Centrum voor Wiskunde en Informatica

REPORT *RAPPORT*

On the structure of the space of product-form models

N. Bayer, R.J. Boucherie

Probability, Networks and Algorithms (PNA)

PNA-R9806 July 31, 1998

Report PNA-R9806
ISSN 1386-3711

CWI
P.O. Box 94079
1090 GB Amsterdam
The Netherlands

On the Structure of the Space of Product-Form Models

Nimrod Bayer
Department of Electrical and Computer Engineering
Ben-Gurion University of the Negev, Beer-Sheva 84105, Israel

Richard J. Boucherie
Department of Econometrics
Universiteit van Amsterdam, Roetersstraat 11, 1018 WB Amsterdam, The Netherlands
and
CWI
P.O. Box 94079, 1090 GB Amsterdam, The Netherlands

ABSTRACT

This paper deals with Markovian models which are defined on a finite-dimensional discrete state space, and possess a stationary state distribution of a product-form. We view the space of such models as a mathematical object, and explore its structure. We focus on models on an orthant $\mathcal{Z}_+^n$, which are homogeneous within subsets of $\mathcal{Z}_+^n$ called walls, and permit only state transitions whose $\| \ \|_\infty$-length is 1. The main finding is that the space of such models exhibits a *decoupling principle*: In order to produce a given product-form distribution, the transition rates on distinct walls of the same dimension can be selected without mutual interference. The selection space of distinct models which share a given product-form state distribution is accounted for.

In addition, we consider models which are homogeneous throughout a finite-dimensional grid $\mathcal{Z}^n$, now without a fixed restriction on the length of the transitions. We characterize the collection of product-form measures which are invariant for a model of this kind. For such models with bounded transitions we prove, using Choquet's theorem, that the only possible invariant measures are product-form measures and their combinations.

1991 Mathematics Subject Classification: 60K25, 90B22.

Keywords and Phrases: Markov chain, product-form, model design.

Note: Nimrod Bayer started being involved in this work during a stay at CWI Amsterdam, supported by the European grant BRA-QMIPS of CEC DG XIII.
The research of R. J. Boucherie has been made possible by a fellowship from the Royal Netherlands Academy of Arts and Sciences.

1. Introduction

The Setting

Product-form models on a discrete state space have played a salient role in the study of stochastic networks since its beginning. See the classical book by Whittle (1986) [8], and the more recent (1993) book by van Dijk [2]. For a now classical contribution not covered in these treatises—the discovery of product-form networks with negative customers, see Gelenbe (1991) [3]. Results such as Gelenbe's have re-established theoretical interest in product-form models, and this line of research has led to various results regarding models with modified boundary rates. See, e.g., Miyazawa and Taylor (1997) [4], and references therein, for recent contributions in this area. Product-form models emerge also in a continuous state space setting. See Williams (1995) [9, Theorem 3.5] for a treatment of questions analogous to those we are about to treat. While new classes of discrete product-form stochastic networks have continued to be discovered, it has not been clear whether these are rare contingencies, or representatives of a rich space of product-form models. The prospects for systematic product-form model design depend on this question.

In this paper we investigate the space of product-form models from an abstract point of view. The word *model* designates here some setup of parameters. The models considered do not carry any system or network semantics. Rather, they directly determine the transition rates of a continuous-time Markov chain. (Such a collection of rates is called *generator*, or *Q-matrix*). Models will be rendered in this paper as a functions, as function arrays, and as finite vectors.

We restrict ourselves to models possessing the property of *space-homogeneity*, or simply *homogeneity*. A model possesses this property when parallel transitions within its state space $\mathcal{S}$, or within subsets of $\mathcal{S}$ referred to as *walls*, necessarily have the same rate. (The Markov chain associated with such a model is sometimes called *random walk*. Random walks arise, for example, as the queue processes in networks of single server queues). The product-form distributions considered here are geometric, i.e. of the type

$$\pi(\vec{\mathbf{a}}) = c\prod_{i=1}^{n} q_i^{a_i}, \qquad \vec{\mathbf{a}} = \langle a_1, \dots, a_n \rangle \in \mathcal{S}, \tag{1.1}$$

where c and the q_i are constants.

Findings Summary

The space $\mathcal{M}_n$ of models of homogeneous continuous-time Markov chains on the n-dimensional integer grid $\{\dots, -1, 0, 1, \dots\}^n$ is considered first. While the notion of stationary state distribution is not very relevant for a model $\varphi \in \mathcal{M}_n$, there may exist infinite measures which are invariant to the transition operators associated with φ. Such measures are said here, simply, to be invariant for φ. Only product-form measures which are invariant for φ can serve, when appropriately trimmed and normalized, as candidates for the stationary state distribution of a product-form model on the orthant, with interior transition rates coinciding with those of φ. The question arises: Which models in $\mathcal{M}_n$ have product-form invariant measures? The following answer is oriented to the case $n \geq 2$: Every $\varphi \in \mathcal{M}_n$ whose drift is finite but nonzero has quite a set of invariant product-form measures; the corresponding set of vectors $\vec{\mathbf{q}} = \langle q_1, \ldots, q_n \rangle$ (playing a role as in Eq. (1.1)) is a smooth $(n-1)$-dimensional manifold Q_φ, which is the boundary of a bounded and convex set in $(0, \infty)^n$. The product-form measures corresponding to Q_φ are proved to be the fundamental invariant measures for φ with bounded transitions: Any other measure invariant for φ cannot be but a mixture of these product-form measures.

Another space, $\mathbb{M}_n$, of models on the nonnegative orthant of the n-dimensional grid is then considered. Its definition relies on partitioning the orthant into walls. The wall to which some state $\vec{\mathbf{a}} = \langle a_1, \ldots, a_n \rangle$ belongs is determined by the coordinates i where $a_i = 0$. All walls but one are pieces of boundary. The exception is the whole interior, which is also regarded as a "wall" for the sake of uniformity. The walls are attributed with dimensions, which range from zero to n. There is a single wall of dimension zero: It contains the single point $\langle 0, \ldots, 0 \rangle$, and is referred to as the *corner*. The homogeneity property postulated for $\mathbb{M}_n$ is weaker than for $\mathcal{M}_n$: Homogeneity prevails within each wall, but transition rates assigned to parallel transitions belonging to different walls may differ. Apart of homogeneity, the models in $\mathbb{M}_n$ comply with a further restrictive assumption: A transition from $\vec{\mathbf{a}} = \langle a_1, \ldots, a_n \rangle$ to $\vec{\mathbf{b}} = \langle b_1, \ldots, b_n \rangle$ is possible only if this transition is "short", in the sense that $|a_i - b_i| \leq 1$ holds for every i. Both the homogeneity assumption and the assumption of short transitions are satisfied by many actual stochastic network models with single servers. The former assumption (homogeneity) seems more essential than the latter (short transitions) in facilitating the analysis. The question regarding the potential for relaxing these assumptions is open.

It is known that there exist models in $\mathbb{M}_n$ with a stationary state distribution not of a product-form. But with current knowledge, these distributions are usually hard to characterize, and their behavior is obscure. This is the reason for interest in those models which do enjoy a product-form distribution. Let $\mathbb{M}_n$'s subspace of product-form models be denoted as $\mathbb{P}_n$. Knowledge about the structure of $\mathbb{P}_n$ may enable the systematic selection of models in order to fit them to systems, or in other words—the design of models.

This paper does not provide design procedures for actual situations. But the structure of $\mathbb{P}_n$, as emerges from the analysis, is naturally described through a model selection procedure. The procedure may start in the selection of the interior transition rates, which is tantamount to a selection of a model φ with short transitions from $\mathcal{M}_n$. The next step may then be a selection of a vector $\vec{\mathbf{q}} = \langle q_1, \ldots, q_n \rangle$ from Q_φ. An alternative way is to start from an arbitrary $\vec{\mathbf{q}}$ with $0 < q_i < 1$, $i = 1, \ldots, n$. (We exclude $q_i = 1$ from reasons explained later). For every such $\vec{\mathbf{q}}$, let $\mathbb{P}_{n,\vec{\mathbf{q}}}$ denote the relevant subset of $\mathbb{P}_n$. To select a model from $\mathbb{P}_{n,\vec{\mathbf{q}}}$, perform the following process. First, select the interior transition rates. This selection is subject to a single linear constraint, so the number of degrees of freedom is one less than the number of variables. Next, select the transition rates within the walls of dimension $n-1$. These selections are decoupled from each other, i.e., they are not coupled by any joint constraint. The same rule regarding the number of degrees of freedom applies again, for each of these walls. The procedure so proceeds to walls of ever smaller dimension, with the same rules holding, until the selection is exhausted. While the selection for a wall is decoupled from other walls of the same dimension, it depends on certain earlier selections for walls of higher dimensions. What makes this procedure valid is the fact that the selections so taken are guaranteed not to interfere with each other at the corner: The corner constraint is shown to be redundant. Thus we have the *decoupling principle*, and along with it—a wealth of product-form models.

Organization of the Paper
The rest of this paper includes a section of preliminaries, two sections of findings, and a section of proofs. The section of preliminaries (Section 2) states some conventions, and introduces model spaces and related objects. The two sections of findings (Section 3 and 4), are dedicated to $\mathcal{M}_n$ and to $\mathbb{M}_n$, respectively. All proofs are concentrated in Section 5.

2. Preliminaries

2.1. General Conventions

$\mathcal{R}$ denotes the real numbers. $\mathcal{R}_+$ denotes the nonnegative reals. $\mathcal{Z}$ denotes the integers. $\mathcal{Z}_+$ and $\mathcal{N}$ denote the nonnegative and positive integers, respectively. Define $\mathcal{B} \triangleq \{0,1\}$ and $\mathcal{T} \triangleq \{-1,0,1\}$. The symbols $\vec{\mathbf{1}}$ and $\vec{\mathbf{0}}$ stand for vectors of all 1's and all 0's, with their length implied by the context. Suppose that $\vec{\mathbf{x}} = \langle x_1, \ldots, x_k\rangle$ and $\vec{\mathbf{y}} = \langle y_1, \ldots, y_k\rangle$ are two vectors, and that A is a set of vectors of the same length. Define $\vec{\mathbf{x}}\vec{\mathbf{y}}$ to be the vector $\langle x_1y_1, \ldots, x_ky_k\rangle$; define $\vec{\mathbf{x}}A$ to be the set $\{\vec{\mathbf{x}}\vec{\mathbf{y}}/\,\vec{\mathbf{y}} \in A\}$; and define $\vec{\mathbf{x}}^{\vec{\mathbf{y}}}$ to be the scalar $\prod_{i=1}^{k} x_i^{y_i}$. Interpret $|\vec{\mathbf{x}}|$ as $\langle |x_1|, \ldots, |x_k|\rangle$. The norms $\|\cdot\|_1$ and $\|\cdot\|_\infty$ are defined as usual:

$$\|\vec{\mathbf{x}}\|_1 \triangleq \sum_{i=1}^{k} |x_i|, \qquad \|\vec{\mathbf{x}}\|_\infty \triangleq \max_{i=1,\ldots,k} |x_i|. \tag{2.1}$$

Relations such as $\vec{\mathbf{x}} \le \vec{\mathbf{y}}$ or $\vec{\mathbf{x}} < \vec{\mathbf{y}}$ are interpreted in the componentwise sense. A nonnegative vector $\vec{\mathbf{x}}$ is said here to be *majorized* by another nonnegative vector $\vec{\mathbf{y}}$ if $\vec{\mathbf{x}} \le \vec{\mathbf{y}}$ as well as $\|\vec{\mathbf{x}}\|_1 < \|\vec{\mathbf{y}}\|_1$ hold; we write $\vec{\mathbf{x}} \prec \vec{\mathbf{y}}$.

2.2. Models and Related Objects

The objects with which we deal in this paper are state spaces and subsets thereof, classes of state transitions, models and model spaces, and state space measures.

State Spaces and Subsets Thereof Our state space $\mathcal{S}$ will either be an n-dimensional grid $\mathcal{Z}^n$, or its orthant $\mathcal{Z}_+^n$. For subsets of $\mathcal{S}$ we use the following notion of dimension, applying in this discrete context only.

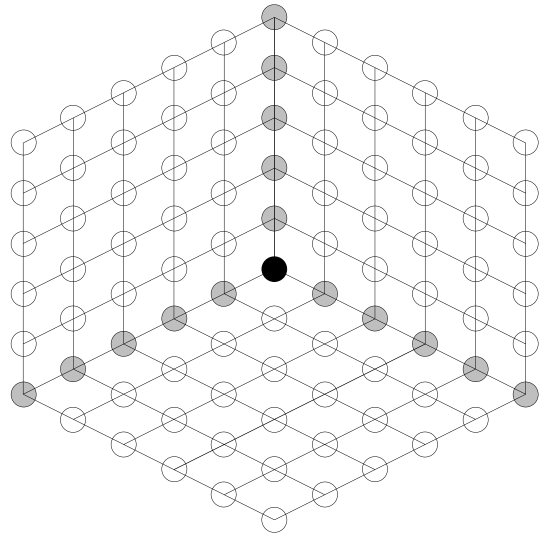

Figure 1: The boundary walls of $\mathcal{Z}_+^3$. The corner, the one-dimensional walls, and the two-dimensional walls are represented by the dark ball, the grey balls, and the white balls, respectively.

Definition 2.1. *The dimension of $A \subset \mathcal{Z}^n$ is less than or equal to k if there exist some $\vec{\mathbf{x}}_1, \ldots, \vec{\mathbf{x}}_k \in \mathcal{Z}^n$, and an $\vec{\mathbf{a}} \in A$, such that every $\vec{\mathbf{b}} \in A$ admits a representation $\vec{\mathbf{b}} = \vec{\mathbf{a}} + \sum_{i=1}^{k} m_i \vec{\mathbf{x}}_i$ with $m_1, \ldots, m_k \in \mathcal{Z}$.*

We now introduce the partitioning of $\mathcal{Z}_+^n$ into walls. Define

$$\mathcal{W}_{n,\vec{\mathbf{w}}} \stackrel{\triangle}{=} \vec{\mathbf{w}}\mathcal{N}^n, \qquad \vec{\mathbf{w}} \in \mathcal{B}^n, \tag{2.2}$$

recalling the $\vec{\mathbf{x}}A$ convention from the previous subsection.

Observation 2.1. *$\mathcal{Z}_+^n$ is the disjoint union $\bigcup_{\vec{\mathbf{w}} \in \mathcal{B}^n} \mathcal{W}_{n,\vec{\mathbf{w}}}$.*

The reason for referring to the $\mathcal{W}_{n,\vec{\mathbf{w}}}$ as the walls of $\mathcal{Z}_+^n$ is that all of them except one are parts of its boundary. The exception is $\mathcal{W}_{n,\vec{\mathbf{1}}} = \mathcal{N}^n$, which constitutes the interior of $\mathcal{Z}_+^n$. The dimension of $\mathcal{W}_{n,\vec{\mathbf{w}}}$ is clearly $\|\vec{\mathbf{w}}\|_1$. The walls thus have various dimensions. For example, the walls of $\mathcal{Z}_+^3$ include the zero-dimensional wall $\mathcal{W}_{3,\langle 0,0,0\rangle}$ (the corner), the one-dimensional walls $\mathcal{W}_{3,\langle 1,0,0\rangle}$, $\mathcal{W}_{3,\langle 0,1,0\rangle}$, and $\mathcal{W}_{3,\langle 0,0,1\rangle}$, the two-dimensional walls $\mathcal{W}_{3,\langle 0,1,1\rangle}$, $\mathcal{W}_{3,\langle 1,0,1\rangle}$, and $\mathcal{W}_{3,\langle 1,1,0\rangle}$, and the three-dimensional interior "wall" $\mathcal{W}_{3,\langle 1,1,1\rangle}$. See Figure 1.

Classes of State Transitions We single out some classes of transitions between state space points. The class of all possible transitions in $\mathcal{Z}^n$ is

$$\mathcal{D}_n \stackrel{\triangle}{=} \left\{ \vec{\mathbf{b}} - \vec{\mathbf{a}} \middle/ \vec{\mathbf{a}}, \vec{\mathbf{b}} \in \mathcal{Z}^n,\ \vec{\mathbf{a}} \neq \vec{\mathbf{b}} \right\} = \mathcal{Z}^n \setminus \left\{ \vec{\mathbf{0}} \right\}.$$

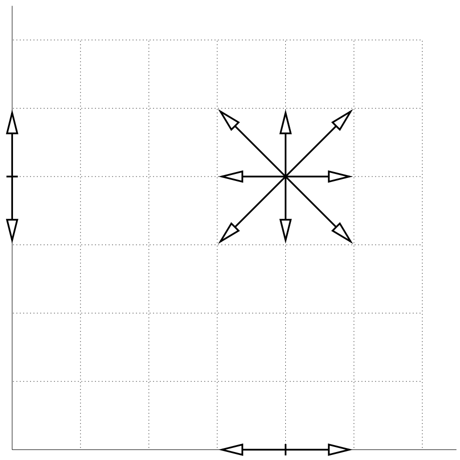

Figure 2: The transition classes $\mathcal{D}_{2,\langle 1,1\rangle}$, $\mathcal{D}_{2,\langle 0,1\rangle}$, and $\mathcal{D}_{2,\langle 1,0\rangle}$ ($\mathcal{D}_{2,\langle 0,0\rangle}$ is empty).

The classes of short transitions within the walls of $\mathcal{Z}_+^n$ are

$$\mathcal{D}_{n,\vec{\mathbf{w}}} \triangleq \left\{ \vec{\mathbf{b}} - \vec{\mathbf{a}} \middle/ \vec{\mathbf{a}}, \vec{\mathbf{b}} \in \mathcal{W}_{n,\vec{\mathbf{w}}},\ \left\| \vec{\mathbf{b}} - \vec{\mathbf{a}} \right\|_\infty = 1 \right\}, \qquad \vec{\mathbf{w}} \in \mathcal{B}^n. \tag{2.3}$$

See an illustration of the transition classes $\mathcal{D}_{2,\vec{\mathbf{w}}}$ in Figure 2.

Observation 2.2. $\mathcal{D}_{n,\vec{\mathbf{w}}} = \vec{\mathbf{w}}\mathcal{T}^n \setminus \left\{\vec{\mathbf{0}}\right\}$ *holds, so the element count* $|\mathcal{D}_{n,\vec{\mathbf{w}}}|$ *is* $3^{\|\vec{\mathbf{w}}\|_1} - 1$. *The overall count* $\sum_{\vec{\mathbf{w}}\in\mathcal{B}^n} |\mathcal{D}_{n,\vec{\mathbf{w}}}|$ *is* $4^n - 2^n$.

Models and Model Spaces Let $\overline{\mathcal{S}^2}$ denote the set of pairs of state space points. Namely, $\overline{\mathcal{S}^2} \triangleq \mathcal{S}^2 \setminus \{(\vec{\mathbf{a}}, \vec{\mathbf{a}})/\, \vec{\mathbf{a}} \in \mathcal{S}\}$. A model, in our context, is a function $\varphi^* : \overline{\mathcal{S}^2} \mapsto \mathcal{R}_+$ satisfying

1. "Communicativity": For every $\left(\vec{\mathbf{a}}, \vec{\mathbf{b}}\right) \in \overline{\mathcal{S}^2}$ there exists a finite sequence $\vec{\mathbf{a}} = \vec{\mathbf{a}}_1, \dots, \vec{\mathbf{a}}_k = \vec{\mathbf{b}}$ of states such that $\prod_{i=1}^{k-1} \varphi^*(\vec{\mathbf{a}}_i, \vec{\mathbf{a}}_{i+1}) > 0$.
2. "Non-instantaneity": For every $\vec{\mathbf{a}} \in \mathcal{S}$, the sum $\sum_{\vec{\mathbf{b}}\in\mathcal{S}} \varphi^*\left(\vec{\mathbf{a}}, \vec{\mathbf{b}}\right)$ is finite.

A value $\varphi^*\left(\vec{\mathbf{a}}, \vec{\mathbf{b}}\right)$ represents the transition rate from $\vec{\mathbf{a}}$ to $\vec{\mathbf{b}}$ of a communicative and non-instantaneous continuous-time Markov chain. Let $\mathcal{M}_n^*$ be the space of all models on $\mathcal{S} = \mathcal{Z}^n$ which possess the following homogeneity property:

$$\vec{\mathbf{b}}_1 - \vec{\mathbf{a}}_1 = \vec{\mathbf{b}}_2 - \vec{\mathbf{a}}_2 \quad \Rightarrow \quad \varphi^*\left(\vec{\mathbf{a}}_1, \vec{\mathbf{b}}_1\right) = \varphi^*\left(\vec{\mathbf{a}}_2, \vec{\mathbf{b}}_2\right), \qquad \vec{\mathbf{a}}_1, \vec{\mathbf{b}}_1, \vec{\mathbf{a}}_2, \vec{\mathbf{b}}_2 \in \mathcal{Z}^n. \tag{2.4}$$

Given some $\varphi^* \in \mathcal{M}_n^*$, let $\varphi : \mathcal{D}_n \mapsto \mathcal{R}_+$ be the function which corresponds to φ^* in the obvious way, by which $\varphi\left(\vec{\mathbf{d}}\right)$ represents the transition rate in the direction $\vec{\mathbf{d}} \in \mathcal{D}_n$: $\varphi\left(\vec{\mathbf{d}}\right)$ is set to the value $\varphi^*\left(\vec{\mathbf{0}}, \vec{\mathbf{d}}\right)$; the inverse mapping sets $\varphi^*\left(\vec{\mathbf{a}}, \vec{\mathbf{b}}\right)$ to $\varphi\left(\vec{\mathbf{b}} - \vec{\mathbf{a}}\right)$. Let $\mathcal{M}_n$ be the space of all such φ's which correspond to members of $\mathcal{M}_n^*$. Let the subclass $\overline{\mathcal{M}}_n$ contain those $\varphi \in \mathcal{M}_n$ whose support in $\mathcal{D}_n$ is finite. Namely, the models in $\overline{\mathcal{M}}_n$ are those with bounded transitions.

Let $\mathbb{M}_n^*$ be the space of all models on $\mathcal{S} = \mathcal{Z}_+^n$ possessing the following two properties. The first property is homogeneity, though weaker than for $\mathcal{M}_n^*$: The implication in Eq. (2.4) applies here only when both $\vec{\mathbf{a}}_1$ and $\vec{\mathbf{a}}_2$, or both $\vec{\mathbf{b}}_1$ and $\vec{\mathbf{b}}_2$, belong to the same wall. The second property is permitting short transitions only:

$$\left\|\vec{\mathbf{b}} - \vec{\mathbf{a}}\right\|_\infty \neq 1 \quad \Rightarrow \quad \varphi^*\left(\vec{\mathbf{a}}, \vec{\mathbf{b}}\right) = 0, \qquad \vec{\mathbf{a}}, \vec{\mathbf{b}} \in \mathcal{Z}_+^n.$$

Given some $\varphi^* \in \mathbb{M}_n^*$, let

$$\varphi = \{\varphi_{\vec{\mathbf{w}}} : \mathcal{D}_{n,\vec{\mathbf{w}}} \mapsto \mathcal{R}_+\}_{\vec{\mathbf{w}} \in \mathcal{B}^n \setminus \{\vec{\mathbf{0}}\}} \tag{2.5}$$

be the function array which corresponds to φ^* in the following way. We shall state both the mapping $\varphi^* \mapsto \varphi$ and the inverse mapping $\varphi \mapsto \varphi^*$. It will not be difficult to see, with the aid of two examples, that the two mappings are proper, and that one is indeed the inverse of the other. First, the mapping $\varphi^* \mapsto \varphi$: For arbitrary $\vec{\mathbf{w}} \in \mathcal{B}^n \setminus \left\{\vec{\mathbf{0}}\right\}$ and $\vec{\mathbf{d}} \in \mathcal{D}_{n,\vec{\mathbf{w}}}$, set $\varphi_{\vec{\mathbf{w}}}\left(\vec{\mathbf{d}}\right)$ to be the value of any $\varphi^*\left(\vec{\mathbf{a}}, \vec{\mathbf{b}}\right)$ with $\vec{\mathbf{a}}$ and $\vec{\mathbf{b}}$ in $\mathcal{W}_{n,\vec{\mathbf{w}}}$ and satisfying $\vec{\mathbf{b}} - \vec{\mathbf{a}} = \vec{\mathbf{d}}$. Now the mapping $\varphi \mapsto \varphi^*$: For arbitrary $\vec{\mathbf{a}}, \vec{\mathbf{b}} \in \mathcal{Z}_+^n$ satisfying $\left\|\vec{\mathbf{b}} - \vec{\mathbf{a}}\right\|_\infty = 1$, with $\vec{\mathbf{a}}$ belonging, say, to $\mathcal{W}_{n,\vec{\mathbf{w}}}$ and $\vec{\mathbf{b}}$ belonging, say, to $\mathcal{W}_{n,\vec{\mathbf{v}}}$, set the value of $\varphi^*\left(\vec{\mathbf{a}}, \vec{\mathbf{b}}\right)$ to be $\varphi_{\max\{\vec{\mathbf{w}},\vec{\mathbf{v}}\}}\left(\vec{\mathbf{b}} - \vec{\mathbf{a}}\right)$; the maximum is taken componentwise. Let $\mathbb{M}_n$ be the space of all function arrays of the type (2.5), which so correspond to members of $\mathbb{M}_n^*$. Note that $\varphi_{\vec{\mathbf{w}}}$ represents the transition rates within the wall $\mathcal{W}_{n,\vec{\mathbf{w}}}$.

Example 2.1. Let us demonstrate the construction by computing $\varphi^*\left(\vec{\mathbf{a}}, \vec{\mathbf{b}}\right)$, with $\varphi^* \in \mathbb{M}_2^*$, $\vec{\mathbf{a}} = \langle 1, 0\rangle$, and $\vec{\mathbf{b}} = \langle 0, 1\rangle$, from the corresponding $\varphi \in \mathbb{M}_2$. The indices of the walls $\mathcal{W}_{2,\vec{\mathbf{w}}}$ and $\mathcal{W}_{2,\vec{\mathbf{v}}}$ to which the points $\vec{\mathbf{a}}$ and $\vec{\mathbf{b}}$ belong happen to be $\vec{\mathbf{w}} = \vec{\mathbf{a}}$ and $\vec{\mathbf{v}} = \vec{\mathbf{b}}$. That is because $\vec{\mathbf{a}}$ and $\vec{\mathbf{b}}$ have been chosen adjacent to the corner. We have $\max\{\vec{\mathbf{w}}, \vec{\mathbf{v}}\} = \langle 1, 1\rangle$ and $\vec{\mathbf{b}} - \vec{\mathbf{a}} = \langle -1, 1\rangle$, so

$$\varphi^*\left(\vec{\mathbf{a}}, \vec{\mathbf{b}}\right) = \varphi_{\langle 1,1\rangle}(\langle -1, 1\rangle). \tag{2.6}$$

Thus, in spite of the fact that both $\vec{\mathbf{a}}$ and $\vec{\mathbf{b}}$ belong to the boundary of the state space, $\varphi^*\left(\vec{\mathbf{a}}, \vec{\mathbf{b}}\right)$ is computed from $\varphi_{\langle 1,1\rangle}$, which expresses the transition rates in the interior.

The reason is that $\vec{\mathbf{a}}$ and $\vec{\mathbf{b}}$ belong to distinct walls, and the transition between them can be regarded as passing through the interior. This transition is parallel, for instance, to the transition from $\langle 2,0\rangle$ to $\langle 1,1\rangle$, and the latter is parallel, say, to the transition from $\langle 2,1\rangle$ to $\langle 1,2\rangle$, whose both ends are interior points. Eq. (2.6) is thus mandated by the (weak) homogeneity assumption. ●

Example 2.2. Let us remain in $\mathbb{M}_2^*$ and in its counterpart $\mathbb{M}_2$. The transition from $\vec{\mathbf{a}} = \langle 5,0\rangle$ to $\vec{\mathbf{b}} = \langle 6,0\rangle$ is parallel to the transition from $\vec{\mathbf{a}}' = \langle 5,3\rangle$ to $\vec{\mathbf{b}}' = \langle 6,3\rangle$. But since the former lies within $\mathcal{W}_{2,\langle 1,0\rangle}$, while the latter lies within $\mathcal{W}_{2,\langle 1,1\rangle}$, homogeneity does not apply. The values $\varphi^*\left(\vec{\mathbf{a}},\vec{\mathbf{b}}\right)$ and $\varphi^*\left(\vec{\mathbf{a}}',\vec{\mathbf{b}}'\right)$ are not forced to be equal. These values are given by $\varphi_{\langle 1,0\rangle}(\langle 1,0\rangle)$ and $\varphi_{\langle 1,1\rangle}(\langle 1,0\rangle)$, respectively. ●

In view of Observation 2.2, $\mathbb{M}_n$ is essentially $\mathcal{R}_+^{4^n-2^n}$, excluding those elements which do not correspond to communicative Markov chains. The members of $\mathcal{M}_n$ and $\mathbb{M}_n$ too, like those of $\mathcal{M}_n^*$ and $\mathbb{M}_n^*$, will be called models.

State Space Measures When speaking of a measure, say μ, we always mean, unless explicitly stating otherwise, that μ is a nonnegative measure on $\left(\mathcal{S}, 2^{\mathcal{S}}\right)$, with $\mu(\mathcal{S}) > 0$. The state space $\mathcal{S}$ can be either $\mathcal{Z}^n$ or $\mathcal{Z}_+^n$. Such a measure is specified through singletons. We write $\mu(\vec{\mathbf{a}})$ as a shorthand for $\mu(\{\vec{\mathbf{a}}\})$. μ is said to be of a *product-form* if there exists a vector $\vec{\mathbf{q}} \in (0,\infty)^n$ satisfying $\mu(\vec{\mathbf{a}}) = \mu(\vec{\mathbf{0}}) \cdot \vec{\mathbf{q}}^{\vec{\mathbf{a}}}$ for every $\vec{\mathbf{a}} \in \mathcal{S}$; recall the $\vec{\mathbf{x}}^{\vec{\mathbf{y}}}$ convention from Subsection 2.1. For every $\vec{\mathbf{q}} \in (0,\infty)^n$, let $\pi_{\vec{\mathbf{q}}}$ denote the corresponding product-form measure with $\pi_{\vec{\mathbf{q}}}(\vec{\mathbf{0}}) = 1$. A measure μ is said to be *invariant* for a model φ from $\mathcal{M}_n$ or from $\mathbb{M}_n$ if it satisfies

$$\mu(\vec{\mathbf{a}}) \sum_{\vec{\mathbf{b}}\in\mathcal{S}} \varphi^*\left(\vec{\mathbf{a}},\vec{\mathbf{b}}\right) = \sum_{\vec{\mathbf{b}}\in\mathcal{S}} \mu\left(\vec{\mathbf{b}}\right)\varphi^*\left(\vec{\mathbf{b}},\vec{\mathbf{a}}\right), \qquad \vec{\mathbf{a}} \in \mathcal{S}, \tag{2.7}$$

where φ^* is the $\mathcal{M}_n^*$ or $\mathbb{M}_n^*$ counterpart of φ. Under the Markov chain semantics of φ^*, Eq. (2.7) is the steady-state version of Kolmogorov's forward equation, but allowing solutions with $\mu(\mathcal{S}) = \infty$. This equation is also known as the *global balance equation.* The communicativity postulate implies

Observation 2.3. *If μ is invariant for some model, then $\mu(\vec{\mathbf{a}})$ is positive for every $\vec{\mathbf{a}} \in \mathcal{S}$.*

3. Measures Invariant for Models in $\mathcal{M}_n$

Given a model $\varphi \in \mathcal{M}_n$, let Q_φ denote the set of vectors $\vec{\mathbf{q}} \in (0,\infty)^n$ such that $\pi_{\vec{\mathbf{q}}}$ is invariant for φ. By rewriting Eq. (2.7) in terms of φ itself, we reach

Observation 3.1. (Q_φ's identification). *Q_φ is the set of $\vec{\mathbf{q}}$'s solving the equation $\tilde{\varphi}(\vec{\mathbf{q}}) = \tilde{\varphi}(\vec{\mathbf{1}})$, where $\tilde{\varphi} : (0,\infty)^n \mapsto \mathcal{R}_+$ is the generating function defined through $\tilde{\varphi}(\vec{\mathbf{s}}) \triangleq \sum_{\vec{\mathbf{d}}\in\mathcal{D}_n} \varphi(-\vec{\mathbf{d}})\vec{\mathbf{s}}^{\vec{\mathbf{d}}}$.*

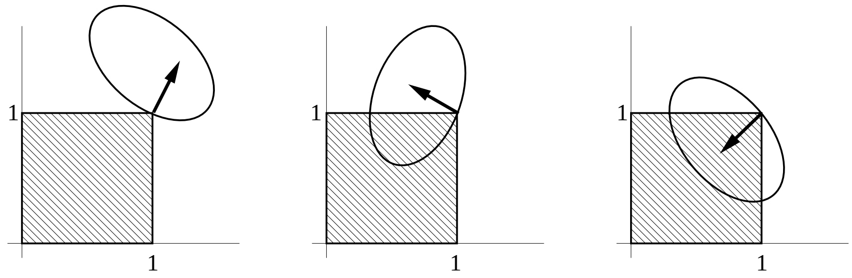

Figure 3: A schematic illustration for Proposition 3.1. The oval represents Q_φ, the box represents $\mathcal{C}$, and the arrow represents the drift.

Having introduced $\tilde{\varphi}$, we are ready for

Definition 3.1. *The quantity* $\sum_{\vec{\mathbf{d}}\in\mathcal{D}_n}\varphi(\vec{\mathbf{d}})\vec{\mathbf{d}}$, *which is equal to* $-\nabla\tilde{\varphi}(\vec{\mathbf{1}})$ *when convergent, is called the* drift *of* φ.

In order to study Q_φ, let us list a few properties of $\tilde{\varphi}$. Its domain of convergence, dom$\tilde{\varphi}$, includes the point $\vec{\mathbf{s}}=\vec{\mathbf{1}}$ due to the "non-instantaneity" property. Being a sum of convex functions, $\tilde{\varphi}$ is convex. Moreover, it is strictly convex since, by the "communicativity" property, for every $i=1,\dots,n$ there exists at least one $\vec{\mathbf{d}}\in\mathcal{D}_n$ with $d_i<0$ such that $\varphi(\vec{\mathbf{d}})>0$. For the same reason, the positive powers are also present. $\tilde{\varphi}$ thus has a unique minimum. When letting $\vec{\mathbf{s}}$ follows any straight line away from this minimum, including in a direction towards the boundary of $(0,\infty)^n$, the value of $\tilde{\varphi}(\vec{\mathbf{s}})$ goes to infinity. If dom$\tilde{\varphi}$ has a nonempty interior, which happens when the drift is convergent, then the gradient $\nabla\tilde{\varphi}$ is defined and finite throughout this interior. When letting $\vec{\mathbf{s}}$ approach a boundary point of dom$\tilde{\varphi}$, along any path in the interior of dom$\tilde{\varphi}$, the value of $\|\nabla\tilde{\varphi}(\vec{\mathbf{s}})\|_1$ goes to infinity. All these facts lead to

Observation 3.2. (Q_φ's properties). *Q_φ is the boundary of a bounded and convex level set of $\tilde{\varphi}$. The point $\vec{\mathbf{1}}$ is always in Q_φ. It is the sole point iff every component of the drift of φ is either zero or non-convergent. When $n\geq 2$ and the drift is convergent, Q_φ is an $(n-1)$-dimensional smooth manifold in $(0,\infty)^n$, and every point of Q_φ is an extreme point.*

Let $\mathcal{C}\triangleq(0,1]^n\setminus\left\{\vec{\mathbf{1}}\right\}$. This cube is the set of $\vec{\mathbf{q}}$'s whose corresponding $\pi_{\vec{\mathbf{q}}}$'s on $\mathcal{Z}_+^n$ are finite.

Proposition 3.1. (The relation between Q_φ and $\mathcal{C}$). *The following three cases may hold when $n\geq 2$ and the drift of φ is convergent (see Figure 3):*

Case (a): The drift is nonnegative (componentwise). In this case $Q_\varphi\cap\mathcal{C}$ is empty.

Case (b): The drift is neither nonnegative nor negative. In this case $Q_\varphi \cap \mathcal{C}$ *is nonempty, and* $Q_\varphi \setminus \left\{\vec{\mathbf{1}}\right\}$ *has a cluster point at* $\vec{\mathbf{1}}$.

Case (c): The drift is negative. In this case $Q_\varphi \cap \mathcal{C}$ *is nonempty, and its distance from* $\vec{\mathbf{1}}$ *is positive.*

Mixtures of product-form measures from $\{\pi_{\vec{\mathbf{q}}}\}_{\vec{\mathbf{q}} \in Q_\varphi}$ satisfy Eq. (2.7), and are thus invariant for φ as well. The converse statement would have been that every measure invariant for φ is either a member of $\{\pi_{\vec{\mathbf{q}}}\}_{\vec{\mathbf{q}} \in Q_\varphi}$, or can be represented as such a mixture. We prove this claim for models in $\overline{\mathcal{M}}_n$ (this subclass was introduced in Subsection 2.2).

Theorem 3.1. (Representation of measures invariant for models in $\overline{\mathcal{M}}_n$). *Let* $\varphi \in \overline{\mathcal{M}}_n$, *and let* μ *be a measure invariant for* φ. *Then there exists a unique Borel probability measure* ζ *on* $(0,\infty)^n$ *such that* $[\mu(\vec{\mathbf{0}})^{-1}]\mu = \int_{Q_\varphi} \pi_{\vec{\mathbf{q}}} d\zeta(\vec{\mathbf{q}})$.

4. The Structure of $\mathbb{P}_n$

Recall from Section 1 that $\mathbb{P}_{n,\vec{\mathbf{q}}}$, with $\vec{\mathbf{q}} \in (0,1)^n$, is the subspace of models from $\mathbb{M}_n$ for which $\pi_{\vec{\mathbf{q}}}$ is invariant. (Note that for every such model, $\pi_{\vec{\mathbf{q}}}$ is the unique invariant probability measure). The $\vec{\mathbf{q}} \in \mathcal{C} \setminus (0,1)^n$ are excluded from this discussion since, as will turn out later, their $\mathbb{P}_{n,\vec{\mathbf{q}}}$ are singular and seem less rich. Recall also that $\mathbb{M}_n$, and thus $\mathbb{P}_{n,\vec{\mathbf{q}}}$, can be viewed as subsets of $\mathcal{R}_+^{4^n-2^n}$. We will characterize $\mathbb{P}_{n,\vec{\mathbf{q}}}$ as the intersection between $\mathbb{M}_n$ and the solution space of a homogeneous linear system

$$\mathbf{A}\vec{\mathbf{x}} = \vec{\mathbf{0}}, \tag{4.1}$$

with $\mathbf{A}$ having $4^n - 2^n$ columns. The sequel gives this characterization, while concentrating on the special features of $\mathbf{A}$ which lead to the decoupling principle. The existence of nonnegative solutions, necessary for the intersection with $\mathbb{M}_n$ to be nonempty, is addressed immediately after the characterization.

Definition 4.1. ("hierarchically partitioned matrix"). *An* $m \times k$ *real matrix* $\mathbf{A} = (a_{i,j})$, *with* $m \leq k$, *will be referred to as a "hierarchically partitioned matrix" if there exists a partial order "*$\ll$*" on* $\{1,\ldots,m\}$, *and a partitioning of* $\{1,\ldots,k\}$ *into* m *nonempty partitions* $P_1,\ldots,P_m$, *such that*

1. *A component* $a_{i,j}$, *with* $j \in P_\ell$, *say, can be nonzero only if* $i = \ell$ *or if* $i \ll \ell$.
2. *For every* $i = 1,\ldots,m$ *there exists at least one* $j \in P_i$ *such that* $a_{i,j} \neq 0$.

The structure suggested by this definition is block triangular, up to a permutation of the columns, yet more sparse. Observe that if the $\mathbf{A}$ of Eq. (4.1) is hierarchically partitioned, then the solution space of (4.1) admits the following recursive characterization.

For every $i = 1, \ldots, m$, the portion $\langle x_j \rangle_{j \in P_i}$ of the vectors $\vec{\mathbf{x}}$ in the solution space is the hyperplane

$$\sum_{j \in P_i} a_{i,j} x_j = - \sum_{j \in \bigcup_{\{\ell / i \ll \ell\}} P_\ell} a_{i,j} x_j; \tag{4.2}$$

here the whole right hand side is regarded as a constant, adopting a point of view which defines the lower portions in terms of the higher ones. Thus, the dimension of a portion $\langle x_j \rangle_{j \in P_i}$, conditional on all higher portions, is $|P_i| - 1$. Suppose that i_1 and i_2 are such that neither $i_1 \ll i_2$ nor $i_2 \ll i_1$ holds. Then, conditional on all portions higher than any of them, the portions belonging to i_1 and to i_2 are decoupled from each other. In the context of the characterization of $\mathbb{P}_{n,\vec{\mathbf{q}}}$, the last property will be referred to as the decoupling principle.

Theorem 4.1. (Characterization of $\mathbb{P}_{n,\vec{\mathbf{q}}}$**).** $\mathbb{P}_{n,\vec{\mathbf{q}}}$ *is the intersection between* $\mathbb{M}_n$ *and the solution space of a homogeneous linear system of the type (4.1), with* $\mathbf{A}$ *being hierarchically partitioned. The rows of the matrix correspond to the walls* $\mathcal{W}_{n,\vec{\mathbf{w}}}$ *of* $\mathcal{Z}_+^n$*, excluding the corner* $\mathcal{W}_{n,\vec{\mathbf{0}}}$*. The variables, i.e. the model elements* $\varphi_{\vec{\mathbf{v}}}(\vec{\mathbf{d}})$ *with* $\vec{\mathbf{v}} \in \mathcal{B}^n \setminus \left\{\vec{\mathbf{0}}\right\}$ *and* $\vec{\mathbf{d}} \in \mathcal{D}_{n,\vec{\mathbf{v}}}$*, are partitioned according to* $\vec{\mathbf{v}}$*. The partial order among partitions is the majorization* $\prec$ *(recall Subsection 2.1). The matrix element serving as the coefficient of* $\varphi_{\vec{\mathbf{v}}}(\vec{\mathbf{d}})$ *in the row contributed by* $\mathcal{W}_{n,\vec{\mathbf{w}}}$*, with* $\vec{\mathbf{w}}$ *being equal to or majorized by* $\vec{\mathbf{v}}$*, is expressed as follows using indicators of conditions:*

$$1_{\left\{-\vec{\mathbf{d}} \leq \vec{\mathbf{1}} - 2(\vec{\mathbf{v}} - \vec{\mathbf{w}})\right\}} - \vec{\mathbf{q}}^{-\vec{\mathbf{d}}} 1_{\left\{\vec{\mathbf{d}} \leq \vec{\mathbf{1}} - 2(\vec{\mathbf{v}} - \vec{\mathbf{w}})\right\}}. \tag{4.3}$$

Having given the characterization, we now address the existence of nonnegative solutions. In a recursive representation of $\mathbb{P}_{n,\vec{\mathbf{q}}}$, of the type discussed in connection with Eq. (4.2), the coefficients at the left hand side are derived from Eq. (4.3) with $\vec{\mathbf{w}} = \vec{\mathbf{v}}$ holding. These coefficients appear in pairs $\left(1 - \vec{\mathbf{q}}^{-\vec{\mathbf{d}}}, 1 - \vec{\mathbf{q}}^{\vec{\mathbf{d}}}\right)$, due to the symmetry of the transition classes $\mathcal{D}_{n,\vec{\mathbf{w}}}$ (see Eq. (2.3)). The restriction $\vec{\mathbf{q}} \in (0,1)^n$ is essential to ensure the existence of at least one such pair whose members are nonzero, for each hyperplane. The opposite signs of the pair members imply that each hyperplane has an unbounded intersection with its pertinent nonnegative orthant.

Theorem 4.1 merely echos the global balance equation system (2.7), with one equation representing each wall $\mathcal{W}_{n,\vec{\mathbf{w}}}$, $\vec{\mathbf{w}} \in \mathcal{B}^n \setminus \left\{\vec{\mathbf{0}}\right\}$. What makes the theorem less banal is the absence of an equation for the corner $\mathcal{W}_{n,\vec{\mathbf{0}}}$. That equation is alive and kicking, and is not degenerated at all. Unless eliminated, the matrix could not have complied with the requirements of Definition 4.1. That is because the row contributed by $\mathcal{W}_{n,\vec{\mathbf{0}}}$ could not have been associated with any nonempty partition of variable indices—$\mathcal{D}_{n,\vec{\mathbf{0}}}$ is empty. It is the redundancy of the corner equation which is responsible for the decoupling principle. This redundancy, which is not very visible, is thus the reason for the richness of $\mathbb{P}_n$.

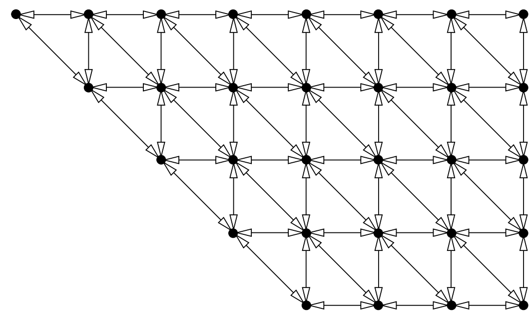

Figure 4: A two-dimensional state space with a slanted wall, creating a blunt corner and a sharp corner.

Remark 4.1. (State spaces with multiple corners). The decoupling principle can be generalized for state spaces of the form $\mathcal{S} = Z_1 \times Z_2 \times \ldots \times Z_n$, where every Z_i is a finite or an infinite succession of integers. ●

Remark 4.2. (Two-dimensional state spaces with slanted walls). State spaces with slanted walls are easy to describe in the two-dimensional case. See Figure 4. The decoupling principle can be generalized for the two corner types shown in the figure. The proof is at the same time a generalization and a specialization of the proof of Theorem 4.1, and is omitted. ●

5. Proofs

Proof of Proposition 3.1

Exclude the case $-\nabla\tilde{\varphi}(\vec{\mathbf{1}}) = \vec{\mathbf{0}}$, which has been commented upon earlier. The essence of the proposition is expressed in Lemma 5.1 below. The lemma is elementary, and is given without a proof. The association between the proposition and the lemma is drawn after the formulation of the latter. The definitions of a *cone* and of a *supporting half-space* are available, e.g., in Rockafellar (1970) [5], pages 13 and 99, respectively.

Lemma 5.1. *Let $A \subset \mathcal{R}^k$ be nonempty, bounded in $\|\cdot\|_1$, convex, containing the point $\vec{\mathbf{0}}$ in its boundary ∂A, and having a unique supporting half-space H with $\vec{\mathbf{0}} \in \partial H$. Also, let $K \subset \mathcal{R}^k$ be a nonempty and convex cone not containing $\vec{\mathbf{0}}$. The following three cases may hold (see a suggestive illustration in Figure 5):*

Case (a): $H \cap K$ is empty. In this case $\partial A \cap K$ is empty.

Case (b): $H \cap K$ is nonempty, but K is not contained in the interior of H. In this case $K \cap \partial A$ has a cluster point at $\vec{\mathbf{0}}$.

Case (c): K is contained in the interior of H. In this case $K \cap \partial A$ is nonempty, and its distance from $\vec{\mathbf{0}}$ is positive.

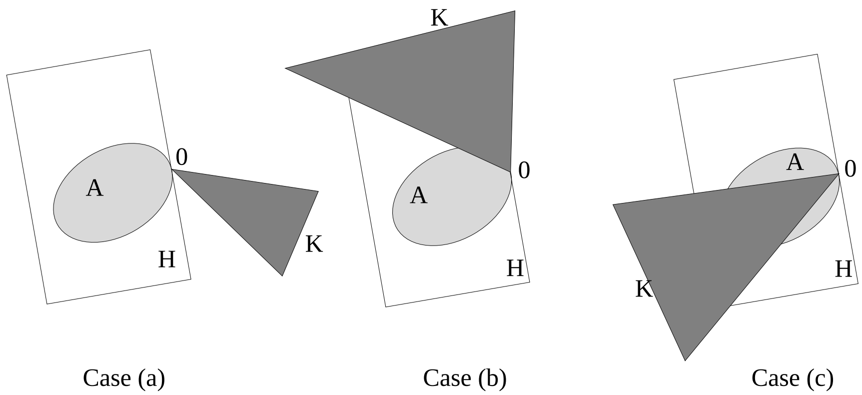

Figure 5: A suggestive illustration for Lemma 5.1.

Proposition 3.1 reduces into Lemma 5.1 via the following association. The role of ∂A is played by $Q_\varphi - \vec{\mathbf{1}}$ (defined as $\{\vec{\mathbf{q}} - \vec{\mathbf{1}} / \vec{\mathbf{q}} \in Q_\varphi\}$). The role of K is played by the nonpositive orthant, excluding $\vec{\mathbf{0}}$. This cone can be viewed as an extension of $\mathcal{C} - \vec{\mathbf{1}}$. In fact, the extension does not have any influence, since Q_φ is confined to $(0,\infty)^n$. The set Q_φ has a unique supporting half-plane H' with $\vec{\mathbf{1}} \in \partial H'$, due to Observation 3.2. The role of H is played by $H' - \vec{\mathbf{1}}$. It is known that

$$H' = \left\{ \vec{\mathbf{x}} \in \mathcal{R}^n \Big/ \left(\vec{\mathbf{x}} - \vec{\mathbf{1}} \right) \cdot \nabla \tilde{\varphi}(\vec{\mathbf{1}}) \leq 0 \right\}.$$

Hence the translation of the conditions on $H \cap K$ into conditions on the drift.

Proof of Theorem 3.1

The uniqueness of ζ is evident from the nature of the $\pi_{\vec{\mathbf{q}}}$: Different combinations of product-form measures cannot give identical aggregated measures. The proof of its existence calls for applying Choquet's theorem. This is proved to be possible, with the theorem applied for a normed space of measures on $\left(\mathcal{Z}^n, 2^{\mathcal{Z}^n}\right)$. The norm is ℓ_1, augmented by geometrical weights. The course of proof is as follows. First, a suitable version of Choquet's theorem is given. Then, the major part of the proof is dedicated to establishing the setting for applying the theorem. Finally, Choquet's theorem is invoked, and the conclusion is adjusted to fit the original setting of Theorem 3.1.

Theorem 5.1. (Choquet's theorem, adapted for a separable normed space). *Let $(Y, \|\ \|)$ be a separable normed vector space. Let $K \subset Y$ be compact and convex. Then for every $x \in K$ there exists a Borel probability measure γ, concentrated on the set of extreme points $E \stackrel{\triangle}{=} \mathrm{ext} K$, such that $x = \int_E y \, d\gamma(y)$.*

In a more general formulation of the theorem, Y is an abstract topological vector space in which the dual space Y^* separates points, and γ is supported on the closure of E.

See, for example, Rudin (1973) [7, p. 85, exercise 25]. But when K is separable and metric, such γ exists on E itself [ibid., p. 376, solution of exercise 25]. Also, when Y is a normed space then it is locally convex. In such a case it is guaranteed that Y^* separates points [ibid., p. 59, the corollary].

We now set out on establishing the setting for the application of Theorem 5.1. Let Υ denote the vector space of *signed* measures on $\left(\mathcal{Z}^n, 2^{\mathcal{Z}^n}\right)$. For every $p > 0$, let the function $\eta_p : \Upsilon \mapsto \mathcal{R}_+ \cup \{\infty\}$ be defined through

$$\eta_p(\nu) \triangleq \sum_{\vec{\mathbf{a}} \in \mathcal{Z}^n} (p\vec{\mathbf{1}})^{-|\vec{\mathbf{a}}|} |\nu(\vec{\mathbf{a}})|, \qquad \nu \in \Upsilon,$$

and let

$$\Upsilon_p \triangleq \{\nu \in \Upsilon \,/\, \eta_p(\nu) < \infty\}.$$

We shall employ normed vector spaces of the type (Υ_p, η_p). Observe that they are separable. We shall also employ the following type of subsets of Υ, whose members are normalized in some sense, and have "bounded growth". For every $r \geq 1$, let $\Delta_r \subset \Upsilon$ contain those ν which satisfy the following two conditions:

1. $\nu(\vec{\mathbf{0}}) = 1$.
2. $\left\|\vec{\mathbf{a}} - \vec{\mathbf{b}}\right\|_1 = 1 \;\Rightarrow\; r^{-1} \leq \dfrac{\nu(\vec{\mathbf{a}})}{\nu(\vec{\mathbf{b}})} \leq r.$

Lemma 5.2. *If $p > r$ then Δ_r is compact in (Υ_p, η_p).*

Proof. First of all, observe that $p > r \Rightarrow \Delta_r \subset \Upsilon_p$. The compactness claim can be reduced into a sequential compactness claim, due to the Borel-Lebesgue theorem (see, for example, Royden (1988) [6, p. 155]). The latter claim can be further reduced into the following one: Every sequence in Δ_r contains a subsequence which converges in the pointwise sense, i.e. at every point of $\mathcal{Z}^n$. The pointwise convergence would imply convergence in the norm η_p, since the major contribution to the norm is due to neighbourhoods of $\vec{\mathbf{0}}$, uniformly throughout Δ_r. Indeed, the requested subsequence can be extracted by diagonalization—see, for example, Chung (1974) [1, p. 84]. ∎

We now turn our attention to measures which are invariant for φ. Let $\Im_\varphi$ denote the set of state space measures ν (of the type of Subsection 2.2, not signed measures) which are invariant for φ and satisfy $\nu(\vec{\mathbf{0}}) = 1$. It can be seen directly that $\Im_\varphi$ is convex. We would like to express the invariance for φ in terms of operators on Υ. For every $\vec{\mathbf{d}} \in \mathcal{D}_n$, define the *shift operator* $\mathtt{S}_{\vec{\mathbf{d}}} : \Upsilon \mapsto \Upsilon$ through

$$\left(\mathtt{S}_{\vec{\mathbf{d}}}\nu\right)(\vec{\mathbf{a}}) = \nu(\vec{\mathbf{a}} - \vec{\mathbf{d}}), \qquad \begin{array}{l} \nu \in \Upsilon, \\ \vec{\mathbf{a}} \in \mathcal{Z}^n. \end{array}$$

By considering the global balance equation (2.7), as rewritten in terms of φ itself and divided by $\mu(\vec{\mathbf{a}})$, we arrive at

Observation 5.1. *The invariance of the elements of* $\Im_\varphi$ *for* φ *is tantamount to an invariance to some operator* $\sum_{\vec{\mathbf{d}}\in\mathcal{D}_n} c_{\vec{\mathbf{d}}}\mathsf{S}_{\vec{\mathbf{d}}}$*; the coefficients* $c_{\vec{\mathbf{d}}}$ *take values in* $[0,1)$*, their sum is* 1*, and only finitely many of them are nonzero. Due to communicativity, for every* $\vec{\mathbf{d}}\in\mathcal{D}_n$ *there exists an operator of the same type, but with* $c_{\vec{\mathbf{d}}}>0$*, to which the members of* $\Im_\varphi$ *are invariant as well.*

This leads to a key fact:

Lemma 5.3. $\Im_\varphi$ *is contained in some* Δ_r*.*

Proof. It is not difficult to see that the claim holds true for every

$$r \geq \max_{\{\vec{\mathbf{d}}\in\mathcal{D}_n/\|\vec{\mathbf{d}}\|_1=1\}} c_{\vec{\mathbf{d}}}^{\,-1},$$

where the $c_{\vec{\mathbf{d}}}$ are any positive coefficients of the type discussed in Observation 5.1. ∎

Lemma 5.4. $\Im_\varphi$ *is compact in some* (Υ_p,η_p)*.*

Proof. Choose some r such that $\Im_\varphi\subset\Delta_r$ (see Lemma 5.3), and some $p>r$. In view of Lemma 5.2, the compactness of $\Im_\varphi$ in (Υ_p,η_p) will be established if we verify that $\Im_\varphi$ is closed in (Υ_p,η_p). The following is to be verified. Let all the elements of a sequence $\{\nu_i\}_{i=1}^\infty\subset\Upsilon_p$ satisfy $\nu_i(\vec{\mathbf{0}})=1$, and be invariant to an operator $\sum_{\vec{\mathbf{d}}\in\mathcal{D}_n} c_{\vec{\mathbf{d}}}\mathsf{S}_{\vec{\mathbf{d}}}$ of the type of Observation 5.1. Suppose that the sequence converges in η_p to some $\nu\in\Upsilon_p$, i.e. $\eta_p(\nu-\nu_i)\to 0$. Then ν must also satisfy $\nu(\vec{\mathbf{0}})=1$, and be invariant to the same operator. So far the target. The fulfillment of $\nu(\vec{\mathbf{0}})=1$ follows from the fact that convergence in η_p obviously implies pointwise convergence. To verify the invariance, we check that $\eta_p\big(\nu-\sum_{\vec{\mathbf{d}}\in\mathcal{D}_n} c_{\vec{\mathbf{d}}}\mathsf{S}_{\vec{\mathbf{d}}}\nu\big)=0$. That is accomplished by applying η_p, and then $\limsup_{i\to\infty}$, on

$$\nu-\sum_{\vec{\mathbf{d}}\in\mathcal{D}_n} c_{\vec{\mathbf{d}}}\mathsf{S}_{\vec{\mathbf{d}}}\nu=(\nu-\nu_i)-\sum_{\vec{\mathbf{d}}\in\mathcal{D}_n} c_{\vec{\mathbf{d}}}\mathsf{S}_{\vec{\mathbf{d}}}(\nu-\nu_i),\qquad i=1,2\ldots,$$

while using

$$\eta_p\big(\mathsf{S}_{\vec{\mathbf{d}}}(\nu-\nu_i)\big)\leq\Big(p\vec{\mathbf{1}}\Big)^{|\vec{\mathbf{d}}|}\eta_p(\nu-\nu_i).$$

∎

Consider now the extreme points of $\Im_\varphi$.

Lemma 5.5. $\mathrm{ext}\Im_\varphi\subset\{\pi_{\vec{\mathbf{q}}}\}_{\vec{\mathbf{q}}\in Q_\varphi}$ *holds.*

Putting Lemma 5.5 in different words, there exists an $Q'_\varphi\subset Q_\varphi$ such that $\mathrm{ext}\Im_\varphi=\{\pi_{\vec{\mathbf{q}}}\}_{\vec{\mathbf{q}}\in Q'_\varphi}$. The equality $Q'_\varphi=Q_\varphi$ will follow from Theorem 3.1 post factum.

Proof of Lemma 5.5. Pick some $\nu \in \text{ext}\Im_{\varphi}$. The claim will follow by verifying that $\mathtt{S}_{\vec{\mathbf{d}}}\nu$ and ν are equal, up to a multiplicative factor, for every $\vec{\mathbf{d}} \in \mathcal{D}_n$. Observation 5.1 implies that there exists a coefficient $c_{\vec{\mathbf{d}}} > 0$ and a nonnegative $\nu' \in \Upsilon$ such that

$$\nu = c_{\vec{\mathbf{d}}}\mathtt{S}_{\vec{\mathbf{d}}}\nu + \nu'. \tag{5.1}$$

But $\mathtt{S}_{\vec{\mathbf{d}}}\nu$ too is invariant for φ (apply the operator and use commutativity). Hence, the right hand side of Eq. (5.1) can be rendered, through appropriate re-normalization, as a convex combination of two elements of $\Im_{\varphi}$. Both of them, with the one proportional to $\mathtt{S}_{\vec{\mathbf{d}}}\nu$ in particular, must be equal to ν, by the hypothesis $\nu \in \text{ext}\Im_{\varphi}$. ∎

Let σ_p denote the Borel σ-algebra on (Υ_p, η_p). We are ready to invoke Theorem 5.1, and draw the following conclusion: There exists a probability measure ζ' on some (Υ_p, σ_p), such that $\{\pi_{\vec{\mathbf{q}}}\}_{\vec{\mathbf{q}} \in Q'_{\varphi}} \in \sigma_p$ and

$$[\mu(\vec{\mathbf{0}})^{-1}]\mu = \int_{\{\pi_{\vec{\mathbf{q}}}\}_{\vec{\mathbf{q}} \in Q'_{\varphi}}} \pi_{\vec{\mathbf{q}}} d\zeta'(\pi_{\vec{\mathbf{q}}}).$$

This conclusion needs a slight adjustment to the original setting of Theorem 3.1, where the integration is performed on Q_{φ} itself. Let σ'_p denote the restriction of σ_p to $\{\pi_{\vec{\mathbf{q}}}\}_{\vec{\mathbf{q}} \in Q'_{\varphi}}$. The required adjustment is enabled by the following observation: The sources of the members of σ'_p, with respect to the mapping $\vec{\mathbf{q}} \mapsto \pi_{\vec{\mathbf{q}}}$, cannot lie outside the Borel σ-algebra on $(0, \infty)^n$.

Proof of Theorem 4.1

This proof should comprise two ingredients:

1. Verification that the matrix whose elements are defined in the theorem indeed satisfies the requirements of Definition 4.1.

2. Verification that $\mathbb{P}_{n,\vec{\mathbf{q}}}$ is indeed the intersection between $\mathbb{M}_n$ and the solution space of (4.1), with the matrix defined in the theorem.

The first ingredient has been addressed by the very formulation of the theorem, and by the ensuing discussion about nonnegative solutions. The second ingredient will be fulfilled by first showing validity for a matrix consisting of the declared $\mathbf{A}$ and an additional row, and then verifying that the additional row is in fact redundant.

The global balance equation (2.7), with a fixed μ, becomes an equation in φ^*. The space $\mathbb{P}_{n,\vec{\mathbf{q}}}$ is the set of all $\varphi \in \mathbb{M}_n$ whose corresponding φ^* satisfy Eq. (2.7) with $\mu = \pi_{\vec{\mathbf{q}}}$. This equation system (we now consider each contribution by some $\vec{\mathbf{a}} \in \mathcal{S}$ as one equation) should be rewritten in terms of φ itself. Due to the structure of $\mathbb{M}_n$, the collection of distinct equations corresponds to the collection of state space walls. In order to write down these equations in φ, some further state transition classes should be introduced. Let

$$\mathcal{D}_{n,\vec{\mathbf{w}}|\vec{\mathbf{v}}} \triangleq \left\{\vec{\mathbf{b}} - \vec{\mathbf{a}} \middle/ \vec{\mathbf{a}} \in \mathcal{W}_{n,\vec{\mathbf{v}}}, \vec{\mathbf{b}} \in \mathcal{W}_{n,\vec{\mathbf{w}}},\ \left\|\vec{\mathbf{b}} - \vec{\mathbf{a}}\right\|_{\infty} = 1\right\}, \qquad \vec{\mathbf{w}}, \vec{\mathbf{v}} \in \mathcal{B}^n.$$

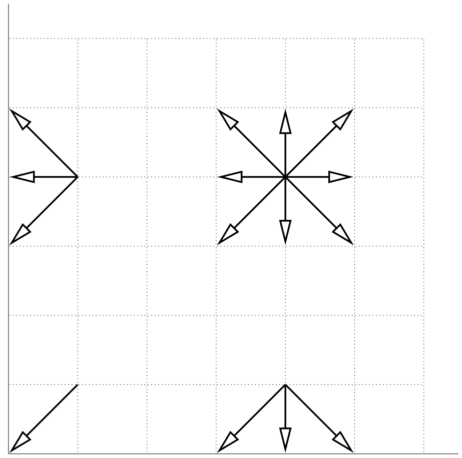

Figure 6: The transition classes $\mathcal{D}_{2,\langle 1,1\rangle|\langle 1,1\rangle}$ $(= \mathcal{D}_{2,\langle 1,1\rangle})$, $\mathcal{D}_{2,\langle 1,0\rangle|\langle 1,1\rangle}$, $\mathcal{D}_{2,\langle 0,1\rangle|\langle 1,1\rangle}$, and $\mathcal{D}_{2,\langle 0,0\rangle|\langle 1,1\rangle}$ (compare with Figure 2).

(Compare with Eq. (2.3)). See an illustration of the classes $\mathcal{D}_{2,\vec{\mathbf{w}}|\vec{\mathbf{1}}}$ in Figure 6. Observe that the overall set of short transitions into any state of $\mathcal{W}_{n,\vec{\mathbf{w}}}$ is

$$\bigcup_{\{\vec{\mathbf{v}}\in\mathcal{B}^n/\vec{\mathbf{v}}\geq\vec{\mathbf{w}}\}} \mathcal{D}_{n,\vec{\mathbf{w}}|\vec{\mathbf{v}}}.$$

Moreover, the above union is disjoint. Observe also that the following characterization holds:

$$\vec{\mathbf{v}} \geq \vec{\mathbf{w}} \;\Rightarrow\; \mathcal{D}_{n,\vec{\mathbf{w}}|\vec{\mathbf{v}}} = \left\{ \vec{\mathbf{d}} \in \mathcal{D}_{n,\vec{\mathbf{w}}} \middle/ \vec{\mathbf{d}} \leq \vec{\mathbf{1}} - 2(\vec{\mathbf{v}} - \vec{\mathbf{w}}) \right\}. \tag{5.2}$$

The rewriting of Eq. (2.7) in terms of the function array (2.5) is based upon the correspondence between φ^* and φ, as defined in Subsection 2.2. In the process, both sides of Eq. (2.7) are divided by $\mu(\vec{\mathbf{a}}) = \vec{\mathbf{q}}^{\vec{\mathbf{a}}}$, and terms are collected according to state transitions. The following equation system results:

$$\sum_{\vec{\mathbf{v}}\geq\vec{\mathbf{w}}} \sum_{\vec{\mathbf{d}}\in\mathcal{D}_{n,\vec{\mathbf{w}}|\vec{\mathbf{v}}}} \left[\varphi_{\vec{\mathbf{v}}}(-\vec{\mathbf{d}}) - \vec{\mathbf{q}}^{-\vec{\mathbf{d}}}\varphi_{\vec{\mathbf{v}}}(\vec{\mathbf{d}})\right] = 0, \qquad \vec{\mathbf{w}} \in \mathcal{B}^n. \tag{5.3}$$

In view of Eq. (5.2), the equation system (5.3) almost matches the description in the theorem. The only discrepancy is in the presence of an equation for $\vec{\mathbf{w}} = \vec{\mathbf{0}}$ in (5.3). (Note that the equation with $\vec{\mathbf{w}} = \vec{\mathbf{0}}$ is proper: The undefined $\varphi_{\vec{\mathbf{0}}}$ does not actually appear, since $\mathcal{D}_{n,\vec{\mathbf{0}}|\vec{\mathbf{0}}}$ is empty). The redundancy of this equation is to be verified. To this end, we show that there exist real numbers $\{g_{\vec{\mathbf{w}}}\}_{\vec{\mathbf{w}}\in\mathcal{B}^n}$, with $g_{\vec{\mathbf{0}}} = 1$, such that the weighted sum of the equations in (5.3), with these numbers serving as the weights, is zero. The numbers we use are $g_{\vec{\mathbf{w}}} = \vec{\mathbf{r}}^{-\vec{\mathbf{w}}}$, where $\vec{\mathbf{r}} = \langle r_1, \dots, r_n \rangle$ is given through

$$r_i = q_i^{-1} - 1, \qquad i = 1, \dots, n. \tag{5.4}$$

Let $h_{\vec{\mathbf{w}},\vec{\mathbf{v}},\vec{\mathbf{d}}}$ denote the coefficient belonging to the variable $\varphi_{\vec{\mathbf{v}}}(\vec{\mathbf{d}})$ in the equation contributed by $\vec{\mathbf{w}}$. From Eqs. (5.2) and (5.3) we have

$$h_{\vec{\mathbf{w}},\vec{\mathbf{v}},\vec{\mathbf{d}}} = \begin{cases} 1_{\left\{-\vec{\mathbf{d}} \le \vec{\mathbf{1}} - 2(\vec{\mathbf{v}} - \vec{\mathbf{w}})\right\}} - \vec{\mathbf{q}}^{-\vec{\mathbf{d}}} 1_{\left\{\vec{\mathbf{d}} \le \vec{\mathbf{1}} - 2(\vec{\mathbf{v}} - \vec{\mathbf{w}})\right\}} & \text{if } \vec{\mathbf{w}} \le \vec{\mathbf{v}}, \\ 0 & \text{otherwise}, \end{cases} \qquad \begin{array}{l} \vec{\mathbf{w}} \in \mathcal{B}^n, \\ \vec{\mathbf{v}} \in \mathcal{B}^n \setminus \left\{\vec{\mathbf{0}}\right\}, \\ \vec{\mathbf{d}} \in \vec{\mathbf{v}}\mathcal{T}^n \setminus \left\{\vec{\mathbf{0}}\right\}. \end{array}$$

Our target is to verify that for each $\varphi_{\vec{\mathbf{v}}}(\vec{\mathbf{d}})$, the weighted sum of coefficients is zero. Namely, we have to verify that

$$\sum_{\vec{\mathbf{w}} \in \mathcal{B}^n} g_{\vec{\mathbf{w}}}\, h_{\vec{\mathbf{w}},\vec{\mathbf{v}},\vec{\mathbf{d}}} = 0, \qquad \begin{array}{l} \vec{\mathbf{v}} \in \mathcal{B}^n \setminus \left\{\vec{\mathbf{0}}\right\}, \\ \vec{\mathbf{d}} \in \vec{\mathbf{v}}\mathcal{T}^n \setminus \left\{\vec{\mathbf{0}}\right\}. \end{array}$$

The above target equation is converted by substitution and a slight manipulation into

$$\vec{\mathbf{q}}^{-\vec{\mathbf{d}}} \sum_{\left\{\vec{\mathbf{w}} \in \mathcal{B}^n / \vec{\mathbf{v}} - \frac{1}{2}\left(\vec{\mathbf{1}} - \vec{\mathbf{d}}\right) \le \vec{\mathbf{w}} \le \vec{\mathbf{v}}\right\}} \vec{\mathbf{r}}^{-\vec{\mathbf{w}}} = \sum_{\left\{\vec{\mathbf{w}} \in \mathcal{B}^n / \vec{\mathbf{v}} - \frac{1}{2}\left(\vec{\mathbf{1}} + \vec{\mathbf{d}}\right) \le \vec{\mathbf{w}} \le \vec{\mathbf{v}}\right\}} \vec{\mathbf{r}}^{-\vec{\mathbf{w}}}, \qquad \begin{array}{l} \vec{\mathbf{v}} \in \mathcal{B}^n \setminus \left\{\vec{\mathbf{0}}\right\}, \\ \vec{\mathbf{d}} \in \vec{\mathbf{v}}\mathcal{T}^n \setminus \left\{\vec{\mathbf{0}}\right\}. \end{array}$$

Fix some $\vec{\mathbf{v}} = \langle v_1, \dots, v_n \rangle \in \mathcal{B}^n \setminus \left\{\vec{\mathbf{0}}\right\}$ and some $\vec{\mathbf{d}} = \langle d_1, \dots, d_i \rangle \in \vec{\mathbf{v}}\mathcal{T}^n \setminus \left\{\vec{\mathbf{0}}\right\}$. Suppose that $v_i = 0$ holds for some $i = 1, \dots, n$. Then d_i must also be zero. Likewise, w_i must be zero for every $\vec{\mathbf{w}} = \langle w_1, \dots, w_n \rangle$ participating in any of the two summations. A coordinate i with $v_i = 0$ can thus be ignored. Hence, no generality will be lost if we focus on $\vec{\mathbf{v}} = \vec{\mathbf{1}}$. The target now reduces into verifying that

$$\vec{\mathbf{q}}^{-\vec{\mathbf{d}}} \sum_{\left\{\vec{\mathbf{w}} \in \mathcal{B}^n / \vec{\mathbf{w}} \ge \frac{1}{2}\left(\vec{\mathbf{1}} + \vec{\mathbf{d}}\right)\right\}} \vec{\mathbf{r}}^{-\vec{\mathbf{w}}} = \sum_{\left\{\vec{\mathbf{w}} \in \mathcal{B}^n / \vec{\mathbf{w}} \ge \frac{1}{2}\left(\vec{\mathbf{1}} - \vec{\mathbf{d}}\right)\right\}} \vec{\mathbf{r}}^{-\vec{\mathbf{w}}}, \qquad \vec{\mathbf{d}} \in \mathcal{T}^n \setminus \left\{\vec{\mathbf{0}}\right\}. \tag{5.5}$$

Fix again an arbitrary $\vec{\mathbf{d}} = \langle d_1, \dots, d_i \rangle$, this time from $\mathcal{T}^n \setminus \left\{\vec{\mathbf{0}}\right\}$. Designate the index sets

$$I_t \stackrel{\triangle}{=} \{i = 1, \dots, n \,/\, d_i = t\}, \qquad t \in \mathcal{T}.$$

Adopt the following convention: Given a vector $\vec{\mathbf{x}} = \langle x_1, \dots, x_k \rangle$ and a partial index set $I \subset \{1, \dots, k\}$, let $\vec{\mathbf{x}}_I$ be the vector of length $|I|$ obtained by restriction. A member $\vec{\mathbf{w}}$ of the summation set at the left hand side of Eq. (5.5) admits the following form:

$\vec{\mathbf{w}}_{I_{-1}}$ can take any value in $\mathcal{B}^{|I_{-1}|}$, while $\vec{\mathbf{w}}_{I_0 \cup I_1}$ must be $\vec{\mathbf{1}}$. Similarly, the form of a member $\vec{\mathbf{w}}$ of the summation set at the right hand side is as follows: $\vec{\mathbf{w}}_{I_{-1} \cup I_0}$ must be $\vec{\mathbf{1}}$, while $\vec{\mathbf{w}}_{I_1}$ can take any value in $\mathcal{B}^{|I_1|}$. By decomposing all the vectors involved in Eq. (5.5) into their I_{-1}, I_0, and I_1 parts, and performing a slight rearrangement, (5.5) further reduces into

$$\vec{\mathbf{q}}_{I_{-1}}^{\vec{\mathbf{1}}} \sum_{\vec{\mathbf{w}} \in \mathcal{B}^{|I_{-1}|}} \vec{\mathbf{r}}_{I_{-1}}^{\,\vec{\mathbf{1}}-\vec{\mathbf{w}}} \;=\; \vec{\mathbf{q}}_{I_1}^{\vec{\mathbf{1}}} \sum_{\vec{\mathbf{w}} \in \mathcal{B}^{|I_1|}} \vec{\mathbf{r}}_{I_1}^{\,\vec{\mathbf{1}}-\vec{\mathbf{w}}}.$$

The last equation indeed holds true, as both sides are equal to 1: Recall Eq. (5.4), and apply the following identity, whose verification by induction on k is immediate:

$$\sum_{\vec{\mathbf{w}} \in \mathcal{B}^k} \vec{\mathbf{y}}^{\vec{\mathbf{w}}} = \left(\vec{\mathbf{y}} + \vec{\mathbf{1}}\right)^{\vec{\mathbf{1}}}, \qquad \vec{\mathbf{y}} \in \mathcal{R}^k.$$

The case where I_{-1} or I_1 are empty requires some attention, but does not require separate treatment if the following convention is adhered to. Let $\mathcal{B}^0$ contain a single element—the "zero-length vector". When raised to the power of itself, this "vector" gives 1—the conventional value of an empty product.

Acknowledgement. The thrust for this work was created at a period when we where both affiliated to the group of Professor Onno J. Boxma at CWI, Amsterdam. We are indebted to him for enabling and encouraging this work, and for his comments.

References

1. Chung, K. L. (1974). *A Course in Probability Theory*, second edition. Academic Press, New York.
2. van Dijk, N. M. (1993). *Queueing Networks and Product Forms.* Wiley, New York.
3. Gelenbe, E. (1991). Product-form queueing networks with negative and positive customers. *Journal of Applied Probability* **28** 656–663.
4. Miyazawa, M., P. G. Taylor (1997). Geometric product-form distribution for a queueing network with nonstandard batch arrivals and batch transfers. *Advances in Applied Probability*, **29** 523–544.
5. Rockafellar, T. R. (1970). *Convex Analysis.* Princeton University Press.
6. Royden, H. L. (1988). *Real Analysis*, third edition. Macmillan, New York.
7. Rudin, W. (1973). *Functional Analysis.* McGraw-Hill, New York.
8. Whittle, P. (1986). *Systems in Stochastic Equilibrium.* Wiley, New York.
9. Williams, R. J. (1995). Semimartingale reflecting Brownian motions in the orthant. In F. P. Kelly and R. J. Williams, editors, *Stochastic Networks*, pages 125–138. Springer, New York.