European Journal of Operational Research 232 (2014) 359–374

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Stochastics and Statistics

A stochastic multiscale model for electricity generation capacity expansion

Panos Parpas [a,*], Mort Webster [b]

[a] Department of Computing, Imperial College London, 180 Queen's Gate, London SW7 2AZ, United Kingdom
[b] Engineering Systems Division, Massachusetts Institute of Technology, 77 Massachusetts Ave, Building E40-235, Cambridge, MA 02139, United States

ARTICLE INFO

Article history:
Received 14 January 2012
Accepted 16 July 2013
Available online 25 July 2013

Keywords:
(I) Control
(I) OR in energy
(I) Markov processes

ABSTRACT

Long-term planning for electric power systems, or capacity expansion, has traditionally been modeled using simplified models or heuristics to approximate the short-term dynamics. However, current trends such as increasing penetration of intermittent renewable generation and increased demand response requires a coupling of both the long and short term dynamics. We present an efficient method for coupling multiple temporal scales using the framework of singular perturbation theory for the control of Markov processes in continuous time. We show that the uncertainties that exist in many energy planning problems, in particular load demand uncertainty and uncertainties in generation availability, can be captured with a multiscale model. We then use a dimensionality reduction technique, which is valid if the scale separation present in the model is large enough, to derive a computationally tractable model. We show that both wind data and electricity demand data do exhibit sufficient scale separation. A numerical example using real data and a finite difference approximation of the Hamilton–Jacobi–Bellman equation is used to illustrate the proposed method. We compare the results of our approximate model with those of the exact model. We also show that the proposed approximation outperforms a commonly used heuristic used in capacity expansion models.

1. Introduction

The general problem of capacity expansion under uncertainty has been extensively studied both as a stochastic optimal control problem as well as a multistage stochastic programming problem. In many ways it is a prototypical example of an optimal control problem; as a result, it has been studied since the late 1950s (Luss, 1982). For electric power systems, long-term investment (capacity expansion) and short-term operations (generation dispatch and unit commitment) were traditionally treated as decoupled decisions, and numerical models of long-term planning used highly simplified models and heuristics to represent the short-term dynamics.

However, the environment in which generation capacity expansion decisions are being made is becoming increasingly complex. This complexity is driven in part by increasing pressure placed on the electricity industry to address the problem of meeting the projected growth in demand in a sustainable manner, including increased reliance on intermittent renewable generation and increased demand–response mechanisms. The variability on the short time scale has important implications for the optimal portfolio of technologies that should be built in the long-run. For example, more intermittent generation will require other dispatchable technologies such as natural gas generation that can ramp up or down quickly. The conventional simplifications in long-term models will not capture this effect and will lead to suboptimal investment strategies (Palmintier & Webster, 2011). Moreover the deregulation of the electricity industry means that utilities cannot pass on the risks of investment decisions to customers. Consequently advanced models are needed in order to capture the complexities of the new decision making environment. The model of the full system, explicitly resolving both short-term (e.g., hourly) and long-term (e.g., annual or decadal) time scales along with the stochastic processes associated with each, would be computationally intractable for any system of realistic size.

In this paper, we present an efficient method for coupling the multiple temporal scales in the capacity expansion problem using the framework of singular perturbation theory for the control of continuous time problems. We demonstrate that for power systems the relevant stochastic processes are highly structured in ways that can be exploited. In particular load demand uncertainty and uncertainties in generation availability can be accurately modeled using weakly connected Markov processes. We take advantage of the properties of weakly connected processes in order to perform dimensionality reduction on the original model and therefore allows useful computation to be performed.

* Corresponding author. Tel.: +44 (0)20 7594 8367; fax: +44 (0)20 7581 8932.
E-mail addresses: p.parpas@imperial.ac.uk (P. Parpas), mort@mit.edu (M. Webster).

0377-2217/$ - see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.ejor.2013.07.022

To make operational the uncertainty structures present in this class of problems we make use of the tools from singular perturbation theory for Finite State Markov Processes (FSMPs) in continuous time. In some respects some models already take advantage of this structure. For example, the widely used MARKAL model (Seebregts, Goldstein, & Smekens, 2001) uses the concept of a "load block" to overcome the onerous requirement of optimizing over all possible loads. Similarly (Palmintier & Webster, 2011) simplify an integrated unit commitment and capacity expansion model by aggregating different power generators together. These types of aggregation approaches can be useful in practice. However, it is also important to understand why heuristics work, when they fail, and what can be done instead. For example, it is not clear how to extend the concept of a "typical" load to handle wind intermittency, or demand elasticity (a major objective of demand response programs). Instead we use perturbation methods to derive an "aggregate" model based on the assumption that the fast processes in our system (e.g wind, demand uncertainty) follow their stationary distribution. The computational complexity of the aggregate model is much less then the exact model and based on initial numerical experiments the error associated with the solution is much less then the existing heuristics used in MARKAL.

The contributions of this paper can be summarized as follows:

1. We formulate the problem of energy planning over multiple scales as a stochastic optimal control problem with weakly interacting FSMPs. We then extend and adapt some existing results from the literature of singular perturbation theory to derive an approximate problem that is computationally more attractive than the original problem. We also establish the conditions under which the approximate problem will yield the same value function as the original problem.
2. We formalize the heuristics of widely used models such as MARKAL. Existing models are based on the assumption that the Markov processes that describe the stochastic intraday dynamics of power systems are regularly perturbed. This assumption is not supported by the data. We show how to relax this assumption, using standard results from singular perturbation theory.
3. We demonstrate the application of this approach using empirical observations. As our approach is based on perturbation theory, we need to make assumptions about the existence of sufficient scale separation. As will be shown in this paper, there is sufficient evidence to suggest that such scale separation is present in the data. We expect some of the statistical techniques we use to be useful in other problem classes as well.

The rest of this paper is structured as follows: in the next section we discuss related literature and outline in more detail the contributions of this paper. In Section 3 we introduce our capacity expansion model and discuss some of its properties. In Section 3.1 we reinforce the arguments that motivated this paper by looking at some real data. In Section 3.2 we review perturbation theory in the context of multiscale Markov processes and link our assumptions with the empirical observations of Section 3.1. We also show that existing models do not capture the correct asymptotic behavior of the uncertainties present in the intraday scale. In Section 4 we introduce an aggregate problem, and show that asymptotically the value function of the approximate problem converges to the value function of the original problem. We also establish the same result for the approximate optimal control. In Section 5 we illustrate how the proposed approach could be implemented in practice. We compare the results of the approximate model with those of the exact model. We also show that the proposed approximation outperforms a commonly used heuristic used in large models.

2. Related literature and contributions

Capacity expansion problems have generated a large amount of literature. This is mainly because expansion problems are applicable to many areas and also because they are a good testbed for new modeling ideas. We will only discuss models that address the effect of the different time scales. Even though some of the work discussed below does not address capacity expansion directly, we believe that the most relevant papers to this work are the ones that address multiple scales since their ideas could be used in a capacity planning problem.

In Sen, Yu, and Genc (2006) the problem of incorporating different scales for addressing risk management problems such as buying and selling forward contracts for fuel are addressed in conjunction with the intraday unit commitment problem. Their model does not address intraday effects from intermittent sources. They formulate the problem as a multistage stochastic programming problem. The resulting large scale mixed integer linear programming problem is solved using a nested column decomposition algorithm. In Epe et al. (2009) the authors do address the problem of dealing with the intermittency of wind. Again a multistage stochastic programming approach is taken and the resulting large scale optimization problem is solved using a recombining tree methodology. In Pritchard, Philpott, and Neame (2005) the operation of a hydro-electric reservoir is addressed. Their problem also has multiple scales since the supply of power occurs intraday, in hourly intervals, but the management of the reservoir occurs over monthly scales. The problem is formulated as a dynamic programming problem. By approximating the decision to have some desirable properties the different scales can be decomposed. In Powell, George, Lamont, and Stewart (2009) the intermittency of wind and solar are addressed by approximate dynamic programming. We refer to the review article in Wallace and Fleten (2003, chap. 10) and the recent book (Weber, 2005) for a more complete overview of stochastic programming approaches as well as approaches based on dynamic programming. What all these papers have in common is that they address the existence of multiple scales using some sort of algorithmic framework. They either use a decomposition algorithm, approximate dynamic programming, or find some way to relax the non-anticipativity constraints in order to make the problem tractable. In this paper we take a different approach from the work described above. We posit that the uncertainties that exist in energy planning problems, in particular load demand uncertainty and uncertainties due to the power source, are highly structured. This structure can be exploited to a great advantage that provides insight into the nature of the problem. It also allows us to construct a reduced order model that is based on model primitives. The main aim of this paper is not to produce more accurate numbers by incorporating intraday variations, but rather define the representational structures that will allow qualitative and quantitative reasoning about the effects of different time scales.

The main results for multiscale FSMPs are summarized in two excellent books (Sethi & Zhang, 1994; Yin & Zhang, 1998). Therefore we only comment on the relations between the analysis here and the literature. In Jiang and Sethi (1991) a similar model to ours is proposed. The model is motivated by a manufacturing system with machines that have failure events that occur on different time scales. However, their model does not address the issue of capacity expansion and demand is deterministic. In Zhang, Yin, and Boukas (1997) another model similar to ours is proposed and is again studied in the context of manufacturing systems. Their model also has fixed capacity and deterministic demand. Even though the problem of capacity expansion with multiple scales has been studied (see chap. 10 in Yin & Zhang (1998)) we address capacity expansion for Markov Chains with weak interactions and multiple scales. This

key difference means that the approximate capacity expansion model described in chap. 10 of Sethi and Zhang (1994) is deterministic, whereas the dimensionality reduction method applied in this paper leads to a stochastic problem. Moreover, all the papers described above are concerned mainly with manufacturing systems, and in order to carry the results over to the power system sector we had to devise a new framework. This is because in order to capture some of the basic features of the energy planning models we need a model that can handle a Markov chain with weak interactions both on capacity (to model wind intermittency for example) as well as on the demand dynamics (to model the demand load dynamics). We also need to handle uncertainties that are not changing fast, such as the long term demand growth for electricity. Finally we extend some of the existing results to the case where general convex constraints are present. These constraints are needed in order to model startup costs, minimum uptimes, and so on. The technical extensions that are needed in order to extend the existing results are discussed in more detail in the beginning of Section 3.3.1. Apart from the theoretical contributions we also discuss the practical implementation of the theory together with the necessary statistical methodology in order for the theory to be implementable in practice.

3. A capacity expansion model

Before we delve into the details of the model, we first motivate our construction by looking at the characteristics of hourly electricity demand data (Section 3.1). In Section 5.1 we show that wind speed time series have the same structure, and thus are amenable to the same techniques. To keep the paper short we will only discuss demand data. In Section 3.2 we show that the assumptions made by existing power systems models can be explained in terms of perturbation theory. The exact, but computationally intractable, model is introduced in Section 3.3. The mathematical theory of perturbation theory, motivated in Section 3.2, will be used as the basis for the approximate, but tractable, model introduced in Section 4.

3.1. Load demand characteristics

Hourly electricity load data time series are driven by strong deterministic cyclical patterns (e.g., time of day, seasonal effects, etc.). By far the most widely used method to analyze such data sets is a two step process. The first step involves identifying the deterministic component of the data series. By subtracting the deterministic component from the original time series, any purely temporal components can be removed. The second step is to fit some stochastic process to the remaining component (we call this the stochastic component). There are a number of ways to perform each of these steps, and the exact method we used, together with our data sources are described in Section 5.1. Fig. 1(a) illustrates the transition matrix associated with the Markov chain of the stochastic component. The transition probabilities were obtained by identifying the maximum likelihood Markov Chain of de-trended hourly load data. The states in the Markov Chain represent the amount by which the hourly load is above or below the deterministic periodic component (for ease of exposition the axis shows the state number rather than the actual state). What is immediately obvious from this figure is that apart from a few cases the Markov Chain will tend to stay around the same state. Of course the variations around the state, as well as far away from the diagonal are important since they will influence the amount of reserves. Fig. 1(b) shows the result of the same process but with using a Markov Chain with only twenty states. Using typical states, optimizing over these and summing the results, essentially assumes a transition matrix as in Fig. 1(c). Note that there is no possibility of transition between clusters of states. These clusters of states are called typical states. However using typical states will always underestimate the costs of running such a system (this follows from Jensen's inequality and will be true for convex models, such as the one studied in this paper). This means that a certain amount of guesswork will always be required in order to find the level of reserves, or any other quantity of interest. Therefore, it seems that assuming the problem to be reducible to a few typical states oversimplifies the problem.

At the same time it is obvious from the figure that the data does not completely lack structure. To see the difference, a process with no structure is shown in Fig. 1(d). Therefore assuming a completely general structure is also not appropriate. Methods that do not take advantage of this highly specialized feature of the problem are implicitly assuming a structure such as the one shown in Fig. 1(d). The aim of this paper is to show how to take advantage of the approximate block diagonal structure of Fig. 1(b) without oversimplifying to a structure like that in 1(c). This provides an insight both into the characteristics of the problem but also leads to a tractable problem. It is worth pointing out that these conclusions are fairly consistent in electricity demand data and have been known for a long time. For example, it was shown in Pirrong and Jermakyan (2008), that electricity loads have a very fast mean reverting component. Most of the literature in this direction uses mean reverting diffusion processes to capture the fast load dynamics. However, it turns out that diffusion processes do not fit our wind data very well. For this reason we use FSMPs. A further possibility is to use diffusions with regime switching (Yin, 2009). Such an approach will be more general than the setup of this paper and it is a promising direction for future work. A review of load demand modeling in the context of multiple scales, is discussed in Ilic and Liu (1996). Any stochastic process can be used to capture this specific structure in energy data. The aim of this paper is to make explicit the implications of this structure to coupled models for capacity expansion and unit commitment.

3.2. Perturbation theory for the long term planning of power systems

The purpose of this section is twofold. The first aim is to establish some notation and terminology that will be used in the rest of the paper. The second aim is to show that the assumptions made by existing power systems models can be explained in terms of perturbation theory. We will argue that existing models assume that the Markov Processes representing the short term uncertainties in the system have generators that are regularly perturbed. However, the regular perturbation assumption is not supported by the data. More importantly the regular perturbation assumption gives results that are qualitatively wrong. Of course this result about the generators of Markov Processes has been known for decades, see for example (Phillips & Kokotovic, 1981). Our aim is to show how the theory of perturbed Markov Processes can be used to capture the assumptions of existing models, highlight the limitations of these assumptions, and propose a solution. The solution is based on the standard technique of singularly perturbing the ODE that governs the time evolution of the Markov generator by a changing the units of time. Finally, note that we are focusing on long term planning models (e.g. capacity expansion models). We will not discuss the literature of how perturbation theory is applied in other power system models. We refer the interested reader to Kokotović, Khalil, and O'reilly (1999) for a review of that literature.

The stochastic processes we consider are continuous time Markov Processes that have a finite number of states. This class of processes are also known as piecewise deterministic. We adopt the terminology of Sethi and Zhang (1994) and call a process that is expected to change a few times (if at all) per unit time a slow process.

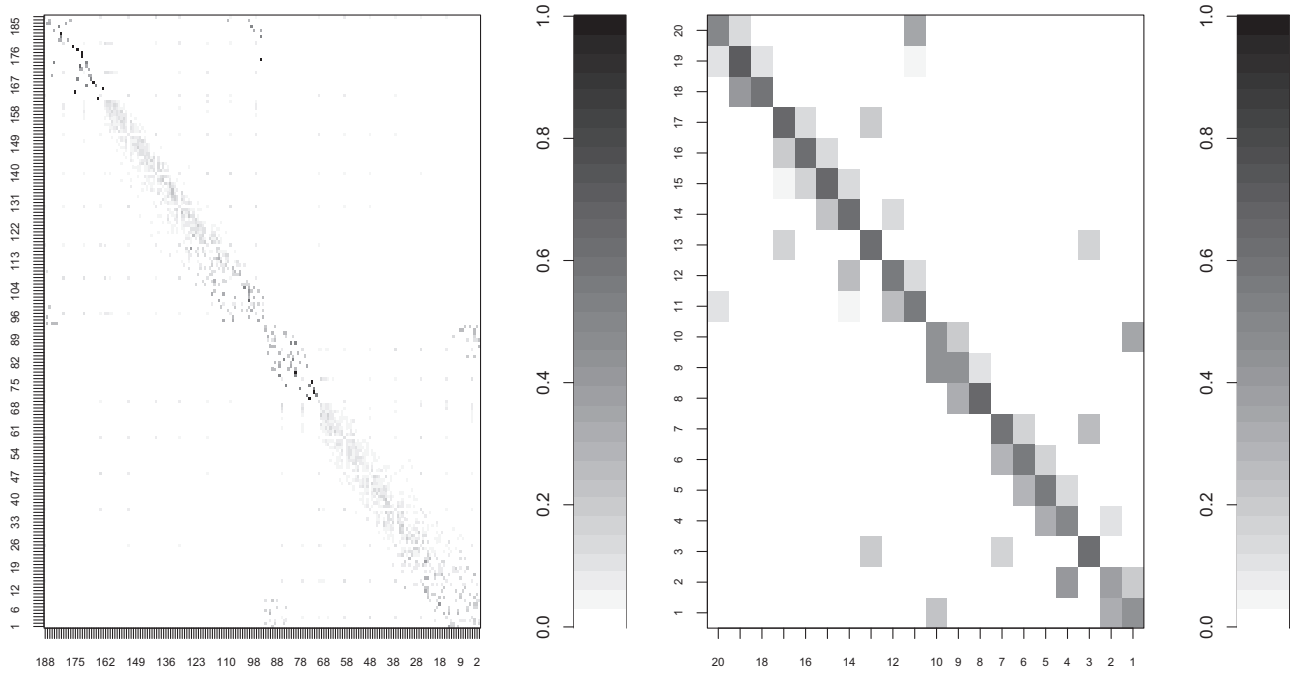

(a) A Markov chain with 200 states

(b) A Markov chain with 20 states

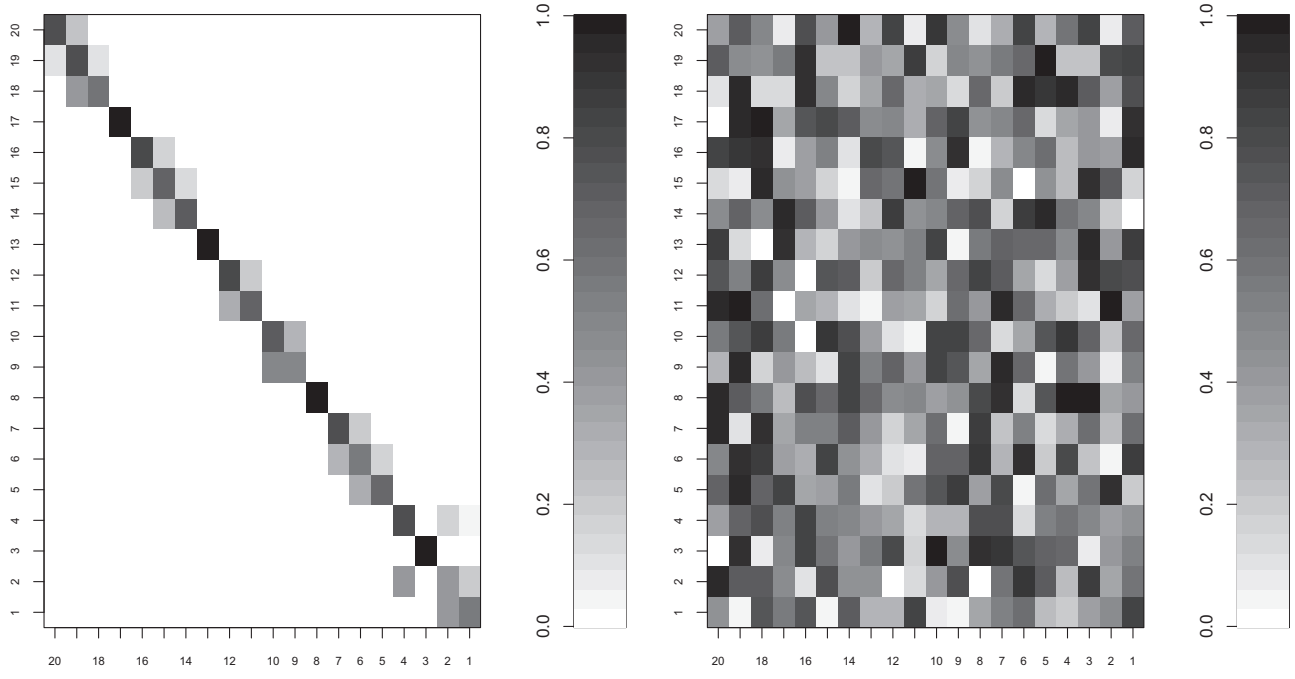

(c) Assumptions in existing models

(d) A structure free transition matrix

Fig. 1. Markov transition matrices. For ease of exposition the labels refer to the index of the state and not its numerical value. (a) and (b) are obtained from a maximum likelihood fit to hourly time series. (c) is the structure implicitly assumed when considering typical states. (d) was obtained by a maximum likelihood fit to a white noise process.

A fast process is one that is expected to make many transitions per unit time. We will use $\eta(t,\epsilon)$ to denote the fast process, and $\xi(t)$ for the slow process. The ϵ in the notation above is the perturbation parameter whose meaning will be explained next.

One way to capture the structure of the probability transition matrices in Fig. 1(a and b) is by assuming that the Markov generator of the fast process has a block diagonal structure. Under the latter assumption the probability transition matrix of the fast process will satisfy the following equation,

$$\begin{aligned} &\frac{dP^{\eta}}{d\tau} = P^{\eta}(\tau)[Q + \epsilon W], \\ &P(0) = I, \end{aligned} \tag{3.1}$$

where $Q = \mathrm{diag}(Q_1,\ldots,Q_m)$. The block diagonal component Q will be used to model the variation within aggregate states, and W will be used to model the transition between aggregate states. Each matrix $\{Q_i\}_{i=1}^{m}$, will be assumed to satisfy the standard conditions for a Markov generator matrix. We will further assume that each Q_i is

irreducible. The parameter $\epsilon > 0$ is used to capture the empirical observation that transitions inside clusters of states happen more frequently than transitions between clusters. To see this, compare Fig. 1(a and b) with Fig. 1(d). In this paper we use the terms aggregate states and clusters of states interchangeably. We note that we have not made any assumptions so far. We finish this section by illustrating how (3.1) is approximated in existing power systems models.

Existing models such as TIMES and MARKAL (e.g., Seebregts et al., 2001) assume that $\epsilon = 0$ in (3.1). This assumption is equivalent to taking the limit $\epsilon \to 0$ in (3.1). The consequence of this assumption is to eliminate the effect of W. For very short term models (e.g., the models used for intraday management of existing generators) this is not a bad approximation. However, for long term planning models this is not a good assumption. In particular, as a consequence of this assumption the models assume that once the process starts in a certain cluster of states it cannot escape from that cluster. This significantly reduces the complexity of the stochastic optimal control problem since it means that we can optimize the system over different clusters and then add up the results. However, the regular perturbation assumption described above does not produce the correct qualitative results. For example, the equilibrium density of the Markov Process in Fig. 1(b) is given in Fig. 2(a) (which is independent of ϵ). As can be seen in this figure the transitions between clusters can occur in equilibrium. Whereas, if the process is assumed to be regularly perturbed, then the equilibrium density shown in Fig. 2(b) does not allow for such transitions.

In order to capture the correct asymptotic behavior one can introduce a singularity by the change of variables $t = \epsilon\tau$ (Phillips & Kokotovic, 1981). With this change (3.1) becomes,

$$\frac{dP}{dt} = P(t)\left[\frac{1}{\epsilon}Q + W\right], \qquad P(0) = I. \tag{3.2}$$

The effect of the change of variables is to "stretch" time. Events that could be expected to happen after a long time interval e.g. after $O(1/\epsilon)$ units of time, will be expected to happen after $O(1)$ units of time. This means that fast processes become even faster. At the limit $\epsilon \to 0$, the fast process jump an infinite number of times within a finite interval. We can therefore expect that the fast process has made so many transitions (per unit time), that it would have reached its equilibrium. The exact details are discussed in Section 4.

3.3. The exact model

In this section, we introduce the stochastic optimal control problem that we will study in the rest of the paper. As was explained in Section 2, this class of models is solved with a large time step (typically five years). However when a large amount of intermittent power is available, a large time step will introduce significant errors. In order to address this issue, we couple the operation and capacity decisions in the model described below. The model is, in general, intractable. Section 4 uses the empirical observations of the previous section and some of the mathematical properties of the exact model discussed below to propose an approximation that is accurate when ϵ is small i.e. in the presence of both fast and slow dynamics. The approximate model in Section 4 takes advantage of the different time scales present in the model, and derives an equivalent model that is tractable.

The Decision Maker (DM) can invest in each of the generating facilities in order to increase capacity. To represent new plants, some of the plants are allowed to have zero capacity to start with. The cumulative investment in the ith plant is denoted by y_i and given by,

$$\frac{dy_i(t)}{dt} = \pi_i(t) \quad y_i(0) = y_i, \quad i = 1, \ldots, N.$$

where π_i is the rate of investment chosen by the DM. Once total investment in plant i reaches a predetermined level K_i, then new

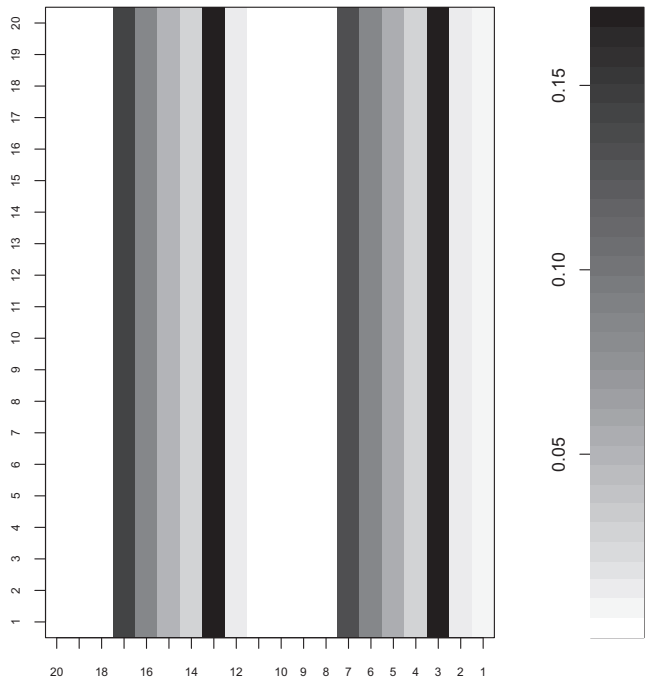

(a) Equilibrium density of ML Markov Chain, for any $\epsilon > 0$.

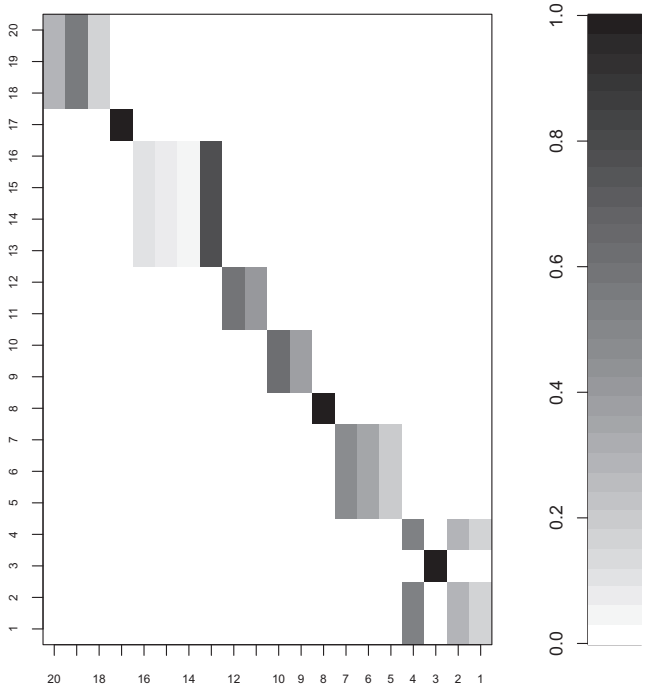

(b) Equilibrium density of ML Markov Chain assuming it is regularly perturbed.

Fig. 2. (a) The equilibrium density of the Maximum Likelihood (ML) Markov Chain is independent of ϵ. (b) Assuming the process is regularly perturbed gives an incorrect asymptotic behavior.

capacity becomes available. By imposing appropriate upper bounds on the rate of investment, π_i, we can ensure that new capacity cannot appear overnight. Each plant can be expanded once, and the level of investment required is assumed to be deterministic. Both of these assumptions can be relaxed, but with some increase in the complexity of the analysis.

The total available capacity in the ith plant is assumed to be a Markov process, and it is given by,

$$\eta^i(\epsilon,t)=\begin{cases}\eta_1^i(\epsilon,t) & t\leqslant\tau_i\\ \eta_2^i(\epsilon,t-\tau_i) & t>\tau_i.\end{cases} \tag{3.3}$$

The stopping time τ_i is given by,

$$\tau_i=\inf\{t|y_i(t)=K_i\}\wedge T, \tag{3.4}$$

where T denotes the terminal time. We assume that the decision maker can use any of the N generating plants to meet demand; the state variable $x(t)$ represents the amount of energy not served up to time t and it is defined as follows,

$$\frac{dx(t)}{dt}=\sum_{i=1}^{N}u_i(t)+\eta^0(t,\epsilon)-z(t)\quad x(0)=x, \tag{3.5}$$

where u_i represents the output from the ith plant, and is controlled by the DM. The finite state Markov Process, η^0 is a fast process that represents the intraday variations in hourly loads. The last term in (3.5), $z(t)$, represents the demand and it is assumed to have the following dynamics,

$$\frac{dz(t)}{dt}=f(z(t),\xi^0(t))\quad z(0)=z. \tag{3.6}$$

In (3.6) ξ^0 is assumed to be a finite state Markov process, but this process changes at a slower rate. The latter process is used to represent uncertainties such as the long term demand trend for electricity. Information about long term demand arrives at slower rate. It is not possible to model this class of uncertainties as fast processes. Throughout the paper, we use the notation η and ξ, for the fast, and slow, respectively, scale dynamics.

Motivated by the discussion in Section 3.2, the Markov process $\eta(t,\epsilon)=[\eta^0(t,\epsilon),\ldots,\eta^N(t,\epsilon)]$ is assumed to have a generator given by,

$$\frac{dP^{\eta}(t)}{dt}=P^{\eta}(t)\left[\frac{1}{\epsilon}Q+W\right], \tag{3.7}$$

where $Q=\operatorname{diag}(Q_1,\ldots,Q_l)$. Furthermore both W and $\{Q_i\}_{i=1}^{l}$ are assumed to be irreducible. The generator for the ξ variables can be more general since we will not perform any approximation on this process. The dynamics of ξ are given by,

$$\frac{dP^{\xi}(t)}{dt}=P^{\xi}(t)R.$$

A policy $(u(t),\pi(t))$ will be called admissible if it satisfies the dynamics specified above (including the stochastic bounds on available capacity). Moreover an admissible policy needs to be non-anticipative and therefore must be adapted to the filtration $\mathcal{F}_t$. The set of all $\mathcal{F}_t$-adapted processes is denoted by $\mathcal{A}_t^{\epsilon}$. In Section 5 we describe a numerical method based on the finite difference approximation technique described in Kushner and Dupuis (2001) to obtain admissible policies.

In the exact model the DM chooses the generation $u(t)$ and investment $\pi(t)$ to minimize the discounted sum of operations and investments costs $G_i(.)$ (of technology i) plus a terminal cost $\Phi(X_T)$, subject to system equations for unmet demand $x(t)$, capacity available $y(t)$, and demand $z(t)$ and constraints on maximum potential capacity and maximum generation:

$$J_s^{\epsilon}(w,\eta,\xi;u,\pi)\triangleq E\left\{\sum_{i=1}^{N}\int_s^T e^{-\rho(t-s)}G_i(x_s,\pi_i(s),u_i(s),\xi(s))ds+e^{-\rho(T-s)}\Phi(X_T)\right\}$$

$$\begin{aligned}v_s(w,\eta,\xi)=\min\quad & J_s^{\epsilon}(w,\eta,\xi;u,\pi)\\ & \frac{dx(t)}{dt}=\sum_{i}^{N}u_i(t)-z(t)+\eta^0(t,\epsilon)\quad x(s)=x\\ & \frac{dz(t)}{dt}=f(z(t),\xi^0(t))\quad z(s)=z\\ & \frac{dy(t)}{dt}=\pi(t),\quad y(s)=y\qquad(\mathcal{H}_{\epsilon})\\ & y_i(t)\leqslant K_i\\ & 0\leqslant u_i(t)\leqslant\eta_i(t,\epsilon)\\ & (u(t),\pi(t))\in\mathcal{A}_t^{\epsilon}\cap\mathcal{C}(t)\end{aligned}$$

For the purposes of this paper we will assume that the controls need to satisfy additional convex constraints. These will be represented by the set $\mathcal{C}(t)$. These constraints can be used to model startup costs, minimum operation times, reserves, etc. Using the approximations described in Weber (2005), these features can be incorporated within the proposed framework with a corresponding increase in computational complexity.

The real valued function G_i captures the operation and investment costs associated with technology i. We assume that G_i is also a function of x (i.e. the amount of energy not served). In order to avoid uninteresting solutions we assume that a high penalty is imposed for not meeting demand. This is captured by having a terminal condition function $\Phi(x)$. This function penalizes the total amount of energy not served during the planning period. In order to simplify the notation in subsequent sections we will use the following notation whenever convenient: $w(t)=[x(t),y(t),z(t)]$, $c(t)=[u(t),\pi(t)]$ and $\beta^{\epsilon}(t)=[\eta(t,\epsilon),\xi(t)]$. Using this notation, the dynamics of $(\mathcal{H}_{\epsilon})$ are given by,

$$\frac{dw(t)}{dt}=g(w(t),c(t),\beta^{\epsilon}(t)), \tag{3.8}$$

where $g(w,c,\beta^{\epsilon})\triangleq g(x,y,z,u,\pi,\eta,\xi)=\left[\sum_i^N u_i-z+\eta^0,f(z,\xi),\pi\right]$. Note that we are using capital G to denote one period instantaneous costs, and lower case g to denote the state evolution equations of our model. We use $G(\cdot)\triangleq\sum_{i=1}^{N}G_i(\cdot)$, when the distinction between the different technologies is irrelevant.

3.3.1. Analysis of the value function

The aggregated problem is obtained asymptotically as ϵ goes to zero in $(\mathcal{H}_{\epsilon})$. Motivated by applications in manufacturing systems, capacity expansion models with fast and slow dynamics have been considered before (see e.g. Sethi & Zhang, 1994 for a review). As was mentioned in Section 2 the problem of capacity expansion with Markov Chains that have weak interactions over multiple scales has not been addressed before. Capacity expansion models introduce a technical complication of having to deal with the random time when the ith plant gets expanded (see (3.4)). The advantage of solving this model with the weakly interacting Markov Chain assumption is that the reduced order model is still stochastic, whereas the dimensionality reduction methods used in other capacity expansion models (see chap. 10 of Sethi & Zhang (1994)) lead to a deterministic approximation. Moreover for the application we consider in this paper, we wish to handle the case where the controls need to satisfy additional convex constraints. In order to deal with the additional technical complications we study the dynamics of our model using the results from Arnold (1998). The approach developed in Section 2.2 of Arnold (1998) allows us to study the

dynamics of (3.8) pathwise. This means that for each fixed path $\omega \in \Omega$, (3.8) can be analyzed as a generalized deterministic ordinary differential equation. This approach allows an easy way to construct various estimates about the value function, and the optimal controls. The theory of viscosity solutions is used in order to establish the connection between the value function and the Hamilton–Jacobi–Bellman (HJB) equation. The results concerning the properties of the value function, and its limiting properties are extensions to the results from Zhang et al. (1997) and Sethi and Zhang (1994) in the sense described above.

As β^ϵ is cadlag the control $c(t) = c(t,\omega)$ is also cadlag. The control will be called admissible if:

- The process $c(\cdot)$ defined on $\mathcal{T} \times \Omega$ is $\mathcal{F}_t$ progressively measurable; i.e., the map $(t,\omega) \to c(t,\omega)$ is $\mathcal{B}[0,t] \otimes \mathcal{F}_t$ measurable for each $t \in \mathcal{T}$.
- The process satisfies,

$$E \int_{t1}^{t2} |c(s)| ds < \infty.$$

We will call the control feasible if, in addition to being admissible, it also satisfies the additional convex constraints given in $(\mathcal{H}_\epsilon)$. Note that we are only concerned with Markov controls. It can be shown, using similar arguments as in Øksendal and Sulem (2007), that Markov controls are sufficient for the problem studied in this paper.

We will use the shift operator θ (see e.g Arnold, 1998) to denote the following:

$$\frac{dw(t)}{dt} = g(w(t), c(t), \theta_t\omega) = g(w(t), c(t,\omega), \beta^\epsilon(t,\omega)).$$

The ω-dependance of β will be dropped whenever possible. Similar notation is used for the control. The process $\beta^\epsilon(t)$ has, by construction, a finite number of discontinuities. Therefore the solution to (3.8) should be considered in a generalized sense and not interpreted as a classical solution. If the following holds,

$$\psi(t,c,\omega)(w) = w + \int_0^t g(\psi_s, \theta_s c, \theta_s\omega) ds, \tag{3.9}$$

we say that ψ solves (3.8). Such a solution is also called a solution in the sense of Caratheodory. If the solution happens to be differentiable with respect to t, then the solution is also a classical solution. As β^ϵ is in general discontinuous, we have to deal with generalized solutions. The significance of using the formalism of generalized solutions is that we can show that the solutions of (3.8) satisfy a co-cycle property (Arnold, 1998). In particular, it can be shown that:

$$\psi(t+s,c,\omega)(w) = \psi(t,\theta_s c,\theta_s\omega) \circ \psi(s,c,\omega)(w). \tag{3.10}$$

The usefulness of the co-cycle property is in studying properties of the value function. For example, if τ denotes the first jump time of the process β^ϵ after time *0*, then we can construct a solution ψ for time $\tau + s$ by using the deterministic solution of ψ up to time time τ and then "restarting" the stochastic solution at the post-jump location and solving the equation forward up to time $\tau + s$. This type of manipulations allows to very easily obtain bounds and establish continuity properties of the value function. We make the following assumptions concerning the dynamics of g.

Assumption 3.1. The function g is uniformly Lipschitz continuous, and satisfies a linear growth condition in w.

Under conditions more general than the one given above, Theorem 2.2.1 in Arnold (1998) can be used to establish the existence and uniqueness of the solutions to (3.8).

Assumption 3.2. G and Φ have a modulus of continuity given by γ_G and γ_Φ respectively. This means that the following holds,

$$|G(x,c,\beta) - G(y,c,\beta)| \leqslant \gamma_G(|x-y|), \quad \forall c, \beta$$
$$|\Phi(x) - \Phi(y)| \leqslant \gamma_\Phi(|x-y|).$$

Assumption 3.3. There exist constants M_G and M_Φ such that:

$$|G(x,c,\beta)| \leqslant M_G,$$
$$|\Phi(x)| \leqslant M_\Phi.$$

Below we show that the value function is equi-continuous in w (Theorem 3.4), and that, therefore, the Arzella-Ascolli theorem can be used to establish the existence of the limit $\lim_{\epsilon\to 0} v^\epsilon$ (Corollary 3.5). Then in Theorem 3.8 we show that this limit does not depend on the particular state η, but rather the set of aggregate states η belongs to.

Theorem 3.4. *For any fixed β and $t \in \mathbb{T}$, the value function $v_t^\epsilon(w,\beta)$ is equi-continuous, and equi-bounded in w.*

Proof. The value function is equi-continuous in w if,

$$\sup_{|w-\widehat{w}|\leqslant\delta} |v_t^\epsilon(w,\beta) - v_t^\epsilon(\widehat{w},\beta)| \to 0 \text{ as } \delta \to 0. \tag{3.11}$$

Note that Assumption 3.1 together with Gronwall's inequality implies that (3.9) satisfies,

$$|\psi(t,c,\omega)(x) - \psi(t,c,\omega)(y)| \leqslant e^{Lt}|x-y|,$$

where L is the Lipschitz constant. Using the observation above and the co-cycle property in (3.10), it can be shown that (3.11) holds in the same way as the estimates in the proof of Proposition 3.1 in Bardi and Capuzzo-Dolcetta (1997) are obtained.

To show that a function is equi-bounded we need to show that for any ϵ the following holds,

$$\sup_w |v_t^\epsilon(w,\beta)| < +\infty. \tag{3.12}$$

The above is a simple consequence of Assumption 3.3. □

Corollary 3.5. *For any fixed t and β^ϵ, the sequence $\{v_t^\epsilon(w,\beta^\epsilon)\}_{\epsilon>0}$ has a convergent subsequence.*

Proof. This is a consequence of the Arzella-Ascoli Theorem, see c.f. (Rudin, 1991). □

Capacity expansion models introduce the possibility of changing the states of the Markov Process. This property of the model introduces additional technical complications in the analysis of the value function. The following Lemma will be used to control the stopping time in (3.4) i.e. the first time a particular plant gets expanded.

Lemma 3.6. *Let $\{\tau_\epsilon\}$ be a sequence of feasible stopping times defined in (3.4). Suppose that,*

$$\lim_{\epsilon\to 0} \tau_\epsilon = t, \quad \text{Pa.s.}. \tag{3.13}$$

Then for some feasible control $\bar{c}$,

$$\lim_{\epsilon\to 0} E[|v_{\tau_\epsilon}^\epsilon(\psi(\tau_\epsilon,\bar{c},\omega)(w_t), \beta^\epsilon(\tau_\epsilon)) - v_t^\epsilon(w_t,\beta^\epsilon(\tau_\epsilon))| \; |\mathcal{F}_t] = 0. \tag{3.14}$$

Proof. This proof is similar to the argument given in the proof of Theorem 3.4, only a sketch is given here. Given any $\epsilon_1 > 0$ let $\bar{c}$ be a feasible control such that,

$$v_t^\epsilon(\widehat{w},\beta) \geqslant E\left[\int_t^T e^{-\rho(s-t)}G(\psi(s,\bar{c},\omega)(\widehat{w}),\bar{c}_s,\beta_s^\epsilon)ds + e^{-\rho(T-t)}(\Phi(\psi(T,\bar{c},\omega)(\widehat{w})))|\mathcal{F}_t\right] - \frac{\epsilon_1}{3}.$$

Then the limit defined in (3.14) can be bounded with standard arguments, to show that if ϵ is small enough then (3.13) is true. □

The HJB equation associated with the value function of $(\mathcal{H}_\epsilon)$ is given by,

$$\begin{aligned} &\rho v_t^\epsilon(w,\beta^\epsilon) = \mathcal{H}^\epsilon(v_t^\epsilon, \nabla_w v_t^\epsilon),\\ &v_T^\epsilon(w,\beta^\epsilon) = \Phi(w), \end{aligned} \tag{3.15}$$

where $\mathcal{H}^\epsilon(v_t^\epsilon, \nabla_w v_t^\epsilon)$ is defined as follows,

$$\min\left\{\sum_{i=1}^N G_i(x,\pi^i,u^i,\xi) + \partial_x v_t^\epsilon\left(\sum_{i=1}^N u_i - z - \eta^0\right) + \nabla_y v_t^\epsilon \cdot \pi\right\} + \partial_t v_t^\epsilon + \partial_z v_t^\epsilon f(z,\xi) + \mathcal{L}^\epsilon v_t^\epsilon[\eta,\xi].$$

The linear operator $\mathcal{L}^\epsilon$ used above is defined as follows,

$$\mathcal{L}^\epsilon f[\eta,\xi] = \mathcal{L}_\eta^\epsilon f[\eta,\xi] + \mathcal{L}_\xi f[\eta,\xi],$$

where,

$$\mathcal{L}_\eta^\epsilon f[\eta,\xi] = \sum_{\widehat{\eta}\neq\eta}\left(w_{\widehat{\eta}\eta} + \frac{q_{\widehat{\eta}\eta}}{\epsilon}\right)(f(\widehat{\eta},\xi) - f(\eta,\xi)),$$

$$\mathcal{L}_\xi f[\eta,\xi] = \sum_{\widehat{\xi}\neq\xi} r_{\widehat{\xi}\xi}(f(\eta,\widehat{\xi}) - f(\eta,\xi)).$$

Above we made use of the standard convention of denoting the (i,j)th entry of a matrix A with the notation a_{ij}. Note that the presence of $\mathcal{L}^\epsilon$ is entirely due to the stochastic nature of the problem. In deterministic optimal control problems this term does not appear.

The value function defined above may not be sufficiently smooth for the derivatives defined above to exist. The theory of viscosity solutions has been developed in order to address this situation. The link between (3.15) and $(\mathcal{H}_\epsilon)$ is well known. We refer to Øksendal and Sulem (2007) for a recent review concerning the control of more general jump processes than the ones considered here. The book (Yin & Zhang, 1998) also contains an overview of viscosity solutions. We will need the properties of viscosity solution later and so we define them below.

Let D^+v (D^-v) denote the superdifferential (subdifferential) of v with respect to w. These two differentials are defined as follows,

$$D^+v = \left\{p \middle| \limsup_{\widehat{w}\to w} \frac{v_t(\widehat{w},\beta) - v_t(w,\beta) - p\cdot(\widehat{w}-w)}{|\widehat{w}-w|} \leqslant 0\right\},$$

$$D^-v = \left\{q \middle| \liminf_{\widehat{w}\to w} \frac{v_t(\widehat{w},\beta) - v_t(w,\beta) - q\cdot(\widehat{w}-w)}{|\widehat{w}-w|} \geqslant 0\right\}.$$

We will say that a continuous function u is a viscosity subsolution if,

$$\rho u_t(w,\beta) - \mathcal{H}^\epsilon(u,p) \leqslant 0 \quad \forall p \in D^+u,$$

and a viscocity supersolution if,

$$\rho u_t(w,\beta) - \mathcal{H}^\epsilon(u,q) \geqslant 0 \quad \forall q \in D^-u.$$

The solution will be called a viscosity solution if it is both a subsolution and supersolution. The following result can be proved using standard techniques (see for example Yin & Zhang, 1998).

Theorem 3.7. *Under Assumptions 3.1, 3.2, 3.3 the value function of $(\mathcal{H}_\epsilon)$ is the unique viscosity solution of (3.15).*

We next show that if ϵ is small enough then the value function does not depend on the exact state η but rather the cluster of states η belongs to.

Theorem 3.8. *Suppose that η belongs to the kth set of aggregate states, $\eta \in \mathcal{M}^k$. Then for any convergent subsequence of $\{v_t^\epsilon\}_\epsilon$ we must have.*

$$v_t^\epsilon(w,\eta,\xi) \to v_t(w,k,\xi) \quad \forall \eta \in \mathcal{M}_k \tag{3.16}$$

Proof. Let τ_i be defined as follows,

$$\tau_i = \inf\{s | y_i(s) = K_i,\ s > t\}.$$

Let τ^y be defined as the stopping time when the first plant gets expanded after time t,

$$\tau^y = \min_i\{\tau_i\}.$$

Also let $\tau^{\epsilon,\eta}$, and $\tau^{\epsilon,\xi}$ denote the jump times of η and ξ respectively. We first show that if ϵ is small enough then $\tau^{\epsilon,\eta} < \tau^y$ and $\tau^{\epsilon,\eta} < \tau^{\epsilon,\xi}$ w.p 1. Let $\delta > 0$ be arbitrary. Then,

$$\begin{aligned} P[\tau^{\epsilon,\eta} - t \in (0,\delta) | t, \eta(\epsilon,t) = \eta] &= -\int_0^\delta e^{\int_t^{t+s} q_{\eta,\eta} + \frac{1}{\epsilon}w_{\eta,\eta}dr}(q_{\eta,\eta} + \frac{1}{\epsilon}w_{\eta,\eta})ds\\ &= 1 - e^{(q_{\eta,\eta}+\frac{1}{\epsilon}w_{\eta,\eta})\delta} \to 1. \end{aligned}$$

Therefore without loss of generality, we can assume that for ϵ small enough $\eta(t,\epsilon)$ jumps before $\xi(t)$ and before the next plant reaches level K w.p. 1. The rest of the proof is a straightforward extension of Lemma 9.6 from Yin and Zhang (1998) to the finite horizon case. □

4. The aggregate model

In this section we derive the limiting problem for $\mathcal{H}_\epsilon$ as $\epsilon\downarrow 0$. We show that the asymptotic value function of $\mathcal{H}_\epsilon$ corresponds to a value function of an aggregate model. We also show how to construct asymptotically optimal controls from the aggregate model.

For the application considered in this paper, it is sufficient to assume that each $\eta^i(\epsilon,t), i = 1,\ldots,N$, has the same generator both before and after it has been expanded. Thus only the states may be different, the transition probabilities remain the same. It is possible to handle the more general case with some additional increase in the notation but the argument remains the same.

The unique stationary distribution associated with the kth aggregate state satisfies the following system of linear equations,

$$\begin{aligned} &\lambda^k Q_k = 0,\\ &\sum_{j\in\mathcal{M}^k}\lambda_j^k = 1, \end{aligned} \tag{4.1}$$

where $\mathcal{M}^k$ denotes the states in the kth aggregate state. The generator Q_k is used to model the transition rates between the kth cluster of states. The existence and uniqueness of the solution to (4.1) follows from our assumption that each Q_k is irreducible.

The reduced order modeling framework adopted in this paper is illustrated in Fig. 3(a). We start from the bottom layer that captures all the intraday variations (e.g., the possibility of changing wind speeds every 5 minutes). This information is combined into aggregate states (the middle layer of Fig. 3(a)). Finally, some other probabilities are devised that control the transition between aggregate states. Our next task is to define the dynamics of this aggregate probability transition function. This can be done rigorously (see c.f. chap. 7 in Yin & Zhang (1998)). Since these dynamics play

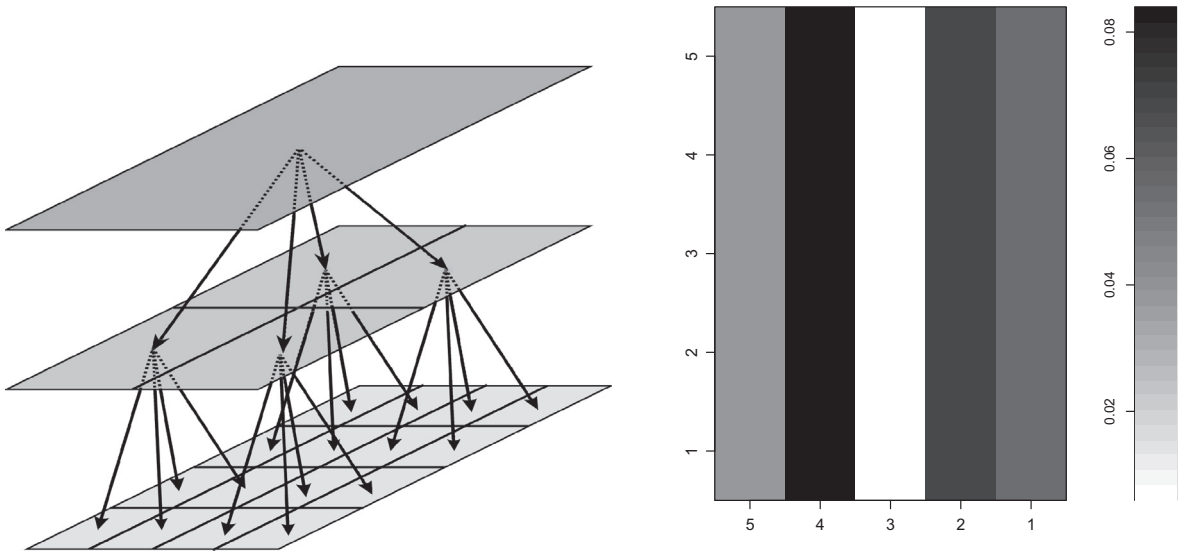

(a) Information structure of the aggregate model

(b) Equilibrium density of the aggregated Markov Chain

Fig. 3. The aggregated model depicted in (b) gives the correct long term behavior (compare (b) with Fig. 2(a)).

a central role in our analysis we sketch below how this derivation can be performed. To this end, let $\widehat{I} = \text{diag}(\mathbf{1}_{|\mathcal{M}^0|}, \ldots, \mathbf{1}_{|\mathcal{M}^l|})$. Where $|\mathcal{M}^k|$ denotes the number of states in the kth aggregate state, and $\mathbf{1}_n$ is a vector of ones of dimension n. The generator associated with the aggregated states dynamics is given by,

$$\overline{\Lambda} = \text{diag}(\lambda^1, \ldots, \lambda^l) W \widehat{I},$$

where W is associated with the transitions between aggregate sets of states in (3.2). Let $\bar{\eta}$ be the Markov process generated by $\overline{\Lambda}$. Its transition matrix satisfies the following,

$$\frac{d\overline{P}(t)}{dt} = \overline{P}(t)\overline{\Lambda}. \tag{4.2}$$

The preceding equation specifies the probability transition function of moving from one set of aggregate states to another. Next we give a heuristic argument of how (4.2) can be derived as the limit of (3.7). Let $\overline{P}_t(k)$ denote the probability of being at the kth cluster of aggregate states at time t. If $\epsilon > 0$ is small enough then,

$$P_t(j) \approx \overline{P}_t(k)\lambda^k_j, \quad j \in \mathcal{M}^k,$$

where P_t is given by (3.7). Using the equation above and (4.1) we obtain,

$$\left(\frac{d\overline{P}_t(1)}{dt}\lambda^1, \ldots, \frac{d\overline{P}_t(l)}{dt}\lambda^l\right) = \left(\sum_{k=1}^{l}\overline{P}_t(k)\lambda^k W^{k,1}, \ldots, \sum_{k=1}^{l}\overline{P}_t(k)\lambda^k W^{k,l}\right), \tag{4.3}$$

where $W^{i,j}$ are the sub-matrices of dimension $\mathbb{R}^{|\mathcal{M}^i|\times|\mathcal{M}^j|}$ obtained from (3.7). The expression in (4.3) can be further reorganized as follows,

$$\frac{d\overline{P}_t(m)}{dt} = \frac{d\overline{P}_t(m)}{dt}(\lambda^m \mathbf{1}_{|\mathcal{M}^m|}) = \sum_{k=1}^{l}\lambda^k\overline{P}_t(k)W^{k,m}\mathbf{1}_{|\mathcal{M}^m|} = \sum_{k=1}^{l}\overline{P}_t(k)\overline{\Lambda}^{k,m},$$

where $\overline{\Lambda}^{k,m} = \lambda^k W^{k,m}\mathbf{1}_{|\mathcal{M}^m|}$. Vectorizing the expression above, we obtain (4.2). To relate this to the discussion of Section 3.2, Fig. 3(b) plots the equilibrium density of $\overline{P}$ defined above. The intuition is that some states become so fast that they become transient. This allows us to aggregate the fast Markov Processes as depicted in Fig. 3(a). As a result the Markov Chain has fewer states (allowing reduction in computational complexity) but still has the same qualitative behavior. To see the latter point, compare Fig. 3(b) with the exact result in Fig. 2(a), and the current approach used in long term planning models Fig. 2(b). In the exact density, there is the possibility of transition between states. This possibility is assumed away in current models in order to reduce computational complexity. The approach proposed in this paper allows for the possibility of transition between clusters of states with a modest increase in computational complexity.

Remark 4.1. For later use we define, $\bar{\eta}^\epsilon(t) = k$, if $\eta^\epsilon(t) \in \mathcal{M}^k$. It was shown in Theorem 2.3 in Zhang and Yin (1997) that,

$$\lim_{\epsilon\to 0} d(\bar{\eta}^\epsilon(t), \bar{\eta}(t)) = 0,$$

where $d(\cdot,\cdot)$ denotes a distance metric on the Skorokhod topology.

Suppose that at time s the system is in the kth aggregate state. The objective function of the aggregate problem is given by,

$$J_s(w,k,\xi;u,\pi) \triangleq E\left\{\int_s^T \sum_{j\in\mathcal{M}^{\bar{\eta}(t)}} \lambda^{\bar{\eta}(t)}_j \left(\sum_{i=1}^{N} e^{-\rho(t-s)} G_i(x(s)), \pi_i(s), u_i^{\bar{\eta}(t),j}(s), \xi(s)\right) ds + e^{-\rho(T-s)}\Phi(X_T)\right\}$$

The aggregate optimal control problem is as follows.

$$\begin{aligned}
\bar{v}_s(w,k,\xi) = \min \quad & J_s(w,k,\xi;u,\pi) \\
& \frac{dx(t)}{dt} = \sum_{j\in\mathcal{M}^{\bar{\eta}(t)}} \lambda^{\bar{\eta}(t)}_j \left(\sum_i^N u_i^{\bar{\eta}(t),j}(s) - z(t) + \eta_0^{\bar{\eta}(t),j}\right), \\
& x(s) = x \\
& \frac{dz(t)}{dt} = f(z(t), \xi^0(t)), \quad z(s) = z \\
& \frac{dy(t)}{dt} = \sum_{j\in\mathcal{M}^{\bar{\eta}(t)}} \lambda^{\bar{\eta}(t)}_j \pi(t), \quad y(s) = y \quad \mathcal{H} \\
& y_i(t) \leqslant K_i, \\
& u_i^{\bar{\eta}(t),j}(t) \in \mathcal{C}_i(\bar{\eta}(t), j), \\
& \bar{\eta}(s) = k.
\end{aligned}$$

Note that the value function is no longer a function of the fast process η, but rather of the current aggregate state k the fast process belongs to. The set $\mathcal{C}(i,j)$ is the feasible set associated with the ith state inside the jth aggregate set of states for the kth plant. In the case with no startup costs then this will be simply given by:

$$0 \leqslant u_i^{\bar{\eta}(t)j}(t) \leqslant \eta_i^{\bar{\eta}(t)j}.$$

For later use we denote the state equations above as follows:

$$\frac{d\bar{w}(t)}{dt} = g(\bar{w},\xi,\eta,\pi,u), \quad \bar{w}(s) = w. \tag{4.4}$$

The constraints on the controls of the aggregate problem will be denoted by $\overline{\mathcal{C}}$.

4.1. Asymptotic analysis of the value function

The main results of this Section are given in this and the next subsection. A similar analysis that covers infinite horizon models, and when the states of the Markov Processes remain constant (i.e., no possibility of capacity expansion) appeared in Zhang et al. (1997), Yin and Zhang (1998), Sethi and Zhang (1994). As noted in the end of Section 2 and the beginning of Section 3.3.1 we extend the aforementioned results to deal with the finite horizon case, the possibility of expanding capacity, and we allow for both weakly interacting processes (η^ϵ) and slow processes (ξ) in the same model.

The convergence proof of v_t^ϵ to $\bar{v}_t$ is given below. As the proof is somewhat technical, we first provide a conceptual explanation of how the exact and aggregate model are related in order to build intuition. We know from the theory of stochastic optimal control that the optimal controls of $(\mathcal{H}_\epsilon)$ are functions of w, ξ and η^ϵ. Suppose that we replaced the full distribution of η^ϵ with the aggregate distribution given by (4.2). The states of the "new" approximate distribution are no longer the states of η^ϵ but the cluster of states that η^ϵ belongs to. Thus the value function will no longer be a function of η^ϵ, but it will be a function of the cluster of states that η^ϵ belongs to. The exact state of η^ϵ can therefore be integrated out from the value function. Since we are optimizing over the set of closed feedback controls, we also need to integrate out the exact state from the controls. The rest of this Section provides the proof of how "integrating out" the exact state works in theory. In Section 5 we put the theory into practice.

The value function of $(\mathcal{H})$ can be shown to be the unique viscosity solution of the following HJB equation,

$$\rho\bar{v}_t(w,k,\xi) = \min\left\{\sum_{j\in\mathcal{M}^k}\lambda_j^k\left(\sum_i G_i(x,\pi^i,u_i^j,\xi) + \partial_x\bar{v}_t\left(\sum_i u_i^j - z - \eta_0^{kj}\right) + \nabla_y\bar{v}_t\cdot\pi\right)\right\} + \partial_z\bar{v}_t f(z,\xi) + \partial_t\bar{v}_t + \mathcal{L}\bar{v}_t[k,\xi], \tag{4.5}$$

where

$$\mathcal{L}_\xi\bar{v}_t[k,\xi] = \mathcal{L}_{\bar{\xi}}\bar{v}_t[k,\xi] + \overline{\mathcal{L}}_\eta\bar{v}_t[k,\xi],$$

and,

$$\mathcal{L}_{\bar{\xi}}\bar{v}_t[k,\xi] = \sum_{\widehat{\xi}\neq\xi} r_{\widehat{\xi}\xi}(\bar{v}_t(w,k,\widehat{\xi}) - \bar{v}_t(w,k,\xi)),$$

$$\overline{\mathcal{L}}_\eta\bar{v}_t[k,\xi] = \sum_{i\neq k}\bar{\lambda}_{ik}(\bar{v}_t(w,i,\xi) - \bar{v}_t(w,k,\xi)).$$

The main result concerning the asymptotic behavior of the value function is given below.

Theorem 4.2. *Suppose that $\eta \in \mathcal{M}_k$, then*

$$\lim_{\epsilon\to 0} v_t^\epsilon(w,\eta,\xi) = \bar{v}_t(w,k,\xi) \quad \forall\eta\in\mathcal{M}_k$$

Where $\bar{v}_t(w,k,\xi)$ is the viscosity solution of $(\mathcal{H})$.

Proof. It was shown in Corollary 3.5 that there exists a subsequence of $\{v_t^\epsilon\}$ that converges. For such a converging subsequence we know, from Theorem 3.8, that $v_t^\epsilon(w,\eta,\xi)\to v_t(w,k,\xi)$. The basic idea behind this proof is to show that $v_t(w,k,\xi)$ is a viscosity solution to the aggregated HJB Eq. (4.5). The uniqueness of the viscosity solution will then imply the result, $v_t(w,k,\xi) = \bar{v}_t(w,k,\xi)$.

For the given fixed k, and for any $\bar{w}$, of appropriate dimension, let

$$\varphi_t(w,k,\xi) = v_t(\bar{w},k,\xi) + \bar{\gamma}\cdot(w-\bar{w}) - \frac{1}{2}|w-\bar{w}|^2, \tag{4.6}$$

where $\bar{\gamma}\in D^- v(\bar{w},k,\xi)$. Then,

$$v_t(\bar{w},k,\xi) - \varphi_t(\bar{w},k,\xi) = \min_w v_t(w,k,\xi) - \varphi_t(w,k,\xi).$$

Due to the definition of φ in (4.6) the above minimum is strict. Next, choose $w_{j,k}^\epsilon$, so that for each $\eta_j\in\mathcal{M}^k$, the following holds,

$$v_t^\epsilon(w_{j,k}^\epsilon,\eta,\xi) - \varphi_t(w_{j,k}^\epsilon,k,\xi) = \min_w v_t^\epsilon(w,\eta,\xi) - \varphi_t(w,k,\xi) \quad \forall\eta_j\in\mathcal{M}^k.$$

It follows that,

$$\lim_{\epsilon\to 0} w_{j,k}^\epsilon = \bar{w}.$$

As v_t^ϵ is a viscosity supersolution to (3.15), we have for any feasible (u,π),

$$\begin{aligned}\rho v_t^\epsilon(w_{j,k}^\epsilon,\eta^{j,k},\xi) \geqslant{}& \sum_i G_i(x_{j,k}^\epsilon,\pi^i,u^i,\xi) + \partial_x\varphi_t(w_{j,k}^\epsilon,\eta^{kj},\xi)\\ &\times\left(\sum_i u_i - z - \eta_0^{kj}\right) + \nabla_y\varphi_t(w_{j,k}^\epsilon,\eta^{kj},\xi)\cdot\pi\\ &+ \partial_z\varphi_t(w_{j,k}^\epsilon,\eta^{kj},\xi)f(z,\xi) + \mathcal{L}^\epsilon v_t^\epsilon(w_{j,k}^\epsilon,\eta^{kj},\xi)[\eta^{kj},\xi].\end{aligned}$$

It follows that,

$$\rho\sum_{j\in\mathcal{M}^k}\lambda_j^k v_t^\epsilon(w_{j,k}^\epsilon,\eta^{kj},\xi) \geqslant \sum_{j\in\mathcal{M}^k}\lambda_j^k\left(\sum_i G_i(x_{j,k}^\epsilon,\pi^i,u^i,\xi) + \partial_x\varphi_t\left(\sum_i u_i - z - \eta_0^{kj}\right) + \nabla_y\varphi_t\cdot\pi + \partial_z\varphi_t f(z,\xi) + \mathcal{L}^\epsilon v_t^\epsilon[\eta^{kj},\xi]\right).$$

Note that,

$$\sum_{j\in\mathcal{M}^k}\lambda_j^k\mathcal{L}^\epsilon v_t^\epsilon[\eta^{kj},\xi] = \sum_{j\in\mathcal{M}^k}\lambda_j^k\mathcal{L}_{\bar{\xi}}v_t^\epsilon[\eta^{kj},\xi] + \sum_{j\in\mathcal{M}^k}\lambda_j^k\mathcal{L}_\eta^\epsilon v_t^\epsilon[\eta^{kj},\xi]$$

The last term in the preceding equation can be decomposed as follows,

$$\begin{aligned}\sum_{j\in\mathcal{M}^k}\lambda_j^k\mathcal{L}_\eta^\epsilon v_t^\epsilon[\eta^{kj},\xi] &= \sum_{j\in\mathcal{M}^k}\lambda_j^k\sum_{m\neq j}\left(w_{m,j(k)} + \frac{q_{m,j(k)}}{\epsilon}\right)\left(v_t^\epsilon(w_{j,k}^\epsilon,\eta^m,\xi) - v_t^\epsilon\left(w_{j,k}^\epsilon,\eta^{kj},\xi\right)\right)\\ &= \sum_{j\in\mathcal{M}^k}\lambda_j^k\sum_{m\neq j}w_{m,j(k)}\left(v_t^\epsilon(w_{j,k}^\epsilon,\eta^m,\xi) - v_t^\epsilon\left(w_{j,k}^\epsilon,\eta^{kj},\xi\right)\right)\\ &\quad + \sum_{j\in\mathcal{M}^k}\lambda_j^k\sum_{m\neq j}\frac{q_{m,j}^k}{\epsilon}\left(v_t^\epsilon\left(w_{j,k}^\epsilon,\eta^m,\xi\right) - v_t^\epsilon\left(w_{j,k}^\epsilon,\eta^{kj},\xi\right)\right)\end{aligned}$$

Let $I_1(\epsilon)$ and $I_2(\epsilon)$ denote the first and second term, respectively, in the preceding equation. Then it follows that,

$$\begin{aligned}\lim_{\epsilon\to 0} I_1(\epsilon) &= \sum_{j\in\mathcal{M}^k}\lambda_j^k\sum_{\substack{m\neq j\\ m\notin\mathcal{M}^k}}(v_t(\bar{w},\hat{k},\xi) - v_t(\bar{w},k,\xi))1_{\{\eta^m\in\widehat{\mathcal{M}^k}\}}\\ &= \sum_{m\neq k}\bar{\lambda}_{mk}(v_t(\bar{w},m,\xi) - v_t(\bar{w},k,\xi)) = \overline{\mathcal{L}}v_t[k,\xi]\end{aligned} \tag{4.7}$$

Using the fact that,

$$v_t^\epsilon(w_{j,k}^\epsilon,\eta^{kj},\xi) - \varphi_t(w_{j,k}^\epsilon,k,\xi) \leqslant v_t^\epsilon(w_{i,k}^\epsilon,\eta^{kj},\xi) - \varphi_t(w_{i,k}^\epsilon,k,\xi) \quad \forall i \in \mathcal{M}^k.$$

The second term, $I_2(\epsilon)$, can be bounded as follows,

$$\begin{aligned}
&\frac{1}{\epsilon}\sum_{j\in\mathcal{M}^k}\lambda_j^k\sum_{m\neq j}q_{m,j}^k\left(v_t^\epsilon\left(w_{j,k}^\epsilon,\eta^{k,m},\xi\right)-v_t^\epsilon\left(w_{j,k}^\epsilon,\eta^{k,j},\xi\right)\right)\\
&=\frac{1}{\epsilon}\sum_{j\in\mathcal{M}^k}\lambda_j^k\sum_{m\neq j}q_{m,j}^k\left(\left(v_t^\epsilon\left(w_{j,k}^\epsilon,\eta^{k,m},\xi\right)-\varphi_t\left(w_{j,k}^\epsilon,k,\xi\right)\right)\right.\\
&\quad\left.-\left(v_t^\epsilon\left(w_{j,k}^\epsilon,\eta^{k,j},\xi\right)\right)-\varphi_t(w_{j,k}^\epsilon,k,\xi)\right)\\
&\geqslant\frac{1}{\epsilon}\sum_{j\in\mathcal{M}^k}\lambda_j^k\sum_{m\neq j}q_{m,j}^k\left(\left(v_t^\epsilon\left(w_{m,k}^\epsilon,\eta^{k,m},\xi\right)-\varphi_t\left(w_{m,k}^\epsilon,k,\xi\right)\right)\right.\\
&\quad\left.-\left(v_t^\epsilon(w_{j,k}^\epsilon,\eta^{k,j},\xi))-\varphi_t\left(w_{j,k}^\epsilon,k,\xi\right)\right)\right.\\
&=\frac{1}{\epsilon}\sum_{m\in\mathcal{M}^k}\left(\left(v_t^\epsilon(w_{m,k}^\epsilon,\eta^{k,m},\xi)-\varphi_t(w_{m,k}^\epsilon,k,\xi)\right)\sum_{j\in\mathcal{M}^k}\lambda_j^k q_{m,j}^k-\frac{1}{\epsilon}\lambda_j^k(v_t^\epsilon(w_{j,k}^\epsilon,\eta^{k,j},\xi))\right.\\
&\quad\left.-\varphi_t(w_{j,k}^\epsilon,k,\xi)\right)\sum_{m\in\mathcal{M}^k}q_{m,j}^k=0.
\end{aligned}$$

Therefore,

$$\begin{aligned}
\rho\sum_{j\in\mathcal{M}^k}\lambda_j^k v_t^\epsilon(w_{j,k}^\epsilon,\eta^{k,j},\xi)\geqslant\sum_{j\in\mathcal{M}^k}\lambda_j^k\left(\sum_i G_i(x_{j,k}^\epsilon,\pi^i,u^i,\xi)+\partial_x v_t^\epsilon\left(\sum_i u_i-z-\eta_0^{k,j}\right)\right.\\
\left.+\partial_t v_t^\epsilon\nabla_y v_t^\epsilon\cdot\pi+\partial_z v_t^\epsilon f(z,\xi)+\sum_{j\in\mathcal{M}^k}\lambda_j^k\mathcal{L}_\xi v_t^\epsilon[\eta^{k,j},\xi]+I_1(\epsilon)\right).
\end{aligned}$$

Taking the limit in the preceding equation we find,

$$\begin{aligned}
\rho\sum_{j\in\mathcal{M}^k}\lambda_j^k v_t(\bar{w},\eta^{k,j},\xi)\geqslant\sum_{j\in\mathcal{M}^k}\lambda_j^k\left(\sum_i G_i(\bar{x},\pi^i,u^i,\xi)+\partial_x\varphi_t\left(\sum_i u_i-z-\eta_0^{k,j}\right)\right.\\
\left.+\partial_t v_t\nabla_y\varphi_t\cdot\pi+\partial_z\varphi_t f(z,\xi)+\sum_{j\in\mathcal{M}^k}\lambda_j^k\mathcal{L}_\xi v_t[\eta^{k,j},\xi]+\bar{\mathcal{L}}_\eta v_t[k,\xi]\right).
\end{aligned}$$

Which shows that v_t is a viscosity supersolution of (4.5).

The proof can be completed by showing that v_t is also viscosity subsolution of (4.5). This argument is symmetric to the one given above so we omit the details. □

4.2. Asymptotic analysis of the optimal controls

We conclude this section by showing that if ϵ is small enough then the optimal control of the aggregate model can be used to construct an optimal control for the original problem.

Consider the following approximate control,

$$(\pi^k,u^k)\in\operatorname{argmin}_{\pi,u}\left\{\sum_{j\in\mathcal{M}^k}\lambda_j^k\left(\sum_i G_i(x,\pi^i,u^{i,j},\xi)+\partial_x\bar{v}_t\left(\sum_i u^{i,j}\right)+\nabla_y\bar{v}_t\cdot\pi\right)\right\}. \tag{4.8}$$

Using the above consider the following dynamics,

$$\begin{aligned}
\frac{dw^\epsilon}{dt}&=\sum_{k=1}^{I}\sum_{j\in\mathcal{M}^k}\mathbb{1}_{\{\eta^\epsilon(t)=\eta^{k,j}\}}g(w^\epsilon,\xi,\eta,\pi^k,u^{k,j})\quad w^\epsilon=w_0,\\
\frac{d\bar{w}^\epsilon}{dt}&=\sum_{k=1}^{I}\sum_{j\in\mathcal{M}^k}\mathbb{1}_{\{\bar{\eta}^\epsilon(t)=k\}}\lambda_j^k g(\bar{w}^\epsilon,\xi,\eta,\pi^k,u^{k,j})\quad \bar{w}^\epsilon=w_0.
\end{aligned} \tag{4.9}$$

For the definition of $\bar{\eta}^\epsilon(t)$ see Remark 4.1. As in Section 3.3.1 we use the concept of a generalized solution for the two ODEs above. Using the same notation as before ψ_{w^ϵ} and $\psi_{\bar{w}^\epsilon}$ will be used to denote the solutions (in the sense of Caratheodory) of the two equations above. As before, Assumption 3.3 and Theorem 2.2.1 in Arnold (1998) can be used to establish the existence and uniqueness of the solutions of (4.9).

Theorem 4.3. *Let (π^k,u^k) be defined as in (4.8), then*

(a) $$\lim_{\epsilon\to 0}E[|\psi_{w^\epsilon}-\psi_{\bar{w}^\epsilon}|]=0$$

(b) $$\lim_{\epsilon\to 0}E[|\psi_{\bar{w}}-\psi_{\bar{w}^\epsilon}|]=0$$

(c) $$\lim_{\epsilon\to 0}|J_t^\epsilon(w,\eta,\xi;u,\pi)-v_t^\epsilon(w,\eta,\xi)|=0$$

Proof.

(a) Let,

$$I_0(s)=E[|\psi_{w^\epsilon}(s,\xi,\eta,\pi,\omega)(w)-\psi_{\bar{w}^\epsilon}(s,\xi,\eta,\pi,\omega)(w)|]$$

$$\begin{aligned}
I_1(s,k,j)=g(\psi_{w^\epsilon}(s,\xi,\eta,\pi,\omega)(w),\xi,\eta,\pi^k,u^{k,j})-g(\psi_{\bar{w}^\epsilon}(s,\xi,\eta,\pi,\omega)\\
\times(w),\xi,\eta,\pi^k,u^{k,j}).
\end{aligned}$$

Then,

$$\begin{aligned}
E[I_0(t)]\leqslant E\left[\left|\sum_{k=1}^{I}\sum_{j\in\mathcal{M}^k}\int_0^t I_1(s,k,j)\mathbb{1}_{\{\eta^\epsilon(s)=\eta^{k,j}\}}ds\right|\right]\\
+E\left[\left|\sum_{k=1}^{I}\sum_{j\in\mathcal{M}^k}\int_0^t g(\psi_{\bar{w}^\epsilon}(s,\xi,\eta,\pi,\omega)(w),\xi,\eta,\pi^k,u^{k,j})\left(\mathbb{1}_{\{\eta^\epsilon(s)=\eta^{k,j}\}}-\mathbb{1}_{\{\bar{\eta}(s)=k\}}\lambda_j^k\right)ds\right|\right].
\end{aligned} \tag{4.10}$$

Using Theorem 2.11 in Arnold (1998) the following estimate can be used for the second term above,

$$\begin{aligned}
I_2(\epsilon,t)=E\left[\left|\sum_{k=1}^{I}\sum_{j\in\mathcal{M}^k}g(\psi_{\bar{w}^\epsilon}(s,\xi,\eta,\pi,\omega)(w),\xi,\eta,\pi^k,u^{k,j})\int_0^t\left(\mathbb{1}_{\{\eta(s)=\eta^{k,j}\}}-\mathbb{1}_{\{\bar{\eta}(s)=k\}}\lambda_j^k\right)ds\right.\right.\\
\left.\left.+\int_0^t\frac{dg}{dt}\left(\int_0^s(\mathbb{1}_{\{\eta(\tau)=\eta^{k,j}\}}-\mathbb{1}_{\{\bar{\eta}(\tau)=k\}}\lambda_j^k)d\tau\right)dt\right|\right].
\end{aligned} \tag{4.11}$$

It follows from Theorem 2.2 in Zhang and Yin (1997) that as $\epsilon\to 0$ we must have,

$$E\left[\left|\left(\mathbb{1}_{\{\eta(t)=\eta^{k,j}\}}-\mathbb{1}_{\{\bar{\eta}(t)=k\}}\lambda_j^k\right)\right|\right]\to 0.$$

The result above together with the boundedness assumptions on g imply that $I_2(\epsilon,t)\to 0$. Using the assumption that g is Lipschitz on w then we have,

$$I_0(t)\leqslant I_2(\epsilon,t)+\widehat{K}_g\int_0^t I_0(s)ds.$$

Then Gronwall's inequality implies,

$$I_0(t)\leqslant I_2(\epsilon,t)+\widehat{K}_g\int_0^t I_2(\epsilon,s)e^{\int_0^s I_0(r)dr}ds.$$

from which (a) follows.

(b) is an easy consequence of Remark 4.1.

(c) Given (a) and (b) above, the result can be proved in the same way as Theorem 9.11 in Yin and Zhang (1998).

□

5. Numerical experiments

In this section we describe a numerical implementation of the theory presented above. Our aim is to compare the aggregate with the exact model. For this reason we will only consider a small problem so that the exact problem is tractable. The problem we consider requires the solution of a five dimensional non-linear PDE, still a non-trivial problem by any means. We do however show that using the aggregate model we can reduce the computation time without introducing any significant deterioration in the accuracy of the solution. In fact, the solution times are comparable to those of a three dimensional problem. Thus a significant gain in computational time is achieved. We also compare the obtained solution with the simple heuristic commonly used in energy planning models in order to deal with uncertainty.

5.1. Data

We consider a model where the decision maker has access to three generating technologies: a coal fired plant, a gas fired plant, and a wind farm. For simplicity, in this illustration, we assume that expansion is only possible for the wind farm. Moreover, the only sources of uncertainty are in hourly load demand and in wind speeds. These assumptions lead to a five dimensional HJB in the form of (3.15) for the exact model, and (4.5) for the approximate model. A numerical scheme based on a finite difference approximation is described in Section 5.2 below. The five dimensions are: cumulative level of energy nonserved (x), deterministic demand level (z), cumulative investment in wind capacity (y), load uncertainty (η^0), and wind speed uncertainty (η^1). This stylized model is not sufficient to draw any conclusions regarding the potential for each technology in the energy mix. However, the model is sufficient for the goals of this paper; i.e., to propose suitable structural representations of the fast scale dynamics that lead to a tractable model. In this example we start with a five dimensional problem, and show how to compute the optimal value function using a formulation of reduced dimensionality.

For fuel prices (coal and natural gas) we used the data from the EIA website (EIA, 2009). We obtained wind data from the NCDC (NCDC, 2009) website. The wind data contain the wind speeds at an hourly time interval over the years 2007–2008, in the Buzzard's bay area in south Massachusetts. We used typical efficiency rates for the coal, and gas fired plants from the IEA study (IEA, 2005). We used a typical wind generator based on the NEG Micon 1000/54 (1000 kilowatt rated power, 54 meter diameter rotor), with a cut wind speed of 4 meter/second, rated windspeed 14 meter/second and furling wind speed 25 meter/second. The data for the wind turbine are from Masters (2004). Finally, hourly load data for 2007 were used from the PJM midatlantic region (PJM.E). These are available from the PJM website.

As was mentioned in Section 3.1, we used a two step process to decompose the load demand data into a periodic deterministic component, and a stochastic component. The continuous time load demand data is modeled as:

$$H(t) = D(t) + \eta^0(t), \quad t \in [0, T]. \tag{5.1}$$

The deterministic component is given by,

$$D(t) = a + bt + \sum_{j=0}^{N} c_j \cos(2\pi\phi_j t + l_j) \quad t \in [0, T]. \tag{5.2}$$

Thus the deterministic component accounts for a linear trend in demand. The seasonal fluctuations are super imposed around this linear trend. In our numerical experiments we used $N = 5$. The deterministic harmonics of (5.2) are obtained from the peaks of the Fourier transform of the hourly load data. In a more realistic model the parameters of (5.2) could also be uncertain. The residual component is given by,

$$r(t) = A(t) - D(t),$$

where $A(t)$ represents the actual data (8760 points). We then fit a maximum likelihood Markov Chain to the residual. The resulting transition matrix of the Markov Chain is shown in Fig. 1(b). We followed a similar procedure for the wind data. The resulting Markov

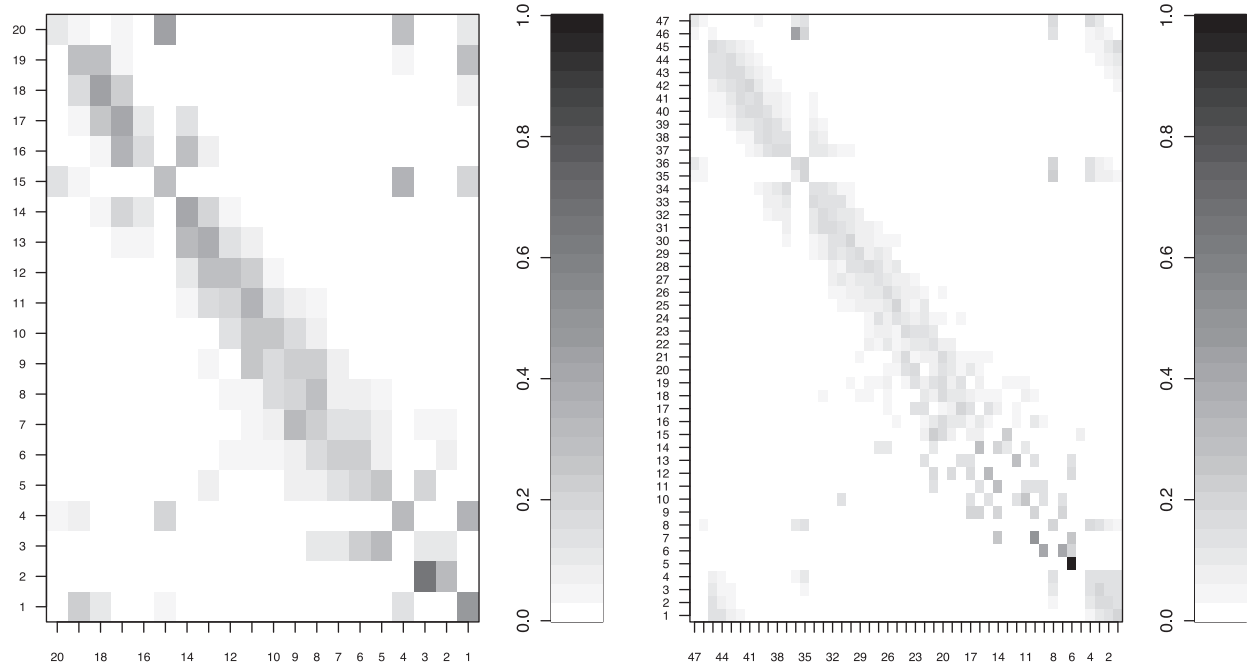

(a) 20 states (b) 50 states

Fig. 4. Markov transition matrices for wind uncertainty. (a) and (b) are obtained from a maximum likelihood fit to hourly time series.

Chain for wind uncertainty is shown in Fig. 4. Again, we observe a structure similar to that of the load demand data. At least at an intuitive level, one would expect that the approximations introduced in this paper, based on weakly interacting Markov processes, should perform well for this class of problems.

The next step is to construct the set of clustered states given the transition matrices estimated above. A procedure to do this is outlined in Phillips and Kokotovic (1981). However, for the application considered here we need to identify the set of aggregate states, and ensure the generator associated with each cluster generates a Markov process with a unique stationary distribution. Moreover, the off-diagonal matrix W must also be a valid generator. The off-diagonal matrix need not admit a unique stationary distribution. Our clustering procedure is based on a variation of spectral clustering (Nadler, Lafon, Coifman, & Kevrekidis, 2005). We use spectral clustering because it tends to cluster points that are on the same manifold (Nadler et al., 2005). This is a desirable property for the class of problems considered here. We want to aggregate points that are likely to have the same value function. This is not a standard application of spectral clustering since we do not define a kernel function. The reason we do not need to define a kernel function is that the original data comes from a Markov Chain. In spectral clustering, the first step is to define a kernel and then define a stochastic matrix given the kernel and the data.

Using the transition matrix estimated above, we calculate its eigenvalues and eigenvectors. We used the standard transformation $(P+P^T)/2$ to make it symmetric. We found no changes in the cluster composition by making this transformation. The sign of the second largest eigenvalue can be used to decompose the set of states into two clusters. Iterating this procedure gives us the set of clusters.

Let $\widehat{P}$ be the maximum likelihood transition matrix generated in the way described above. We will denote the orthogonal set of eigenvectors and eigenvalues of $\widehat{P}$ by $\{\psi_k, \phi_k, \lambda_k\}$ so that,

$$\widehat{P}_{i,j} = \sum_{k=1}^{m} \psi_k(i)\phi_k(j)\lambda_k.$$

Then the maximum likelihood generator can be obtained from,

$$\widehat{Q}_{i,j} = \sum_{k=1}^{m} \frac{1}{\delta t}\psi_k(i)\phi_k(j)\log\lambda_k,$$

where δt is a small step size. We used $\delta t = 1/8760$ in our numerical simulations. The method described above transforms the ML transition matrix into the ML generator. In some practical cases it is known to pose numerical and theoretical issues (Crommelin & Vanden-Eijnden, 2009). We found no numerical issues associated with the procedure above. In the future this part of the proposed methodology can be made more robust by using the quadratic model proposed in Crommelin and Vanden-Eijnden (2009). For the data set we used in our numerical experiments, the results would be identical.

The final step in the estimation procedure is to enforce the block diagonal assumption. This step is necessary because we found that there are small numerical differences between the estimated generators and the assumed block diagonal structures. The effect of enforcing this structure can be easily quantified for the example application considered in this paper by solving the exact model with the original maximum likelihood generators. We found that no significant discrepancies between the solutions obtained with the two generators. In order to ensure that our assumptions hold, we solve the following optimization problem.

$$\begin{aligned}
\min_{Q,W,\lambda} \quad & \sum_k \sum_{i,j} \left(W_{i,j}^{k,k} + \tfrac{1}{\epsilon} Q_{i,j}^k - \widehat{Q}_{i,j}^{k,k}\right)^2 + \sum_{k\neq m}\sum_{i,j}\left(W_{i,j}^{k,m} - \widehat{Q}_{i,j}^{k,m}\right)^2 \\
& \lambda_k^T Q^k = 0, \quad k = 1,\ldots,l \\
& Q^k 1 = 0, \quad k = 1,\ldots,l \\
& W1 = 0, \\
& \lambda_k^T 1 = 1, \quad \lambda_k \geqslant 0,
\end{aligned} \tag{5.3}$$

where $\widehat{Q}_{i,j}^{k,m}$ represents the rate of transition from the ith state in the kth cluster to the jth state in the mth cluster of the estimated maximum likelihood generator. Similar notation is used for W and Q. The λ_k represents the stationary distribution of the kth cluster. In (5.3), and throughout our numerical simulations with the "exact" model we used $\epsilon = 0.1$. We have empirically observed that this parameter value is a good fit to the data. The objective ensures that the estimated generator is as close as possible to the original generator. The constraints enforce the specific structure assumed in the paper. We empirically found that the constraints in (5.3) make the problem difficult to solve numerically. However, using our assumption that the block diagonal matrices must be irreducible and the change of variables given by:

$$R_j^k = \lambda_j Q_j^k,$$

where Q_j^k is the jth column of Q^k, the problem is transformed into the following well behaved optimization problem,

$$\begin{aligned}
\min_{R,W,\lambda} \quad & \sum_k \sum_{i,j} \left(\lambda_i W_{i,j}^{k,k} + \tfrac{1}{\epsilon} R_{i,j}^k - \lambda_i \widehat{Q}_{i,j}^{k,k}\right)^2 + \sum_{k\neq m}\sum_{i,j}\left(W_{i,j}^{k,m} - \widehat{Q}_{i,j}^{k,m}\right)^2 \\
& R^k 1 = 0, \quad k = 1,\ldots,l \\
& 1^T R^k = 0, \quad k = 1,\ldots,l \\
& W1 = 0, \\
& \lambda_k^T 1 = 1, \quad \lambda_k \geqslant 0.
\end{aligned}$$

As mentioned earlier, we found that the procedure outlined above gives good results. The matrices obtained are sufficiently close to the original maximum likelihood generator. However, in the future we plan to compare this method with other alternatives. For example, one could perform the joint estimation of the maximum likelihood generator and impose additional constraints on the structure of the matrix.

5.2. Finite difference approximation

In order to solve both the exact and the aggregate model, an upwind finite difference was used. The theoretical properties of similar schemes have been considered in Kushner and Dupuis (2001). Similar arguments can be used for the analysis of the algorithm described below. Here we just sketch out the necessary ingredients for such an analysis. The algorithm is described using the aggregate model in (4.5).

Let $[\Delta t, \Delta w] = [\Delta t, \Delta x, \Delta z, \{\Delta y_i\}_{i=1}^n]$ denote the finite difference discretization intervals along each direction. The partial derivative with respect to the state w will be approximated as follows:

$$\partial_{w_i} v_{t-\Delta t}(w) = \begin{cases} \frac{\bar{v}_t(w_i^+) - \bar{v}_t(w)}{\Delta w_i} + O(\Delta w_i) & \text{if } g_i(w) \geqslant 0 \\ \frac{\bar{v}_t(w) - \bar{v}_t(w_i^-)}{\Delta w_i} + O(\Delta w_i) & \text{if } g_i(w) < 0. \end{cases}$$

where g is defined by (4.4). By w_i^+, w_i^- we mean that the ith dimension in w is shifted by a positive, respectively negative amount Δw_i. For simplicity an explicit time stepping scheme was used. Of course, many other numerical algorithms can be used to solve the same problem. However, for the purposes of this paper the scheme below is sufficient.

Using the approximations outlined above we obtain the following,

$$\rho\bar{v}_t^{\Delta}(w,m,\xi) = \min_{(u,\pi,y)\in\bar{C}} \sum_{j\in\mathcal{M}^m} \lambda_j^m \left(\sum_i G_i(w,\pi^i,u^{ij},\xi) + \sum \left(\frac{\bar{v}_{t+\Delta t}^{\Delta}(w_i^+)}{\Delta w_i} \mathbb{1}_{\{g(w)\geqslant 0\}} + \frac{\bar{v}_{t+\Delta t}^{\Delta}(w_i^-)}{\Delta w_i} \mathbb{1}_{\{g(w)<0\}} - \frac{\bar{v}_{t+\Delta t}^{\Delta}(w)}{\Delta w_i} \right) |g(w)| + \frac{\bar{v}_{t+\Delta t}^{\Delta}(w) - \bar{v}_t^{\Delta}(w)}{\Delta t} \right) + \sum_{\widehat{\xi}} r_{\widehat{\xi}\xi}\, \bar{v}_{t+\Delta t}^{\Delta}(w,m,\widehat{\xi}) + \sum_i \bar{\lambda}_{i,m}\, \bar{v}_{t+\Delta t}^{\Delta}(w,i,\xi) + r_{\xi\xi}\, \bar{v}_t^{\Delta} + \bar{\lambda}_{m,m}\, \bar{v}_t^{\Delta}(w,m,\xi).$$

Let,

$$I^{\Delta}(t) = \left(\rho - r_{\xi\xi} - \bar{\lambda}_{m,m} + \frac{1}{\Delta t} \right)^{-1}.$$

Then using the last two expressions we obtain the following,

$$\mathcal{B}^{\Delta}[\bar{v}_{t+\Delta t}] = \min_{(u,\pi,y)\in\bar{C}} I^{\Delta}(t) \left\{ \sum_{j\in\mathcal{M}^m} \lambda_j^m \left(\sum_i G_i(w,\pi^i,u^{ij},\xi) + \sum \left(\frac{\bar{v}_{t+\Delta t}^{\Delta}(w_i^+)}{\Delta w_i} \mathbb{1}_{\{g(w)\geqslant 0\}} + \frac{\bar{v}_{t+\Delta t}^{\Delta}(w_i^-)}{\Delta w_i} \mathbb{1}_{\{g(w)<0\}} - \frac{\bar{v}_{t+\Delta t}^{\Delta}(w)}{\Delta w_i} \right) |g(w)| \right) + \sum_{\widehat{\xi}} r_{\widehat{\xi}\xi}\, \bar{v}_{t+\Delta t}^{\Delta}(w,m,\widehat{\xi}) + \sum_i \bar{\lambda}_{i,m}\, \bar{v}_{t+\Delta t}^{\Delta}(w,i,\xi) \right\}.$$

Thus the finite difference iteration can be written as,

$$\bar{v}_t^{\Delta}(w,m,\xi) = \mathcal{B}^{\Delta}[\bar{v}_{t+\Delta t}(w,m,\xi)], \tag{5.4}$$

with the boundary condition specified in Section 4.1. The infinite horizon finite difference approximation of (4.5) can be derived in an identical manner. It can easily be shown that it has a unique fixed point.

5.3. Comparison of the different models

Below we report on some computational experience with the numerical procedure outlined above. The algorithm was implemented in C++, and run on a modern desktop computer. Motivated by the contraction property above, our first numerical test is with respect to the speed of convergence of this map. This test has the advantage of eliminating any purely temporal effects. Both the exact and the aggregate problem were started from the same point of all zeros. Wind and load uncertainty were clustered into four aggregate states using the procedure outlined above. The computation times are shown in Table 1, along with the total number of iterations. In Table 2 the error is measured as follows:

$$\left\| \frac{\bar{v}(w) - v^{\epsilon}(w)}{v^{\epsilon}(w)} \right\|_{\infty}.$$

The value function for the heuristic method was obtained by optimizing for the set of points in the cluster and summing over the results. These preliminary results show that the error from the aggregate model is much smaller than the simple heuristic. The computational times for the aggregate and heuristic models are identical. What is strikingly obvious from Table 1 is how much faster the aggregate model converges to the fixed point. The reason for this is twofold. Firstly, we are only looping over aggregate states and not the complete state space so each iteration is faster. Secondly, as it can be seen when the iterative scheme is applied to the full model we observe what is sometimes referred to as the smoothing property. This property means that at initial iterations the high frequency component of the error is removed, while the low frequency remains the same. As a result, at the first few iterations we see a big drop in the error, but progress is much slower in the following iterations. The aggregate model does not have this property, partly because it iterates over a coarser grid. Clearly there exist links between multigrid methods Briggs and McCormick (2000), and the methods introduced in this paper. In future work we plan to investigate the implications of this connection.

Table 1
Solution times.

Model	Time (seconds)	Iterations
Full	3717.26	77
Aggregate	5.63	2

Table 2
Error.

Model	Error (%)
Aggregate	1.74
Heuristic	14.33

As a second test we solved the finite horizon problem with four hour steps for one year. Using the approximate problem the solution is obtained in about 90 minutes. We estimated that it would take approximately 2200 minutes to solve the exact model.[1] Thus for the simple model we tested in this paper, the proposed approach is an order of magnitude faster than solving the exact model. Unfortunately, the error is only guaranteed to be small if ϵ is small enough. It is currently not known how to derive useful error bounds given the input data and a specific ϵ. However, for the specific class of problems we discussed in this paper the approximations seem to be good enough, and they represent an improvement over the currently used heuristics.

5.4. Case study: large scale integration of wind power into an existing power system

The model we propose in this paper is particularly suited to address questions regarding the integration of technologies with fast dynamics (such as wind power) into an existing power system. An additional advantage of our model is that we can easily investigate the effect of lead times in capacity expansion decisions. This is done by imposing the constraint in (3.3) so that capacity cannot appear overnight. The aim of our case study is to investigate the impact of large scale integration of wind power to capacity expansion decisions of conventional thermal generators under the assumptions that investment decisions are irreversible and with realistic lead times.

In order to maximize the impact of wind power one possibility is to dispatch it first and then meet the residual demand with conventional power generators. This is not the only possibility, but it is a policy that is widely used or seriously considered in many countries. Our aim in this section is to analyze the implications of this policy to the capacity expansion decisions of other technologies that will be used to meet residual demand. Our analysis is based on two observations. Firstly residual demand will tend to have higher variance than the original demand. For example, if demand and wind power generation are independent then the variance of the residual demand will equal the variance of the original demand plus the variance of wind power availability. This observation is backed up by the data we used to calibrate our model. Our second observation is that the impact of this policy will be more significant on thermal generators that have relatively low start-up costs and high ramping rates. Based on these two observations we

[1] Due to the amount of CPU time required to solve the exact model we just solved the model for a single day and extrapolated the solution times for the whole year.

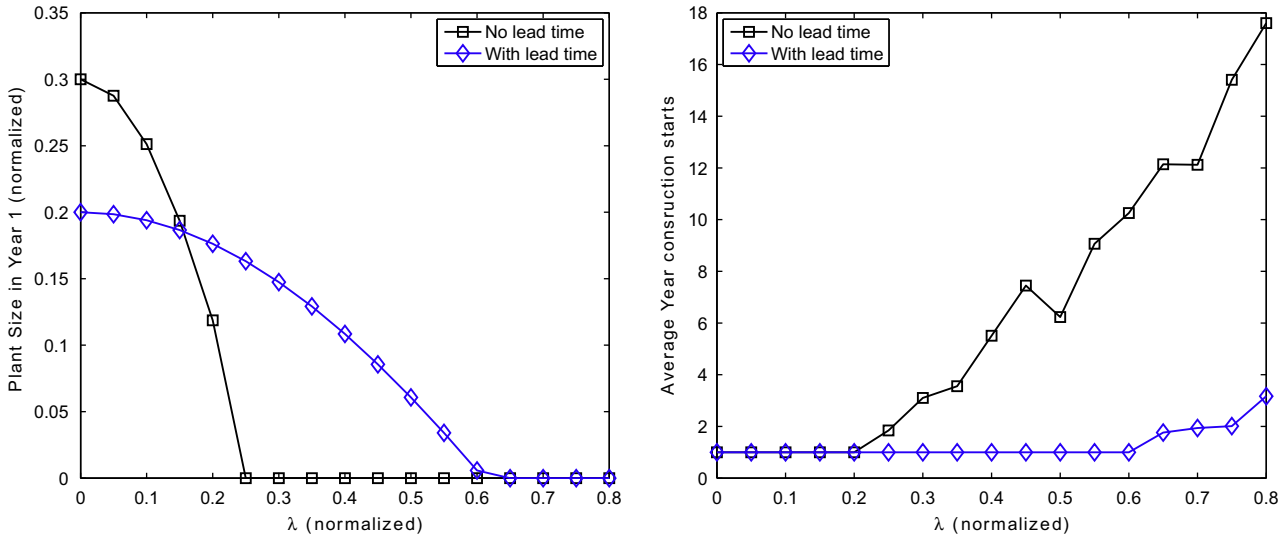

(a) Plant size in Year 1 (normalized) (b) Average year construction starts

Fig. 5. The effect of increasing wind penetration (in our model this means increasing λ).

consider a model where the objective of the decision maker is to satisfy residual demand by expanding capacity in a gas generator.

There are three sources of uncertainty: fuel price, capital cost and residual demand uncertainty. All sources of uncertainty were modeled using the continuous time Markov chain approach and dataset described in the previous sections. Residual demand (D_r) is derived by subtracting the two calibrated Markov processes:

$$D_r(t) = H(t) - \lambda W(t),$$

where H is the demand, λ is a fixed scalar denoting the installed capacity in wind, and W denotes the availability of wind power. Below we report results with normalized values of λ. Each finite state Markov Chain has 100 states. When $\lambda = 0$ all the demand has to be satisfied using gas, and when $\lambda = 1$ the amount of installed wind power capacity is equal to maximum demand (but not necessarily available). The state variables of our model are given by: capacity currently available, amount invested in new capacity, and the state of the gas generator. We used a time step of four hours over twenty years, and a discount rate of 5%. The value function associated with this model has six dimensions, and cannot be solved using an exact method. However the aggregate model described in Section 4 has three state variables and can therefore still be solved using the exact algorithm described in Section 5.2.

In our first experiment we did not impose the constraint in (3.3). This means that capacity expansions have no lead time so that additional capacity can be constructed and be available in the next time period. Note that the investment decisions are still non-anticipative, irreversible and lumpy (i.e. a certain amount of investment needs to be reached before additional capacity becomes available). As can be seen in Fig. 5(a) (the results labeled 'No lead time') as wind penetration increases the size of the gas generator installed in Year 1 decreases. However after a certain critical threshold (around 22% of peak demand with our parameters) construction in year 1 does not take place. After Year 1 it is difficult to see what capacity decisions are optimal. In order to get some idea of the optimal policy after year 1 we use the following procedure. We first construct the value function using the backward induction defined in (5.2). Given the value function approximation, we then estimate the optimal policy at every state. We then use a non-parametric regression method described in Parpas and Webster (2013) to estimate the control law. We then perform forward Monte Carlo simulations using the optimal control law. During the simulations we keep track of the first time construction starts. The results are shown in Fig. 5(b). Note that the results in Fig. 5(b) are noisy due to the forward simulation sampling and the approximate nature of the methodology used to construct the optimal control law. Despite these approximations errors the trend is quite clear. As the percentage of demand that can be met with wind is increased the investment in gas is postponed. With respect to the size of each plant we obtained a result similar to the Year 1 results shown in Fig. 5(a). In summary, the effect of increasing wind penetration seems to be delayed investment in gas technology and a smaller gas plant is built. For the latter result it is clear that as the proportion of demand that is satisfied from wind is increased then investment in gas falls. By increasing λ, the uncertainty of the residual demand increases while the mean of the residual demand decreases. In our model, we impose a high penalty for unsatisfied demand (50 $/kilowatt hour). But there is also a penalty for building unused capacity since expansion costs are irreversible. These two costs are competing. However, when there are no lead times, there is an asymmetry since if demand is greater than expected then, additional capacity can be quickly installed. This means that additional capacity can be delayed. This asymmetry is removed once realistic lead times are included. In Fig. 5(a) the plots labeled 'With lead time' were constructed by imposing the constraint in (3.3). With our choice of parameters the results with lead times required a construction time of at least 5 years (by limiting the amount that could be invested in each time period). We find that when lead times are taken into consideration then construction starts earlier but the size of the plant is smaller than in the case when no lead times are modeled.

6. Conclusions

We introduced a stochastic multiscale model that can be used for generation capacity expansion planning in the power systems sector. This class of models has stochastic dynamics that evolve over multiple temporal and spatial scales. Due to the computational challenges involved with simulating and optimizing integrated multiscale models, energy system models have been hitherto treated with each scale decoupled and interactions between scales ignored. However, with the introduction of new generation technologies (e.g., intermittent wind generation, dynamic demand response management, plug-in hybrid vehicles etc.), combined with the need for optimal utilization of existing and new generation capacity, it is becoming more and more important to

consider integrated models. We have empirically shown that some key uncertainties in this area exhibit structures that are amenable to multiscale mathematical techniques. We introduced a multiscale model and showed how the theory of singularly perturbed Markov processes can be used in order to reduce the dimensionality of the problem and therefore make it computationally tractable. We also discussed the application of the theory using real data and showed that the multiscale model can outperform heuristics that are currently in use while requiring a comparable amount of computational effort. Given the current need for advanced modeling techniques, together with a multitude of challenging theoretical and applied issues in this area, energy systems modeling can become a fruitful application area for stochastic multiscale methods.

Acknowledgement

The authors wish to acknowledge the two anonymous referees for their helpful comments that led to substantial improvements of the paper. The work of the first author was partially supported by a FP7 Marie Curie Career Integration Grant (PCIG11-GA-2012-321698 SOC-MP-ES) and by the Cyprus Program at MIT Energy Initiative. The work of the second author was supported by the U.S. National Science Foundation Grant No. 1128147 and by the US Department of Energy Office of Science, Biological and Environmental Research Program, Integrated Assessment Research Program, Grant Nos. DE-SC0005171 and DE-SC0003906.

References

Arnold, L. (1998). *Random dynamical systems*. *Springer monographs in mathematics*. Berlin: Springer-Verlag.

Bardi, M., & Capuzzo-Dolcetta, I. (1997). *Optimal control and viscosity solutions of Hamilton–Jacobi–Bellman equations*. Boston, MA: Birkhäuser Boston Inc..

Briggs, W., & McCormick, S. (2000). A multigrid tutorial. *Society for industrial mathematics*.

Crommelin, D., & Vanden-Eijnden, E. (2009). Data-based inference of generators for Markov jump processes using convex optimization. *Multiscale Modeling & Simulation, 7*, 1751.

EIA (2009). U.S. Energy Information Administration. <http://www.eia.doe.gov/>.

Epe, A., Küchler, C., Römisch, W., Vigerske, S., Wagner, H., Weber, C., et al. (2009). Optimization of dispersed energy supply – stochastic programming with recombining scenario trees. *Optimization in the Energy Industry*, 347.

IEA, N. (2005). *Projected costs of generating electricity 2005 update*.

Ilic, M., & Liu, S. (1996). *Hierarchical power systems control: its value in a changing industry*. Springer.

Jiang, J., & Sethi, S. (1991). A state aggregation approach to manufacturing systems having machine states with weak and strong interactions. *Operations Research, 39*(6), 970–978.

Kokotović, P., Khalil, H., & O'reilly, J. (1999). Singular perturbation methods in control: analysis and design. *Society for industrial mathematics*.

Kushner, H., & Dupuis, P. (2001). *Numerical methods for stochastic control problems in continuous time*. Springer Verlag.

Luss, H. (1982). Operations research and capacity expansion problems: A survey. *Operations Research, 30*(5), 907–947.

Masters, G. (2004). *Renewable and efficient electric power systems*. Wiley, Interscience.

Nadler, B., Lafon, S., Coifman, R., & Kevrekidis, I. (2005). *Diffusion maps, spectral clustering and eigenfunctions of fokker-planck operators*. Arxiv preprint math/0506090.

NCDC (2009). National Climatic Data Center. <http://www.ncdc.noaa.gov/oa/ncdc.html>.

Øksendal, B., & Sulem, A. (2007). Applied stochastic control of jump diffusions. *second ed*. Berlin: Universitext. Springer.

Palmintier, B., & Webster, M. (2011). Impact of unit commitment constraints on generation expansion planning with renewables. In *Proc. of 2011 IEEE power and energy society general meeting*, Detroit, MI.

Parpas, P., & Webster, M. (2013). A stochastic minimum principle and an adaptive pathwise algorithm for stochastic optimal control. *Automatica* (0).

Phillips, R., & Kokotovic, P. (1981). A singular perturbation approach to modeling and control of Markov chains. *IEEE Transactions on Automatic Control, 26*(5), 1087–1094.

Pirrong, C., & Jermakyan, M. (2008). The price of power: The valuation of power and weather derivatives. *Journal of Banking & Finance, 32*(12), 2520–2529.

Powell, W., George, A., Lamont, A., & Stewart, J. (2009). *SMART: A stochastic multiscale model for the analysis of energy resources, technology and policy*. <http://joc.journal.informs.org/content/24/4/665.abstract.

Pritchard, G., Philpott, A. B., & Neame, P. J. (2005). Hydroelectric reservoir optimization in a pool market. *Mathmatical Programming, 103*(3, Ser. A), 445–461.

Rudin, W. (1991). *Functional analysis. International series in pure and applied mathematics* (second ed.). New York: McGraw-Hill Inc..

Seebregts, A., Goldstein, G., & Smekens, K. (2001). Energy/environmental modeling with the MARKAL family of models. In *Proc. int. conf. on operations research (OR 2001)*.

Sen, S., Yu, L., & Genc, T. (2006). A stochastic programming approach to power portfolio optimization. *Operations Research, 54*(1), 55–72.

Sethi, S. P., & Zhang, Q. (1994). *Hierarchical decision making in stochastic manufacturing systems. Systems & control: foundations & applications*. Boston, MA: Birkhäuser Boston Inc..

Wallace, S., & Fleten, S. (2003). Stochastic programming in energy. In A. Ruszczynski, & A. Shapiro (Eds), *Handbooks in OR & MS*. Elsevier Science.

Weber, C. (2005). *Uncertainty in the electric power industry: Methods and models for decision support*. Springer Verlag.

Yin, G. (2009). Asymptotic expansions of option price under regime-switching diffusions with a fast-varying switching process. *Asymptotic Analysis, 65*(3), 203–222.

Yin, G. G., & Zhang, Q. (1998). *Continuous-time Markov chains and applications. Applications of Mathematics (New York)* (vol. 37). New York: Springer-Verlag (A singular perturbation approach).

Zhang, Q., & Yin, G. (1997). Structural properties of markov chains with weak and strong interactions. *Stochastic Processes and their Applications, 70*(2), 181–197.

Zhang, Q., Yin, G., & Boukas, E. (1997). Controlled Markov chains with weak and strong interactions: asymptotic optimality and applications to manufacturing. *Journal of Optimization Theory and Applications, 94*(1), 169–194.