qquad t \geq 0.$$

In particular, we examine the case where $\epsilon \to 0$, keeping the concentrations constant - the so called *thermodynamic limit*. Specifying the volume-normalized mass-action jump rates $\lambda_k^\epsilon : \epsilon\mathbb{N}_0^d \to [0,\infty)$ by

$$\lambda_k^\epsilon(x) = \alpha_k \epsilon^{\|y_k\|_1 - 1} \prod_{i=1}^{d} \binom{x_i/\epsilon}{y_{ki}} y_{ki}!,$$

the stochastic process $(X_t^\epsilon : t \geq 0)$ then satisfies the stochastic equation

$$X_t^\epsilon = X_0^\epsilon + \sum_{k \in \mathcal{R}} Y_k \left(1/\epsilon \int_0^t \Lambda_k^\epsilon(X_s^\epsilon)\, ds \right) \epsilon \xi_k,$$

where $\Lambda_k^\epsilon = \epsilon \lambda_k^\epsilon$ and Y_k are independent and identically distributed unit-rate Poisson processes (Anderson and Kurtz, 2015; Ethier and Kurtz, 1986; Norris, 2009). This stochastic equation is referred to as a *random time change representation* of the classically scaled process. If reaction $y_k \to y_k'$ occurs at time t, then the new state is $X_t^\epsilon = X_{t-}^\epsilon + \epsilon \xi_k$, where X_{t-}^ϵ denotes the previous state. Thus $(X_t^\epsilon : t \geq 0)$ is a continuous time pure jump Markov process with generator

$$L_\epsilon f(x) = \frac{1}{\epsilon} \sum_{k \in \mathcal{R}} \Lambda_k^\epsilon(x)(f(x + \epsilon \xi_k) - f(x)), \tag{4.2.4}$$

for all bounded $f : (\epsilon\mathbb{N}_0)^d \to \mathbb{R}$ (Anderson and Kurtz, 2015). Note that Λ_k^ϵ is independent of ϵ to first order and $\Lambda_k^\epsilon \to \lambda_k$ for $\epsilon \to 0$. Let $D^\epsilon \subseteq \epsilon\mathbb{N}_0^d$ denote the state space upon which the process $(X_t^\epsilon : t \geq 0)$ takes its values. We may formally embed this process into $D \subseteq [0,\infty)^d$ by allowing X_0^ϵ to be any point in D and update with the jump rates

$$\lambda_k^\epsilon(x) = \alpha_k \epsilon^{\|y_k\|_1 - 1} \prod_{i=1}^{d} \binom{\lfloor x_i/\epsilon \rfloor}{y_{ki}} y_{ki}!,$$

for the same set of jump directions ξ_k. Note that if ξ_k is a possible jump direction then, for all $1 \leq i \leq d$,

$$x_i + \epsilon \xi_{ki} \geq \epsilon \lfloor x_i/\epsilon \rfloor + \epsilon \xi_{ki} \geq \epsilon(y_{ki} + \xi_{ki}) \geq 0,$$

hence the embedded process is well defined on D.

4.2.2 Quasi-stationary Distribution

The stochastic description is more detailed than the deterministic one (Qian, 2011). In particular, absorbing sets, or extinctions, corresponding to the loss of one or several species may be captured. These kinds of systems appear naturally in many models in ecology, epidemiology and chemical processes (Nåsell, 2011; Méléard and Villemonais, 2012). Let $D \subseteq \mathbb{R}_0^d$ be a closed subset. Given a reaction network with corresponding embedded stochastic process $(X_t^\epsilon : t \geq 0)$, we may divide the state space D into the endorsed set $D_E^\epsilon \subseteq \mathbb{R}_0^d$ and

absorbing set $D_A^\epsilon \subset \mathbb{R}_0^d$ (Hansen and Wiuf, 2018a). Note that the absorbing set is just a "fattening" of the absorbing set for the corresponding non-embedded process. Further, in this continuous setting, D_E^ϵ has a continuum of endorsed sets. For each $\epsilon > 0$, we restrict attention to the unique endorsed set, E^ϵ, containing X_0^ϵ and assume for simplicity, that this set is irreducible for each $\epsilon > 0$. The corresponding absorbing set is denoted A^ϵ. The next lemma shows that this decomposition in turn will define a corresponding absorbing set for the deterministic dynamical system. Let $\{\epsilon_n\}_{n\geq 1}$ be a sequence converging to 0.

Lemma 4.2.5. *The state space can be written $D = D_0 \sqcup D_1 \subseteq [0,\infty)^d$ where*

(i) $D_0 = \lim_{n\to\infty} D_A^{\epsilon_n}$ *is a closed subset of D;*

(ii) $D_1 = \lim_{n\to\infty} D_E^{\epsilon_n}$ *is an open subset of D;*

(iii) D_0 *and* D_1 *are positively φ-invariant;*

(iv) D_0 *is absorbing for the random perturbations,*

$$p^\epsilon(t,x,D_1) = 0 \qquad \forall \epsilon > 0, t > 0, x \in D_0.$$

Proof. First note that if $\epsilon_2 < \epsilon_1$ and $\lambda_k^{\epsilon_1}(x) > 0$ then $\lambda_k^{\epsilon_2}(x) > 0$. By definition of endorsed and absorbing sets, we then have the monotone sequences of sets

$$D_A^{\epsilon_1} \supset D_A^{\epsilon_2} \supset \ldots, \qquad D_E^{\epsilon_1} \subset D_E^{\epsilon_2} \subset \ldots.$$

Clearly, it follows that $D_0 \subseteq \partial[0,\infty)^d$. Indeed, let $x \in (0,\infty)^d$. Then as $\epsilon_n \to 0$, there exists an n large enough such that all reactions may fire, that is $\lambda_k(x) > 0$ for all $k \in \mathcal{R}$. In particular, $x \notin D_A^{\epsilon_n}$. As $D_0 \subset D_A^{\epsilon_n}$ we conclude that $x \notin D_0$. In fact, we claim that

$$\lim_{n\to\infty} D_A^{\epsilon_n} = \bigcup_{1\leq i\leq d} span_{[0,\infty)}\{e_i \colon p_i(D_A^1) = [0,\infty)\},$$

where p_i is the projection onto the ith axis. To see this, let $\epsilon > 0$ be given and suppose first that $x \in span_{(0,\infty)}\{e_i\}$ for some $1 \leq i \leq d$ satisfying $p_i(D_A^\epsilon) \neq [0,\infty)$. As above, there exists an ϵ_1 sufficiently small such that $\lambda_k^{\epsilon_1}(x) > 0$ for all $k \in \mathcal{R}$. Thus $x \in D_E^{\epsilon_1}$. In particular $x \notin D_0$. On the other hand, if $p_i(D_A^\epsilon) = [0,\infty)$ then x has no path to a point where all reactions may fire. As x was arbitrary, the same conclusion holds for all points x/ϵ and any $\epsilon > 0$. Thus $x \in D_0$. As the union of rays from origo is closed this proves (i). The disjointness of absorbing and endorsed sets then immediately implies (ii).

Suppose there exists $\epsilon > 0, t > 0$ and $x \in D_0$ such that $p^\epsilon(t,x,D_1) > 0$. Then there is a path from $x \in D_0$ to $y \in D_1$. However, as $x \in D_A^\epsilon$ for all $\epsilon > 0$ and $y \in D_E^\epsilon$ this contradicts the fact that D_A^ϵ is absorbing. Thus, D_0 is indeed absorbing for the random perturbations, proving (iv).

To prove (iii), suppose for contradiction that $\varphi_t(D_0) \subseteq D_0$ does not hold for all t. Then there exists a $t > 0$ and $y \in \varphi_t(D_0) \cap D_1$, that is, there exists $x \in D_0$ such that $\varphi_t(x) = y \in D_1$. In particular, we have $p^\epsilon(t,x,D_1) > 0$ for ϵ sufficiently small contradicting (iv). Finally, by uniqueness of the orbits of an ODE system, we conclude that D_1 is positive φ-invariant as well. This completes the proof. □

We are interested in examining the process before it is absorbed. In particular, whether the process exhibits stationarity in this time domain. For this purpose, we introduce the quasi-stationary distribution (Collet et al., 2013; Méléard and Villemonais, 2012; van Doorn and Pollett, 2013).

Definition 4.2.6. A probability measure ν on E is called a *quasi-stationary distribution* (QSD) for the process $(X_t \colon t \geq 0)$ absorbed at A, if for every measurable set $B \subseteq E$

$$\mathbb{P}_\nu(X_t \in B \,|\, \tau_A > t) = \nu(B), \qquad t \geq 0,$$

or equivalently, if there exists $\theta \in (0, \infty)$ such that

$$\int_E p(t, x, B)\nu(dx) = e^{-\theta t}\nu(B), \qquad t \geq 0.$$

We shall refer to the last equality as the *QSD property*.

Note 4.2.7. The equivalence between the two definitions follows from the general property of QSDs that $\mathbb{P}_\nu(t < \tau_A) = e^{-\theta t}$ (Collet et al., 2013). Thus, on the one hand

$$\int_E p(t, x, B)\nu(dx) = \int_E \mathbb{P}_x(X_t \in B)\nu(dx) = \mathbb{P}_\nu(X_t \in B),$$

while on the other

$$\mu(B) = \mathbb{P}_\mu(X_t \in B \,|\, t < \tau_A) = \frac{\mathbb{P}_\mu(X_t \in B, t < \tau_A)}{\mathbb{P}_\mu(t < \tau_A)} = \frac{\mathbb{P}_\mu(X_t \in B)}{e^{-\theta t}},$$

establishing the equivalence. Note furthermore that as $E^\epsilon \subset D_1$ for all $\epsilon > 0$, we may consider all ν_ϵ as measures on the same space, D_1. In particular, it makes sense to look at the weak limit of ν_ϵ.

One may refine the law of large numbers theorem of Kurtz (1972) to a large deviations type result, telling us the rate at which the probability of deviation from the fluid limit decreases with system size (Kurtz, 1978).

Theorem 4.2.8 (Kurtz)**.** For every $T > 0$, there exist $C_1, C_2 > 0$ such that for every $\delta > 0$ and $\epsilon > 0$ sufficiently small,

$$\mathbb{P}_x\left(\sup_{0 \leq t \leq T} |X_t^\epsilon - \varphi_t(x)| \geq \delta\right) \leq C_1 e^{-C_2\delta^2/\epsilon}.$$

Proof. Note that once X_0^ϵ is specified, there exists a $\gamma \in [0, \epsilon]$ such that one may write $\lambda_k^\epsilon(x) = \alpha_k \epsilon^{\|y_k\|_1 - 1} \prod_{i=1}^d \binom{x_i/\epsilon}{y_{ki}} y_{ki}! - \gamma$. Thus, for $\epsilon > 0$ sufficiently small we may consider the rate functions to be Lipschitz continuous and bounded on the compact set

$$K_\delta = \left\{y : \inf_{0 \leq s \leq T} |y - \varphi_s(x)| \leq \delta\right\} \subset D_1,$$

for δ sufficiently small and some finite $T > 0$. As $\mathbb{P}_x\left(\sup_{0\leq t\leq T}|X_t^\epsilon - \varphi_t(x)| \geq \delta\right)$ depends only on $\lambda_k^\epsilon(x)$ for $x \in K_\delta$, we may assume without loss of generality that λ_k^ϵ are uniformly bounded and Lipschitz continuous. The result may now be obtained directly from Theorem 5.3 Shwartz and Weiss (1995), which contains all details. □

This result is the backbone to all what follows. Note that, for any $t \in [0, T]$, we have

$$|X_t^\epsilon - \varphi_t(x)| \leq \sup_{t\in[0,T]}|X_t^\epsilon - \varphi_t(x)|,$$

yielding the inequality

$$\mathbb{P}_x(|X_t^\epsilon - \varphi_t(x)| \geq \delta) \leq \mathbb{P}_x\left(\sup_{t\in[0,T]}|X_t^\epsilon - \varphi_t(x)| \geq \delta\right).$$

Therefore, by taking supremum on both sides,

$$\sup_{t\in[0,T]}\mathbb{P}_x(|X_t^\epsilon - \varphi_t(x)| \geq \delta) \leq \mathbb{P}_x\left(\sup_{t\in[0,T]}|X_t^\epsilon - \varphi_t(x)| \geq \delta\right) \leq C_1 e^{-C_2\delta^2/\epsilon}.$$

It follows by taking supremum over $x \in K$ that

$$\beta_{\delta,K}(\epsilon) = \sup_{t\in[0,T]}\sup_{x\in K} p^\epsilon(t, x, D\backslash N^\delta(\varphi_t(x))) \leq C_1 e^{-C_2\delta^2/\epsilon} \to 0, \tag{4.2.9}$$

for $\epsilon \to 0$. We may therefore consider $\{(X_t^\epsilon \colon t \geq 0)\}_{\epsilon>0}$ as a random perturbation of the semi-flow $\{\varphi_t\}$ arising from the deterministic description of the underlying network.

4.3 Assuming a Positive Attractor

In this section, we derive some preliminary results based solely on the assumption of the existence of a positive attractor and that the dynamics of $\varphi_t(x)$ is not too wild. In particular, we make the very weak assumption that for each $x \in D_1$, either $\lim_{t\to\infty}\varphi_t(x) \in D_0$ or $\inf_{t\geq 0} d(\varphi_t(x), D_0) > 0$. Recall the following terminology from dynamical systems theory (Katok and Hasselblatt, 1995; Brin and Stuck, 2002).

Definition 4.3.1. A compact set Γ is an *attractor* for the semi-flow $\{\varphi_t\}_{t\geq 0}$ if there exists an open set U containing Γ such that $\bigcap_{t\geq 0}\varphi_t(U) = \Gamma$ and for any open set $V \supset \Gamma$, there exists $T > 0$ such that $\varphi_t(\overline{U}) \subset V$ for all $t \geq T$. We shall use the term *positive attractor* if furthermore $\Gamma \subset D_1$. The set U is referred to as a *fundamental neighborhood* of Γ.

4.3.1 Mean Time to Extinction

Upon existence of a positive attractor, we may find the asymptotic behavior of the mean time to extinction starting from a QSD. The proof follows (Marmet, 2013).

Theorem 4.3.2. Assume that the semi-flow $\{\varphi_t\}_{t\geq 0}$ admits an attractor $\Gamma \subset D_1$. Then, starting according to the QSD, ν_ϵ, the probability of being absorbed by time $t > 0$ is $\mathcal{O}(\epsilon e^{-\gamma/\epsilon})$ while the mean time to extinction is $\mathcal{O}(\epsilon e^{c/\epsilon})$, where $\gamma, c > 0$.

Proof. Let $V \subset D_1$. For $\epsilon > 0$ sufficiently small, we have $V \subset D_E^\epsilon$. By the QSD property, it follows for $t \geq 0$ that

$$e^{-\theta_\epsilon t}\nu_\epsilon(V) = \int_{E^\epsilon} p^\epsilon(t,x,V)\nu_\epsilon(dx) \geq \int_V p^\epsilon(t,x,V)\nu_\epsilon(dx) \geq \inf_{x\in V} p^\epsilon(t,x,V)\nu_\epsilon(V).$$

By irreducibility of E^ϵ, it follows that $\nu_\epsilon(V) > 0$ for all $\epsilon > 0$ thus we find

$$e^{-\theta_\epsilon t} \geq \inf_{x\in V} p^\epsilon(t,x,V).$$

Let $U \subset D_1$ be a fundamental neighborhood of the attractor Γ. Then for all $\eta > 0$ there exist $T > 0$ such that for all $t \geq T$ and $x \in U$ we have $d(\varphi_t(x), \Gamma) < \eta$. Let $\alpha = d(\Gamma, U^c)$ and $\eta < \alpha$, $\delta < \alpha - \eta$. Then for all $x \in U$, we have for given $t > T$

$$\begin{aligned} p^\epsilon(t,x,U^c) &\leq \mathbb{P}_x(d(X_t^\epsilon, \Gamma) > \alpha) \\ &\leq \mathbb{P}_x(d(X_t^\epsilon, \varphi_t(x)) + d(\varphi_t(x), \Gamma) > \alpha) \\ &\leq \mathbb{P}_x(d(X_t^\epsilon, \varphi_t(x)) > \alpha - \eta) \\ &\leq \mathbb{P}_x(d(X_t^\epsilon, \varphi_t(x)) > \delta) \\ &\leq C_1 e^{-C_2\delta^2/\epsilon}, \end{aligned}$$

for ϵ sufficiently small by Theorem 4.2.8. It follows, since for any $t > 0$ there exists $\epsilon > 0$ sufficiently small such that $t/\epsilon > T$,

$$e^{-\theta_\epsilon t/\epsilon} \geq \inf_{x\in U} p^\epsilon(t/\epsilon, x, U) = 1 - \sup_{x\in U} p^\epsilon(t/\epsilon, x, U^c) \geq 1 - C_1 e^{-C_2\delta^2/\epsilon}.$$

Thus, we find that the probability of being absorbed by time t is

$$\begin{aligned} 1 - e^{-\theta_\epsilon t} &\leq 1 - (1 - C_1 e^{-C_2\delta^2/\epsilon})^\epsilon = 1 - \exp(\log((1 - C_1 e^{-C_2\delta^2/\epsilon})^\epsilon)) \\ &= 1 - \exp(\epsilon \log(1 - C_1 e^{-C_2\delta^2/\epsilon})) \\ &= 1 - \exp\left(\epsilon(-C_1 e^{-C_2\delta^2/\epsilon} - o(e^{-C_2\delta^2/\epsilon}))\right) \\ &= 1 - \left(1 + \epsilon(-C_1 e^{-C_2\delta^2/\epsilon} + o(e^{-C_2\delta^2/\epsilon}))\right) \\ &= \epsilon C_1 e^{-C_2\delta^2/\epsilon} + o(\epsilon C_1 e^{-C_2\delta^2/\epsilon}) = \mathcal{O}(\epsilon e^{-\gamma/\epsilon}) \end{aligned}$$

In particular, we see that the mean time to extinction, starting from the QSD is

$$\mathbb{E}_{\nu_\epsilon}(\tau_{A^\epsilon}) = \frac{1}{\theta_\epsilon} \geq \frac{-1}{\epsilon \log(1 - C_1 e^{-C_2\delta^2/\epsilon})} = \mathcal{O}(\epsilon e^{c/\epsilon}), \tag{4.3.3}$$

as desired. □

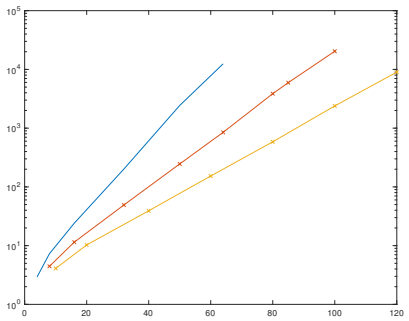

Figure 4.1: Stochastic simulation, using Gillespie's algorithm, of the mean time to extinction for the reaction system $S \xrightarrow{\alpha_1} 2S \xrightarrow{\alpha_2} \emptyset$ for $\alpha_1 = 1$ and $\alpha_2 = 2, 3, 4$.

4.3.2 Invariance of Limiting Measure

Theorem 4.3.4. Suppose the semi-flow $\{\varphi_t\}_{t\geq 0}$ admits an attractor $\Gamma \subset D_1$. Then the set of limit points of $\{\nu_\epsilon\}_{\epsilon>0}$ for the weak* topology is a subset of the set of invariant measures for the semi-flow $\{\varphi_t\}_{t\geq 0}$.

Proof. We follow Marmet (2013) and extend the arguments where necessary. Suppose the sequence ν_ϵ of QSDs converges weakly to a measure ν. We may view this convergence taking place on D. Let $t \geq 0$. By the Portemanteau theorem (Klenke, 2014), it suffices to prove that

$$\lim_{\epsilon\to 0} \left| \int_D f(x)\nu_\epsilon(dx) - \int_D f(\varphi_t(x))\nu_\epsilon(dx) \right| = 0,$$

for every bounded Lipschitz function f. The QSD property implies that for all $t \geq 0$

$$\int_D f(x)\nu^\epsilon(dx) = \mathbb{E}_{\nu^\epsilon}(f(X_t^\epsilon)) = \mathbb{E}_{\nu^\epsilon}(f(X_t^\epsilon)|\tau_{A^\epsilon} > t) = \int_D \mathbb{E}_x\left(f(X_t^\epsilon)|\tau_{A^\epsilon} > t\right)\nu^\epsilon(dx).$$

Thus, for all $t \geq 0$, one may choose $\epsilon > 0$ sufficiently small such that

$$\begin{aligned} I &:= \left| \int_D f(x)\nu^\epsilon(dx) - \int_D f(\varphi_t(x))\nu^\epsilon(dx) \right| \\ &= \left| \int_D \mathbb{E}_x(f(X_t^\epsilon)|\tau_{A^\epsilon} > t)\nu^\epsilon(dx) - \int_D f(\varphi_t(x))\nu^\epsilon(dx) \right| \\ &= \left| \int_D \mathbb{E}_x(f(X_t^\epsilon) - f(\varphi_t(x))|\tau_{A^\epsilon} > t)\nu^\epsilon(dx) \right|, \end{aligned}$$

where the last equality holds due to the assumptions on φ_t. By Theorem 4.2.8 we infer that

$$\mathbb{E}_x\left(\sup_{0\le t\le T}|X_t^\epsilon-\varphi_t(x)|\ge\delta\right)=\int_0^\infty\mathbb{P}_x\left(\sup_{0\le t\le T}|X_t^\epsilon-\varphi_t(x)|\ge\delta\right)d\delta$$
$$\le\int_0^\infty C_1e^{-C_2\delta^2/\epsilon}\,d\delta=\frac{C_1\sqrt{\pi}}{2\sqrt{C_2}}\sqrt{\epsilon}.$$

Further, by (the proof of) Theorem 4.3.2 we see that

$$1-\exp(\theta_\epsilon t)\le\mathcal{O}(\epsilon e^{c/\epsilon}),$$

for $\epsilon>0$ sufficiently small, hence $1-\exp(\theta_\epsilon t)\to 0$ as $\epsilon\to 0$ which implies that $\exp(\theta_\epsilon t)$ is bounded. Let C_L be the Lipschitz constant for f. Then we see that

$$\begin{aligned}I&=\left|\int_D\mathbb{E}_x(f(X_t^\epsilon)-f(\varphi_t(x))|\tau_{A^\epsilon}>t)\nu^\epsilon(dx)\right|\\&\le\left|\int_D\frac{\mathbb{E}_x(f(X_t^\epsilon)-f(\varphi_t(x)))}{\mathbb{P}_x(\tau_{A^\epsilon}>t)}\nu^\epsilon(dx)\right|\\&\le\int_D\frac{C_L\frac{C_1\sqrt{\pi}}{2\sqrt{C_2}}\sqrt{\epsilon}}{e^{-\theta_\epsilon t}}\nu^\epsilon(dx)\le C_3\sqrt{\epsilon}e^{\theta_\epsilon t}\to 0,\end{aligned}$$

for $\epsilon\to 0$. We conclude that

$$\int_D f(x)\nu(dx)=\int_D f(\varphi_t(x))\nu(dx),$$

for all $t\ge 0$ as desired. □

Lemma 4.3.5. *For all $c>0$, there exists an open neighborhood V_c of D_0 such that*

$$\liminf_{\epsilon\to 0}\inf_{x\in V_c}\epsilon\log p^\epsilon(1,x,A^\epsilon)\ge -c$$

Proof. Given $\epsilon>0$, every point $x\in E^\epsilon$ has a path to A^ϵ. Let $x\in E^\epsilon$ be given. There is a fixed set of reactions, $\mathcal{R}_A\subseteq\mathcal{R}$, which takes x to A^ϵ. The number of jumps needed to take x to A^ϵ is then bounded by η_1/ϵ for some $\eta_1\in\mathbb{N}$. Let $k'\in\mathcal{R}$ be the last reaction on the path from x to A^ϵ and let its order be m. We may without loss of generality assume that m is the maximal order of a reaction in $\mathcal{R}_A$. Further, let n be the lowest order of any reaction in the network. Define the following events

$$\mathcal{E}_1=\{X_t^\epsilon\text{ follows the path }x\mapsto A^\epsilon\},$$
$$\mathcal{E}_2=\{X_t^\epsilon\text{ makes }\eta_1/\epsilon\text{ jumps before time }1\}.$$

Then we have the rough bound

$$\epsilon\log p^\epsilon(t,x,A^\epsilon)\ge\epsilon\log\mathbb{P}_x(\mathcal{E}_1\mathcal{E}_2)=\epsilon\log\mathbb{P}_x(\mathcal{E}_1)+\epsilon\log\mathbb{P}_x(\mathcal{E}_2).$$

For $\epsilon > 0$ sufficiently small there exists $\eta_2 < 1$ such that with $\eta = \eta_1\eta_2 < 1$,

$$\mathcal{O}(\epsilon \log \mathbb{P}_x(\mathcal{E}_1)) \geq \mathcal{O}\left(\epsilon \log \prod_{i=1}^{\eta/\epsilon} \frac{\epsilon^{m-1} i^m}{\epsilon^{n-1} i^n}\right) = \mathcal{O}(\epsilon \sum_{i=1}^{\eta/\epsilon} \log(\epsilon i)^{m-n}) = \mathcal{O}(\eta \log(\eta)).$$

We further have for ϵ sufficiently small

$$\begin{aligned}\epsilon \log \mathbb{P}_x(\mathcal{E}_2) &\geq \epsilon \log \left(\prod_{i=1}^{\eta/\epsilon} \left(1 - e^{-C_1(\epsilon i)^m}\right)\right) = \epsilon \left(\sum_{i=1}^{\eta/\epsilon} \log \left(1 - e^{-C_1(i\epsilon)^m}\right)\right) \\ &\geq \epsilon \sum_{i=1}^{\eta/\epsilon} -e^{-C_1(i\epsilon)^m} \geq -\eta e^{\eta^m}.\end{aligned}$$

Thus, we conclude that

$$\epsilon \log p^\epsilon(1, x, A^\epsilon) > -c,$$

for all x in a sufficiently small neighborhood V_c of D_0 as desired. □

4.4 Support of the limiting measure

Assume that there exists a decreasing sequence $\epsilon_n \to 0$ such that for every $n \in \mathbb{N}$, ν_n is a quasi-stationary distribution for p^{ϵ_n} with associated θ_n. Additionally, we assume that ν_n converges weakly to a Borel probability measure ν. Under the assumption of a positive attractor, the limit set of the sequence $(\nu_\epsilon)_{\epsilon>0}$ in the weak* topology consists of invariant measures. Thus, by the Poincaré Recurrence theorem (Walters, 1982), for any $B \subset D$, if $\nu(B) > 0$ then almost all points in B returns infinitely often to B under φ_t, that is almost all points in B are recurrent, $x \in \omega(x)$. Recall that a point $y \in D$ is called an ω-limit point for $x \in D$ is there exists a sequence of times $(t_n)_{n\in\mathbb{N}}$ going to infinity such that $\lim_{n\to\infty} \varphi_{t_n}(x) = y$. The set of all ω-limit points is (Katok and Hasselblatt, 1995),

$$\omega(x) = \bigcap_{T=0}^{\infty} \left(\overline{\bigcup_{t\geq T} \varphi_t(x)}\right). \tag{4.4.1}$$

Letting $BC(\varphi) = \overline{\{x \in D\colon x \in \omega(x)\}}$ denote the Birkhoff center, we conclude that

$$\operatorname{supp} \nu \subseteq BC(\varphi).$$

In this section, we shall refine this result further. In particular, as the QSDs, ν_n have support in D_1, it is natural to ask whether this is true for the limit measure ν as well. Note that in many cases 0 would be a fixed point for the flow thus $0 \in BC(\varphi) \cap D_0$. An answer to this question would therefore indeed be a refinement. To obtain this result, we shall assume that there is only a finite number of positive attractors $\{\Gamma_i\}_{i=1,\dots,\ell}$, and rely on the theory of large deviations.

4.4.1 Large Deviation Property

The theory of *Large Deviations* describes the rate at which rare events occur. More concretely, in our setting, we need a sample path large deviations principle, in which the rare events are given in terms of entire paths lying within certain sets (Dembo and Zeitouni, 1998). On the entire state space, D, we face two main obstacles, which hinders a direct application of standard results - the rates of the jump process are neither bounded nor uniformly Lipschitz continuous. Recently, an application of a general result on mean field interacting particles from Dupuis et al. (2016) applied to reaction networks, was shown to suffice for a large deviations principle (LDP) (Agazzi et al., 2017a,b). However, a pivotal assumption in this work is that the process cannot "get stuck near the boundary" – it thus rules out the possibility of absorbing sets. To circumvent these obstacles, we rely on an assumption on the global dynamics of φ_t on D.

By assumption, there exists a convex compact set $W \subset D_1$ containing all positive attractors $\{\Gamma_i\}_{i=1,\dots,\ell}$. Further, for $\epsilon > 0$ sufficiently small, all reactions may fire in W. Let $D_{0,T}(W)$ denote the space of càdlàg functions $z : [0,T] \to W$ equipped with the topology induced from the Skorohod metric on $D_{0,T}(W)$ given by

$$d_d(z_1, z_2) = \inf_{\tau} \left\{ \max \left(\gamma(\tau), \sup_{0 \le t \le T} |z_1(t) - z_2(\tau(t))| \right) \right\},$$

where τ is a strictly increasing function on $[0,T]$ such that $\tau(0) = 0$, $\tau(T) = T$ and

$$\gamma(\tau) := \sup_{0 \le s \le t \le T} \left| \log \frac{\tau(s) - \tau(t)}{s - t} \right| < \infty,$$

(Ethier and Kurtz, 1986). The induced topology is coarser than the topology of uniform convergence (Dembo and Zeitouni, 1998). Further, let $AC_{0,T}(W)$ denote the subspace of absolutely continuous functions from $[0,T]$ to W.

Theorem 4.4.2. For each $\epsilon > 0$, the sample paths $\{X_t^\epsilon : t \in [0,T]\}$ satisfy the LDP in $D_{0,T}(W)$ with speed $1/\epsilon$ and the good rate function $I : W \times [0,\infty) \times D_{0,T}(W) \to [0,\infty]$ given by

$$I(x,T,z) = \begin{cases} \int_0^T L(z(t), \dot{z}(t))\, dt, & z(0) = x, z \in AC_{0,T}(W) \\ \infty, & \text{otherwise} \end{cases}$$

uniformly for x on compact subsets of W, where L is the Lagrangian

$$L(y,\beta) = \sup_{\theta \in \mathbb{R}^d} \left(\langle \theta, \beta \rangle - \sum_{k \in \mathcal{R}} \lambda_k(y) \left[e^{\langle \theta, \xi_k \rangle} - 1 \right] \right).$$

That is, for any $B \subset D_{0,T}(W)$ and $K \subset W$ compact,

$$\limsup_{\epsilon\to 0} \epsilon \log \sup_{x\in K} \mathbb{P}_x[X^\epsilon_\cdot \in \overline{B}] \leq -\inf_{x\in K} \inf_{z\in \overline{B}} I(x,T,z) \tag{4.4.3}$$

$$\liminf_{\epsilon\to 0} \epsilon \log \inf_{x\in K} \mathbb{P}_x[X^\epsilon_\cdot \in B^\circ] \geq -\sup_{x\in K} \inf_{z\in B^\circ} I(x,T,z). \tag{4.4.4}$$

Proof. In the limit $\epsilon \to 0$, the generator (4.2.4) is identical to the expression (10.1) in Dupuis and Ellis (2011) when letting $a(x) \equiv 0 \in Mat_{d,d}(\mathbb{R})$ and $b(x) = \sum_{k\in\mathcal{R}} \xi_k \lambda_k(x) \in \mathbb{R}^d$ and the family of measures $\mu_x(dy) = \sum_{k\in\mathcal{R}} \delta_{\xi_k}(y)\lambda_k(x)$. We extend the kernel $\mu_x(dy)$ to all of $\mathbb{R}^d$ by letting

$$\sigma_x(dy) = \mu_{p_W(x)}(dy),$$

where p_W is the convex projection onto W. Condition 10.2.2 of Dupuis and Ellis (2011) is then easily satisfied. The support of σ_x is the set $\{\xi_k \colon k \in \mathcal{R}\}$. Letting Σ be the relative interior of the smallest convex cone that contains the support of σ_x, we find

$$\Sigma = ri(conv\{\xi_k \colon k \in \mathcal{R}\}),$$

which is evidently independent of $x \in \mathbb{R}^d$. Finally, as there exists a positive attractor we claim that we must have

$$0 \in \Sigma,$$

hence condition 10.2.4 Dupuis and Ellis (2011) would be satisfied as well. To see this, note that as $\{0\}$ and Σ are convex and $0 \notin \partial\Sigma$, it follows by the hyperplane separation theorem that $0 \notin \Sigma$ exactly when there exists a separating hyperplane. Suppose therefore for contradiction that there exists a separating hyperplane with normal vector γ. We may choose γ such that $\langle \gamma, \xi_k \rangle > 0$ for all $k \in \mathcal{R}$. Let $U \subset D_1$ be a fundamental neighborhood of the attractor K. Then, for all $x \in U$,

$$\langle \gamma, F(x) \rangle = \sum_{k\in\mathcal{R}} \lambda_k(x) \langle \gamma, \xi_k \rangle > 0,$$

that is, the vector field F points towards the same half-space on U which contradicts K being an attractor. We conclude that $0 \in \Sigma$ as desired.

Applying Theorem 10.2.6 of Dupuis and Ellis (2011) we find an LDP for σ on all of $\mathbb{R}^d$ with speed $1/\epsilon$ and good rate function

$$I_\sigma(x,T,z) = \begin{cases} \int_0^T L_\sigma(z(t), z'(t))\, dt, & z \in AC_{0,T}(\mathbb{R}^d) \\ 0, & \text{otherwise.} \end{cases}$$

Now, when $z(\cdot) \subset W$ we have $\sigma_x(dy) = \mu_x(dy)$ and $H_\sigma(z,\theta) = H_\mu(z,\theta)$ hence $L_\sigma(z,u) = L_\mu(z,u)$ and $I_\sigma(x,T,z) = I_\mu(x,T,z)$ which proves the Laplace principle on compacts. Finally, Theorem 1.2.3 Dupuis and Ellis (2011) gives the desired. □

4.4.2 Metastability

Using the LDP, we immediately find that if there exists a positive attractor, then any limiting distribution has support away from D_0. That is, the system is metastable. Indeed, inspired by Marmet (2013); Faure and Schreiber (2014) we obtain the following proposition.

Proposition 4.4.5. *Suppose the semi-flow $\{\varphi_t\}_{t\geq 0}$ admits an attractor $\Gamma \subset D_1$. Then there exists a neighborhood V_0 of D_0 such that $\nu(V_0) = 0$.*

Proof. Let $K \subset D_1$ be a compact neighborhood of the attractor Γ such that $K \subseteq W$. We may by definition find a fundamental neighborhood $U \subset K$ of Γ and $\delta > 0$ such that $N^\delta(\varphi_1(\overline{U})) \subset U$. Note that for $\epsilon > 0$ sufficiently small, $E^{\epsilon_n} \cap U \neq \emptyset$. Thus, with $T \geq 1$,

$$\begin{aligned} e^{-\theta_n}\nu_n(U) &= \int_{D_1} p^{\epsilon_n}(1,x,U)\nu_n(dx) \geq \nu_n(U) \inf_{x\in U} p^{\epsilon_n}(1,x,U) \\ &= \nu_n(U)\left(1 - \sup_{x\in U} p^{\epsilon_n}(1,x,U^c)\right) \geq \nu_n(U)\left(1 - \sup_{x\in U} p^{\epsilon_n}(1,x,N^\delta(\varphi_1(\overline{U}))^c)\right) \\ &\geq \nu_n(U)\left(1 - \sup_{t\in[0,T]}\sup_{x\in K} p^{\epsilon_n}(t,x,N^\delta(\varphi_t(x))^c)\right) = \nu_n(U)(1-\beta_{\delta,K}(\epsilon_n)). \end{aligned}$$

As $\nu_n(U) > 0$ for any n by irreducibility of E^{ϵ_n}, we conclude that

$$1 - \beta_{\delta,K}(\epsilon_n) \leq e^{-\theta_n}. \tag{4.4.6}$$

Now, for any neighborhood V_0 of D_0,

$$\begin{aligned} 1 - \beta_{\delta,K}(\epsilon_n) \leq e^{-\theta_n} &= \int_{E^\epsilon} p^{\epsilon_n}(1,x,E^\epsilon)\nu_n(dx) = \int_{E^\epsilon}(1 - p^{\epsilon_n}(1,x,A^\epsilon))\nu_n(dx) \\ &= \int_{E^\epsilon\setminus V_0}(1 - p^{\epsilon_n}(1,x,A^\epsilon))\nu_n(dx) + \int_{V_0}(1 - p^{\epsilon_n}(1,x,A^\epsilon))\nu_n(dx) \\ &\leq \nu_n(E^\epsilon\setminus V_0) + \nu_n(V_0)\left(1 - \inf_{x\in V_0} p^{\epsilon_n}(1,x,A^\epsilon)\right) \\ &= \nu_n(E^\epsilon) - \nu_n(V_0)\inf_{x\in V_0} p^{\epsilon_n}(1,x,A^\epsilon) = 1 - \nu_n(V_0)\inf_{x\in V_0} p^{\epsilon_n}(1,x,A^\epsilon), \end{aligned}$$

thus we infer

$$\nu_n(V_0) \leq \frac{\beta_{\delta,K}(\epsilon_n)}{\inf_{x\in V_0} p^{\epsilon_n}(1,x,A^\epsilon)}. \tag{4.4.7}$$

As D_1 is invariant under φ_t, we may choose $\delta_0 > 0$ such that

$$\delta_0 < \inf_{t\in[0,T],x\in K} d(\varphi_t(x), D_0).$$

Without loss of generality, we may assume $\sup(d(W, D_0)) < \delta_0$ hence $\varphi_t(x) \in W$ for all

$t \in [0,T]$ with $x \in K$. Define the quantity

$$c(K) := \frac{1}{4} \inf \left(I(x,T,z) \colon x \in K, d(z, \varphi_\cdot(x)) \geq \delta_0 \right) > 0.$$

Positivity comes from the fact that $I(x,T,z) = 0$ iff z solves the differential equation, that is $z(t) = \varphi_t(x)$ for all $t \in [0,T]$. By the LDP from Theorem 4.4.2 we have, for $x \in K$

$$\begin{aligned} \limsup_{\epsilon \to 0} \epsilon \log \mathbb{P}_x(X^\epsilon_\cdot \in D_{0,T}(W), d(X^\epsilon_\cdot, \varphi_\cdot(x)) \geq \delta_0) &\leq - \inf_{x \in K} \inf_{z \in D_{0,T}(W) : d(z(\cdot), \varphi_\cdot(x)) \geq \delta_0} I(x,T,z) \\ &\leq -4c(K). \end{aligned}$$

Thus, for $\epsilon > 0$ sufficiently small and all $x \in K$,

$$\beta_{\delta_0,K}(\epsilon) = \mathbb{P}_x(X^\epsilon_\cdot \in D_{0,T}(W), d(X^\epsilon_\cdot, \varphi(x)) \geq \delta_0) \leq e^{-3c(K)/\epsilon}. \tag{4.4.8}$$

By Lemma 4.3.5, there exists a neighborhood $V_{c(K)}$ of D_0 such that

$$\lim_{\epsilon \to 0} \inf_{x \in V_{c(K)}} \epsilon \log p^\epsilon(1,x,A^\epsilon) \geq -2c(K).$$

Thus, for $\epsilon > 0$ sufficiently small, we have for all $x \in V_{c(K)}$

$$p^\epsilon(1,x,A^\epsilon) \geq e^{-2c(K)/\epsilon},$$

and we conclude from (4.4.7) that

$$\nu\left(V_{c(K)}\right) \leq \lim_{\epsilon \to 0} \frac{\beta_{\delta,K}(\epsilon)}{\inf_{x \in V_{c(K)}} p^\epsilon(1,x,A^\epsilon)} \leq \lim_{\epsilon \to 0} e^{-c(K)/\epsilon} = 0,$$

as desired. □

4.5 The General Main Result

We need a few preliminary results before proving the main theorem. These rest on the large deviation principle of Theorem 4.4.2.

Proposition 4.5.1. *Let K be a compact subset of W. For all $\eta, \delta, \tilde{T} > 0$ there exists $\epsilon_0 > 0$ such that*

$$\mathbb{P}_x(d(z, X^\epsilon_\cdot) < \eta) \geq \exp\left(-\frac{I(x,T,z) + \delta}{\epsilon} \right)$$

for any $x \in K$, $\epsilon \in (0, \epsilon_0)$, $T < \tilde{T}$ and all $z \in D_{[0,T]}(K)$.

Proof. We follow (Kifer, 1988; Faure and Schreiber, 2014; Marmet, 2013) and define the

quantity

$$\Delta^K_{\gamma,T} = \sup\{|I(x,T,z) - I(x',T,z')|\colon x,x' \in K, d(z,z') < \gamma\}.$$

As I is a good rate function, the lower level-sets $\{z \in D_{0,T}(W)\colon I(x,T,z) \le s\}$ are compact for every $s \in (0,\infty)$, $T > 0$ (Dembo and Zeitouni, 1998). In particular, $I(x,T,z)$ is lower semi-continuous in z. Therefore $\liminf_{z'\to z} I(x',T,z') \ge I(x,T,z)$. Equivalently, this may be written $\limsup_{z'\to z} I(x,T,z) - I(x',T'z') \le 0$. As $x \mapsto |x|$ is non-decreasing on $[0,\infty]$, we find that

$$\lim_{\gamma\to 0} \Delta^K_{\gamma,T} = \limsup_{z'\to z} |I(x,T,z) - I(x',T,z')| = 0.$$

Choose $\eta > 0$, $\delta > 0$ and $\tilde{T} > 0$ such that $\gamma < \eta$, $\Delta^K_{\gamma,\tilde{T}} < \delta/2$ and $N^\gamma(K) \subset W$. By the uniform LDP lower bound (4.4.4) we conclude that for $z \in D_{[0,T]}(K)$,

$$\liminf_{\epsilon\to 0} \epsilon \log \inf_{x\in K} \mathbb{P}_x(X^\epsilon_\cdot \in N^\gamma(z)) \ge -\sup_{x'\in K} \inf_{z'\in N^\gamma(z)} I(x',T,z'),$$

where $N^\gamma(z) = \{z' \in D_{[0,T]}(K)\colon d(z,z') < \gamma\}$. We conclude that there exists a function $g\colon (0,\infty) \to (0,\infty)$ such that

$$\lim_{\epsilon\to 0} g(\epsilon) = 0, \qquad \text{and} \qquad \epsilon \log \inf_{x\in K} \mathbb{P}_x(X^\epsilon_\cdot \in N^\gamma(z)) \ge -\sup_{x'\in K} \inf_{z'\in N^\gamma(z)} I(x',T,z') - g(\epsilon).$$

It follows that for any $x \in K$,

$$\begin{aligned}
\mathbb{P}_x(d(z,X^\epsilon_\cdot) < \eta) &\ge \mathbb{P}_x(d(z,X^\epsilon_\cdot) < \gamma) \ge \inf_{x\in K} \mathbb{P}_x(X^\epsilon_\cdot \in N^\gamma(z)) \\
&\ge \exp\left(\frac{-\sup_{x'\in K} \inf_{z'\in N^\gamma(z)} I(x',T,z') - g(\epsilon)}{\epsilon}\right) \\
&= \exp\left(\frac{-I(x,T,z) - \sup_{x'\in K} \inf_{z'\in N^\gamma(z)} (I(x',T,z') - I(x,T,z)) - g(\epsilon)}{\epsilon}\right) \\
&\ge \exp\left(\frac{-I(x,T,z) - \Delta^K_{\gamma,T} - g(\epsilon)}{\epsilon}\right).
\end{aligned}$$

Choosing ϵ sufficiently small so that $g(\epsilon) < \delta/2$, we conclude that

$$\mathbb{P}_x(d(z,X^\epsilon_\cdot) < \eta) \ge \exp\left(-\frac{I(x,T,z) + \delta}{\epsilon}\right),$$

as desired. □

Proposition 4.5.2. *Let $K_1,\dots,K_b$ be compact non-attractors. Given any $\delta > 0$, there exist neighborhoods $V_j \subset N^\delta(K_j)$ of K_j for $1 \le j \le b$ and $\delta_1 \in (0,1)$ such that*

$$\sup_{x\in V_j} \mathbb{P}_x(\tau^\epsilon_{V_j} > h(\epsilon)) \le \zeta(\epsilon), \qquad \lim_{\epsilon\to 0} \zeta(\epsilon) = 0$$

for a function h satisfying $\lim_{\epsilon\to 0} h(\epsilon)\beta_{\delta_1,K}(\epsilon) = 0$.

Proof. The proof is given in Faure and Schreiber (2014); Marmet (2013), however, for completeness we reformulate it to fit our notation. Let $j \in \{1,\dots,b\}$. As K_j is not an attractor, there exists $\eta > 0$ such that $\overline{N^{2\eta}(K_j)} \subset W$ and for any $\gamma > 0$ and any $x \in N^\eta(K_j) \subset U := N^{2\eta}(K_j)$ there exists T^γ and z^γ such that

$$z^\gamma(0) = x, \qquad z^\gamma(T^\gamma) = y^\gamma \notin U, \qquad I(x, T^\gamma, z^\gamma) < \gamma.$$

Let $r > 0$ such that $N^r(U) \subset K$ for some compact set in W. As z^γ starts in U and ends outside U, it must pass through $K\backslash U$. Without loss of generality, we may assume that $y^\gamma \in K\backslash U$ and z^γ is contained in K.

By Proposition 4.5.1 we find that for any $\eta > 0, \delta > 0$ and $\tilde{T} > 0$ there exists $\epsilon_0 > 0$ such that for all $\epsilon \in (0, \epsilon_0)$ and $x \in K$,

$$\mathbb{P}_x(d(z^\gamma, X^\epsilon_\cdot) < \eta) \geq \exp\left(-\frac{I(x, T^\gamma, z^\gamma) + \delta}{\epsilon}\right) \geq \exp\left(-\frac{\gamma + \delta}{\epsilon}\right).$$

Now, if $d(z^\gamma, X^\epsilon_\cdot) < \eta$ then $X^\epsilon_\cdot$ leaves $N^\eta(K_j)$ before T^γ thus

$$\mathbb{P}_x\left(\tau^\epsilon_{N^\eta(K_j)} < T^\gamma\right) \geq \mathbb{P}_x(d(z^\gamma, X^\epsilon) < \eta) \geq \exp\left(-\frac{\gamma + \delta}{\epsilon}\right).$$

In particular, for all $x \in N^\eta(K_j)$ and ϵ sufficiently small, we find by the Markov property

$$\begin{aligned}
\mathbb{P}_x\left(\tau^\epsilon_{N^\eta(K_j)} > e^{2\gamma/\epsilon}\right) &= \mathbb{P}_x\left(\tau^\epsilon_{N^\eta(K_j)} > T^\gamma \frac{e^{2\gamma/\epsilon}}{T^\gamma}\right) \leq \mathbb{P}_x\left(\tau^\epsilon_{N^\eta(K_j)} > T^\gamma\right)^{\frac{e^{2\gamma/\epsilon}}{T^\gamma}} \\
&= \left(1 - \mathbb{P}_x\left(\tau^\epsilon_{N^\eta(K_j)} < T^\gamma\right)\right)^{\frac{e^{2\gamma/\epsilon}}{T^\gamma}} \leq \left(1 - \exp\left(-\frac{\gamma+\delta}{\epsilon}\right)\right)^{\frac{e^{2\gamma/\epsilon}}{T^\gamma}} \\
&= e^{\ln\left(1-\exp(-\frac{\gamma+\delta}{\epsilon})\right)\frac{e^{2\gamma/\epsilon}}{T^\gamma}} < e^{-e^{-\frac{\gamma+\delta}{\epsilon}}\frac{e^{2\gamma/\epsilon}}{T^\gamma}} = e^{-\frac{e^{\frac{\gamma-\delta}{\epsilon}}}{T^\gamma}}.
\end{aligned}$$

Choosing $\delta < \gamma$, the desired follows by taking $h(\epsilon) = e^{2\gamma/\epsilon}$ and $\zeta(\epsilon) = \exp(-e^{-\frac{\gamma-\delta}{\epsilon}}/T^\gamma)$. □

4.5.1 Absorption Preserving Chain Recurrence

In order to prove the main result, we need some theory on chain recurrence and pseudo orbits, which can be traced back to Conley (1978). The generalization to absorption preserving chain recurrence was introduced in Jacobs and Schreiber (2006).

Definition 4.5.3. Let $\{\varphi_t\}_{t>0}$ be a semi-flow given by an ordinary differential equation on $D \subset \mathbb{R}^d_0$ for which D_0 is an invariant set. Given $\delta > 0$, a piecewise continuous path from x to y,

$$x = x_1, \{\varphi_t(x_1) : t \in [0, t_1]\}, \{\varphi_t(x_2), t \in [0, t_2]\} \dots \{\varphi_t(x_k), t \in [0, t_k]\}, x_{k+1} = y$$

which is uniquely defined by the sequence of points $x_1, \ldots, x_{k+1}$ and times $t_1, \ldots, t_k$ such that the following holds

$$\begin{aligned} &d(\varphi_{t_j}(x_j), x_{j+1}) < \delta, \quad \forall j = 1, \ldots, k \\ &t_i \geq T, \quad \forall i = 1, \ldots, k, \\ &x_j \in D_0 \Rightarrow x_{j+1} \in D_0, \quad \forall j = 1, \ldots, k, \end{aligned}$$

is called an absorption preserving (δ, T)-pseudoorbit, or an ap (δ, T)-pseudoorbit for short. If there exists an ap (δ, T)-pseudoorbit from x to y, we will write $x <_{ap,\delta,T} y$. If $x <_{ap,\delta,T} y$ holds for every $\delta > 0$ and $T > 0$, we will write $x <_{ap} y$. If $x <_{ap} y$ and $y <_{ap} x$ we write $x \sim_{ap} y$. If $x \sim_{ap} x$ then x is an ap-chain recurrent point. We denote by $\mathcal{R}_{ap}$ the set of ap-chain recurrent points. The relation $\sim_{ap}$ restricted to this set defines an equivalence relation and the equivalence classes, $[x]_{ap}$ with $x \in \mathcal{R}_{ap}$ are called ap-basic classes. We write $[x]_{ap} <_{ap} [y]_{ap}$ if $x <_{ap} y$. A maximal ap-basic class is called an ap-quasiattractor.

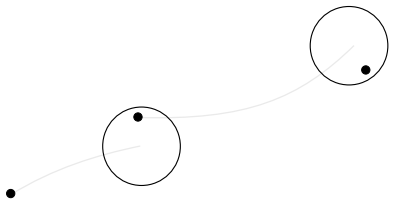

Figure 4.2: Schematic representation of a (δ, T)-pseudoorbit.

In the appendix, the theory of ap-chain-recurrence is developed in detail, under the following dynamical assumption.

Assumption 3. There exists a $T > 0$ and a compact set K_T such that for all $x \in D$ and all $t \geq T$, the semi-flow $\{\varphi_t\}_{t \geq 0}$ satisfies $\varphi_t(x) \in K_T$.

Note that this is slightly different from the requirement for permanence. On one hand it is stronger since we require a *common* time T after which all orbits must be in the compact set K_T. On the other hand, we don't require this set to be strictly inside $\mathbb{R}^d_0$. Note, however, that if the network is strongly endotactic, an assumption on the order of the reactions would be sufficient for this strong dissipative behavior. In general, Assumption 3 holds for a much larger class of networks, including those for which the existence of a unique quasi-stationary distribution is guaranteed (see Hansen and Wiuf, 2018a). The second assumption is of technical nature.

Assumption 4. There is only a finite number of ap-basic classes in D_1: $\{K_i\}_{i=1,\ldots,v}$. These are all closed sets with $\{K_i\}_{i=1,\ldots,\ell}$ being attractors, $\ell \geq 1$, while $\{K_i\}_{i=\ell+1,\ldots,v}$ are not attractors.

With these assumptions, one may arrive at the following corollary crucial to proving the general main theorem (Faure and Schreiber, 2014; Marmet, 2013).

Corollary 4.5.4. Given $\delta, T > 0$ there exists a family $\{V_i\}_{i=1,\dots,v}$ of isolating open neighborhoods of the ap-basic classes $\{K_i\}_{i=1,\dots,v}$, and positive constants δ_1 and T_1 such that

(i) $\overline{N^{\delta_1}(K_i)} \subset V_i$ for $1 \leq i \leq v$;

(ii) any ap (δ_1, T_1)-pseudoorbit starting in V_i remains in V_i for $i = 1, \dots, \ell$;

(iii) if there exists an ap (δ_1, T_1)-pseudoorbit $(x, t_1), \dots, (x_n, t_n), y$ with $x \in N^{\delta_1}(K_i)$ and $y \in N^{\delta_1}(K_{i'})$ such that

$$\varphi_s(x_j) \notin V_i, \qquad \text{for some } j \in \{2, \dots, n-1\}, s \in [0, t_j]$$

then $i \neq i'$ and $K_i <_{ap} K_{i'}$.

(iv) Every ap (δ_1, T_1)-pseudoorbit intersects $N^{\delta}(\mathcal{R}_{ap})$.

Support Concentrated Away From Non-attractors. We are now in a position to prove the general version of the main theorem. Once this is result is obtained, we shall present a more tangible version getting rid of the absorption preserving chain recurrence all together. Inspired by (Faure and Schreiber, 2014), we may prove the following.

Theorem 4.5.5 (Main)**.** Under Assumption 4 the limiting measure ν is supported by the union of positive attractors $\bigcup_{i=1}^{\ell} \Gamma_i$.

Proof. By Assumption 4, there is at least one positive attractor. Thus, we already know from (4.3.3) that $\theta_n \to 0$ for $n \to \infty$, that ν is φ-invariant and has its support in $BC(\varphi) \subseteq \mathcal{R}_{ap}$. Further, it follows from Proposition 4.4.5 that the support of ν is contained in D_1. In particular, $\nu(\mathcal{R}_{ap} \cap D_1) = 1$ hence $\nu(K) = 1$ where $K = \bigcup_{i=1}^{v} K_i$. We must show that for every non ap-quasiattractor, K_j, $j = \ell+1, \dots, v$, there exists an open neighborhood $V_j \supset K_j$ such that $\nu(V_j) = 0$.

Let $j \in \{\ell + 1, \dots, v\}$ be fixed and denote by $b = v - \ell$ the number of non-attractors in D_1. Further, we may choose sequences $\{m_n\}_{n\geq 1}$ and$\{m'_n\}_{n\geq 1}$ such that

$$\lim_{n\to\infty} m_n \beta_{\delta_1, K}(\epsilon_n) = 0, \qquad \lim_{n\to\infty} \frac{m'_n}{m_n} = 0, \qquad \lim_{n\to\infty} \frac{h(\epsilon_n)}{m'_n} = 0 \tag{4.5.6}$$

By the QSD property for μ_n, we have

$$\nu_n(V_j) = \frac{1}{e^{\theta_n t}} \int_{D_1} \mathbb{P}_x(X_t^{\epsilon_n} \in V_j) \nu_n(dx), \qquad t \geq 0.$$

For $j = \{1, \dots, b\}$ let $t_n^j = m_n/j$. Let δ_1 be given such that Corollary 4.5.4 holds. Define the events $\mathcal{E}_n$ and $\mathcal{E}'_n$ by

$$\mathcal{E}_n = \{(X_t^{\epsilon_n})_{t\in[0,m_n]} \text{ is an ap } (\delta_1, T_1)\text{-pseudoorbit}\}$$
$$\mathcal{E}'_n = \{\forall j \in \{\ell + 1, \dots, v\}, \forall s \geq m'_n, X_t^{\epsilon_n} \in N^{\delta_1}(K_j) \Rightarrow X_{t+s}^{\epsilon_n} \notin N^{\delta_1}(K_j)\}.$$

Thus, on the event $\mathcal{E}'_n$, after the first entry to $N^{\delta_1}(K_i)$ the Markov process will have escaped from the set after spending at most m'_n amounts of time there and will never come back. On the event $\mathcal{E}_n \cap \mathcal{E}'_n$, the process $(X_t^{\epsilon_n})_{t\in[0,m_n]}$ is furthermore an ap (δ_1, T_1)-pseudoorbit hence by Corollary 4.5.4 (i)-(ii) it gets trapped inside $\bigcup_{i=1}^{\ell} N^{\delta_1}(K_i)$ if it enters this set. Further, Corollary 4.5.4 (iii) and the property of $\mathcal{E}'_n$, ensures that this pseudo orbit cannot spend more than bm'_n amounts of time in $\bigcup_{i=\ell+1}^{v} N^{\delta_1}(K_i)$. We conclude that, for n sufficiently large so that $m_n > bm'_n$, if $X_{m_n}^{\epsilon_n}$ is in V_j then $X_{t_n^i}^{\epsilon_n} \notin N^{\delta_1}(K)$ for some $i \in \{1, \dots, b\}$. Therefore,

$$\mathbb{P}_x(\{X_{m_n}^{\epsilon_n} \in V_j\} \cap \mathcal{E}_n \cap \mathcal{E}'_n) \leq \sum_{i=1}^{b} \mathbb{P}_x(X_{t_n^i}^{\epsilon_n} \notin N^{\delta_1}(K)).$$

On the other hand, it follows from Corollary 4.5.4 (iii) that on the event $\mathcal{E}_n$, starting from $N^{\delta_1}(K_i)$, the chain cannot enter back into $N^{\delta_1}(K_i)$ once it exited V_i. Let $\tau_V^n = \inf\{t \geq 0\colon X_t^{\epsilon_n} \notin V\}$. Proposition 4.5.2 implies that, as $m'_n > h(\epsilon_n)$ for n sufficiently large,

$$\mathbb{P}_x((\mathcal{E}'_n)^c \cap \mathcal{E}_n) \leq \sum_{i=\ell+1}^{v} \sup_{y\in V_i} \mathbb{P}_y(\tau_{V_i}^n \geq m'_n) \leq \sum_{i=\ell+1}^{v} \sup_{y\in V_i} \mathbb{P}_y(\tau_{V_i}^n \geq h(\epsilon_n)) \leq b\zeta(\epsilon_n).$$

By Lemma 4.8.12 we see, with $\{X_{m_n}^{\epsilon_n} \in V_j\} = A_1$, $\mathcal{E}_n = A_2$ and $\mathcal{E}'_n = A_3$ that

$$\begin{aligned}
\mathbb{P}_x(\{X_{m_n}^{\epsilon_n} \in V_j\}) &\leq \mathbb{P}_x(\mathcal{E}_n^c) + \mathbb{P}_x((\mathcal{E}'_n)^c|\mathcal{E}_n) + \mathbb{P}_x(\{X_{m_n}^{\epsilon_n} \in V_j\}|\mathcal{E}_n \cap \mathcal{E}'_n) \\
&\leq \mathbb{P}_x(\mathcal{E}_n^c) + \mathbb{P}_x((\mathcal{E}'_n)^c \cap \mathcal{E}_n) + \mathbb{P}_x(\{X_{m_n}^{\epsilon_n} \in V_j\} \cap \mathcal{E}_n \cap E'_n) \\
&\leq \beta_{\delta_1,K}(\epsilon_n) + b\zeta(\epsilon_n) + \sum_{i=1}^{b} \mathbb{P}_x(X_{t_n^i}^{\epsilon_n} \notin N^{\delta_1}(K)),
\end{aligned}$$

for n sufficiently large. By the QSD property of ν_n we find

$$\begin{aligned}
\int_{D_1} \mathbb{P}_x(\{X_{m_n}^{\epsilon_n} \in V_j\})\nu_n(dx) &\leq \beta_{\delta_1,K}(\epsilon_n) + b\zeta(\epsilon_n) + \sum_{i=1}^{b} \int_{D_1} \mathbb{P}_x(X_{t_n^i}^{\epsilon_n} \in (N^{\delta_1}(K))^c)\nu_n(dx) \\
&= \beta_{\delta_1,K}(\epsilon_n) + b\zeta(\epsilon_n) + \sum_{i=1}^{b} e^{-\theta_n t_n^i} \nu_n(N^{\delta_1}(K)^c) \\
&\leq \beta_{\delta_1,K}(\epsilon_n) + b\zeta(\epsilon_n) + b\nu_n(N^{\delta_1}(K)^c) \to 0,
\end{aligned}$$

for $n \to \infty$. Here we have used the fact that $\nu_n \to \nu$ with $\nu(N^{\delta_1}(K)^c) = 1 - \nu(N^{\delta_1}(K)) \leq 1 - \nu(K) = 0$. Finally,

$$\lim_{n\to\infty} \nu_n(V_j) = \lim_{n\to\infty} \frac{1}{e^{\theta_n m_n}} \int_{D_1} \mathbb{P}_x(X_{m_n}^{\epsilon_n} \in V_j)\nu_n(dx) = 0,$$

as $\lim_{n\to\infty} e^{\theta_n m_n} = 1$. Indeed, by (4.4.6) and the choice of m_n we find

$$\begin{aligned}\liminf_{n\to\infty} e^{-\theta_n m_n} &\geq \liminf_{n\to\infty}(1-\beta_{\delta_1,K}(\epsilon_n))^{m_n} = \exp\left(\liminf_{n\to\infty} m_n \log(1-\beta_{\delta_1,K}(\epsilon_n))\right)\\ &= \exp\left(-\liminf_{n\to\infty} m_n\beta_{\delta_1,K}(\epsilon_n) + m_n\beta_{\delta_1,K}(\epsilon_n)\frac{\beta_{\delta_1,K}(\epsilon_n)}{2} + \dots\right) = 1.\end{aligned}$$

By the property of weak convergence we then have $\nu(V_j) \leq \liminf_{n\to\infty} \nu_n(V_j) = 0$ and the desired follows - ν is not supported by the non-attractors. □

In the paper Faure and Schreiber (2014), one furthermore has the very natural result that if D_0 is a global attractor, then ν is supported by D_0. However, due to the restriction to a compact set K in our definition of $\beta_{\delta,K}$, the proof is not readily generalizable, although we strongly suspect it to hold. We furthermore note that the problem of determining the support of ν in the case where there is no positive attractor in D_1(and D_0 is not globally attracting) is still open (Schreiber, 2016).

4.5.2 Morse Decompositions

In order to state the main result in a more straightforward fashion which circumvents the theory of ap-basic classes, we shall assert global restrictions on the semi-low φ_t by introducing the notion of Morse decompositions from dynamical systems theory (Conley, 1978; Colonius and Kliemann, 2014).

It follows from Assumption 3 that the semi-flow $\{\varphi_t\}$ on the state space D associated to the deterministic reaction network satisfies that

$$G = \bigcap_{t\geq 0} \varphi_t(D)$$

is a global attractor. We introduce the following definition, which describes the global structure of the flow by ordering invariant subsets of the state space capturing the limit behavior forwards and backwards in time (Colonius and Kliemann, 2014).

Definition 4.5.7. A Morse decomposition of the dynamics of φ_t is a collection of non-empty φ-invariant pairwise disjoint compact sets $\{M_1, \dots, M_m\}$, called Morse sets, such that

- M_i is isolated, $\quad 1 \leq i \leq m$
- for every $x \in G \backslash \bigcup_{i=1}^m M_i$, there exists $i > j$ such that $\omega(x) \subseteq M_i$ and $\alpha(x) \subseteq M_j$.

Note that the Morse sets contain all limit sets, and no cycles between the Morse sets are allowed. A Morse decomposition thus describes the flow through movement from lower to higher indexed Morse sets. A Morse decomposition $\{M_1, \dots, M_m\}$ is called finer than a Morse decomposition $\{M'_1, \dots, M'_{m'}\}$ if for all $j \in \{1, \dots, m'\}$ there is $i \in \{1, \dots, m\}$ with $M_i \subset M'_j$. There need not exist a finest Morse decomposition, but when it does, it is unique Ayala et al. (2006).

Proposition 4.5.8. *There exists a finest Morse decomposition $\{M_i\}_{i=1,\dots,m}$ if and only if $\mathcal{R}_{ap}$ has only finitely many ap-basic classes. In this case, the ap-basic classes in D_1 coincide with the Morse sets $M_i \subset D_1$. In particular, there is a finite number of ap-basic classes in D_1 and each of them is closed.*

Proof. By (Ayala et al., 2006, Theorem 6.4), there exists a finest Morse decomposition $\{M_1, \dots, M_m\}$ if and only if the regular chain recurrent set $\mathcal{R}$ has only finitely many basic classes. As $\mathcal{R}_{ap} \subseteq \mathcal{R}$, there are only finitely many ap-basic classes. Furthermore, as the Morse sets coincide with the basic classes by (Ayala et al., 2006, Theorem 6.4) whenever $M_i \subset D_1$, it must coincide with an ap-basic class. As the Morse sets are compact, they are in particular closed. □

One may now state the main theorem in terms completely eradicated from the complexities of absorption preserving chain recurrence. We note that if the system is gradient or more generally Morse-Smale, then a Morse decomposition exists (Faure and Schreiber, 2014).

Corollary 4.5.9. Let $M_1, \dots M_m$ be the finest Morse decomposition for φ_t such that $M_j, \dots, M_m$ are attractors. If

- $M_i \subset D_0$ or $M_i \subset D_1$,
- $M_i \subset D_1$ for some $i \geq j$.

then any weak*-limit point of $\{\nu^\epsilon\}_{\epsilon>0}$ is φ_t-invariant and is supported by the union of attractors in D_1.

Proof. As there exists a finest Morse decomposition, it follows by Proposition 4.5.8 that there is a finite number of ap-basic classes in D_1 and each of them is closed. Further, as the Morse decomposition contains at least one positive attractor, Assumption 4 is satisfied. The desired now follows from Theorem 4.5.5. □

To examine the existence and structure of connecting orbits between the Morse sets, one may apply the Conley index for which we refer the reader to Conley (1978). Further, the connection between Morse-decompositions and Lyapunov functions could be of particular interest in reaction network theory (Patrao, 2007).

4.6 Examples

In this section, we apply the main result, in the form of Corollary 4.5.9, on reaction networks displaying a range of different dynamical properties.

4.6.1 1D Reaction Networks

Example 4.6.1. The logistic network from the introduction, giving rise to Keizer's paradox,

$$\emptyset \xleftarrow{\alpha_1} S \underset{\alpha_3}{\overset{\alpha_2}{\rightleftarrows}} 2S$$

has the corresponding deterministic flow defined on the state space $D = [0, \infty)$ where

$$\dot{x} = (\alpha_2 - \alpha_1)x - \alpha_3 x^2.$$

The set $D_0 = \{0\}$ is closed and absorbing and $D_1 = (0, \infty)$ is open and invariant under φ. Assuming $\alpha_2 > \alpha_1$, the flow may be written out explicitly,

$$\varphi_t(x) = \frac{\alpha_2 - \alpha_1}{\frac{-\alpha_3 x + \alpha_2 - \alpha_1}{x} e^{-(\alpha_2 - \alpha_1)t} + \alpha_3},$$

for $t \geq 0$. Thus we may explicitly see that Assumption 3 is satisfied. Note that the global attractor is $G = [0, x^*]$ where $x^* = \frac{\alpha_2 - \alpha_1}{\alpha_3}$. It follows that the finest Morse decomposition of the dynamics is given by

$$M_1 = \{0\}, \qquad M_2 = \{x^*\}$$

where $M_2 \subset D_1$ is an attractor. By Proposition 4.4.5, there is a neighborhood V_0 of $\{0\}$ such that all weak* limit points ν of $\{\nu_\epsilon\}_{\epsilon > 0}$ have $\nu(V_0) = 0$. In other words: the stochastic process exhibits quasi-stationary or metastable behavior for large volume size. In addition, Corollary 4.5.9 implies that any weak* limit point of $\{\nu_\epsilon\}_{\epsilon > 0}$ is supported by the attractor M_2. This is indeed what happens numerically, as illustrated in Figure 4.3.

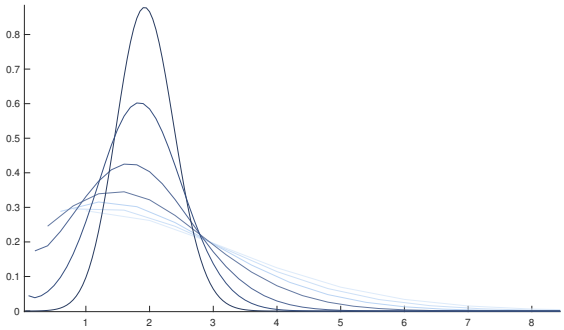

Figure 4.3: Quasi-stationary distributions for the logistic network with $\alpha_1 = 1, \alpha_2 = 3, \alpha_3 = 1$. Darker colors correspond to smaller values of the scaling parameter ϵ.

We have therefore resolved Keizer's paradox in the sense that we may bridge the gap between the deterministic and the stochastic descriptions of the underlying network which predict fundamentally different outcomes. Note that the poisson-form of the distribution should not come as a surprise. Indeed, when $\alpha_2 > \alpha_1$, the weakly reversible deficiency zero network

$$S \underset{\alpha_3}{\overset{\alpha_2-\alpha_1}{\rightleftarrows}} 2S$$

has the same deterministic dynamics as the original network, and this related network has a Poisson distribution with parameter $\frac{\alpha_2-\alpha_1}{\alpha_3}$ (Anderson et al., 2010). △

Example 4.6.2. The previous example might seem obvious, as there are no non-positive attractors. Consider now the reaction network

$$\emptyset \overset{\alpha_1}{\leftarrow} S_1 \quad 2S_1 \underset{\alpha_3}{\overset{\alpha_2}{\rightleftharpoons}} 3S_1.$$

The deterministic rate equation under mass-action is given by

$$\dot{x} = \alpha_2 x^2 - \alpha_1 x - \alpha_3 x^3,$$

defined on the state space $D = [0, \infty)$ where $D_0 = \{0\}$ is closed and absorbing and $D_1 = (0, \infty)$. A standard analysis of the dynamical system reveals that, under the assumption that $\alpha_2^2 > 4\alpha_1\alpha_3$, a finest Morse decomposition is given by

$$M_1 = \left\{ \frac{\alpha_2 - \sqrt{\alpha_2^2 - 4\alpha_1\alpha_3}}{2\alpha_3} \right\}, \quad M_2 = \{0\}, \quad M_3 = \left\{ \frac{\alpha_2 + \sqrt{\alpha_2^2 - 4\alpha_1\alpha_3}}{2\alpha_3} \right\},$$

with $G = [0, M_3]$ being the global attractor. Here, $M_3 \subset D_1$ is the only positive attractor hence by Corollary 4.5.9, the limit of the QSDs, guaranteed to exist and be unique by Hansen and Wiuf (2018a), will be supported by this set. Again, this may be verified numerically as seen in Figure 4.4. △

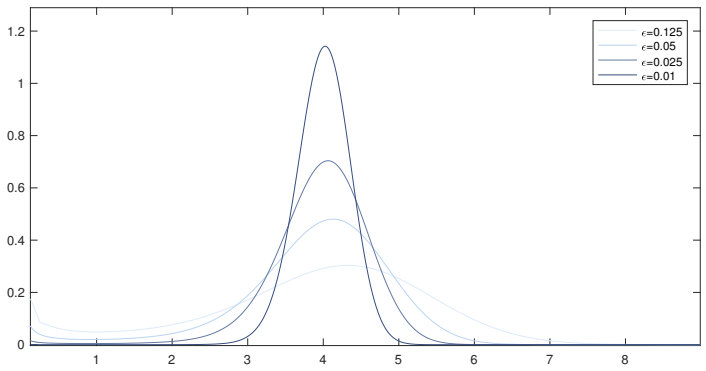

Figure 4.4: $\alpha_1 = 8, \alpha_2 = 6$ and $\alpha_3 = 1$. Thus $M_1 = \{2\}, M_2 = \{0\}$ and $M_3 = \{4\}$.

Example 4.6.3. Consider the reaction network

$$\emptyset \underset{\alpha_2}{\overset{\alpha_1}{\rightleftarrows}} S \to 2S \qquad 3S \underset{\alpha_4}{\overset{\alpha_3}{\rightleftarrows}} 4S$$

One may conclude by Hansen and Wiuf (2018a) that the associated stochastic process has a unique QSD for each $\epsilon > 0$ and all parameter values. The state space is $D = \{0\} \sqcup (0, \infty)$, and one may readily verify that for the specific choice of parameter values $\alpha_1 = 900, \alpha_2 = 320, \alpha_3 = 33, \alpha_4 = 1$, the finest Morse decomposition is

$$M_1 = \{0\}, \qquad M_2 = \{10\}, \qquad M_3 = \{5\}, \qquad M_4 = \{18\}.$$

where M_3 and M_4 are positive attractors. As there are multiple positive attractors, we may conclude by Corollary 4.5.9 that

$$\operatorname{supp} \nu \subseteq \{5\} \cup \{18\}. \tag{4.6.4}$$

Numerically, we obtain the interesting behavior of Figure 4.5. Initially, for from equilibrium, the QSD is unimodal with a mode approaching the positive attractor $\{5\}$ as ϵ is decreased. For even smaller values of ϵ, the QSD becomes multimodal, with a second mode appearing at the other positive attractor $\{18\}$. Continuing to decrease ϵ, we observe a shift and the QSD is eventually only supported at the single positive attractor $\{18\}$. △

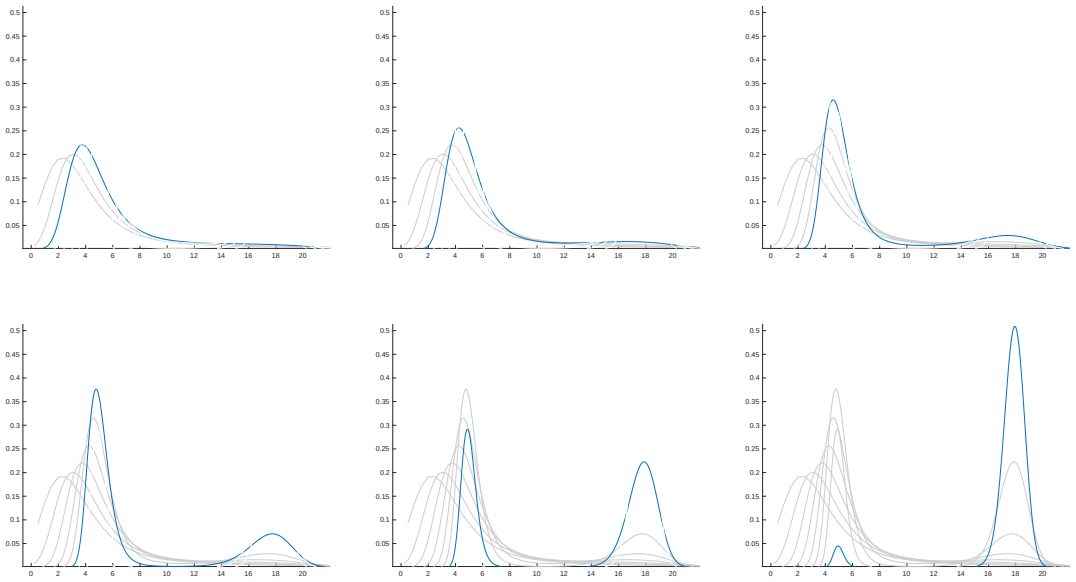

Figure 4.5: Quasi-stationary distributions for $\epsilon = 1/16, 1/32, 1/64, 1/128, 1/256, 1/512$ (blue). Previous QSDs are included for comparison (grey).

4.6.2 2D Reaction Networks

Example 4.6.5. We consider a slight alteration of the Bruxellator (Prigogine and Lefever, 1968), which was used to show the possibility of order originating from oscillations in a system out of equilibrium (Luisi, 2006).

$$\begin{aligned} S_1 &\xrightarrow{\alpha_1} 2S_1 \\ 2S_1 &\xrightarrow{\alpha_2} S_1 + S_2 \\ 3S_1 + S_2 &\xrightarrow{\alpha_3} 4S_1 \\ 2S_1 &\xrightarrow{\alpha_4} \emptyset \end{aligned}$$

We may write the state space $D = \mathbb{R}_0^2$ as $D = D_0 \cup D_1$, where

$$\begin{aligned} D_0 &= \{(x_1, x_2) \in \mathbb{R}_0^2 : x_1 = 0\}, \\ D_1 &= \{(x_1, x_2) \in \mathbb{R}_0^2 : x_1 > 0\}. \end{aligned}$$

Thus D_0 is a closed subset of D, D_0 and D_1 are positively φ-invariant and D_0 is absorbing for the random perturbations,

$$p^\epsilon(t, x, D_1) = 0,$$

for all $\epsilon > 0$, $t > 0$, $x \in D_0$. The deterministic rate equations are given by

$$\begin{aligned} \dot{x_1} &= \alpha_1 x_1 - \alpha_2 x_1^2 + \alpha_3 x_1^3 x_2 - 2\alpha_4 x_1^2, \\ \dot{x_2} &= \alpha_2 x_1^2 - \alpha_3 x_1^3 x_2. \end{aligned}$$

One may readily verify that the equilibrium points are

$$x^* = (0, x_2), \qquad x^{**} = \left(\frac{\alpha_1}{2\alpha_4}, \frac{2\alpha_2\alpha_4}{\alpha_1\alpha_3} \right).$$

where $x_2 \in \mathbb{R}_0$. The linearized system is

$$Df(x) = \begin{pmatrix} \alpha_1 - 2\alpha_2 x_1 + 2\alpha_3 x_1^2 x_2 - 4\alpha_4 x_1 & \alpha_3 x_1^3 \\ 2\alpha_2 x_1 - 3\alpha_3 x_1^2 x_2 & -\alpha_3 x_1^3 \end{pmatrix}$$

Evaluated at the equilibrium points one obtains

$$Df(x^*) = \begin{pmatrix} \alpha_1 & 0 \\ 0 & 0 \end{pmatrix}, \qquad Df(x^{**}) = \begin{pmatrix} -\alpha_1 & \frac{\alpha_3\alpha_1^3}{8\alpha_4^3} \\ -\frac{\alpha_1\alpha_2}{2\alpha_4} & -\frac{\alpha_3\alpha_1^3}{8\alpha_4^3} \end{pmatrix}. \tag{4.6.6}$$

It follows that the second axis consists of repelling stationary points, and that x^{**} is a stable stationary point.

Note that we are not guaranteed by Hansen and Wiuf (2018a) to have a unique QSD in this reaction network. Further, Assumption 3 is not satisfied hence we cannot conclude

that the limit distribution ν has support contained in the positive attractor $(\frac{\alpha_1}{2\alpha_4}, \frac{2\alpha_2\alpha_4}{\alpha_1\alpha_3})$. However, this behavior is exactly what is observed numerically. Figure 4.6 and 4.7 shows QSDs corresponding to diminishing values of the scaling parameter ϵ for the same altered Bruxellator, but with different rate parameters.

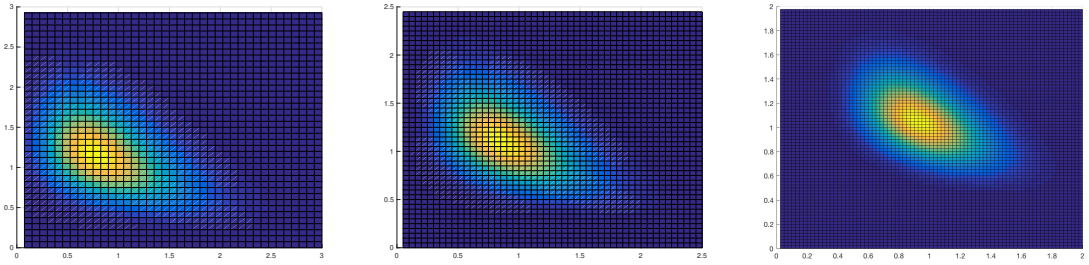

Figure 4.6: $\alpha_1 = 2, \alpha_2 = \alpha_3 = \alpha_4 = 1$. $\epsilon = 0.075, 0.05, 0.025$. Stable spiral.

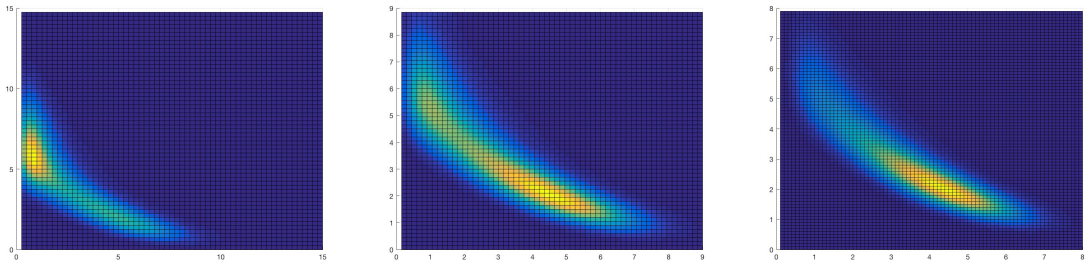

Figure 4.7: $\alpha_1 = \alpha_2 = 8, \alpha_3 = \alpha_4 = 1$. $\epsilon = 0.25, 0.15, 0.1$.

Based on the behavior for birth-death processes, one could take this as another indication of the conjecture that the weaker dissipativity allows a continuum of quasi-stationary distributions. Note furthermore how the behavior of the QSDs ν_ϵ change from large to small values of ϵ. In Figure 4.6 the convergence appears regular while in Figure 4.7, the mode switches from being close to the boundary to being at the stable stationary point. This suggests that quasi-stationarity is only observed for vary large system size. △

Example 4.6.7. Consider the reaction network

$$X \xrightarrow{\alpha_1} 2X \xrightarrow{\alpha_2} 2Y \xrightarrow{\alpha_3} \emptyset$$

The deficiency is $\delta = 4 - 2 - 1 = 1$ and there is a single terminal strong linkage class hence the deficiency one theorem applies and we may conclude that there is no multistationarity. The deterministic rate equations are

$$\dot{x_1} = \alpha_1 x_1 - 2\alpha_2 x_1^2,$$
$$\dot{x_2} = 2\alpha_2 x_1^2 - 2\alpha_3 x_2^2.$$

One may readily find the equillibrium points

$$x^* = (0,0), \qquad x^{**} = \left(\frac{\alpha_1}{2\alpha_2}, \frac{\alpha_1}{2\sqrt{\alpha_2\alpha_3}}\right),$$

and a linearization shows that x^* is unstable while x^{**} is stable. In fact, it is globally stable. It follows that the finest Morse decomposition is

$$M_1 = \{0\}, \qquad M_2 = \{x^{**}\}. \tag{4.6.8}$$

It follows from Hansen and Wiuf (2018a) that for each $\epsilon > 0$ there exists a unique QSD, ν_ϵ for the system hence one may apply Corollary 4.5.9 to conclude that the weak limit, ν, has support on $\delta_{x^{**}}$. Note that in this case, for each $\epsilon > 0$ there are two endorsed sets. △

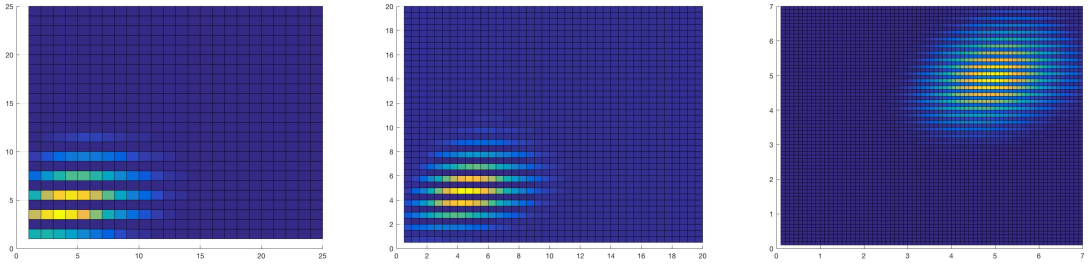

Figure 4.8: $\alpha_1 = 10, \alpha_2 = 1, \alpha_3 = 1$. $x = (5,5)$ is a stable equilibrium point (sink).

Example 4.6.9. To examine what happens for more complicated dynamics, we use the following reaction network. One may readily verify that the sufficient conditions of Hansen and Wiuf (2018a) are verified. Thus, for all $\epsilon > 0$ and any choice of rate parameter, there exists a unique quasi-stationary distribution.

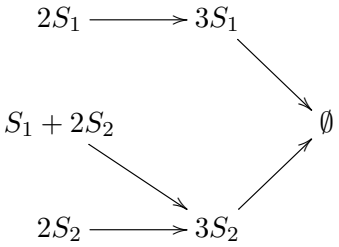

The state space is given by

$$D = \partial\mathbb{R}^2_0 \sqcup (0,\infty)^2. \tag{4.6.10}$$

One may, by a suitable choice of rate parameters obtain the following vector field for the deterministic system. There is a unique positive attractor, x^* as well as an attractor on the second axis, x^{**}.

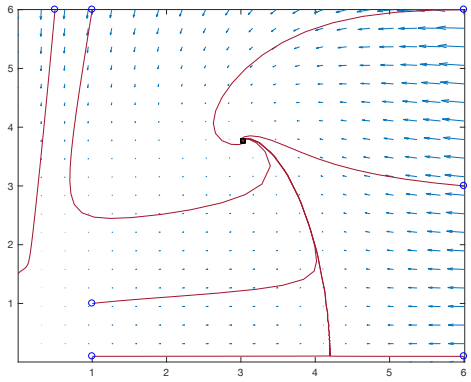

Figure 4.9: The vector field of the deterministic system (blue) with orbits (red) from the initial points (circles).

The minimal Morse-decomposition would therefore consists of the three sets

$$M_1 = \{0\}, \qquad M_2 = \{x^{**}\}, \qquad M_3 = \{x^{*}\}.$$

As there is a unique positive attractor, it follows from Corollary 4.5.9 that $\nu_\epsilon \Rightarrow \nu$ as $\epsilon \to 0$, where ν is supported by M_3. In Figure 4.10, we are able to see this numerically.

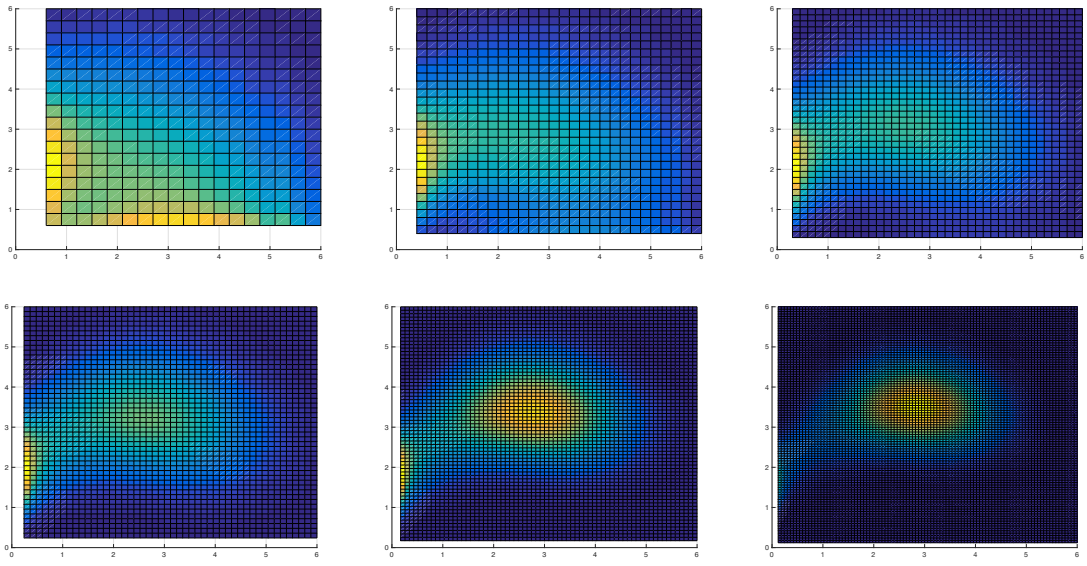

Figure 4.10: $\epsilon = 1/2, 1/4, 1/8, 1/16, 1/32, 1/64$.

Choosing other rate constants, one may obtain an attracting limit cycle. However, we are not able to display the convergence to the limit measure supported on such a limit cycle. The problems arise due to the complexity in finding eigenvalues of a large matrix. Indeed,

as can be seen from Figure 4.11, one would need $\epsilon \approx 10^{-3}$ in order to capture the structure of the limit cycle sufficiently well, which would correspond to a $10^8 \times 10^8$ matrix. △

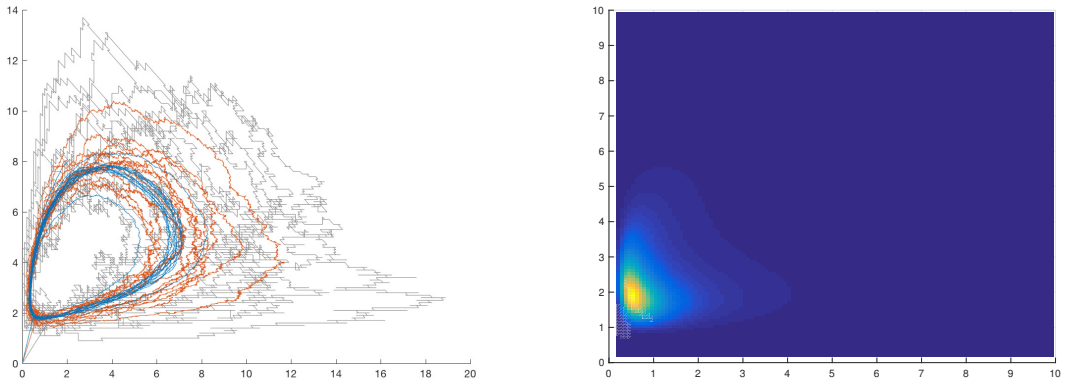

Figure 4.11: Left: Three sample paths for the reaction network with $\epsilon = 0.1$ (gray), $\epsilon = 0.01$ (red) and $\epsilon = 0.001$ (blue). Right: The QSD for $\epsilon = 0.083$.

4.6.3 Competitive systems

The general class of dynamical systems known as competitive systems are defined by having the off diagonal terms of the Jacobian matrix for the vector field F nonpositive. That is, $\partial F_i/\partial x_j \leq 0$ for $j \neq i$, where $\dot{x} = F(x)$ (Hirsch, 1982). The ecological Lotka-Volterra system with three competing species is one such example, and one may view it as a reaction network with underlying reaction graph as given below.

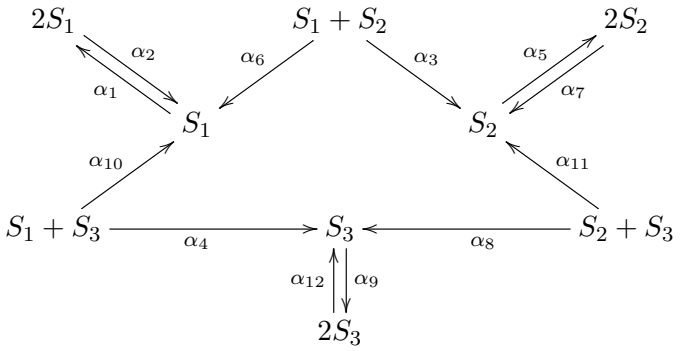

Indeed, the deterministic rate equations following mass-action is given by

$$\begin{aligned}
\dot{x}_1 &= x_1(\alpha_1 - \alpha_2 x_1 - \alpha_3 x_2 - \alpha_4 x_3),\\
\dot{x}_2 &= x_2(\alpha_5 - \alpha_6 x_1 - \alpha_7 x_2 - \alpha_8 x_3),\\
\dot{x}_3 &= x_3(\alpha_9 - \alpha_{10} x_1 - \alpha_{11} x_2 - \alpha_{12} x_3).
\end{aligned}$$

It is straightforward to verify that this system has a unique QSD for any set of rate parameters using Hansen and Wiuf (2018a). If in addition to Assumption 3, the set D_1 contains a unique equilibrium point x^* and $\det(DF(x^*)) < 0$, then as F is analytic in D_1, it follows from Theorem 1.2-1.3 Zhu and Smith (1994) that there is at least one but no more than finitely many periodic orbits, and at least one of these is orbitally asymptotically stable. Furthermore, from general results of Hirsch (1990), the system is Morse-Smale. In particular, by Corollary 4.5.9, the system displays quasi-stationarity and the QSDs ν_ϵ converges weakly to a measure ν supported on a union of limit cycles.

4.7 Conclusion

Using the classical scaling of Kurtz (1970, 1972, 1978), we were able to describe the stochastic reaction system as a random perturbation of the corresponding deterministic reaction system. Under the assumption of a strong dissipative flow, which also guarantees the existence of a unique QSD for all values of the scaling parameter ϵ (Hansen and Wiuf, 2018a), and the existence of a minimal Morse decomposition of the dynamics, we proved that the weak limit ν of the sequence of quasi-stationary measures ν_ϵ has support contained in the union of positive attractors for the deterministic flow. Furthermore, given the existence of a positive attractor, the expected time to extinction scales exponentially with system size ϵ^{-1}, hence for sufficiently small ϵ all such systems display quasi-stationarity. This correspondence between the stochastic and the deterministic descriptions in terms of the quasi-stationary distributions resolves Keizer's paradox (Keizer, 1987; Vellela and Qian, 2007).

4.7.1 Discussion of Freidlin-Wentzell Theory and Outlook

Far from equilibrium, that is for large values of the scaling parameter ϵ, the picture may be radically different as illustrated in the examples. Furthermore, as this regime is what calls for stochastic models in the first place, it would be of much interest to understand the extend to which the limiting measure may be of use in this non-equilibrium domain. Further, as we have seen in Example 4.6.3, when there are more than one positive attractor, the limiting measure is believed generically to only be supported on a single one of these, and one would like to know which of the attractors it is (Schreiber, 2016). To answer these questions, one normally tends to Freidlin-Wentzell theory (Freidlin and Wentzell, 2012), where the large deviation results are bootstrapped into estimates on the behavior of a process over very long time intervals, by splitting time into finite time-intervals and exploiting the Markov property (Shwartz and Weiss, 1995). However, this approach is developed for stationary measures where it rests on a large-deviations principle holding on the entire state space. It is therefore not directly applicable to quasi-stationary distributions where the absorbing set poses challenges by essentially erasing the Markov property. Future work lies in developing a conditional or reflecting large-deviations principle suitable for extending the main results of Freidlin-Wentzell theory (Freidlin and Wentzell, 2012).

Indeed, under the assumptions made in this manuscript we conjecture the natural extensions of Theorem 4.1, 4.2, 4.3 of Freidlin and Wentzell (2012) to hold for quasi-stationary

measures ν_ϵ. To be more specific, define the set of $\{i\}$-graphs over the set of attractors $\{M_1, \ldots, M_m\} \simeq \{1, \ldots, m\}$, denoted $G\{i\}$, to be graphs consisting of arrows $m \to n$ with $m \in \{1, \ldots, m\}\backslash\{i\}$, $n \in \{1, \ldots, m\}$ such that every point $m \in \{1, \ldots, m\}\backslash\{i\}$ is the initial point of exactly one arrow, there are no closed loops and from any point $m \in \{1, \ldots, m\}\backslash\{i\}$ there is a path to $\{i\}$. Define the quantities

$$W(M_i) = \min_{g \in G\{i\}} \sum_{(m \to n) \in g} \tilde{V}(M_n, M_m),$$

where we introduce the so called *quasi-potential*

$$\tilde{V}(M_i, M_j) = \inf_z \left\{ I(x, T, z) \colon x \in M_i, z(T) \in M_j, z(t) \in D_1 \backslash \bigcup_{s \neq i,j} M_s, 0 < t < T \right\},$$

measuring how difficult a transition from the ith Morse set to the jth Morse set is (when not passing through any others). Set $\tilde{V}(M_i, M_j) = \infty$ if no such path exists. Then for any $\gamma > 0$ there exists a $\delta > 0$ such that

$$\exp(-\epsilon(W(M_i) - \min_i W(M_i) + \gamma)) \leq \nu_\epsilon(N^\delta(M_i)) \leq \exp(-\epsilon(W(M_i) - \min_i W(M_i) - \gamma))$$

In particular, as $\epsilon \to 0$ the QSDs ν_ϵ converges weakly to a measure ν whose support is concentrated in a small neighborhood of the union of M_i for which $\min_i W(M_i)$ is attained. In the generic case, there will be just a single such M_i. Similarly, for $x \in D$ define

$$W(x) = \min_{M_1, \ldots, M_m} (W(M_i) + V(M_i, x)).$$

For any sufficiently small neighborhood $N^\delta(x)$ of x there exists $\epsilon_0 > 0$ such that for $\epsilon \leq \epsilon_0$

$$\exp(-\epsilon(W(x) - \min_i W(M_i) + \gamma)) \leq \nu_\epsilon(N^\delta(x))) \leq \exp(-\epsilon(W(x) - \min_i W(M_i) - \gamma))$$

Thus we should see exponential decay away from the support of the limiting distribution ν. The numerical examples considered here supports this conjecture.

Birth-death processes. In the small subclass of reaction networks, that can be represented as birth-death processes, the problem of finding a large deviation principle has been examined (Chan, 1998). However, the extension to Freidlin-Wentzell theory is still open even for this case.

To illustrate, consider once again the logistic network, which in the stochastic case is modeled by a birth-death process. In the deterministic case, there is a stable equilibrium at

$$x^* = \frac{\alpha_2 - \alpha_1}{\alpha_3},$$

provided $\alpha_2 > \alpha_1$, with domain of attraction $(0, \infty)$. In this simple case, one may explicitly find the quasi-potential. Indeed, letting $\lambda(x)$ denote the birth rates and $\mu(x)$ the death rates,

then one has (Pakdaman et al., 2010)

$$\tilde{V}(x^*, x) = \int_{x^*}^{x} \log \frac{\mu(z)}{\lambda(z)}\, dz = \int_{x^*}^{x} \log \frac{\alpha_1 z + \alpha_3 z^2}{\alpha_2 z}\, dz.$$

We note that in the setting of stationary distributions, this function has previously been characterized as the limit of the *non-equilibrium potential* under the classical scaling (Anderson et al., 2015). This correspondence was also demonstrated (but not proved) for models in the quasi-stationary regime, suggesting yet again that an extension should be possible. The classical Freidlin-Wentzell theory applied to the domain $(0, \infty)$ now yields

$$\lim_{\epsilon \to 0} \epsilon \log \mathbb{E}_x(\tau_A) = \frac{\alpha_1 \log(\alpha_1/\alpha_2) - \alpha_1 + \alpha_2}{\alpha_3}.$$

In accordance with Theorem 4.3.2, the expected time of extinction grows exponentially with ϵ. However, here we obtain an analytic expression for the rate.

For the problem of determining which attractor ν is supported on in Example 4.6.3, it is tempting to simply calculate and compare the quasi-potentials $W(18) = \tilde{V}(5, 18)$ and $W(5) = \tilde{V}(18, 5)$ and conclude that the support is on the attractor with the smallest quantity. However, this naive approach may contain a small error due to the lack of a large deviation property all the way to 0. Furthermore, the network considered in Example 4.6.3 is not birth-death yielding a much harder variational problem to find the quasi-potential. For more complex reaction networks, an analytic solution is generally not be possible.

4.8 Appendix

In this appendix, we take the opportunity to explain some of the more technical background material. In particular, we review some of the results on chain recurrence. The purpose is to keep the paper self contained, as equivalent results and their proofs may be found in Kifer (1988); Faure and Schreiber (2014); Marmet (2013); Colonius and Kliemann (2014).

4.8.1 Absorption Preserving Chain Recurrence

To prove the main result, we introduced the following notion of chain recurrence, which we repeat here for ease of reference.

Definition 4.8.1. Let $\{\varphi_t\}_{t\in\mathbb{R}}$ be a flow given by an ordinary differential equation on $\mathcal{X} \subset \mathbb{R}^d_+$ for which M_0 is an invariant set. Given $\delta > 0$, a piecewise continuous path from x to y,

$$x = x_1, \{\varphi_t(x_1) : t \in [0, t_1]\}, \{\varphi_t(x_2), t \in [0, t_2]\} \dots \{\varphi_t(x_k), t \in [0, t_k]\}, x_{k+1} = y$$

which is uniquely defined by the sequence of points $x_1, \dots, x_{k+1}$ and times $t_1, \dots, t_k$ such that the following holds

$$\begin{aligned} &d(\varphi_{t_j}(x_j), x_{j+1}) < \delta, \quad \forall j = 1, \dots, k \\ &t_i \geq T, \quad \forall i = 1, \dots, k, \\ &x_j \in M_0 \Rightarrow x_{j+1} \in M_0, \quad \forall j = 1, \dots, k, \end{aligned}$$

is called an absorption preserving (δ, T)-pseudoorbit, or an ap (δ, T)-pseudoorbit for short. If there exists an ap (δ, T)-pseudoorbit from x to y, we will write $x <_{ap,\delta,T} y$. If $x <_{ap,\delta,T} y$ holds for every $\delta > 0$ and $T > 0$, we will write $x <_{ap} y$. If $x <_{ap} y$ and $y <_{ap} x$ we write $x \sim_{ap} y$. If $x \sim_{ap} x$ then x is an ap-chain recurrent point. We denote by $\mathcal{R}_{ap}$ the set of ap-chain recurrent points. The relation $\sim_{ap}$ restricted to this set defines an equivalence relation and the equivalence classes, $[x]_{ap}$ with $x \in \mathcal{R}_{ap}$ are called ap-basic classes. We write $[x]_{ap} <_{ap} [y]_{ap}$ if $x <_{ap} y$. A maximal ap-basic class is called an ap-quasiattractor.

Lemma 4.8.2. *Let $x \in \mathcal{R}_{ap}$. Then $\overline{[x]}_{ap} \subseteq D_0 \cup [x]_{ap}$. Moreover, $[x]_{ap}$ is positively φ-invariant (actually φ-invariant).*

Proof. Let $y \in \overline{[x]}_{ap}$. Then there exists a sequence $\{y_k\}$ in $[x]_{ap}$ convergning to y. Thus, for $\delta > 0$, there exists k large enough such that $d(y_k, y) < \delta$. Let $T > 0$. As $y_k \in [x]_{ap}$, there exists an ap (δ, T)-pseudoorbit $(x_1, t_1), \dots, (x_n, t_n), y_k$ linking x to y_k. As we have

$$d(\varphi_{t_n}(x_n), y) \leq d(\varphi_{t_n}(x_n), y_k) + d(y_k, y) < 2\delta$$

The sequence $(x_1, t_1), \dots, (x_n, t_n), y$ is an ap $(2\delta, T)$-pseudoorbit linking x to y. Thus $x <_{ap} y$. On the other hand, assume $y \notin M_0$ and consider an ap (δ, T)-pseudoorbit from y_k to x. Then by continuity of φ_{t_1} it follows that for k large enough

$$d(\varphi_{t_1}(y), x_2) \leq d(\varphi_{t_1}(y), \varphi_{t_1}(y_k)) + d(\varphi_{t_1}(y_k), x_2) < 2\delta$$

hence $y <_{ap} x$. Therefore $y \in [x]_{ap}$ and we conclude that $\overline{[x]}_{ap} \subseteq D_0 \cup [x]_{ap}$ as desired.

Let $T, T', \epsilon > 0$. We may without loss of generality assume that $T > T'$. As $\varphi_t(D)$ is bounded for some $t > 0$ sufficiently large by Assumption 3, we may assume that the map $\varphi_{T'}$ is uniformly continuous. Thus, we may find $0 < \delta < \epsilon$ such that $d(x, y) < \delta$ implies $d(\varphi_{T'}(x), \varphi_{T'}(y)) < \epsilon$. We know that there exists an ap (δ, T)-pseudoorbit $(x, t_1), \ldots, (x_n, t_n), x$ linking x to x. Hence $d(\varphi_{t_n+T'}(x_n), \varphi_{T'}(x)) < \epsilon$ and $(x, t_1), \ldots, (x_n, t_n + T'), \varphi_{T'}(x)$ is an ap (ϵ, T)-pseudoorbit linking x to $\varphi_{T'}(x)$.

For the converse, consider an ap (δ, T)-pseudoorbit given by $(x, t_1), \ldots, (x_n, t_n), x$ linking x to x. Note that, as $t_1, t_2 > T$, we have $t_1 + t_2 - T' > T$. Thus

$$(\varphi_{T'}(x), t_1 + t_2 - T'), (x_3, t_3), \ldots, (x_n, t_n), x$$

is an ap (ϵ, T)-pseudoorbit linking $\varphi_{T'}(x)$ to x, as

$$d(\varphi_{T'+t_1+t_2-T'}(x), x_3) = d(\varphi_{t_1+t_2}(x), x_3) \leq d(\varphi_{t_2}(\varphi_{t_1}(x)), \varphi_{t_2}(x_2)) + d(\varphi_{t_2}(x_2), x_3) \leq \epsilon$$

for δ chosen sufficiently small. Thus $x \sim_{ap} \varphi_{T'}(x)$. As $T' > 0$ was arbitrary, we conclude that $\varphi_t([x]_{ap}) \subseteq [x]_{ap}$ for $t > 0$ as desired. □

Lemma 4.8.3. *For $x \in D$, $\omega(x) \cap \mathcal{R}_{ap} \neq \emptyset$*

Proof. If $x \in D_0$ or $\omega(x) \subset D_1$ then the classical result (Conley, 1978, p. 37) for chain recurrence implies $\omega(x) \subset \mathcal{R}_{ap}$. If $x \in D_1$ and $y \in \omega(x) \cap D_0$, then $\omega(y) \subset \mathcal{R}_{ap}$ and since $\omega(y) \subseteq \omega(x)$ we find $\omega(x) \cap \mathcal{R}_{ap} \supseteq \omega(y) \cap \mathcal{R}_{ap} \neq \emptyset$. □

Lemma 4.8.4. *If $[x]_{ap}$ is maximal, then $x <_{ap} z$ if and only if $z \in [x]_{ap}$. In particular, any ap-quasiattractor $[x]_{ap}$ is closed.*

Proof. Let z be such that $x <_{ap} z$. we must prove that $z <_{ap} x$. By lemma 3.3, $\omega(z) \cap \mathcal{R}_{ap} \neq \emptyset$. Thus there exists $z' \in \omega(z) \cap \mathcal{R}_{ap}$ such that $x <_{ap} z <_{ap} z'$. As $z' \in \mathcal{R}_{ap}$, maximality of $[x]_{ap}$ implies that $z' \in [x]_{ap}$. Thus $z <_{ap} x$ as desired. In particular, it follows from the proof of Lemma 4.8.2 that if $y \in \overline{[x]}_{ap}$, then $x <_{ap} y$ hence from the above argument, $y \in [x]_{ap}$ and thus $[x]_{ap}$ is closed. □

Theorem 4.8.5. Let $[x]_{ap}$ be an isolated ap-quasiattractor in D_1. Then $[x]_{ap}$ is an (irreducible) attractor.

Proof. As $[x]_{ap}$ isolated, there exists an open neighborhood $G \supset [x]_{ap}$ whose closure $\overline{G} \subset D_1$ is compact and disjoint from other ap-basic classes. Further, as $[x]_{ap}$ is an ap-quasiattractor, it follows from Lemma 4.8.4 that $[x]_{ap}$ is closed, and by Lemma 4.8.2 that $\varphi_t([x]_{ap}) \subset [x]_{ap}$. As the family $\{\varphi_t, 0 \leq t \leq T\}$ is equicontinuous on $\overline{G}$, there exists another open neighborhood W of $[x]_{ap}$ such that $\overline{W} \subset G$ and $\varphi_t(\overline{W}) \subset G$ for all $t \in [0, T]$.

We claim that there exists $\delta_0 > 0$, $T_0 > 0$ such that any ap (δ_0, T_0)-pseudoorbit starting in a point $y \in [x]$ is completely contained in W. Suppose for contradiction this was not true. Then there would exist sequences of numbers $\delta_n \to 0$, $T_n \to \infty$ and a sequence of ap

(δ_n, T_n)-pseudoorbits starting at $y_n \in [x]$ and end at points $z_n \in \overline{G}\backslash W$. As both sets $[x]$ and $\overline{G}\backslash W$ are compact, we may choose converging subsequences

$$y_{n_i} \to y \in [x], \qquad z_{n_i} \to z \in \overline{G}\backslash W,$$

as $i \to \infty$. Thus, for any $\delta > 0$, $T > 0$ there is an ap (δ, T)-pseudoorbit from y to z hence $y <_{ap} z$. However, as $[x]$ is an ap quasiattractor, we have from Lemma 4.8.4 that $z \in [x]$, which is impossible as $[x] \cap \overline{G}\backslash W = \emptyset$.

We now prove that

$$\varphi_t([x]) = [x] \qquad t \geq 0. \tag{4.8.6}$$

It suffices to get this for $t = 1$ since this by Lemma 4.8.2 would imply that $[x] = \varphi_1([x]) \subset \varphi_t([x])$ for $0 \leq t \leq 1$. Applying this result repeatedly, we would obtain $[x] \subset \varphi_t([x])$ for all $t \geq 0$ as required.

Thus for any $y \in [x]$ we must find $z \in [x]$ such that $\varphi_1(z) = y$. If $[x]$ consists of a fixed point then we are done. Otherwise, we may choose $\tilde{y} \in [x]$, $\tilde{y} \neq y$ such that for any sequences $\delta_n \to 0$, $T_n \to \infty$ there exists a sequence of ap (δ_n, T_n)-pseudoorbits beginning at $\tilde{y}$ and ending at y which contains points $v_n, w_n \in W$ satisfying

$$d(\varphi_1(v_n), w_n) < \delta_n, \qquad \varphi_{t_{k(n)}}(w_n) = y,$$

for some $t_{k(n)} \geq T_n$. By compactness, we may choose a subsequence n_i such that

$$t_{n_i} \to t_0, \qquad v_{n_i} \to v, \qquad w_{n_i} \to w,$$

thus we have $\tilde{y} <_{ap} v <_{ap} w$ which, as $[x]$ is an ap-quasiattractor, implies that $v, w \in [x]$. As $\varphi_{t_0}(w) = y$ and $\varphi_1(v) = w$ we conclude that $\varphi_{1+t_0}(v) = y$. Thus, if we set

$$z := \varphi_{t_0}(v)$$

we see that $z \in [x]$ by Lemma 4.8.2 as $v \in [x]$, and $\varphi_1(z) = \varphi_{1+t_0}(v) = y$ as desired.

Finally, we show that $[x]$ is an attractor. By the above, we know that any ap (δ_0, T_0)-pseudoorbit starting in $[x]$ remains in W. Thus, we infer that

$$z \in U_0 := \{y\colon d(y, [x]) < \delta_0\} \Rightarrow \varphi_t(z) \in W, \tag{4.8.7}$$

for all $t \geq 0$. In the same way, we see that there exists $\delta_1 > 0$, $\delta_1 < \delta_0$ such that

$$z \in U_1 := \{y\colon d(y, [x]) < \delta_1\} \Rightarrow \varphi_t(z) \in U_0,$$

for all $t \geq 0$. Take an arbitrary open set $V \supset [x]$. We must show that $\varphi_t(U_1) \subset V$ provided $t \geq t(V)$ is sufficiently large. Pick an open set $V_1 \supset [x]$ such that $\overline{V}_1 \subset V$. We claim that

$$t(V) := \inf\{t\colon \varphi_t(U_0) \subset V_1\} < \infty.$$

Indeed, otherwise there would exist a sequence $z_n \in U_0$ and numbers $t_n \to \infty$ such that

$\varphi_t(z_n) \notin V_1$ for all $t \in [0, t_n]$. Choose a converging subsequence $z_{n_i} \to z \in \overline{U}_0$. Then $\varphi_t(z) \notin V_1$ for all $t \geq 0$. However, as $\varphi_t(U_0) \subset W$ by (4.8.7) we have $\varphi_t(\overline{U}_0) \subset \overline{W}$ hence $\varphi_t(z) \in \overline{W}$ for all $t \geq 0$. In other words the entire trajectory $\{\varphi_t(z), t \geq 0\}$ is contained in the compact set $\overline{W} \backslash V_1$ and thus admits a limit point and therefore an ap basic class. However, this contradicts the fact that $\overline{W} \backslash V_1$ is disjoint from any ap basic class. Finally, we see that for $t \geq t(V)$,

$$\varphi_t(U_1) = \varphi_{t(V)}\varphi_{t-t(V)}(U_1) \subset \varphi_{t(V)}(U_0) \subset V_1 \subset V.$$

From this and (4.8.6) it follows that $\bigcap_{t\geq 0} \varphi_t(U_0) = [x]$, thus $[x]$ ia an attractor as desired. □

4.8.2 Finiteness of the ap-basic Classes

In this section we prove that if the number of classes is finite, then they are isolated, and we may use the previous subsection to conclude that ap-quasiattractors are attractors.

Assumption 5. The number of ap-basic classes in D_1, $\{K_i\}_{i=1,\dots,v}$, is finite. Moreover, they are closed sets and $\{K_i\}_{i=1,\dots,\ell}$ with $\ell \geq 1$ are ap-quasiattractors while $\{K_i\}_{i=\ell+1,\dots v}$ are not.

The following lemma shows that for δ small and T large, the ap (δ, T)-pseudoorbits respect the partial order.

Lemma 4.8.8. *For every $\theta > 0$ small enough, there exists $\delta_\theta \in (0, \theta)$ and $T_\theta \in (1/\theta, \infty)$ such that, if there exists an ap $(\delta_\theta, T_\theta)$-pseudoorbit $\xi_0, \dots, \xi_n$ satisfying*

$$d(\xi_0, K_i) < \delta_\theta, \qquad d(\xi_n, K_{i'}) < \delta_\theta, \qquad d(\xi_j, K_i) > \theta,$$

for some $i, i' \in \{1, \dots, v\}$ and $j \in \{1, \dots, n\}$, then $i \neq i'$ and $K_i <_{ap} K_{i'}$

Proof. Suppose that for every $\delta, T > 0$, there exists an ap (δ, T)-pseudoorbit $\xi_0, \dots \xi_n$ such that

$$d(\xi_0, K_i) < \delta, \qquad d(\xi_n, K_{i'}) < \delta,$$

then we may construct an ap (δ, T)-pseudoorbit from K_i to $K_{i'}$, hence $K_i <_{ap} K_{i'}$.

Hence, if $K_i \not<_{ap} K_{i'}$ there exists $\tilde{\delta} \geq 0$ and $\tilde{T} \in (0, \infty]$ such that an ap (δ, T)-pseudoorbit verifying

$$d(\xi_0, K_i) < \delta, \qquad d(\xi_n, K_{i'}) < \delta,$$

may only exist if $\delta > \tilde{\delta}$ or $T < \tilde{T}$.

Suppose now that $i = i'$. We will show that there exists $\hat{\delta}$ and $\hat{T}$ such that every ap (δ, T)-pseudoorbit $\xi_0, \dots, \xi_n$ verifying either $\delta < \hat{\delta}$ or $T > \hat{T}$ and

$$d(\xi_0, K_i) < \delta, \qquad d(\xi_n, K_i) < \delta,$$

doesn't contain any point at a distance greater than θ from K_1. Suppose this assertion is false. Then we have real sequences $\delta_l \to 0$ and $T_l \to \infty$ and a sequence of ap (δ_l, T_l)-pseudoorbits $\xi_0^l, \ldots, \xi_{n_l}^l$ satisfying

$$d(\xi_0^l, K_i) < \delta_\theta, \qquad d(\xi_{n_l}^l, K_i) < \delta_\theta, \qquad d(\xi_{j_l}^l, K_i) > \theta.$$

Choosing θ sufficiently small, we have $\overline{N^\theta(K_i)} \subset D_1$, which is a compact set. Without loss of generality, we may then assume that

$$\lim_{l\to\infty} \xi_0^l = x \in K_i, \qquad \lim_{l\to\infty} \xi_{n_l}^l = z \in K_i, \qquad \lim_{l\to\infty} \xi_{j_l}^l = y \in K \backslash N^\theta(K_i)$$

where $K \subset M_1$ is a compact set. We see that $x <_{ap} y <_{ap} z$ thus $y \in K_i$ which contradicts $y \in K\backslash N^\theta(K_i)$. □

Lemma 4.8.9. *For every $\delta_0 > 0$, there exists $T_0 > 0$ such that every ap (δ_0, T_0)-pseudoorbit intersects $N^{\delta_0}(\mathcal{R}_{ap})$*

Proof. Let $x \in D$, $\gamma > 0$ and define

$$T^\gamma(x) = \inf\{t \geq 0 \colon \varphi_t(x) \in N^\gamma(\mathcal{R}_{ap})\}.$$

By Lemma 4.8.3 we have $\omega(x) \cap \mathcal{R}_{ap} \neq \emptyset$ hence $T^\gamma(x) < \infty$. For $\alpha > 0$ denote the level sets of T^γ by

$$L_\alpha^\gamma = \{x \in M \colon T^\gamma(x) \geq \alpha\}.$$

We wish to show that L_α^γ is closed. Let $(x_n)_{n\in\mathbb{N}}$ be a sequence of elements in L_α^γ converging to y. By continuity of φ_t we see that for every $t > 0$, $\lim_{n\to\infty} \varphi_t(x_n) = \varphi_t(y)$. If $t < \alpha$ then $\varphi_t(x_n) \in (N^\gamma(\mathcal{R}_{ap}))^c$ which is closed. In particular, $\varphi_t(y) \notin N^\gamma(\mathcal{R}_{ap})$ for $t < \alpha$ hence $y \in L_\alpha^\gamma$. Thus L_α^γ is closed. By definition, it follows that T^γ is upper-semicontinuous. Thus for any $t > 0$,

$$T^\gamma := \sup_{x\in M} T^\gamma(x) \leq \max_{y\in\overline{\varphi_t(M)}} T^\gamma(x) + t < \infty.$$

Taking $T_0 > T^{\delta_0}$ yields the desired. □

We may now collect the results into the central Corollary used in the paper.

Corollary 4.8.10. Given $\delta, T > 0$ there exists a family $\{V_i\}_{i=1,\ldots,v}$ of isolating open neighborhoods of the ap-basic classes $\{K_i\}_{i=1,\ldots,v}$, and positive constants δ_1 and T_1 such that

1. $\overline{N^{\delta_1}(K_i)} \subset V_i$ for $1 \leq i \leq v$;

2. any ap (δ_1, T_1)-pseudoorbit starting in V_i remains in V_i for $i = 1, \ldots, \ell.$;

3. if there exists an ap (δ_1, T_1)-pseudoorbit $(x, t_1), \ldots, (x_n, t_n), y$ with $x \in N^{\delta_1}(K_i)$ and $y \in N^{\delta_1}(K_{i'})$ such that

$$\varphi_s(x_j) \notin V_i, \qquad \text{for some } j \in \{2, \ldots, n-1\}, s \in [0, t_j],$$

then $i \neq i'$ and $K_i <_{ap} K_{i'}$.

4. Every ap (δ_1, T_1)-pseudoorbit intersects $N^{\delta}(\mathcal{R}_{as})$.

Proof. Choose $\theta \in (0, \delta)$ sufficiently small so that Lemma 4.8.8 holds and let δ_θ be given by this lemma. From Lemma 4.8.2 we see that the ap basic classes K_i are compact for $i = 1, \ldots, v$ hence we may find neighborhoods V_i of K_i such that $\overline{N^\theta(K_i)} \subset V_i$. Further, by Theorem 4.8.5, K_i is an attractor for $i = 1, \ldots, \ell$ thus we may choose the neighborhoods such that $\varphi_t(\overline{V}_i) \subset V_i$ for $i = 1, \ldots, \ell$, $t > 0$. Now choose $\delta_1 \in (0, \delta_\theta) \subset (0, \theta)$ such that δ_1 is less than the δ given by Lemma 4.8.8. □

Corollary 4.8.11. Under Assumption 5, there are only finitely many positive attractors.

Proof. By Corollary 4.8.10, the ap-basic classes are isolated thus by Theorem 4.8.5 the ap-quasiattractors are attractors of which there are finitely many. □

We include a small lemma on conditional probabilities.

Lemma 4.8.12. *For three events A_1, A_2, A_3 one has*

$$\mathbb{P}(A_1) \leq \mathbb{P}(A_3^c | A_2) + \mathbb{P}(A_1 | A_2 \cap A_3) + \mathbb{P}(A_2^c)$$

Proof. We simply calculate, using conditional probabilities,

$$\begin{aligned}
\mathbb{P}(A_1) &= \mathbb{P}(A_1 \cap A_2) + \mathbb{P}(A_1 \cap A_2^c) \\
&= \mathbb{P}(A_1 \cap A_2 \cap A_3) + \mathbb{P}(A_1 \cap A_2 \cap A_3^c) + \mathbb{P}(A_1 \cap A_2^c) \\
&\leq \mathbb{P}(A_1 \mid A_2 \cap A_3)\mathbb{P}(A_2 \cap A_3) + \mathbb{P}(A_3^c \cap A_2) + \mathbb{P}(A_2^c) \\
&\leq \mathbb{P}(A_1 \mid A_2 \cap A_3) + \mathbb{P}(A_3^c | A_2) + \mathbb{P}(A_2^c)
\end{aligned}$$

as desired. □

5

A Description of Quasi-stationary Distributions in Reaction Networks using a Slow Manifold Reduction

Mads Christian Hansen and Carsten Wiuf

University of Copenhagen

Abstract. Stochastic reaction networks often allow a separation of timescales in which case a slow manifold reduction, in the space of total probabilities, approximates the global flow well. When this is the case, we exploit the fact that the quasi-stationary distribution lies in the center manifold to provide an explicit iterative method for describing this distribution. Further, we show conditions under which the stationary distribution of a certain reduced network provide a good approximation for the true quasi-stationary distribution of the full network. As a consequence, knowing the stationary distribution of a reaction network becomes relevant as a first step in describing the quasi-stationary distribution. For the case of one-species reaction network, we provide a method for determining such limit distributions.

Keywords: *Center manifold, quasi-stationary distribution, stationary distribution, one-species reaction network.*

Publication details: Unsubmitted

5.1 Introduction

The term *quasi-stationarity* is used for the phenomenon that although a system is certain to die out eventually, it appears to be stationary over any reasonable time scale (van Doorn and Pollett, 2013). The quasi-stationary distributions, which are stationary distributions when we condition on the event that the process has not reached extinction, captures this variability (Nåsell, 2011). However, although a quasi-stationary distribution exists for a given reaction network, this does not imply that quasi-stationarity is actually observed. Only when the expected time to extinction is much longer than the expected time to reach the quasi-stationary distribution, would the latter be of interest (Qian, 2011). In other words, one needs to have a certain separation of time-scales (Kuehn, 2015).

The class of stochastic dynamical systems known as reaction networks, poses an intuitive framework for modeling a range of natural phenomena where not only the entities under consideration but also their interactions are of transformational form (Veloz and Razeto-Barry, 2017). The set of entities (species) may in principle be of any nature, and specifying not just which ones interact (stoichiometry and reactions) but also quantifying how frequent they interact (kinetics), we obtain the dynamical system of a reaction network. Examples abound in biochemistry, where the language originated. Furthermore, at this cellular level, the copy-numbers of interacting entities may be low (Elowitz et al., 2002) and stochastic effects resulting from these small numbers may be physiologically significant (Cook et al., 2009). In particular, extinction events occur naturally in this domain and the quasi-stationary distribution, rather than the usual stationary distribution, becomes the object of interest. We underline that low copy-numbers occur in many other scientific disciplines where reaction networks may be used, including social (Dittrich and Winter, 2008), economical (Veloz et al., 2014), ecological (Shakil et al., 2015) or epidemiological (Nåsell, 2011) contexts.

Sufficient conditions for the existence and uniqueness of a quasi-stationary distribution in stochastic reaction networks exist (Hansen and Wiuf, 2018a). However, the problem of determining its form analytically is in general very hard, and one has to resort to either approximations or numerical methods. Furthermore, as the state-space of the stochastic process associated to a reaction network is in general countably infinite, even calculating the quasi-stationary distribution numerically poses several challenges (van Doorn and Pollett, 2013). Here, we present a formal way of obtaining the quasi-stationary distribution specifically developed for reaction networks. Indeed, by a slight variation of the approach in Pollett and Roberts (1990), we exploit a connection between a center manifold and the quasi-stationary distribution. By splitting the reaction network into two suitably defined sub-networks, one may apply existing knowledge about the ordinary stationary distribution for such systems, to approximate the quasi-stationary distribution. In fact, one may consider the stationary distribution on a subnetwork as a first approximation to the quasi-stationary distribution on the full network. Hereby a more informative characterization is obtained, compared to a purely numerical approach.

The question of characterizing the stationary distribution therefore becomes of interest. In general, the form of these distributions are only known for mono-molecular or complex balanced networks (Jahnke and Huisinga, 2007; Anderson et al., 2010), although sufficient conditions for their existence are satisfied for a large class of networks (Gupta et al., 2014). It is still unknown precisely how large the class of networks having a stationary distribution is. However, it is believed that all weakly reversible networks is contained herein. In the case of single-species reaction networks, we provide an iterative analytic procedure for calculating the stationary distribution.

To set notation, consider a right-continuous time-homogenous Markov process $(X_t \colon t \geq 0)$ (Rogers and Williams, 2000), that evolves in a domain $D \subseteq \mathbb{R}^d$, wherein there is a set of absorbing states, a "trap", $A \subset D$. The process is absorbed, also referred to as being killed, when it hits the set of absorbing states, implying $X_t \in A$ for all $t \geq \tau_A$, where $\tau_A = \inf\{t \geq 0 : X_t \in A\}$ is the hitting time of A. As we are interested in the process before reaching A, there is no loss of generality in assuming $X_t = X_{t \wedge \tau_A}$. We refer to the complement,

$$E := D \backslash A,$$

as the set of endorsed states. For any probability distribution, μ, on E, we let $\mathbb{P}_\mu$ and $\mathbb{E}_\mu$ be the probability and expectation respectively, associated with the process $(X_t \colon t \geq 0)$, initially distributed with respect to μ. For any $x \in E$, we let $\mathbb{P}_x = \mathbb{P}_{\delta_x}$ and $\mathbb{E}_x = \mathbb{E}_{\delta_x}$. Under suitable conditions, the process hits the absorbing set almost surely (a.s.), that is $\mathbb{P}_x(\tau_A < \infty) = 1$ for all $x \in E$, and we investigate the behavior of the process before being absorbed (Collet et al., 2013).

Definition 5.1.1. A probability measure ν on E is called a quasi-stationary distribution (QSD) for the process $(X_t \colon t \geq 0)$ absorbed at A, if for every measurable set $B \subseteq E$

$$\mathbb{P}_\nu(X_t \in B \,|\, t < \tau_A) = \nu(B), \qquad t \geq 0,$$

or equivalently, if there exists a probability measure μ on E such that

$$\lim_{t \to \infty} \mathbb{P}_\mu(X_t \in B \,|\, t < \tau_A) = \nu(B),$$

in which case we also say that ν is a quasi-limiting distribution.

We refer to Méléard and Villemonais (2012) for a proof of the equivalence of quasi-limiting and quasi-stationary distributions.

The manuscsript is organized as follows. In section 5.2 we provide some background on stochastic reaction network theory and introduce the decomposition of a reaction network with a non-empty absorbing set into the inferred sub-networks, yielding a corresponding decomposition of the rate-matrix. This allows viewing the Kolmogorov forward equation as a linear perturbation. Section 5.3 provides a dynamical systems view on the quasi-stationary distribution. In particular, we apply the center manifold theorem to obtain a geometrical perspective on the QSD. This is in turn exploited to provide an iterative scheme for calcu-

lating the QSD. Section 5.4 is devoted to finding the stationary distribution of one-species reaction networks, where an iterative analytic procedure is provided. Finally, section 5.5 contains a discussion of the assumptions needed to conclude a proper separation of time-scales and thereby of the validity of the center manifold approach.

5.2 Reaction Network Setup

Denote the real numbers by $\mathbb{R}$, the integers by $\mathbb{Z}$, the natural numbers by $\mathbb{N} = \{1, 2, \dots\}$ and the nonnegative integers by $\mathbb{N}_0 = \{0, 1, 2, \dots\}$. Further, for any set, B, let $|B|$ denote its cardinality and denote by $\mathbb{1}_B \colon D \to \{0, 1\}$ the indicator function of a subset $B \subseteq D$.

A reaction network is a triple $\mathcal{N} = (\mathcal{S}, \mathcal{C}, \mathcal{R})$, where $\mathcal{S}$ is a finite ordered set of species[1], $\mathcal{C}$ is a finite set of complexes, consisting of linear combinations over $\mathbb{N}_0$ of the species, and $\mathcal{R} \subset \mathcal{C} \times \mathcal{C}$ is an irreflexive relation on $\mathcal{C}$, referred to as the set of reactions (Anderson and Kurtz, 2015; Feinberg, 1979; Gunawardena, 2003). Furthermore, $\mathcal{R}$ is assumed to be ordered.

We define the dimension of the reaction network, $d = |\mathcal{S}|$. Any species $S_i \in \mathcal{S}$ can be identified with the unit vector $e_i \in \mathbb{N}_0^d$, thus any complex $y \in \mathcal{C}$ can be identified with a vector in $\mathbb{N}_0^d$. It is customary to denote an element $(y_k, y_k') \in \mathcal{R}$ by $y_k \to y_k' \in \mathcal{R}$ in which case we refer to y_k as the source complex and to y_k' as the product complex of reaction k. We may thus write $\mathcal{R} = \{y_k \to y_k' \colon k = 1, \dots, r\}$. Employing a standard, although slight abuse of, notation, we identify $\mathcal{S} = \{S_1, \dots, S_d\}$ with the set $\{1, \dots, d\}$ and $\mathcal{R}$ with $\{1, \dots, r\}$. We write the k'th reaction with the notation

$$\sum_{i \in \mathcal{S}} y_{ki} S_i \to \sum_{i \in \mathcal{S}} y_{ki}' S_i,$$

where $y_{ki} = (y_k)_i$ and $y_{ki}' = (y_k')_i$ are the stoichiometric coefficients associated with the source and product complexes of reaction k, respectively. Define the reaction vectors $\xi_k = y_k' - y_k$ and the stoichiometric matrix

$$\Xi = (\xi_1 \, \xi_2 \, \dots \, \xi_r) \in \mathbb{N}_0^{d \times r}.$$

The order of reaction k is the sum of the stoichiometric coefficients of the source complex, $\sum_{i \in \mathcal{S}} y_{ki}$. Finally, we define the maximum of a vector over the set $\mathcal{R}$, $x = \max_{k \in \mathcal{R}} y_k$, as the entry-wise maximum, $x_i = \max_{k \in \mathcal{R}} y_{ki}$.

A set of reactions $\mathcal{R}$ induces a set of complexes and a set of species, namely the complexes and species that appear in the reactions. We will assume that a reaction network is always given in this way by $\mathcal{R}$, and one may then completely describe a reaction network in terms of its reaction graph, whose nodes are the complexes and whose directed edges are the reactions. This concise description will be employed in the rest of the paper. To avoid trivialities, we assume $\mathcal{R} \neq \emptyset$.

[1] The terminology "species" is standard. One may equally think of them as general entities or agents.

For each reaction we specify an intensity function $\lambda_k \colon \mathbb{N}_0^d \to [0, \infty)$, $k \in \mathcal{R}$, which satisfies the stoichiometric admissibility condition:

$$\lambda_k(x) > 0 \quad \Leftrightarrow \quad x \geq y_k,$$

where we use the usual vector inequality notation; $x \geq y$ if $x_i \geq y_i$ for all $i \in \mathcal{S}$. Thus, reactions are only allowed to take place whenever the copy-numbers of each species in the current state is at least as great as those of the corresponding source complex. A widely used example is stochastic mass action kinetics given by

$$\lambda_k(x) = \alpha_k \prod_{i=1}^{d} y_{ki}! \binom{x}{y_k} = \alpha_k \prod_{i=1}^{d} \frac{x_i!}{(x_i - y_{ki})!},$$

for some reaction rate constants $\alpha_k > 0$ (Anderson and Kurtz, 2015). The idea is that the rate is proportional to the number of distinct subsets of the molecules present that can form the input of the reaction. It reflects the assumption that the system is well-stirred (Anderson and Kurtz, 2015). Other examples include power law kinetics or generalized mass action kinetics (Anderson, 2008; Horn and Jackson, 1972; Müller and Regensburger, 2012). A particular choice of such rate functions constitute a stochastic kinetics $\lambda = (\lambda_1, \ldots, \lambda_r)$ for the reaction network $\mathcal{N}$, and the pair $(\mathcal{N}, \lambda)$ is referred to as a stochastic reaction system, or simply a reaction network with kinetics λ.

We may then specify the stochastic process $(X_t \colon t \geq 0)$ on the state space $D := \mathbb{N}_0^d$ related to the reaction system $(\mathcal{N}, \lambda)$. Let X_t be the vector in $\mathbb{N}_0^d$ whose entries are the species counts at time t. If reaction $y_k \to y_k'$ occurs at time t, then the new state is $X_t = X_{t-} + y_k' - y_k = X_{t-} + \xi_k$, where X_{t-} denotes the previous state. The stochastic process then follows,

$$X_t = X_0 + \sum_{k \in \mathcal{R}} Y_k \left(\int_0^t \lambda_k(X_s)\, ds \right) \xi_k, \tag{5.2.1}$$

where Y_k are independent and identically distributed unit-rate Poisson processes (Anderson and Kurtz, 2015; Ethier and Kurtz, 1986; Norris, 2009). This stochastic equation is referred to as a random time change representation. We assume throughout the paper that the process is non-explosive, so that the process is well defined.

Decomposing the Reaction Network. Let a reaction network with mass-action kinetics $(\mathcal{N}, \lambda)$ be given. Following (Hansen and Wiuf, 2018a) we introduce the region of large copy numbers, where all reactions may take place,

$$R = \{x \in \mathbb{N}_0^d \,|\, \lambda_k(x) > 0 \,\forall\, k \in \mathcal{R}\}.$$

Letting $D_E = \{x \in D \,|\, x \mapsto R\}$, we may decompose the state space into a disjoint union $D = D_E \sqcup D_A$. Here, we shall assume that D_E is irreducible thus there is a single endorsed

set, E, with corresponding absorbing set A and we may write the state space

$$D = A \sqcup E \subseteq \mathbb{N}_0^d.$$

We shall further assume $A \neq \emptyset$ and that E is a countably infinite set of transient states. We define the set of reactions through which the stochastic process may reach the absorbing set,

$$\mathcal{R}_A = \{k \in \mathcal{R} \,|\, \exists x \in E \colon x + \xi_k \in A\},$$

and define the following sub-systems.

Definition 5.2.2. Given a reaction network $\mathcal{N} = (\mathcal{S}, \mathcal{C}, \mathcal{R})$, the *sub-network* with set of reactions $\mathcal{R}' \subseteq \mathcal{R}$ is the reaction network

$$\mathcal{N}_{\mathcal{R}'} = (\mathcal{S}', \mathcal{C}', \mathcal{R}')$$

where $\mathcal{C}' = \{y_k, y_k' \in \mathcal{C} \colon k \in \mathcal{R}'\}$, $\mathcal{S}' = \{S \in \mathcal{S} \colon S \in \operatorname{supp}(\mathcal{C}')\}$. The *endorsed and absorbing networks* are given, respectively, by

$$\mathcal{N}_E = \mathcal{N}_{\mathcal{R}\setminus\mathcal{R}_A} \qquad \mathcal{N}_A = \mathcal{N}_{\mathcal{R}_A}.$$

Without loss of generality, we may identify $D = \{0\} \sqcup \mathbb{N}$ by collapsing A to a single state $\{0\}$. The jump-rate matrix is then given by

$$\tilde{Q} = \begin{pmatrix} 0 & 0' \\ a & Q \end{pmatrix},$$

where by assumption $a \neq 0$, such that absorption is possible. Now, for all $k \in \mathcal{R}$, one may write the reaction rates $\alpha_k = \beta_k \alpha_{k'}$ for some $k' \in \mathcal{R}_A$ and $\beta_k \in (0, \infty)$. Set $\epsilon = \alpha_{k'}$ and let $\tilde{Q}_E$ and $\epsilon\tilde{Q}_A$ denote the rate matrices for the endorsed and absorbing reaction network respectively. Then it is always possible to decompose the full rate-matrix $\tilde{Q}$ into the sum of two conservative rate-matrices

$$\tilde{Q} = \tilde{Q}_E + \epsilon\tilde{Q}_A, \quad \text{where} \quad \tilde{Q}_E = \begin{pmatrix} 0 & 0' \\ 0 & Q_E \end{pmatrix}, \quad \epsilon\tilde{Q}_A = \epsilon \begin{pmatrix} 0 & 0' \\ b & Q_A \end{pmatrix}$$

where Q_E, Q_A, b are constant in ϵ. We assume that E is positive-recurrent for Q_E. Then there exists a unique invariant probability measure, ρ, for Q_E. Thus if we had $\epsilon = 0$, any stationary distribution π on D, would be a mixture of $\phi' = (1, 0, \dots)$ and $\psi' = (0, \rho')$. Hence, writing Kolmogorov's forward equation as

$$\frac{dp'}{dt} = p'(\tilde{Q}_E + \epsilon\tilde{Q}_A), \tag{5.2.3}$$

we can interpret $\epsilon\tilde{Q}_A$ as a perturbation of the linear system (Kuehn, 2015). As ϕ and ψ are the only stationary measures, and thus eigenvectors with corresponding eigenvalue 0, we seek a two-dimensional centre manifold. The aim is to modify this, taking into account the perturbation, to obtain the QSD.

Example 5.2.4. Consider as an example the classical example of Keizer's paradox.

$$\emptyset \xleftarrow{\alpha_1} S \underset{\alpha_3}{\overset{\alpha_2}{\rightleftarrows}} 2S$$

It follows that the full rate matrix is

$$\tilde{Q} = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 \\ \alpha_1 & -\alpha_1 - \alpha_2 & \alpha_2 & 0 & 0 \\ 0 & 2\alpha_1 + 2\alpha_3 & -2\alpha_1 - 2\alpha_2 - 2\alpha_3 & 2\alpha_2 & 0 \\ 0 & 0 & 3\alpha_1 + 6\alpha_6 & -3\alpha_1 - 3\alpha_2 - 6\alpha_3 & 3\alpha_2 \\ 0 & 0 & 0 & \ddots & \ddots \end{pmatrix}$$

The endorsed and absorbing set are

$$E = \mathbb{N}, \qquad A = \{0\},$$

respectively, where E is irreducible. Thus $\mathcal{R}_A = \{S \to \emptyset\}$ and we find the endorsed and absorbing networks

$$\mathcal{N}_E: \quad S \underset{\alpha_3}{\overset{\alpha_2}{\rightleftarrows}} 2S \qquad\qquad \mathcal{N}_A: \quad S \xrightarrow{\alpha_1} \emptyset$$

Treating α_1 as a small parameter ϵ one obtains the decomposition

$$\tilde{Q} = \tilde{Q}_E + \epsilon \tilde{Q}_A$$

with the constant rate matrices given by

$$\tilde{Q}_E = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & -\alpha_2 & \alpha_2 & 0 & 0 \\ 0 & 2\alpha_3 & -2\alpha_2 - 2\alpha_3 & 2\alpha_2 & 0 \\ 0 & 0 & 6\alpha_6 & -3\alpha_2 - 6\alpha_3 & 3\alpha_2 \\ 0 & 0 & 0 & \ddots & \ddots \end{pmatrix} \quad \tilde{Q}_A = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 \\ 1 & -1 & 0 & 0 & 0 \\ 0 & 2 & -2 & 0 & 0 \\ 0 & 0 & 3 & -3 & 0 \\ 0 & 0 & 0 & \ddots & \ddots \end{pmatrix}$$

for the endorsed and absorbing sub-networks, respectively. △

5.3 A Dynamical Systems Perspective

Given an initial distribution μ on $\mathbb{N}_0$, let $p' = (p_0, p_1, \dots)$ denote the probabilities given by

$$p_j(t) = \mathbb{P}_\mu(X_t = j), \qquad j \geq 0, t \geq 0,$$

which then satisfy Kolmogorov's forward equation

$$\frac{dp'}{dt} = p'\tilde{Q}, \qquad p(0) = \mu. \tag{5.3.1}$$

Following (Pollett and Roberts, 1990), define

$$\nu_j(t) := \mathbb{P}_\mu(X_t = j \,|\, t < \tau_A) = \frac{\mathbb{P}_\mu(X_t = j, t < \tau_A)}{\mathbb{P}_\mu(t < \tau_A)} = \frac{p_j(t)}{1 - p_0(t)}, \quad j > 0, t \geq 0$$

It then follows that $\nu(t)$ converges to a QLD, (and thus a QSD), ν for $t \to \infty$, where $\nu' = (\nu_1, \nu_2, \dots)$. Substituting this into the Kolmogorov forward equation we find

$$\frac{dp_0(t)}{dt} = \nu'(t)a(1 - p_0(t)),$$

and, letting $p_E = (p_1, p_2, \dots)$ one obtains

$$\frac{d\nu'}{dt}(1 - p_0(t)) - \nu'(t)\frac{dp_0}{dt} = \frac{dp'_E}{dt} = p'_E Q = \nu'(t)Q(1 - p_0(t)).$$

Thus we find the non-linear Riccati system of differential equations for $\nu(t)$,

$$\frac{d\nu'}{dt} = \nu' Q + (\nu' a)\nu'. \tag{5.3.2}$$

One should therefore expect that potential QLDs are stationary points for (5.3.2). As QLDs are QSDs, we rediscover the condition for being a QSD

$$\nu' Q = -(\nu' a)\nu' = -\theta(\nu)\nu'.$$

5.3.1 Center Manifold

In this section, to avoid subtle problems with the spectral theory of operators, and as we are interested in a numerical scheme for calculating QSDs, we assume that one may consider a sufficiently large truncated state space, to contain the dynamics. Thus we have $D = \{0\} \cup \{1, \dots, N\}$ for some $N \in \mathbb{N}$. We treat ϵ as a parameter and consider systems of the canonical form

$$\begin{aligned} \dot{x} &= Ax + f(x, y, \epsilon), \\ \dot{y} &= By + g(x, y, \epsilon), \\ \dot{\epsilon} &= 0 \end{aligned} \tag{5.3.3}$$

where the functions $f, g \in C^2$ satisfy that $f(0,0,0) = 0, g(0,0,0) = 0$ and $\nabla f(0,0,0) = 0, \nabla g(0,0,0) = 0$. Further, $A \in \mathbb{R}^{c \times c}$ is a matrix having eigenvalues with zero real parts and $B \in \mathbb{R}^{(N+1-c)\times(N+1-c)}$ is a matrix having eigenvalues with negative real parts. We assume that the functions f, g are sufficiently differentiable. Recall that a *center manifold* for (5.3.3) is an invariant manifold which can locally be represented as follows

$$W^c(0) = \{(x, y, \epsilon) \in \mathbb{R}^c \times \mathbb{R}^{N+1-c} \times \mathbb{R} \,|\, y = h(x, \epsilon), |x| < \delta, |\epsilon| < \bar{\delta}, h(0,0) = 0, Dh(0,0) = 0\}$$

for $\delta, \bar{\delta}$ sufficiently small (Wiggins, 2003). The center manifold is also referred to as a *slow manifold*, as the evolution on W^c is slow due to the eigenvalues of A having zero real part

(Pollett and Roberts, 1990). We have the following central theorem (Wiggins, 2003; Pollett and Roberts, 1990)

Theorem 5.3.4 (Center Manifold Theorem)**.** There exists a center manifold for (5.3.3). Further, the dynamics of (5.3.3) restricted to the center manifold is, for u and ϵ sufficiently small, given by the $c+1$-dimensional vector field

$$\begin{aligned} \dot{u} &= Au + f(u, h(u,\epsilon), \epsilon), \\ \dot{\epsilon} &= 0. \end{aligned} \tag{5.3.5}$$

If the zero solution of this system is stable, so is the zero solution of the full system, and there exists a solution of (5.3.5) such that, as $t \to \infty$

$$\begin{aligned} x(t) &= u(t) + \mathcal{O}(\exp(-\gamma t)), \\ y(t) &= h(u(t), \epsilon) + \mathcal{O}(\exp(-\gamma t)), \end{aligned}$$

for some $\gamma > 0$, where $(x, y, \epsilon)'$ is any solution of (5.3.3) with $(x(0), y(0), \epsilon)$ sufficiently close to 0. In particular, the flow on W^c is given by (5.3.5).

To apply the center manifold theorem, our first step is to put the forward equation,

$$\frac{dp'}{dt} = p'(\tilde{Q}_E + \epsilon \tilde{Q}_A),$$

obtained in (5.2.3) into the canonical form of (5.3.3). It is known from the Perron-Frobenius theorem that the eigenvalues of $\tilde{Q}_E$ lie in the negative half-plane, and that 0 is an eigenvalue of multiplicity 2. If $\tilde{Q}_E$ is diagonalizable (this will be the generic case as the set of diagonalizable matrices is dense in the set of all complex matrices), one may obtain

$$U \tilde{Q}_E U^{-1} = \tilde{D} \tag{5.3.6}$$

where $\tilde{D} = diag[0, 0, d_1, d_2, \dots]$ is a diagonal matrix. Let $D = diag[d_1, d_2, \dots]$ and $p' = z'N$ with $z' = (x', y')$. It then follows that

$$\frac{dz'}{dt} = p'(\tilde{Q}_E N^{-1} + \epsilon \tilde{Q}_A U^{-1}) = z'\tilde{D} + \epsilon z' U \tilde{Q}_A U^{-1}, \tag{5.3.7}$$

which we may then write in the canonical form,

$$\begin{aligned} \frac{dx'}{dt} &= f'(x, y, \epsilon), \\ \frac{dy'}{dt} &= y'D + g'(x, y, \epsilon), \\ \frac{d\epsilon}{dt} &= 0 \end{aligned} \tag{5.3.8}$$

It follows from the center manifold theorem, that there exists a center manifold, W^c, consisting of points (x, y, ϵ) such that $y = h(x, \epsilon)$. Solutions to (5.3.8) approach W^c exponentially

fast as $t \to \infty$ and, on W^c, the flow $u' = (u_0, u_1)$ is governed by

$$\frac{du'}{dt} = f(u, h(u, \epsilon), \epsilon).$$

Performing the transformation $p' = z'U$, we obtain a center manifold in the space of total probabilities (Pollett and Roberts, 1990).

The Center Manifold in p-space. To obtain the center manifold, we follow Pollett and Roberts (1990). Here it was argued that one should look for a two-dimensional manifold spanned by the limiting distribution ϕ and another distribution φ assigning positive probability to each state in E. Indeed, these are orthogonal eigenvectors with zero eigenvalues. Further, the span of these vectors intersect the p-space in a line, yielding our effective center manifold. Suppose therefore that one may write $p(t) = f(x(t))$ where $x' = (x_0, x_1)$ is a non-negative vector such that $x_0 + x_1 = 1$ and

$$f(x) = x_0\phi + x_1\varphi,$$

where $\varphi' = (0, m')$ is some probability distribution over E. Further, the exponential convergence implies that the change in weights follows $dx/dt = g(x)$ where

$$g'(x) = x' \begin{pmatrix} 0 & 0 \\ \eta & -\eta \end{pmatrix}, \tag{5.3.9}$$

for $\eta > 0$. Inserting this in the Master equation (5.3.1), we arrive at the equation

$$f'\tilde{Q} = g'\nabla_x f.$$

Substituting for $\tilde{Q}, f$ and g we find

$$x_1(m'a, m'Q) = \left(x_0\phi' + x_1\varphi'\right)\tilde{Q} = f'\tilde{Q} = g'\nabla_x f = x_1 \begin{pmatrix} \eta & -\eta \end{pmatrix} \begin{pmatrix} \phi' \\ \varphi' \end{pmatrix} = x_1(\eta, -\eta m').$$

Hence, in the non-trivial case where $x_1 \neq 0$ we arrive at the equations

$$m'a = \eta, \qquad m'Q = -\mu m',$$

that is

$$m'Q = -(m'a)m',$$

which is exactly the condition for being a QSD (Collet et al., 2013). We conclude that $m = \nu$ and it follows that the center manifold, W^c, is the line in p-space connecting the points ϕ (the limiting distribution) and ν (a quasi-stationary distribution) depending on ϵ (Pollett and Roberts, 1990). Further, on W^c, the flow approximating that of $p(t)$ is

$$p(t) = f(x(t)) = x_0(t) \begin{pmatrix} 1 \\ 0 \end{pmatrix} + x_1(t) \begin{pmatrix} 0 \\ \nu \end{pmatrix} \tag{5.3.10}$$

where $x(t)$ satisfies the differential equation (5.3.9), that is, the probabilities over E decay to 0 exponentially, although at a very slow rate. It was argued in Pollett and Roberts (1990) that η is related to the decay parameter as $\eta \leq \theta^*$.

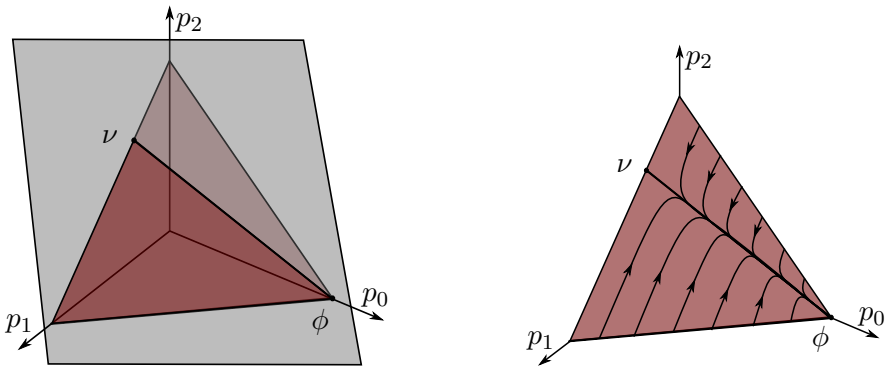

Figure 5.1: Schematic presentation of the center manifold in p-space.

Observability of Quasi-stationarity. To justify the center manifold reduction, one needs to argue that the convergence from any initial distribution onto the slow manifold happens sufficiently quick. For this, one may employ a *spectral gap* analysis. Indeed, if the spectral gap of $\tilde{Q}_E$ is much larger than 0, then we may observe quasi-stationarity for a range of comparatively smaller values of the parameter ϵ (Childs and Keener, 2012). Note that as we have assumed Q_E has a stationary distribution, results connected to the spectrum is more readily available. Indeed, for birth-death processes, still on a truncated state space, one may employ a general procedure whereby a lower bound for the spectral gap may be computed (Granovsky and Zeifman, 1997). It was argued in Childs and Keener (2012) that this procedure can be bootstrapped to the non-truncated case by letting the size of the truncation go to infinity.

An Iterative Scheme for the QSD. One may exploit the connection between the center manifold and the QSD, ν. Indeed, if one can approximate the former through an approximation of f, one will as a byproduct have approximated the QSD. In fact, one may do so to an arbitrarily small degree of error. We follow Pollett and Roberts (1990), where a method avoiding the change of variable is explained. Substituting (5.3.10) into the master equation, we find that

$$f'\tilde{Q}_E + \epsilon f'\tilde{Q}_A = \frac{dp'}{dt} = \frac{df(x(t))'}{dt} = g'\nabla_x f,$$

that is

$$f'\tilde{Q}_E = g'\nabla_x f - \epsilon f'\tilde{Q}_A, \tag{5.3.11}$$

where we may consider the terms on the right hand side small (Pollett and Roberts, 1990). Using that the solution to the Kolmogorov forward equation my be written as an exponential function, which is analytic, we can write f and g in the expansion

$$f(x) = \sum_{n=0}^{\infty} f_n(x), \qquad g(x) = \sum_{n=1}^{\infty} g_n(x),$$

where f_n, g_n are terms of order n in ϵ. Further, one may by Pollett and Roberts (1990) obtain the following iterative scheme

$$\begin{aligned} f_0' \tilde{Q}_E &= 0, \\ f_n' \tilde{Q}_E &= \sum_{m=1}^{\infty} g_m' \nabla_x f_{n-m} - \epsilon f_{n-1}' \tilde{Q}_A, \qquad n \geq 1. \end{aligned} \tag{5.3.12}$$

In the setting of reaction networks introduced in Section 5.2, we then have the following iterative scheme for the quasi-stationary distribution.

Theorem 5.3.13. Suppose the reaction network $(\mathcal{N}, \lambda)$ has a quasi-stationary distribution, ν, and that the corresponding rate matrix allows a slow manifold reduction. Further, assume that the endorsed reaction network $\mathcal{N}_E$ has a stationary distribution ρ. Then

$$\nu = \sum_{n=0}^{\infty} \rho_n \epsilon^n,$$

where

$$\rho_0 = \rho, \qquad \rho_n = (Q_E')^{-1} \left(-Q_A' \rho_{n-1} - \sum_{m=0}^{n-1} (\rho_m' b) \rho_{n-m} \right), \qquad n \geq 1$$

Proof. We prove the statement by induction on n. For $n = 0$ we have the equation $f_0' \tilde{Q}_E = 0$ which has the solution

$$f_0(x) = x_0 \phi + x_1 \rho.$$

With $f_0(x) = x_0 \phi + x_1 \psi$ and $\phi' = (1, 0, \dots)$ and $\psi' = (0, \rho')$, we find

$$\begin{aligned} f_1' \tilde{Q}_E &= g_1' \nabla_x f_0 - \epsilon f_0' \tilde{Q}_A \\ &= g_1' \begin{pmatrix} 1 & 0' \\ 0 & \rho' \end{pmatrix} - \epsilon (x_0 \phi' + x_1 \psi') \begin{pmatrix} 0 & 0' \\ b & Q_A \end{pmatrix} \\ &= g_1' \begin{pmatrix} 1 & 0' \\ 0 & \rho' \end{pmatrix} - \epsilon x_1 (0, \rho') \begin{pmatrix} 0 & 0' \\ b & Q_A \end{pmatrix} \end{aligned} \tag{5.3.14}$$

Multiplying with the eigenvector of $\tilde{Q}_E$, $(1, 0')$ (with eigenvalue 0), we get

$$g_1' \begin{pmatrix} 1 \\ 0 \end{pmatrix} = g_1' \begin{pmatrix} 1 & 0' \\ 0 & \rho' \end{pmatrix} \begin{pmatrix} 1 \\ 0 \end{pmatrix} = \epsilon x_1(0,\rho') \begin{pmatrix} 0 & 0' \\ b & Q_A \end{pmatrix} \begin{pmatrix} 1 \\ 0 \end{pmatrix}$$
$$= \epsilon x_1(0,\rho') \begin{pmatrix} 0 \\ b \end{pmatrix} = \epsilon x_1 \rho' b.$$

And if we multiply by the other eigenvector $(0, 1')$ (with eigenvalue 0) we get, using that the row sums of Q_A is $-b$ and ρ is a probability measure

$$g_1' \begin{pmatrix} 0 \\ 1 \end{pmatrix} = g_1' \begin{pmatrix} 0 \\ \rho' \cdot 1 \end{pmatrix} = g_1' \begin{pmatrix} 1 & 0' \\ 0 & \rho' \end{pmatrix} \begin{pmatrix} 0 \\ 1 \end{pmatrix} = \epsilon x_1(0,\rho') \begin{pmatrix} 0 & 0' \\ b & Q_A \end{pmatrix} \begin{pmatrix} 0 \\ 1 \end{pmatrix}$$
$$= \epsilon x_1(0,\rho') \begin{pmatrix} 0 \\ -b \end{pmatrix} = \epsilon x_1(-\rho' b).$$

We therefore see that

$$g_1' = \epsilon x' \begin{pmatrix} 0 & 0 \\ \rho' b & -\rho' b \end{pmatrix}. \tag{5.3.15}$$

Substituting back into (5.3.14) one finds

$$\begin{aligned} f_1' \tilde{Q}_E &= g_1' \begin{pmatrix} 1 & 0' \\ 0 & \rho' \end{pmatrix} - \epsilon x_1(0,\rho') \begin{pmatrix} 0 & 0' \\ b & Q_A \end{pmatrix} \\ &= \epsilon x' \begin{pmatrix} 0 & 0 \\ \rho' b & -\rho' b \end{pmatrix} \begin{pmatrix} 1 & 0' \\ 0 & \rho' \end{pmatrix} - \epsilon x_1(0,\rho') \begin{pmatrix} 0 & 0' \\ b & Q_A \end{pmatrix} \\ &= \epsilon x_1 \begin{pmatrix} \rho' b & (-\rho' b)\rho' \end{pmatrix} - \epsilon x_1 \begin{pmatrix} \rho' b & \rho' Q_A \end{pmatrix} \\ &= \epsilon x_1 \begin{pmatrix} 0 & -\rho' Q_A - (\rho' b)\rho' \end{pmatrix}. \end{aligned}$$

Setting $\rho_1' = (-\rho' Q_A - (\rho' b)\rho')(Q_E)^{-1}$, we may write this as

$$f_1' = \epsilon x_1 \begin{pmatrix} 0 & \rho_1' \end{pmatrix}$$

in agreement with the desired. We can then find ρ_2 as follows

$$\begin{aligned} f_2' \tilde{Q}_E &= g_1' \nabla_x f_1 + g_2' \nabla_x f_0 - \epsilon f_1' \tilde{Q}_A \\ &= \epsilon^2 x' \begin{pmatrix} 0 & 0 \\ \rho' b & -\rho' b \end{pmatrix} \begin{pmatrix} 0 & 0 \\ 0 & \rho_1' \end{pmatrix} + g_2' \begin{pmatrix} 1 & 0' \\ 0 & \rho' \end{pmatrix} - \epsilon^2 x_1 \begin{pmatrix} 0 & \rho_1' \end{pmatrix} \begin{pmatrix} 0 & 0 \\ b & Q_A \end{pmatrix} \\ &= \epsilon^2 x_1 \begin{pmatrix} 0 & -\rho' b \rho_1' \end{pmatrix} + g_2' \begin{pmatrix} 1 & 0' \\ 0 & \rho' \end{pmatrix} - \epsilon^2 x_1 \begin{pmatrix} \rho_1' b & \rho_1' Q_A \end{pmatrix}. \end{aligned}$$

Multiplying with the eigenvector of $\tilde{Q}_E$, $(1, 0')$, we get

$$g_2' \begin{pmatrix} 1 \\ 0 \end{pmatrix} = \epsilon^2 x_1 \rho_1' b. \tag{5.3.16}$$

And multiplying by the other eigenvector $(0, 1')$ one finds, using that the row sums of Q_A is $-b$ and ρ is a probability measure and $\|\rho_1\|_1 = 0$,

$$g_2' \begin{pmatrix} 0 \\ 1 \end{pmatrix} = g_2' \begin{pmatrix} 1 & 0' \\ 0 & \rho' \end{pmatrix} \begin{pmatrix} 0 \\ 1 \end{pmatrix} = \epsilon^2 x_1 \begin{pmatrix} \rho_1' b & \rho_1' Q_A \end{pmatrix} \begin{pmatrix} 0 \\ 1 \end{pmatrix} = -\epsilon^2 x_1 \rho_1' b,$$

thus

$$g_2' = \epsilon^2 x' \begin{pmatrix} 0 & 0 \\ \rho_1' b & -\rho_1' b \end{pmatrix}. \tag{5.3.17}$$

We get

$$\begin{aligned} f_2' \tilde{Q}_E &= \epsilon^2 x_1 \begin{pmatrix} 0 & -\rho' b \rho_1' \end{pmatrix} + g_2' \begin{pmatrix} 1 & 0' \\ 0 & \rho' \end{pmatrix} - \epsilon^2 x_1 \begin{pmatrix} \rho_1' b & \rho_1' Q_A \end{pmatrix} \\ &= \epsilon^2 x_1 \begin{pmatrix} 0 & -\rho' b \rho_1' \end{pmatrix} + \epsilon^2 x' \begin{pmatrix} 0 & 0 \\ \rho_1' b & -\rho_1' b \end{pmatrix} \begin{pmatrix} 1 & 0' \\ 0 & \rho' \end{pmatrix} - \epsilon^2 x_1 \begin{pmatrix} \rho_1' b & \rho_1' Q_A \end{pmatrix} \\ &= \epsilon^2 x_1 \begin{pmatrix} 0 & -\rho_1' Q_A - \rho_1' b \rho' - \rho' b \rho_1' \end{pmatrix}. \end{aligned}$$

Setting $\rho_2' = (-\rho_1' Q_A - \rho_1' b \rho' - \rho' b \rho_1')(Q_E')^{-1}$ it follows that one may write

$$f_2' = \epsilon^2 x_1 \begin{pmatrix} 0 & \rho_2' \end{pmatrix},$$

as desired. In general suppose that the statement of the proposition holds for some $n \geq 2$. Then in particular $f_n' = \epsilon^n x_1 \begin{pmatrix} 0 & \rho_n' \end{pmatrix}$ hence

$$\epsilon f_n' \tilde{Q}_A = \epsilon^{n+1} x_1 \begin{pmatrix} 0 & \rho_n' \end{pmatrix} \tilde{Q}_A, \tag{5.3.18}$$

and multiplying with the two eigenvectors $(1, 0')$ and $(0, 1')$ respectively we infer that

$$g_{n+1}' = \epsilon^{n+1} x' \begin{pmatrix} 0 & 0 \\ \rho_n' b & -\rho_n' b \end{pmatrix}. \tag{5.3.19}$$

From the iterative scheme (5.3.12) one then finds

$$\begin{aligned} f_{n+1}' \tilde{Q}_E &= \sum_{m=1}^{\infty} g_m' \nabla_x f_{n-m+1} - \epsilon f_n' \tilde{Q}_A \\ &= \sum_{m=1}^{\infty} \epsilon^m x' \begin{pmatrix} 0 & 0 \\ \rho_{m-1}' b & -\rho_{m-1}' b \end{pmatrix} \nabla_x \epsilon^{n-m+1} x_1 \begin{pmatrix} 0 & \rho_{n-m+1}' \end{pmatrix} - \epsilon f_n' \tilde{Q}_A \\ &= \epsilon^{n+1} \left(\sum_{m=1}^{\infty} x_1 \begin{pmatrix} \rho_{m-1}' b & -\rho_{m-1}' b \end{pmatrix} \begin{pmatrix} \mathbb{1}_{n+1}(m) & 0' \\ 0 & \rho_{n-m+1}' \end{pmatrix} - x_1 \begin{pmatrix} 0 & \rho_n' \end{pmatrix} \tilde{Q}_A \right) \\ &= \epsilon^{n+1} x_1 \begin{pmatrix} 0 & -\sum_{m=0}^{\infty} (\rho_m' b) \rho_{n-m} - \rho_n' Q_A \end{pmatrix}. \end{aligned}$$

Thus we conclude that

$$\rho'_{n+1} = \left(-\rho'_n Q_A - \sum_{m=0}^{\infty} (\rho'_m b)\rho_{n-m} \right) (Q_E)^{-1}, \tag{5.3.20}$$

which is the desired upon transposing. □

One may immediately obtain the following very natural result, when the leakage, by means of ϵ, tends to zero.

Corollary 5.3.21. Let ν be the quasi-stationary distribution for $\mathcal{N}$ and ρ the stationary distribution for $\mathcal{N}_E$. Then $\|\nu - \rho\|_{TV} \to 0$ as $\epsilon \to 0$.

The way we have constructed the division of the original reaction network, $\mathcal{N}$, into the smaller sub-networks $\mathcal{N}_E$ and $\mathcal{N}_A$ allows us to exploit some of the existing results in stochastic reaction network theory. In particular, using the result of Anderson et al. (2010), we find that if the endorsed network $\mathcal{N}_E$ is weakly reversible and has deficiency zero, then the quasi-stationary distribution for $\mathcal{N}$ is approximated by a product-form poisson distribution.

Example 5.3.22. Returning to the logistic network giving rise to Keizer's paradox,

$$\emptyset \xleftarrow{\alpha_1} S \underset{\alpha_3}{\overset{\alpha_2}{\rightleftarrows}} 2S$$

The endorsed network, $\mathcal{N}_E$, is given by

$$S \underset{\alpha_3}{\overset{\alpha_2}{\rightleftarrows}} 2S$$

which is weakly reversible and has deficiency 0. Thus it has a poisson-form distribution,

$$\rho_0(x) = \frac{1}{Z}\frac{c^x}{x!}e^{-c}, \qquad x \in \mathbb{N}$$

where $c = \alpha_2/\alpha_3$ is the complex balanced equilibrium.

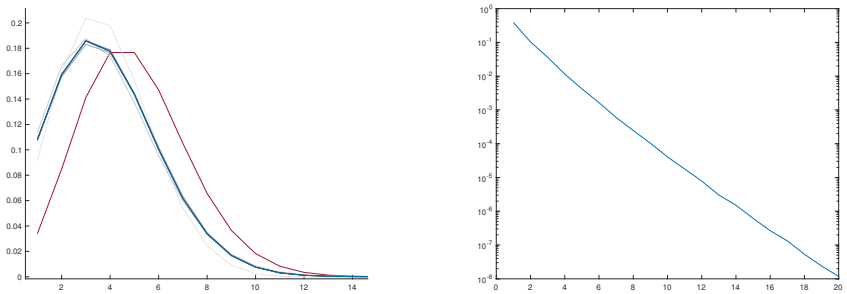

Figure 5.2: Left: The first 6 approximations of the QSD. Right: The error compared to exact eigenvalue computation in 1-norm. There is an exponential convergence, $\varepsilon \approx \mathcal{O}(e^{-1.3n})$. Here $\alpha_1 = 1, \alpha_2 = 5, \alpha_3 = 1$.

This is a first approximation to the QSD ν by Theorem 5.3.13. Calculating further terms in the sum, one may approximate ν to any degree desired, see Figure 5.2.

Example 5.3.23. Consider the semi-reversible Lotka-Volterra network,

$$S_1 \underset{\alpha_2}{\overset{\alpha_1}{\rightleftharpoons}} 2S_1$$
$$S_1 + S_2 \underset{\alpha_4}{\overset{\alpha_3}{\rightleftharpoons}} 2S_2$$
$$S_2 \overset{\alpha_5}{\to} \emptyset$$

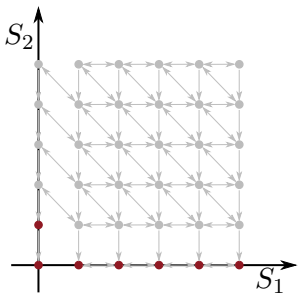

Only through the reaction $S_2 \to \emptyset$ may the absorbing set be reached hence the endorsed and absorbing networks are given by

$\mathcal{N}_E$:
$$S_1 \underset{\alpha_2}{\overset{\alpha_1}{\rightleftharpoons}} 2S_1$$
$$S_1 + S_2 \underset{\alpha_4}{\overset{\alpha_3}{\rightleftharpoons}} 2S_2$$

$\mathcal{N}_A$: $S_2 \xrightarrow{\alpha_5} \emptyset$

The endorsed network is weakly reversible and has deficiency 0. Furthermore, the endorsed set is irreducible. We conclude that for sufficiently small values of the rate α_5, the quasi-stationary distribution for the semi-reversible Lotka-Volterra network is approximated by the product form poisson distribution

$$\rho(x) = \frac{1}{Z} \frac{(\frac{\alpha_1}{\alpha_2})^{x_1} (\frac{\alpha_1\alpha_3}{\alpha_2\alpha_4})^{x_2}}{x_1! x_2!} e^{-\left(\frac{\alpha_1}{\alpha_2} + \frac{\alpha_1\alpha_3}{\alpha_2\alpha_4}\right)}, \qquad x \in E.$$

where Z is a normalizing constant.

5.4 Limit Distributions for One-species Reaction Networks

When the endorsed network is neither mono-molecular nor complex balanced, not much is known about the form of the stationary distribution. Here, we present a scheme for calculating it analytically, in the case of one-species mass-action reaction networks. These may be seen as generalizations of birth-death processes. The first step is to have sufficient conditions for the existence of a unique limit distribution. Define the total rate out of the highest order complex, $y_{\max}$, to be

$$\alpha^+ = \sum_{k \in \mathcal{R}: y = y_{\max}} \alpha_k \xi_k$$

Then we have the following proposition.

Proposition 5.4.1. *If $\alpha^+ < 0$ and each endorsed set contains a unique minimal irreducible class, then the process is exponentially ergodic. In particular, it has a unique stationary distribution.*

Proof. By Gupta et al. (2014) and Hansen and Wiuf (2018a) it suffices to find constants $c_1, c_2 > 0$ such that the following negative drift condition holds for all $x \in \mathbb{N}_0$

$$\sum_{k \in \mathcal{R}} \lambda_k(x)(V(x + \xi_k) - V(x)) \leq c_1 - c_2 V(x), \tag{5.4.2}$$

where V is a Lyapunov function. Simply take $V : \mathbb{N}_0 \to \mathbb{R}_0$ to be the identity $V(x) = x$. Then $V(x) \to \infty$ as $x \to \infty$ as required. The statement then reduces to

$$F(x) := \sum_{k \in \mathcal{R}} \lambda_k(x)\xi_k \leq c_1 - c_2 x, \tag{5.4.3}$$

for all $x \in \mathbb{N}_0$. As $\alpha^+ < 0$ it follows that $F(x) < 0$ for x sufficiently large. If $F(x) < 0$ for all $x \in \mathbb{N}_0$ then the process would not be bounded in $\mathbb{N}_0$, thus we conclude that there exists a zero of F. Let $x^* - 1$ be the largest zero of F and define

$$c_1 = \max\{F(x) : 0 \leq x \leq x^*\} - x^* F'(x^*), \qquad c_2 = -F'(x^*). \tag{5.4.4}$$

Note that since $x^* - 1$ is the largest zero and $\alpha^+ < 0$ we must have $F'(x^*) < 0$. Therefore, $c_1, c_2 > 0$ as required. Now, for $x \leq x^*$ we see that

$$F(x) \leq \max\{F(x) : 0 \leq x \leq x^*\} - (x^* - x)F'(x^*) = c_1 - c_2 x,$$

and for $x > x^*$ we find, as F is concave on this segment,

$$\begin{aligned} F(x) &\leq F(x^*) + F'(x^*)(x - x^*) \leq F'(x^*)x - x^* F'(x^*) \\ &\leq -x^* F'(x^*) + \max\{F(x) : 0 \leq x \leq x^*\} + F'(x^*)x = c_1 - c_2 x. \end{aligned} \tag{5.4.5}$$

We conclude that the process is exponentially ergodic. □

Corollary 5.4.6. A weakly reversible one-species network with mass-action kinetics has a unique non-degenerate stationary distribution.

Proof. As the network is weakly reversible, the highest order complex can have no reactions in the positive direction. In particular $\alpha^+ < 0$ and there exists a unique stationary distribution. Similarly, by weak reversibility, the complex of lowest order must have a reaction in the positive direction and the stationary distribution cannot be degenerate. □

The question of whether weak reversibility is a sufficient condition for a reaction network of any dimension to have a stationary distribution is an active area of research, and it is generally believed that this should be the case.

5.4.1 Solving the Master Equation

To find the stationary distribution in one-species reaction networks, we devise a procedure to solve the Master equation. Let Ω be the set of jump sizes,

$$\Omega = \{\omega \in \mathbb{N} \,|\, \exists k \in \mathcal{R} \colon |\xi_k| = \omega\}.$$

For $\omega \in \Omega$ divide the set of reactions into positive and negative jumps with jump-size ω

$$\mathcal{R}_-^\omega = \{k \in \mathcal{R} \,|\, \xi_k = -\omega\}, \qquad \mathcal{R}_+^\omega = \{k \in \mathcal{R} \,|\, \xi_k = \omega\}.$$

Finally, define the following quantities coming from the mass-action rates

$$a_x^\omega = \sum_{k \in \mathcal{R}_-^\omega} \alpha_k x^{\underline{y_k}}, \qquad b_x^\omega = \sum_{k \in \mathcal{R}_+^\omega} \alpha_k x^{\underline{y_k}}.$$

where $x^{\underline{y}}$ denotes the falling factorial. Note that when $\Omega = \{1\}$, these quantities reduce to variable death and birth rates respectively. Now, letting $p_x(t) = \mathbb{P}(X(t) = x)$, We may write Kolmogorov's forward equation, also referred to as the master equation, as

$$\dot{p}_x(t) = \sum_{k \in \mathcal{R}} \lambda_k(x - \xi_k) p_{x-\xi_k}(t) - \sum_{k \in \mathcal{R}} \lambda_k(x) p_x(t). \tag{5.4.7}$$

The stationary distribution, π, is found at equilibrium hence

$$\begin{aligned}
0 &= \sum_{k \in \mathcal{R}} \lambda_k(x - \xi_k)\pi_{x-\xi_k} - \sum_{k \in \mathcal{R}} \lambda_k(x)\pi_x \\
&= \sum_{\omega \in \Omega} \sum_{k \in \mathcal{R}_-^\omega} \alpha_k (x+\omega)^{\underline{y_k}} \pi_{x+\omega}(t) + \sum_{\omega \in \Omega} \sum_{k \in \mathcal{R}_+^\omega} \alpha_k (x-\omega)^{\underline{y_k}} \pi_{x-\omega} \\
&\qquad - \sum_{\omega \in \Omega} \sum_{k \in \mathcal{R}_-^\omega} \alpha_k x^{\underline{y_k}} \pi_x - \sum_{\omega \in \Omega} \sum_{k \in \mathcal{R}_+^\omega} \alpha_k x^{\underline{y_k}} \pi_x \\
&= \sum_{\omega \in \Omega} a_{x+\omega}^\omega \pi_{x+\omega} + \sum_{\omega \in \Omega} b_{x-\omega}^\omega \pi_{x-\omega} - \sum_{\omega \in \Omega} a_x^\omega \pi_x - \sum_{\omega \in \Omega} b_x^\omega \pi_x.
\end{aligned}$$

We arrive at the following form of the master equation

$$\sum_{\omega \in \Omega} \pi_{x-\omega} b_{x-\omega}^\omega - \pi_x \left(\sum_{\omega \in \Omega} a_x^\omega + b_x^\omega \right) + \sum_{\omega \in \Omega} \pi_{x+\omega} a_{x+\omega}^\omega = 0, \qquad x \geq 0. \tag{5.4.8}$$

Note that the sums are finite since the set of jump-sizes Ω is bounded.

Proposition 5.4.9. *One may write the Master equations as*

$$\sum_{i=1}^{\infty} \pi_{x-i} \left(\sum_{\omega \geq i} b_{x-i}^\omega \right) = \sum_{i=1}^{\infty} \sum_{\omega \geq i} \pi_{x+\omega-i} a_{x+\omega-i}^\omega, \qquad x \geq 0. \tag{5.4.10}$$

Proof. Inserting $x = 0$, the left hand side is readily seen to be zero, while the fact that $a_x^\omega = 0$ for $x < \omega$ implies that the right hand side is zero. Now, suppose (5.4.10) holds for some $x \geq 0$. One may write this as

$$\pi_x \sum_{\delta>0} a_x^\omega = \sum_{i=1}^{\infty} \pi_{x-i} \left(\sum_{\omega \geq i} b_{x-i}^\omega \right) - \sum_{i=1}^{\infty} \sum_{\omega > i} \pi_{x+\omega-i} a_{x+\omega-i}^\omega. \tag{5.4.11}$$

From the Master equation (5.4.8) it follows that

$$\sum_{\omega>0} \pi_{x+\omega} a_{x+\omega}^\omega = \pi_n \sum_{\omega>0} a_x^\omega + \pi_x \sum_{\omega>0} b_x^\omega - \sum_{\omega>0} \pi_{x-\omega} b_{x-\omega}^\omega.$$

Inserting (5.4.11) and collecting the terms containing a on the left hand side and the terms containing b on the right hand side one obtains

$$\sum_{\omega>0} \pi_{x+\omega} a_{x+\omega}^\omega + \sum_{i=1}^{\infty} \sum_{\omega>i} \pi_{x+\omega-i} a_{x+\omega-i}^\omega = \sum_{i=1}^{\infty} \pi_{x-i} \left(\sum_{\omega \geq i} b_{x-i}^\omega \right) + \pi_x \sum_{\omega>0} b_x^\omega - \sum_{\omega>0} \pi_{x-\omega} b_{x-\omega}^\omega$$

which upon reshuffling of the indices, we may write as

$$\sum_{i=1}^{\infty} \pi_{x+1-i} \left(\sum_{\omega \geq i} b_{x+1-i}^\omega \right) - \sum_{i=1}^{\infty} \sum_{\omega \geq i} \pi_{x+1+\omega-i} a_{x+1+\omega-i}^\omega = 0.$$

The required now follows by the principle of simple induction. □

Example 5.4.12.

$$S \rightleftarrows 2S \longrightarrow 3S$$

The set of jump-sizes is $\Omega = \{1, 2\}$, yielding the master equation

$$\pi_{x-1} b_{x-1}^1 - \pi_x (a_x^2 + b_x^1) + \pi_{x+2} a_{x+2}^2 = 0, \qquad x \geq 0$$

and we conclude that for $x \geq 0$,

$$\pi_{x+1} a_{x+1}^2 = \pi_{x-1} b_{x-1}^1 - \pi_x a_x^2,$$

by Proposition 5.4.9. △

5.4.2 The Stationary Distribution

Let λ_* be the lowest order of a product complex in the network, and let $\lambda^* + 1$ be the minimal order of a reaction with largest possible jump, ω_*, in negative direction. We note that $\lambda_* \leq \lambda^*$. Indeed, suppose for contradiction that $\lambda_* \geq \lambda^* + 1$. As there is a reaction in the negative direction from $\lambda^* + 1$, the product complex is smaller than λ_* which is a contradiction. The next proposition shows that we can express π_x in terms of $\pi_{\lambda_*}, \ldots, \pi_{\lambda^*}$.

Proposition 5.4.13. *Let $\omega_*, \omega^* \in \Omega$ be the size of the largest jump in the negative and positive directions respectively. Then*

$$\pi_x a_x^{\omega_*} = \pi_{\lambda_*}\gamma_x^{(\lambda_*)} + \cdots + \pi_{\lambda^*}\gamma_x^{(\lambda^*)} \tag{5.4.14}$$

where for $x \geq \omega^* + \lambda^*$

$$\gamma_{x+\omega_*}^{(j)} = \sum_{i=0}^{\infty}\gamma_{x-i}^{(j)}\left(\sum_{\omega>i}\frac{b_{x-i}^{\omega}}{a_{x-i}^{\omega_*}}\right) - \sum_{i=1}^{\infty}\sum_{\omega>i}\gamma_{x+\omega-i}^{(j)}\frac{a_{x+\omega-i}^{\omega}}{a_{x+\omega-i}^{\omega_*}} - \sum_{\omega\neq\omega_*}\gamma_{x+\omega}^{(j)}\frac{a_{x+\omega}^{\omega}}{a_{x+\omega}^{\omega_*}}$$

for $j = \lambda_*, \ldots, \lambda^*$.

Proof. For $x \leq \lambda^*$ the statement (5.4.14) is trivial. If $x > \lambda^*$ then in particular, we have $a_x^{\omega_*} \neq 0$. Let $x = \lambda^* + 1$. From the master equation (as $x - \omega_* \geq 0$) we have

$$\pi_x a_x^{\omega_*} = -\sum_{\omega\in\Omega}\pi_{x-\omega_*-\omega}b_{x-\omega}^{\omega} + \pi_{x-\omega_*}\left(\sum_{\omega\in\Omega}a_{x-\omega_*}^{\omega} + b_{x-\omega_*}^{\omega}\right) - \sum_{\omega\neq\omega_*}\pi_{x-\omega_*+\omega}a_{x-\omega_*+\omega}^{\omega},$$

hence dividing by $a_x^{\omega_*}$ it follows that π_x is a linear combination of $\pi_0, \ldots, \pi_{x-1}$, that is, $\pi_0, \ldots, \pi_{\lambda^*}$. We may proceed inductively in the same manner, allowing us to write

$$\pi_x = \sum_{j=0}^{\lambda^*}\pi_j\gamma_x^{(j)},$$

for some value of $\gamma_x^{(j)}$, $j = 0, \ldots, \lambda^*$, as desired. Now, the state space of the process is confined to $\{\lambda_*, \lambda_* + 1 \ldots\}$ since no reaction in negative direction can happen from λ_* – otherwise there would be a product complex of lower order. Therefore, we may view it as $\pi_0 = \cdots = \pi_{\lambda_*-1} = 0$, yielding

$$\pi_x = \sum_{j=\lambda_*}^{\lambda^*}\pi_j\gamma_x^{(j)},$$

as required. Finally, since $x > \lambda^*$ was arbitrary, the result (5.4.14) follows. By Proposition 5.4.9 we have, for $x \geq 0$

$$\begin{aligned}\sum_{i=1}^{\infty}\pi_{x-i}\left(\sum_{\omega\geq i}b_{x-i}^{\omega}\right) &= \sum_{i=1}^{\infty}\sum_{\omega\geq i}\pi_{x+\omega-i}a_{x+\omega-i}^{\omega} = \sum_{i=2}^{\infty}\sum_{\omega\geq i}\pi_{x+\omega-i}a_{x+\omega-i}^{\omega} + \sum_{\omega\geq i}\pi_{x+\omega-1}a_{x+\omega-1}^{\omega} \\ &= \sum_{i=2}^{\infty}\sum_{\omega\geq i}\pi_{x+\omega-i}a_{x+\omega-i}^{\omega} + \sum_{\omega\neq\omega_*}\pi_{x+\omega-1}a_{x+\omega-1}^{\omega} + \pi_{x+\omega_*-1}a_{x+\omega_*-1}^{\omega_*}.\end{aligned}$$

Shifting indices, we may rearrange this expression,

$$\begin{aligned}\pi_{x+\omega_*}a^{\omega_*}_{x+\omega_*} &= \sum_{i=1}^{\infty}\pi_{x+1-i}\left(\sum_{\omega\geq i}b^{\omega}_{x+1-i}\right) - \sum_{i=2}^{\infty}\sum_{\omega\geq i}\pi_{x+1+\omega-i}a^{\omega}_{x+1+\omega-i} - \sum_{\omega\neq\omega_*}\pi_{x+\omega}a^{\omega}_{x+\omega}\\ &= \sum_{i=0}^{\infty}\pi_{x-i}\left(\sum_{\omega\geq i+1}b^{\omega}_{x-i}\right) - \sum_{i=1}^{\infty}\sum_{\omega\geq i+1}\pi_{x+\omega-i}a^{\omega}_{x+\omega-i} - \sum_{\omega\neq\omega_*}\pi_{x+\omega}a^{\omega}_{x+\omega}\\ &= \sum_{i=0}^{\infty}\pi_{x-i}\left(\sum_{\omega>i}b^{\omega}_{x-i}\right) - \sum_{i=1}^{\infty}\sum_{\omega>i}\pi_{x+\omega-i}a^{\omega}_{x+\omega-i} - \sum_{\omega\neq\omega_*}\pi_{x+\omega}a^{\omega}_{x+\omega}.\end{aligned}$$

Note that all indices of π are smaller on the right-hand side than on the left-hand side. Further, as $a^{\omega_*}_{x-i}=0$ and $\omega>i$ holds if and only if $x-i\leq\lambda^*$ and $\omega>i$, that is if and only if $x-\lambda^*<\omega$, we infer that the fraction $\frac{b^{\omega}_{x-i}}{a^{\omega_*}_{x-i}}$ is well defined whenever $x\geq\omega^*+\lambda^*$. Applying the representation (5.4.14), letting $\Lambda=\{\lambda_*,\dots,\lambda^*\}$, we obtain

$$\begin{aligned}\pi_{x+\omega_*}a^{\omega_*}_{x+\omega_*} =& \sum_{i=0}^{\infty}\sum_{j\in\Lambda}\pi_j\gamma^{(j)}_{x-i}\left(\sum_{\omega>i}\frac{b^{\omega}_{x-i}}{a^{\omega_*}_{x-i}}\right) - \sum_{i=1}^{\infty}\sum_{\omega>i}\sum_{j\in\Lambda}\pi_j\gamma^{(j)}_{x+\omega-i}\frac{a^{\omega}_{x+\omega-i}}{a^{\omega_*}_{x+\omega-i}}\\ &- \sum_{\omega\neq\omega_*}\sum_{j\in\Lambda}\pi_j\gamma^{(j)}_{x+\omega}\frac{a^{\omega}_{x+\omega}}{a^{\omega_*}_{x+\omega}}\\ =& \sum_{j\in\Lambda}\pi_j\left(\sum_{i=0}^{\infty}\gamma^{(j)}_{x-i}\left(\sum_{\omega>i}\frac{b^{\omega}_{x-i}}{a^{\omega_*}_{x-i}}\right) - \sum_{i=1}^{\infty}\sum_{\omega>i}\gamma^{(j)}_{x+\omega-i}\frac{a^{\omega}_{x+\omega-i}}{a^{\omega_*}_{x+\omega-i}} - \sum_{\omega\neq\omega_*}\gamma^{(j)}_{x+\omega}\frac{a^{\omega}_{x+\omega}}{a^{\omega_*}_{x+\omega}}\right).\end{aligned}$$

Thus, $\gamma^{(j)}_{x+\omega_*}$ has the desired expression for all $x\geq\omega^*+\lambda^*$. □

Example 5.4.15.

$$A \rightleftarrows 2A \longrightarrow 3A$$

We see that

$$\omega^*=2,\qquad \omega^*=1,\qquad \lambda_*=1,\qquad \lambda^*=2.$$

By Proposition 5.4.13 we can write π_x as a combination of π_1,π_2

$$\pi_x a^2_x=\pi_1\gamma^{(1)}_x+\pi_2\gamma^{(2)}_x$$

where for $j=1,2$

$$\gamma^{(j)}_n=\gamma^{(j)}_{n-2}\frac{b^1_{n-2}}{a^2_{n-2}}-\gamma^{(j)}_{n-1},\qquad n\geq 5$$

By Proposition 5.4.9 we furthermore have $\gamma_1^{(1)} = 1, \gamma_2^{(1)} = 0, \gamma_1^{(2)} = 0, \gamma_2^{(2)} = 1$, $\gamma_3^{(1)} = b_1^1 = \alpha_1$, $\gamma_3^{(2)} = -a_2^2 = 0$, $\gamma_4^{(1)} = -b_1^1 = -\alpha_1$ and $\gamma_4^{(2)} = b_2^1 + a_2^2 = b_2^1 = 2(\alpha_1 + \alpha_2)$. △

Example 5.4.16.

$$A \underset{\alpha_2}{\overset{\alpha_1}{\rightleftarrows}} 3A$$

It follows that

$$\omega_* = 2, \qquad \omega^* = 2, \qquad \lambda_* = 1, \qquad \lambda^* = 2.$$

By Proposition 5.4.13 we can write π_x as a combination of π_1, π_2,

$$\pi_x a_x^2 = \pi_1 \gamma_x^{(1)} + \pi_2 \gamma_x^{(2)},$$

where for $j = 1, 2$

$$\gamma_x^{(j)} = \gamma_{x-3}^{(j)} \frac{b_{x-3}^2}{a_{x-3}^2} + \gamma_{x-2}^{(j)} \frac{b_{x-2}^2}{a_{x-2}^2} - \gamma_{x-1}^{(j)}, \qquad x \geq 6.$$

Using Proposition 5.4.9 we can write the first terms $\gamma_1^{(1)} = \gamma_1^{(2)} = \gamma_2^{(1)} = \gamma_2^{(2)} = 0$, $\gamma_3^{(1)} = b_1^2 = \alpha_1$, $\gamma_3^{(2)} = 0$, $\gamma_4^{(1)} = 0$, $\gamma_4^{(2)} = b_2^2 = 2\alpha_1$, $\gamma_5^{(1)} = b_1^2 \frac{b_3^2}{a_3^2}$, $\gamma_5^{(2)} = 0$. It follows that one may obtain the expressions

$$\gamma_x^{(1)} = \begin{cases} \gamma_{x-2}^{(1)} \frac{b_{x-2}^2}{a_{x-2}^2} & x \text{ odd} \\ 0 & x \text{ even} \end{cases}, \qquad \gamma_x^{(2)} = \begin{cases} 0 & x \text{ odd} \\ \gamma_{x-2}^{(2)} \frac{b_{x-2}^2}{a_{x-2}^2} & x \text{ even} \end{cases}$$

for $x \geq 5$. This in turn yields the representation

$$\pi_x a_x^2 = \begin{cases} \pi_1 \gamma_{x-2}^{(1)} \frac{b_{x-2}^2}{a_{x-2}^2} & x \text{ odd} \\ \pi_2 \gamma_{x-2}^{(2)} \frac{b_{x-2}^2}{a_{x-2}^2} & x \text{ even} \end{cases}$$

for $x \geq 5$. In fact, we can even write it explicitly in this case, for $x \geq 3$,

$$\pi_x = \begin{cases} \pi_1 \frac{b_1^2}{a_x^2} \prod_{i=1}^{(x-3)/2} \frac{b_{x-2i}^2}{a_{x-2i}^2}, & x \text{ odd} \\ \pi_2 \frac{b_2^2}{a_x^2} \prod_{i=1}^{(x-4)/2} \frac{b_{x-2i}^2}{a_{x-2i}^2}, & x \text{ even} \end{cases}$$

Of course, as the network is weakly reversible and has deficiency zero, we know for this particular case that the stationary distribution has a poisson form for each of the irreducible classes (Anderson et al., 2010). △

It follows by Proposition 5.4.13 that we may characterize the entire distribution once $\pi_{\lambda_*}, \ldots, \pi_{\lambda^*}$ are known. In some cases, one way to obtain these is through the probability generating function, $G(z) = \sum_{x \in E} \pi_x z^x$, which for one-species systems may be characterized through an ODE of a particular simple form.

Lemma 5.4.17. *For simplicity, assume $E = \mathbb{N}_0$. The differential equation for the probability generating function, $G(z)$, of the stationary distribution of a one-species reaction network is given by*

$$\sum_{k\in\mathcal{R}} \alpha_k (z^{y_k+\xi_k} - z^{y_k}) G^{(y_k)}(z) = 0,$$

where $G^{(n)}(z)$ denotes the nth derivative of G.

Proof. The master equation (5.4.7) at equilibrium takes the form

$$\sum_{k\in\mathcal{R}} \lambda_k(x-\xi_k)\pi_{x-\xi_k} - \sum_{k\in\mathcal{R}} \lambda_k(x)\pi_x = 0.$$

Multiplying through by z^x and summing over $x \in \mathbb{N}_0$ we obtain

$$\sum_{k\in\mathcal{R}}\sum_{x=0}^{\infty} \lambda_k(x-\xi_k)\pi_{x-\xi_k} z^x - \sum_{k\in\mathcal{R}}\sum_{x=0}^{\infty} \lambda_k(x)\pi_x z^x = 0.$$

Using the form of mass action kinetics, and rearranging terms,

$$\sum_{k\in\mathcal{R}} \alpha_k z^{\xi_k+y_k} \sum_{x=0}^{\infty} x^{\underline{y_k}} \pi_x z^{x-y_k} - \sum_{k\in\mathcal{R}} \alpha_k z^{y_k} \sum_{x=0}^{\infty} x^{\underline{y_k}} \pi_x z^{x-y_k} = 0.$$

As $G(z) = \sum_{x=0}^{\infty} \pi_x z^x$ we may write this as

$$\sum_{k\in\mathcal{R}} \alpha_k z^{\xi_k+y_k} G^{(y_k)}(z) - \sum_{k\in\mathcal{R}} \alpha_k z^{y_k} G^{(y_k)}(z) = 0,$$

whence upon collecting terms yields the desired. □

Example 5.4.18. To Illustrate the approach, consider the following reaction network

$$0 \underset{\alpha_3}{\overset{\alpha_1}{\rightleftarrows}} A \xrightarrow{\alpha_2} 2A$$

Applying Lemma 5.4.17, we can write down the governing differential equation of the probability generating function,

$$\alpha_3 G''(z)(1-z^2) - \alpha_2 G'(z)(z-z^2) - \alpha_1 G(z)(1-z) = 0.$$

Using that $G(1) = 1$, one may obtain the solution

$$G(z) = \frac{{}_1F_1\left(\frac{\alpha_1}{\alpha_2}, \frac{\alpha_2}{\alpha_3}, \frac{\alpha_2(z+1)}{\alpha_3}\right)}{{}_1F_1\left(\frac{\alpha_1}{\alpha_2}, \frac{\alpha_2}{\alpha_3}, \frac{2\alpha_2}{\alpha_3}\right)},$$

where ${}_qF_p$ is the generalized hypergeometric function (Riley et al., 2006). Now, using Proposition 5.4.9 we find

$$a^2_{x+1}\pi_{x+1} = \pi_{x-1}b^1_{x-1} - \pi_x a^2_x, \qquad x \geq 0.$$

Thus,

$$\pi_2 = \pi_0 \frac{b^1_0}{a^2_2} = \pi_0 \frac{\alpha_1}{2\alpha_3}, \qquad \pi_3 = \frac{\pi_0 b^1_0 + \pi_1 b^1_1}{a^2_3} = \pi_0 \frac{-\alpha_1}{6\alpha_3} + \pi_1 \frac{\alpha_2 + \alpha_1}{6\alpha_3},$$

and from Proposition 5.4.13 we get

$$\pi_x = \pi_0 \frac{\gamma^{(0)}_x}{a^2_x} + \pi_1 \frac{\gamma^{(1)}_x}{a^2_x},$$

where $\gamma^{(0)}_2 = \alpha_1, \gamma^{(1)}_2 = 0, \gamma^{(0)}_3 = -\alpha_1, \gamma^{(1)}_3 = \alpha_2 + \alpha_1$ and for $x \geq 4$,

$$\gamma^{(j)}_x = \gamma^{(j)}_{x-2}\frac{b^1_{x-2}}{a^2_{x-2}} - \gamma^{(j)}_{x-1} = \gamma^{(j)}_{x-2}\frac{\alpha_1 + \alpha_2(x-2)}{\alpha_3(x-2)(x-3)} - \gamma^{(j)}_{x-1},$$

for $j = 0, 1$. Thus, knowing π_0 and π_1 is sufficient to characterize the limit distribution, and these are given by $G(0)$ and $G'(0)$ respectively.

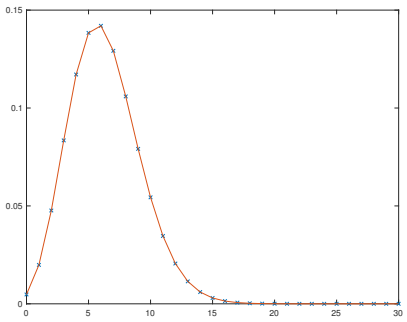

Figure 5.3: The stationary distribution π calculated in two ways. Blue cross is explicit while orange line is through the recursions. They are identical. $\alpha = 10$, $\beta = 40$, $\gamma = 1$.

In this particular example, it is in fact possible to calculate the explicit stationary distribution which is given as follows

$$\pi_x = \frac{1}{x!}\left(\frac{\alpha_2}{\alpha_3}\right)^x \frac{\Gamma(\alpha_2/\alpha_3)}{\Gamma(\alpha_1/\alpha_2)}\frac{\Gamma(x+\alpha_1/\alpha_2)}{\Gamma(x+\alpha_2/\alpha_3)}\frac{{}_1F_1\left(x+\frac{\alpha_1}{\alpha_2}, x+\frac{\alpha_2}{\alpha_3}, \frac{\alpha_2}{\alpha_3}\right)}{{}_1F_1\left(\frac{\alpha_1}{\alpha_2}, \frac{\alpha_2}{\alpha_3}, \frac{2\alpha_2}{\alpha_3}\right)}, \qquad x \geq 0.$$

Note that if $\alpha_1/\alpha_2 = \alpha_2/\alpha_3$ then π has is $Pois(\alpha_2/\alpha_3)$. △

Example 5.4.19. Consider the following reaction network, which may exhibit multistationarity for certain parameter values.

$$0 \underset{\alpha_2}{\overset{\alpha_1}{\rightleftarrows}} A \qquad 2A \underset{\alpha_4}{\overset{\alpha_3}{\rightleftarrows}} 3A$$

The governing differential equation for the probability generating function may be found by Lemma 5.4.17. It is given by

$$\alpha_4 G'''(z) - \alpha_3 G''(z) + \alpha_2 G'(z)\frac{1}{z^2} - \alpha_1 G'(z)\frac{1}{z^2} = 0.$$

One may then find the general (real) solution,

$$G(z) = C \cdot {}_2F_2\left(\left[\frac{-\alpha_3 + \sqrt{\alpha_3^2 - 4\alpha_3\alpha_1}}{2\alpha_3}, \frac{-\alpha_3 - \sqrt{\alpha_3^2 - 4\alpha_3\alpha_1}}{2\alpha_3}\right], \left[\frac{-\alpha_4 + \sqrt{\alpha_4^2 - 4\alpha_4\alpha_2}}{2\alpha_4}, \frac{-\alpha_4 - \sqrt{\alpha_4^2 - 4\alpha_4\alpha_2}}{2\alpha_4}\right], \frac{\alpha_3 z}{\alpha_4}\right),$$

where C is the normalizing constant found by letting $G(1) = 1$. Note that if $\alpha_1 = \alpha_2$ and $\alpha_3 = \alpha_4$ then we obtain $G(z) = e^{z-1}$ which is the probability generating function for $Pois(1)$, and $x^* = 1$ is the only stable equilibrium. An explicit expression for π_x is not readily available. However, note that we immediately have

$$\pi_0 = G(0) = C.$$

Further, from Proposition 5.4.9 and 5.4.13 we find that

$$\pi_1 = \pi_0 \frac{\alpha_1}{\alpha_2}, \qquad \pi_x = \pi_0 \frac{\gamma_x^{(0)}}{a_x^1},$$

where $\gamma_x^{(0)} = \gamma_{x-1}^{(0)} \frac{b_{x-1}^1}{a_{x-1}^1}$ for $x \geq 2$. Thus, we have completely characterized the stationary distribution. As an example, taking the parameter values

$$\alpha_1 = 40, \quad \alpha_2 = 62, \quad \alpha_3 = 23, \quad \alpha_4 = 1,$$

one obtains the following bimodal stationary distribution. △

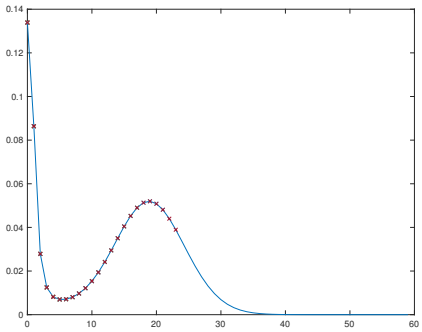

Figure 5.4: The stationary distribution, π, calculated in two ways. Explicitly by differentiating the generating function and by the iterative method. $\alpha_4 = 1, \alpha_3 = 23, \alpha_2 = 62, \alpha_1 = 40$.

5.5 Discussion

We have devised a procedure by which one may split a reaction network, $\mathcal{N}$, into two subnetworks; the endorsed network $\mathcal{N}_E$, and the absorbing network $\mathcal{N}_A$ containing all reactions making absorption possible. Under the further assumption that there exists a stationary measure for the stochastic process associated to $\mathcal{N}_E$ a characterization of observability of quasi-stationary behavior was discussed. Indeed, when the spectral gap of the generator $\tilde{Q}_E$ is large compared to the parameter value ϵ, quasi-stationarity is observed. In this case, the total probabilities $p(t)$ will rapidly approach the center manifold, on which a very slow convergence to the degenerate stationary measure takes place. Exploiting the relationship between this slow manifold and the quasi-stationary distribution, we found a numerical procedure to inductively describe this QSD.

A Numerical Procedure for One-species Reaction Networks. When the reaction network is sufficiently complex, there may not exist an analytical solution to the probability generating function. In such cases, one might find an estimate of the relative sizes of $\pi_{\lambda_*}, \dots \pi_{\lambda^*}$, which in turn determines π_x up to an unknown scaling constant. This constant may then be found numerically by calculating the first N probabilities π_x, $x = \lambda_*, \dots, N$.

Here, we sketch how on may proceed in practice. As π is a probability distribution, it follows that $\lim_{x\to\infty} \pi_x = 0$. Thus, form Proposition 5.4.13 we have the approximate linear system

$$G_x\pi := \begin{pmatrix} \gamma_x^{(\lambda_*)} & \cdots & \gamma_x^{(\lambda^*)} \\ \vdots & & \vdots \\ \gamma_{x+\lambda^*-\lambda_*+1}^{(\lambda_*)} & \cdots & \gamma_{x+\lambda^*-\lambda_*+1}^{(\lambda^*)} \end{pmatrix} \pi \approx 0 \tag{5.5.1}$$

with en error that can be made arbitrarily small by choosing x sufficiently large. These $\lambda^* - \lambda_* + 1$ equations define hyperplanes intersecting approximately in a straight line, giving

the relative sizes of $\pi_{\lambda_*}, \ldots, \pi_{\lambda^*}$. To obtain this line, we look only at the last $\lambda^* - \lambda_*$ equations, which when put in reduced row echelon form yields

$$\begin{pmatrix} 1 & 0 & \ldots & 0 & -r_{\lambda_*} \\ 0 & 1 & \ldots & 0 & -r_{\lambda_*+1} \\ \vdots & \vdots & \ldots & \vdots & \vdots \\ 0 & 0 & \ldots & 1 & -r_{\lambda^*-1} \end{pmatrix} \pi = 0$$

Thus the stationary distribution approximately satisfies

$$\begin{pmatrix} \pi_{\lambda_*} \\ \vdots \\ \pi_{\lambda^*-1} \\ \pi_{\lambda^*} \end{pmatrix} \approx t \begin{pmatrix} r_{\lambda_*} \\ \vdots \\ r_{\lambda^*-1} \\ 1 \end{pmatrix} \tag{5.5.2}$$

for some $t \in \mathbb{R}$. Now, choosing some $\hat{t} > 0$ at random, we find the measure $\hat{\pi}$ through (5.5.2) and apply Proposition 5.4.13, to numerically obtain $\hat{\pi}_x$, $x = \lambda_*, \ldots, N$ for some sufficiently large N. Then

$$t \approx \frac{1}{\sum_{x=\lambda_*}^{N} \hat{\pi}_x}. \tag{5.5.3}$$

with a precision increasing with N.

Example 5.5.4. To illustrate the numerical scheme, consider the following reaction network

$$A \underset{\alpha_3}{\overset{\alpha_1}{\rightleftarrows}} 2A \xrightarrow{\alpha_2} 4A$$

We find the constants

$$\omega_* = 3, \qquad \omega^* = 2, \qquad \lambda_* = 1, \qquad \lambda^* = 3.$$

By Proposition 5.4.13 we can write π_x as a combination of $\pi_1, \ldots, \pi_3$,

$$\pi_x a_x^3 = \pi_1 \gamma_x^{(1)} + \pi_2 \gamma_x^{(2)} + \pi_3 \gamma_x^{(3)},$$

where for $j = 1, 2, 3$

$$\gamma_x^{(j)} = \gamma_{x-4}^{(j)} \frac{b_{x-4}^2}{a_{x-4}^3} + \gamma_{x-3}^{(j)} \frac{b_{x-3}^1 + b_{x-3}^2}{a_{x-3}^3} - \gamma_{x-2}^{(j)} - \gamma_{x-1}^{(j)}, \qquad x \geq 8.$$

Suppose we have the rate constants $\alpha_1 = 10, \alpha_2 = 1, \alpha_3 = 1$. Choosing $x = 300, 301, 302$, we

find the approximative linear system

$$\begin{pmatrix} -2.037 & -41.0209 & 54.9297 \\ 44.1816 & -4.0756 & -45.8043 \\ -42.1489 & 45.0969 & -9.1260 \end{pmatrix} \pi \approx 0$$

These three equations define planes intersecting approximately in a straight line. Looking only at the two last equations, we get using reduced row echelon form

$$\begin{pmatrix} 1 & 0 & -1.15497291951308 \\ 0 & 1 & -1.28183618136202 \end{pmatrix} \pi = 0 \tag{5.5.5}$$

Thus the distribution approximately satisfies

$$\begin{pmatrix} \pi_1 \\ \pi_2 \\ \pi_3 \end{pmatrix} \approx t \begin{pmatrix} 1.15497 \\ 1.28183 \\ 1 \end{pmatrix} \tag{5.5.6}$$

for some $t > 0$. An approximate value for t can now easily be found by truncating the space at a sufficiently large point. △

To provide more insight to the behavior of the sequences $(\gamma_x^{(j)})_{x\in\mathbb{N}}$ for $j = \lambda_*, \dots, \lambda^*$, one may construct the sequences of vectors for $n = 0, \lambda^* - \lambda_*$,

$$\Gamma_x^n = \begin{pmatrix} \gamma_{n+x(\lambda^*-\lambda_*+1)}^{(\lambda_*)} & \cdots & \gamma_{n+x(\lambda^*-\lambda_*+1)}^{(\lambda^*)} \end{pmatrix}' \tag{5.5.7}$$

which by Proposition 5.4.13 may be seen as the normal vectors to the line defining the stationary distribution, π. In general, there need not be a convergence of the individual $\gamma_x^{(j)}$ nor of Γ_x^n or indeed their normed counterparts. However, considering Example 5.5.4, we find three such sequences, whose behavior is depicted on Figure 5.5, and we observe a very fast convergence to a common plane, rendering π well defined.

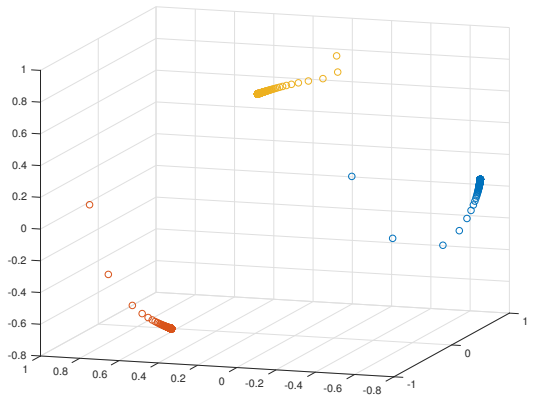

Figure 5.5: $\alpha_1 = 10, \alpha_2 = 10, \alpha_3 = 1$. The three sequences $(\Gamma_x^j / \|\Gamma_x^j\|)_{x\in\mathbb{N}}$

A proof of the convergence of Γ_x^n to a common plane for $x \to \infty$ would not only make the numerical procedure accurate, but also provide another route to prove existence of stationary measures in general one-species reaction networks. There is strong numerical evidence in favor of this convergence.

Future research lies in obtaining good bounds for the spectral gaps, of generators associated with general reaction networks, whereby whole parameter regions for ϵ guaranteeing quasi-stationarity, may be characterized. Furthermore, sufficient conditions suitable for reaction networks guaranteeing an extension of the application of the center-manifold theorem to infinite dimensions, whereby exact analytical results on quasi-stationary distributions on countably infinite state space may be obtsined, is being investigated.

6 Perspective

In this chapter we recall some of the questions raised by the previous chapters and discuss general future directions of study in the intersection of stochastic reaction network theory and quasi-stationarity.

In the paper by Hansen and Wiuf (2018a) we found sufficient conditions for both existence and uniqueness of a quasi-stationary distribution for a given reaction network. The question then becomes whether one can find weaker conditions guaranteeing only existence. Based on the behavior of birth-death processes one should expect that Assumption 2 with $\eta = 0$ would be sufficient for the existence of a QSD, however this problem remains open (van Doorn, 1991; van Doorn and Pollett, 2013). The conditions obtained by Hansen and Wiuf (2018a) are both graphical and dynamical in the sense that they rely both on the reaction vectors and the kinetics. In the spirit of reaction network theory, a more direct graphical condition implying the existence and uniqueness of QSDs is sought for. In particular for the simplest case of mass-action kinetics. However, the lack of complex balance in the reaction networks under consideration poses a great obstacle. Strongly endotactic networks (Gopalkrishnan et al., 2014), which is currently the most well studied class of reaction networks allowing extinction events, do not entail sufficient regularity in the stochastic case. Indeed, some are known to be explosive (Anderson et al., 2018a).

Through the work of Kurtz (1970, 1978) we established a correspondence between the stochastic and the deterministic descriptions of the same underlying reaction network, in terms of the quasi-stationary distributions. Again, even though strongly endotactic reaction networks are both persistent and permanent, which implies that the system possesses all the requirements needed for quasi-stationarity, the need for a strong dissipative system reemerged in the conditions for the results of Hansen and Schreiber (2018). We believe that if D_0 is a global attractor for the semi-flow φ_t then the weak* limit, ν, of the quasi-stationary distributions ν_ϵ is supported on D_0. As we only require the existence and not uniqueness of the QSDs ν_ϵ for each $\epsilon > 0$, the question emerges of finding the weakest possible assumptions on the flow for the results to hold.

A full extension of Freidlin-Wentzell theory to the quasi-stationary domain would allow one, through the quasi-potential, to examine the transition times between positive attractors and describe the extinction time in terms of rate parameters. Furthermore, it would extend the results of Hansen and Schreiber (2018) by answering the question, in the generic case, of which of the positive attractors the limit measure ν would be concentrated on, as well as explain the exponential decay of ν_ϵ away from the attractor, as verified by numerical experiments (Freidlin and Wentzell, 2012; Shwartz and Weiss, 1995). However, this program requires a conditioned or reflecting large deviations principle (Chan, 1998). Work towards this goal has been initiated by the author.

In computing the quasi-potential, connections to WKB theory (Assaf and Meerson, 2017; Bender and Orszag, 1999) seems promising, as it reduces the variational problem to a Hamiltonian system. For this reason it is increasingly used in applications (Nieddu et al., 2017). However, possible issues may arise from divergence of the WKB expansion (Agazzi et al., 2017a), and the theory may not be verified with full mathematical rigor (Chazottes et al., 2017). An examination of the method applied to reaction networks could nevertheless be of great value in pointing towards hypothesis that can later be proven.

Although characterization of the deterministic flow is considerably more tangible than the stochastic counterpart, the problem of determining a Morse decomposition still poses great challenges for higher dimensional reaction networks (Hansen and Schreiber, 2018). It would be interesting to see whether there is something general to be inferred from the graphical structure to the Morse decomposition. In particular, as this relates to the development of new Lyapunov functions, which in itself is one of the hottest topics of modern reaction network theory (Al-Radhawi and Angeli, 2016).

When the limit of QSDs concentrate on a particular attractor, its type appears from numerical examinations to have great qualitative impact on the QSDs corresponding to the system away from equilibrium (Hansen and Schreiber, 2018). Whether one can find general quantitative correspondences or improvements of approximations through such a theory is another topic to investigate. In particular, with the new developments of diffusion approximations in reaction network theory (Bibbona and Sirovich, 2017; Anderson et al., 2018b), one should examine how this fits with the normal approximations to QSDs found in the litterature (van Doorn and Pollett, 2013; Nåsell, 2011). This is closely connected to the question of the behavior near bifurcation points. Consider for example the network

$$\emptyset \overset{\alpha_1}{\leftarrow} S_1 \overset{\alpha_2}{\rightarrow} 2S_1 \underset{\alpha_4}{\overset{\alpha_3}{\rightleftharpoons}} 3S_1$$

Treating the variable $\lambda = 4\alpha_4(\alpha_2 - \alpha_1) + \alpha_3^2$ as the bifurcation parameter, one may verify that there is a supercritical saddle-node bifurcation at $\lambda = 0$. Fix the rates $\alpha_1 = 100.2, \alpha_2 = 0.1, \alpha_4 = 1$. Close to the bifurcation, we see a "ghost" where a new fixed point is emerging. Further, after the bifurcation, the QSD is still monotone, implying that the existence of a stable fixed point is not sufficient for the QSD to have a nearby mode. Moving further away from the bifurcation point the correspondence between the positive stable fixed point and the mode of the QSD is restored. This is illustrated in Figure 6.1

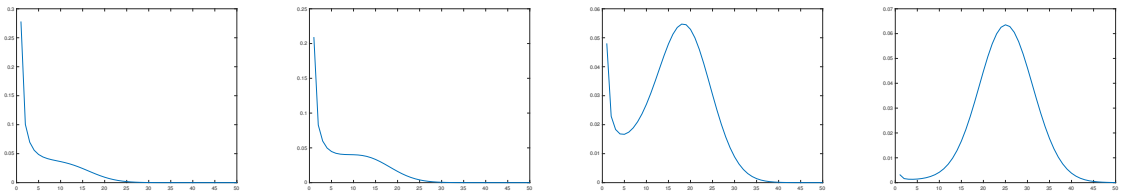

Figure 6.1: $\lambda = -0.4, \lambda = 40.6, \lambda = 224.6, \lambda = 499.6$.

In turn, when considering the classical scaling, the speed of convergence may be dependent on how close the system is to a bifurcation point. General scaling laws would be of great interest.

The standard way of numerically calculating quasi-stationary distributions is by finding the eigenvector corresponding to the second smallest eigenvalue of the generator. This is computationally heavy, $\mathcal{O}(n^3)$, posing problems for large systems. In particular, if the dimension of the reaction network is large, or the scaling ϵ is small, this procedure quickly becomes infeasible. The method proposed by Hansen and Wiuf (2018b) relies on inverting a matrix, hence this is also $\mathcal{O}(n^3)$, although for some regimes computationally more stable (Pollett and Roberts, 1990). A way of simulating the quasi-stationary distribution for reaction networks could be advantageous, allowing one to visualize the weak* convergence to more interesting attractors such as limit cycles.

In the paper by Hansen and Wiuf (2018a) we found that there is an exponential convergence towards the QSD, however, one would like estimates on the convergence rate, γ, since it tells one how likely it is to observe a process in quasi-stationarity. In general, there may be multiple QSDs and hence multiple domains of attraction. In turn, the convergence rate may depend greatly upon which domain of attraction the initial distribution belongs to. Clarifying these domains is thus another case of interest, and only few current results are available (Villemonais, 2015).

When there is just one unique QSD, the spectral gap is closely connected to the convergence rate as demonstrated by Hansen and Wiuf (2018b). Sufficient conditions for the application of the center manifold theorem in infinite dimensions would greatly improve the theoretical understanding of the quasi-stationary distributions from a dynamical viewpoint. In particular, it is the aim of future research to characterize parameter regions in which, for any given reaction network, the system displays quasi-stationarity.

Bibliography

A. Agazzi, A. Dembo, and J.-P. Eckmann. Large Deviations for Markov Jump Models of Chemichal Reaction Networks. *arXiv*, 2017a.

A. Agazzi, A. Dembo, and J.-P. Eckmann. On the Geometry of Chemical Reaction Networks: Lyapunov Function and Large Deviations. *arXiv*, 2017b.

M. A. Al-Radhawi and D. Angeli. New Approach to the Stability of Chemical Reaction Networks: Piecewise Linear in Rates Lyapunov Functions. *IEEE Trans. Automa. Control*, 61:76–89, 2016.

D. F. Anderson. Global Asymptotic Atability for a Class of Nonlinear Chemical Reactions. *SIAM J. Appl. Math.*, 68(5):1464–1476, 2008.

D. F. Anderson and D. Cappelletti. Discrepancies between extinction events and boundary equilibria in reaction networks. *arXiv*, September 2018.

D. F. Anderson and T. G. Kurtz. *Stochastic Analysis of Biochemical Systems*. Mathematical Biosciences Institute Lecture Series. Springer, 2015.

D. F. Anderson, G. Craciun, and T. G. Kurtz. Product-form Stationary Distributions for Deficiency zero Chemical Reaction Networks. *Bulletin of Mathematical Biology*, 72(8): 1947–1970, 2010.

D. F. Anderson, G. A. Enciso, and M. D. Johnston. Stochastic analysis of biochemical reaction networks with absolute concentration robustness. *Journal of the royal society*, 2014.

D. F. Anderson, G. Craciun, M. Gopalkrishnan, and C. Wiuf. Lyapunov functions, stationary distributions, and non-equilibrium potential for reaction networks. *Bulletin of Mathematical Biology*, 77(9):1744–67, September 2015.

D. F. Anderson, D. Cappelletti, J. Kim, and T. Nguyen. Tier structure of strongly endotactic reaction networks. *arXiv*, 2018a.

D. F. Anderson, D. J. Higham, S. C. Leite, and R. J. Williams. Constrained langevin equations and (bio)chemical reaction networks. 2018b.

J. C. Anderson, E. J. Clarke, A. P. Arkin, and C. A. Voigt. Environmentally controlled invasion of cancer cells by engineered bacteria. *Journal of Molecular Biology*, 355(4): 619–627, 2006.

W. J. Anderson. *Continuous-Time Markov Chains: An Applications-Oriented Approach*. Probability and its applications. Springer, 1991.

D. Angeli. A tutorial on chemical reaction network dynamics. *European Journal of Control*, 15(3-4):398–406, 2009.

V. S. Anishchenko, V. Astakhov, A. Neiman, T. Vadivasova, and L. Schimansky-Geier. *Nonlinear Dynamics of Chaotic and Stochastic Systems*. Springer Series in Synergetics. Springer, 2nd edition, 2007.

A. Arkin, J. Ross, and H. H. McAdams. Stochastic kinetic analysis of developmental pathway bifurcation in phage λ-infected escherichia coli cells. *Genetics*, 149:1633–1648, 1998.

M. Assaf and B. Meerson. WKB theory of large deviations in stochastic populations. *J. Phys. A: Math. Theor.*, 50, 2017.

J. Ayala, P. Corbin, K. McConville, W. Kliemann, and J. Peters. Morse decompositions, attractors and chain recurrence. *Proyecciones Journal of Mathematics*, 25(1):79–109, 2006.

A. D. Barbour and P. K. Pollett. Total variation approximation for quasi-equilibrium distributions. *Journal of Applied Probability*, 2010.

N. Barkai and S. Leibler. Biological rhythms: Circadian clocks limited by noise. *Nature*, 403:p.267, 2000.

M. S. Bartlett. *Stochastic Population Models in Ecology and Epidemiology*. Methuen, 1960.

P. M. . M. Beccuti, A. Horvath, T. Jaki, R. Sirovich, and E. Bibbona. A review of the deterministic and diffusion approximations for stochastic chemical reaction networks. *arXiv*, 2018.

C. M. Bender and S. A. Orszag. *Advanced Mathematical Methods for Scientists and Engineers I. Asymptotic Methods and Perturbation Theory*. Springer, 1999.

E. Bibbona and R. Sirovich. Strong approximation of density dependent Markov chains on bounded domains. *arXiv*, 2017.

J. Blanchet, P. Glynn, and S. Zheng. Analysis of a stochastic approximation algorithm for computing quasi-stationary distributions. *Adv. Appl. Prob.*, 48:792–811, 2016.

L. A. Breyer and A. G. Hart. Approximations of quasi-stationary distributions for Markov chains. *Mathematical and Computer Modelling*, 31:69–76, 2000.

M. Brin and G. Stuck. *Introduction to Dynamical Systems*. Cambridge University Press, 2002.

D. Cappelletti. *Limits for Stochastic Reaction Networks*. PhD thesis, University of Copenhagen, October 2015.

D. Cappelletti and C. Wiuf. Elimination of intermediate species in multiscale stochastic reaction networks. *The Annals of Applied Probability*, 26(5):2915–2958, 2016.

J. A. Cavender. Quasistationary distributions of birth-death processes. *Adv. Appl. Prob.*, 10:570–586, 1978.

N. Champagnat and D. Villemonais. Quasi-stationary distribution for multi-dimensional birth and death processes conditioned to survival of all coordinates. *arXiv*, 2015.

N. Champagnat and D. Villemonais. Exponential convergence to quasi-stationary distribution and q-process. *Probability Theory and Related Fields*, 164:243–286, 2016.

N. Champagnat and D. Villemonais. Lyapunov criteria for uniform convergence of conditional distributions of absorbed Markov processes. *arXiv*, April 2017.

N. Champagnat and D. Villemonais. General criteria for the study of quasi-stationarity. *arXiv*, 2018.

T. Chan. Large deviations and quasi-stationarity for density-dependent birth-death processes. *J. Austral. Math. Soc. Ser. B*, 40:238–256, 1998.

J. R. Chazottes, P. Collet, and S. Méléard. On time scales and quasi-stationary distributions for multitype birth-and-death processes. *arXiv*, 2017.

P. Childs and J. P. Keener. Slow manifold reduction of a stochastic chemical reaction: Exploring keizers paradox. *Discrete and Continuous Dynamical Systems Series B*, 17(6): 1775–1794, 2012.

G. Cohen, T. Mora, and O. Moreno, editors. *Applied Algebra, Algebraic Algorithms and Error-Correcting Codes*. Number 673 in Lecture Notes in Computer Science. Springer, 1993.

P. Collet, S. Martinez, and J. S. Martin. *Quasi-stationary distributions. Markov Chains, Diffusions and Dynamical Systems*. Probability and its applications. Springer, 2013.

F. Colonius and W. Kliemann. *Dynamical Systems and Linear Algebra*, volume 158 of *Graduate Studies in Mathematics*. American Mathematical Society, 2014.

C. Conley. *Isolated Invariant Sets and the Morse Index*. Number 38 in Conference Board of the Mathematical Sciences. American Mathematical Society, 1978.

M. Cook, D. Soloveichik, E. Winfree, and J. Bruck. *Programmability of Chemical Reaction Networks, Chapter in Algorithmic Bioprocesses*. Springer, 2009.

S. L. Cotter. Constrained approximation of effective generators for multiscale stochastic reaction networks and application to conditioned path sampling. *Journal of Computational Physics*, 323:265–282, 2016.

G. Craciun. Toric differential inclusions and a proof of the global attractor conjecture. *arXiv*, 2015.

G. Craciun, F. Nazarov, and C. Pantea. Persistence and permanence of mass-action and power-law dynamical systems. *SIAM J. Appl. Math.*, 73(1):305–329, 2013.

D. Dadush. *Lattices, Convexity & Algorithms*. Lecture Notes, NYU, 2013.

A. Dembo and O. Zeitouni. *Large Deviations Techniques and Applications*. Number 38 in Applications of Mathematics. Springer, 2nd edition, 1998.

P. Dittrich and L. Winter. Chemical organizations in a toy model of the political system. *Adv. Complex Syst.*, 11:609–627, 2008.

P. Dupuis and R. S. Ellis. *A Weak Convergence Approach to the Theory of Large Deviations*. Wiley Series in Probability and Statistics. Wiley, 2011.

P. Dupuis, K. Ramanan, and W. Wu. Large deviation principle for finite-state mean field interacting particle systems. *arXiv*, 2016.

R. Durrett. *Probability: Theory and Examples*. Duxbury Press, 2nd edition, 1996.

M. B. Elowitz, A. J. Levine, E. D. Siggia, and P. S. Swain. Stochastic gene expression in a single cell. *Science*, 297(5584):1183, 2002.

F. E. Emery, editor. *Systems Thinking*. Penguin, 1969.

G. A. Enciso. Transient absolute robustness in stochastic biochemical networks. *Journal of the royal society*, 13, 2016.

R. Erban, S. J. Chapman, and P. K. Maini. A practical guide to stochastic simulations of reaction-diffusion processes. *arXiv*, 2007.

P. Erdi and J. Toth. *Mathematical models of chemical reactions*. Manchester University Press, 1989.

S. N. Ethier and T. G. Kurtz. *Markov Processes*. Wiley Series in Probability and Mathematical Statistics. Wiley, 1986.

M. Faure and S. J. Schreiber. Quasi-stationary distributions for randomly perturbed systems. *The Annals of Applied Probability*, 24(2):553–598, 2014.

M. Feinberg. Complex balancing in general kinetic systems. *Arch. Ration. Mech. Anal.*, 49: 187–194, 1972.

M. Feinberg. Lectures on chemical reaction networks. delivered at the Mathematics Research Center, Univ. Wisc.-Madison., 1979.

M. Feinberg. The existence and uniqueness of steady states for a class of chemical reaction networks. *Arch. Ration. Mech. Anal.*, 132:311–370, 1995.

M. Feinberg. *The Princeton Companion to Applied Mathematics*. Princeton University Press, 2015.

P. A. Ferrari, H. Kesten, S. Martinez, and S. Picco. Existence of quasistationary distributions. a renewal dynamical approach. *Annals of Probability*, 23:501–521, 1995.

M. I. Freidlin and A. D. Wentzell. *Random Perturbations of Dynamical Systems*. Number 260 in Grundlehren der mathematischen Wissenschafften. Springer, 3rd edition, 2012.

D. Gillespie. Exact stochastic simulation of coupled chemical reactions. *Journal of Physical Chemistry*, 81(25):2340–2361, 1977.

M. Gopalkrishnan, E. Miller, and A. Shiu. A geometric approach to the global attractor conjecture. *SIAM J. Appl. Dyn. Syst.*, 13(2):758–797, 2014.

B. L. Granovsky and A. I. Zeifman. The decay function of nonhomogeneous birth-death processes with application to mean-field models. *Stoch. Process. Appl.*, 72:105–120, 1997.

A. Griffin. *Quasi-Stationary Distributions for Evolving Epidemic Models: Simulation and Characterisation*. PhD thesis, Mathematics Institute, The University of Warwick, 2016.

A. Griffin, P. A. Jenkins, G. O. Roberts, and S. E. F. Spencer. Simulation from quasi-stationary distributions on reducible state spaces. *Advances in Applied Probability*, 49(3): 960–980, 2017.

P. Groisman and M. Jonckheere. Simulation of quasi-stationary distributions on countable spaces. *arXiv*, 2012.

J. Gunawardena. Chemical reaction network theory for in-silico biologists, 2003. URL `http://vcp.med.harvard.edu/papers/crnt.pdf`.

A. Gupta and M. Khammash. Determining the long-term behavior of cell populations: A new procedure for detecting ergodicity in large stochastic reaction networks. *arXiv*, 2013.

A. Gupta, C. Briat, and M. Khammash. A scalable computational framework for establishing long-term behavior of stochastic reaction networks. *PLOS Computational Biology*, 10(6), 2014.

H. Haken. *Synergetics; An Introduction*. Springer, 1983.

M. C. Hansen and S. J. Schreiber. Quasi-stationary distributions for randomly perturbed reaction networks. *Unpublished*, 2018.

M. C. Hansen and C. Wiuf. Existence of a unique quasi-stationary distribution for stochastic reaction networks. *arXiv*, 2018a.

M. C. Hansen and C. Wiuf. A description of quasi-stationary distributions in reaction networks using a slow manifold reduction. *Unpublished*, 2018b.

M. W. Hirsch. Systems of differential equations that are competitive or cooperative. i: Limit sets. *SIAM J. Math. Anal.*, 13(2), 1982.

M. W. Hirsch. Systems of differential equations that are competitive or cooperative. iv: Structural stability in three-dimensional systems. *SIAM J. Math. Anal.*, 21(5):1225–1234, 1990.

J. H. Holland. *Complexity; A very short introduction*. Oxford University Press, 2014.

F. Horn. Necessary and sufficient conditions for complex balancing in chemical kinetics. *Arch. Ration. Mech. Anal.*, 49:172–186, 1972.

F. J. M. Horn and R. Jackson. General mass action kinetics. *Arch. Ration. Mech. Anal.*, 47: 81–116, 1972.

Y. Ilyashenko. Centennial history of Hilbert's 16th problem. *Bull. of the American Mathematical Society*, 39(3):301–354, 2002.

F. Jacobs and S. J. Schreiber. Random perturbations of dynamical systems with absorbing states. *SIAM J. Appl. Dyn. Syst.*, 5:293–312, 2006.

T. Jahnke and W. Huisinga. Solving the chemical master equation for monomolecular reaction systems analytically. *J. Math. Biol.*, 54(1):1–26, 2007.

M. D. Johnston, C. Pantea, and P. Donnell. A computational approach to persistence, permanence, and endotacticity of biochemical reaction systems. *Journal of Mathematical Biology*, 72(1-2):467–498, 2016.

M. D. Johnston, D. F. Anderson, G. Craciun, and R. Brijder. Conditions for extinction events in chemical reaction networks with discrete state spaces. *Journal of Mathematical Biology*, pages 1–24, 2017.

B. Joshi and A. Shiu. A survey of methods for deciding whether a reaction network is multistationary. *Math. Model. Nat. Phenom.*, 10(5):47–67, 2015.

O. Kallenberg. *Foundations of Modern Probability*. Probability and its applications. Springer, 2nd edition, 2001.

H.-W. Kang and T. G. Kurtz. Separation of time-scales and model reduction for stochastic reaction networks. *The Annals of Applied Probability*, 23(2):529–583, 2013.

S. Karlin and J. L. McGregor. The differential equations of birth-death processes and the stieltjes moment problem. *Trans. Amer. Math. Soc.*, 85:489–546, 1957.

A. Katok and B. Hasselblatt. *Introduction to the Modern Theory of Dynamical Systems*. Cambridge University Press, 1995.

M. Ke, Z. Fang, and C. Gao. Generalized Gibbs' free energy for stability analysis of chemical reaction networks. *arXiv*, August 2017.

J. Keizer. *Statistical Thermodynamics of Nonequilibrium Processes*. Springer, 1987.

W. O. Kermack and A. G. McKendrick. A contribution to the mathematical theory of epidemics. *Proceedings of the Royal Society A*, 115(772), 1927.

Y. Kifer. *Random Perturbations of Dynamical Systems*, volume 16 of *Progress in Probability and Statistics*. Birkhäuser, 1988.

R. K. Kittappa. A representation of the solution of the nth order linear difference equation with variable coefficients. *Linear Algebra and its applications*, 193:211–222, 1993.

A. Klenke. *Probability Theory, A Comprehensive Course*. Universitext. Springer, 2nd edition, 2014.

C. Kuehn. *Multiple Time Scale Dynamics*. Applied Mathematical Sciences. Springer, 2015.

T. G. Kurtz. Solutions of ordinary differential equations as limits of pure jump Markov processes. *J. Appl. Prob.*, 7:49–58, 1970.

T. G. Kurtz. The relationship between stochastic and deterministic models for chemical reactions. *The journal of chemical physics*, 57(7):2976–2978, 1972.

T. G. Kurtz. Strong approximation theorems for density dependent Markov chains. *Stochastic processes and their applications*, 6(3):223–240, 1978.

A. J. Lotka. *Elements of Physical Biology*. Williams and Witkins, Baltimore, 1925.

P. L. Luisi. *The Emergence of Life; From Chemical Origins to Synthetic Biology*. Cambridge University Press, 2006.

B. Marmet. Quasi-stationary distributions for stochastic approximation algorithms with constant step size. *arXiv*, 2013.

S. Martinez, J. S. Martin, and D. Villemonais. Existence and uniqueness of a quasi-stationary distribution for markov processes with fast return from infinity. *J. Appl. Prob.*, 51:756–768, 2014.

S. Méléard and D. Villemonais. Quasi-stationary distributions and population processes. *Probability Surveys*, 9:340–410, 2012.

S. P. Meyn and R. L. Tweedie. Stability of markovian processes III: Foster-Lyapunov criteria for continuous-time processes. *Advances in Applied Probability*, 25(3):518–548, 1993.

S. P. Meyn and R. L. Tweedie. *Markov Chains and Stochastic Stability*. Cambridge University Press, 2nd edition, 2009.

D. S. Mitrinovic, J. Pecaric, and A. M. Fink. *Inequalities Involving Functions and Their Integrals and Derivatives*. Mathematics and Its Applications (East European Series). Springer, 1991.

S. Müller and G. Regensburger. Generalized mass action systems: Complex balancing equilibria and sign vectors of the stoichiometric and kinetic-order subspaces. *SIAM J. Appl. Math.*, 72:1926–1947, 2012.

B. Munsky, B. Trinh, and M. Khammash. Listening to the noise: random fluctuations reveal gene network parameters. *Molecular Systems Biology*, 5(318), 2009.

I. Nåsell. *Extinction and Quasi-stationarity in the Stochastic logistic SIS model*. Lecture Notes in Mathematics. Springer, 2011.

M. E. J. Newman. *Networks; An introduction*. Oxford University Press, 2010.

G. T. Nieddu, L. Billings, J. H. Kaufman, E. Forgoston, and S. Bianco. Extinction pathways and outbreak vulnerability in a stochastic Ebola model. *J. R. Soc. Interface*, 14, 2017.

J. R. Norris. *Markov Chains*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2009.

B. Øksendal. *Stochastic Differetial Equations*. Springer, 6. edition, 2010.

P. D. O'neill. Constructing population processes with specified quasi-stationary distributions. *Stochastic Models*, 23(3):439–449, 2007.

K. Pakdaman, M. Thieullen, and G. Wainrib. Diffusion approximation of birth-death processes: Comparison in terms of large deviations and exit points. *Statistics and Probability Letters*, 80:1121–1127, 2010.

M. Patrao. Morse decomposition of semiflows on topological spaces. *Journal of Dynamics and Differential Equations*, 19(1), 2007.

L. Paulevé, G. Craciun, and H. Koeppl. Dynamical properties of discrete reaction networks. *J. Math. Biol.*, 69:55–72, 2014.

D. Pollard. Total variation distance between measures, 2005. URL `http://www.stat.yale.edu/~pollard/Courses/607.spring05/handouts/Totalvar.pdf`.

P. K. Pollett. *Diffusion approximations for ecological models*. 2001.

P. K. Pollett. Quasi-stationary distributions: A bibliography, 2015.

P. K. Pollett and A. J. Roberts. A description of the long-term behaviour of absorbing continuous-time markov chains using a centre manifold. *Advances in Applied Probability*, 22(1):111–128, 1990.

I. Prigogine and R. Lefever. Symmetry breaking instabilities in dissipative systems. *The journal of chemical physics*, 48:1695–1700, 1968.

H. Qian. Nonlinear stochastic dynamics of mesoscopic homogeneous biochemical reaction systems - an analytical theory. *Nonlinearity*, 24, 2011.

E. Renshaw. *Stochastic Population Processes; Analysis, Approximations and Simulations*. Oxford University Press, 2011.

K. F. Riley, M. P. Hobson, and S. J. Bence. *Mathematical Methods for Physics and Engineering*. Cambridge University Press, 3rd edition, 2006.

L. C. G. Rogers and D. Williams. *Diffusions, Markov Processes and Martingales*, volume 1 of *Cambridge Mathematical Library*. Cambridge University Press, 2000.

S. J. Schreiber. Criteria for C^r robust permanence. *Journal of Differential Equations*, 162: 400–426, 2000.

S. J. Schreiber. Coexistence in the face of uncertainty. *arXiv*, 2016.

M. Shakil, H. A. Wahab, M. Naeem, S. Bhatti, and M. Shahzad. The modeling of predator-prey interactions. *Network Biology*, 5(2):71–81, 2015.

S. W. Shin, C. Thachuk, and E. Winfree. Verifying chemical reaction network implementations: A pathway decomposition approach. *Theoretical Computer Science*, 2017.

G. Shinar and M. Feinberg. Structural sources of robustness in biochemical reaction networks. *Science*, 327(5971):1389–1391, 2010.

A. Shwartz and A. Weiss. *Large Deviations for Performance Analysis.* Stochastic Modeling Series. Chapman & Hall, 1995.

M. E. Taylor. *Partial Differential Equations 1.* Springer, 2011.

A. M. Turing. The chemical basis of morphogenesis. *Philosophical Transactions of the Royal Society of London*, 237(641):37–72, 1952.

E. A. van Doorn. Quasi-stationary distributions and convergence to quasi-stationarity of birth-death processes. *Advances in Applied Probability*, 23:683–700, 1991.

E. A. van Doorn and P. K. Pollett. Quasi-stationary distributions for reducible absorbing Markov chains in discrete time. *Markov Processes and Relat. Fields*, 15:191–204, 2009.

E. A. van Doorn and P. K. Pollett. Quasi-stationary distributions for descrete-state models. *European Journal of Operational Research*, 230:1–14, 2013.

J. H. van Lint and R. M. Wilson. *A course in Combinatorics.* Cambridge University Press, 2nd edition, 2006.

M. Vellela and H. Qian. A quasistationary analysis of a stochastic chemical reaction: Keizer's paradox. *Bulletin of Mathematical Biology*, 69:1727–1746, 2007.

T. Veloz and P. Razeto-Barry. Reaction networks as a language for systemic modeling: Fundamentals and examples. *Systems*, 5(11), 2017.

T. Veloz, P. Razeto-Barry, P. Dittrich, and A. Fajardo. Reaction networks and evolutionary game theory. *Journal of Mathematical Biology*, 68:181–206, 2014.

D. Vere-Jones. Some limit theorems for evanescent processes. *Austral. J. Statist.*, 11:67–78, 1969.

D. Villemonais. Minimal quasi-stationary distributions. *arXiv*, 2015.

V. Volterra. Fluctuations in the abundance of a species considered mathematically. *Nature*, 118:558–560, 1926.

P. Walters. *An Introduction to Ergodic Theory.* Number 79 in Graduate texts in Mathematics. Springer, 1982.

S. Watanabe. On discontinuous additive functionals and Lévy measures of a Markov process. *Japan J. Math.*, 34:53–70, 1964.

W. Weidlich and G. Haag. *Concepts and models of a quantitative sociology: the dynamics of interacting populations.* Springer, 2012.

S. Wiggins. *Introduction to Applied Nonlinear Dynamical Systems and Chaos*. Texts in Applied Mathematics. Springer, 2nd edition, 2003.

D. J. Wilkinson. *Stochastic Modelling for Systems Biology*. Chapman & Hall/CRC, 2nd edition, 2012.

C. Wiuf and E. Feliu. Power-law kinetics and determinant criteria for the preclusion of multistationarity in networks of interacting species. *SIAM J. Appl. Dyn. Syst.*, 12(4): 1685–1721, 2013.

H.-R. Zhu and H. L. Smith. Stable periodic orbits for a class of three dimensional competitive systems. *Journal of Differential Equations*, 110:143–156, 1994.

Index