Efficiency-Risk Tradeoffs in Electricity Markets with Dynamic Demand Response

Qingqing Huang, Mardavij Roozbehani, and Munther A. Dahleh

***Abstract*— In order to study the impact of dynamic demand response in the future smart grid, we examine in an abstract framework, how a tradeoff between efficiency and risk arises under different market architectures. We first examine the system performance under non-cooperative and cooperative market architectures. The statistics of the stationary aggregate demand processes show that, although the non-cooperative load scheduling scheme leads to an efficiency loss, the stationary distribution of the corresponding aggregate demand process has a smaller tail, resulting in less frequent aggregate demand spikes. Cooperative dynamic demand response, on the other hand, makes the market place more efficient at the cost of increased risk of aggregate demand spikes. The market architecture determines the locus of the system performance with respect to the tradeoff curve. We also investigate how a properly designed real-time electricity pricing mechanism can help the system operator achieve a target tradeoff between efficiency and risk in a non-cooperative market. We further provide a convex characterization of the Pareto front of system performance measures, which serves as a benchmark of the tradeoffs for the system operator to evaluate the pricing rules.**

I. Introduction

The distributed renewable energy sources pose new challenges for the design and operation of the future smart grid. On the supply side, the intermittency of the renewable sources introduces exogenous supply shocks. On the demand side, one challenge is system modeling and system design when large or perhaps small consumers may be able to actively respond to the real-time prices or other incentive signals [1], [2], [3]. A considerable amount of the consumer response will take the form of scheduling flexible loads, for example, electrical vehicle charging, building heating, and industrial processing. Several recent paper have examined optimization and control policies for scheduling such flexible electric loads, either under the assumption that the prices are exogenous, or that the players are small and therefore, price taking [4], [5], [6], [7], [8]. In the cases where the endogenously determined price is increasing in the instantaneous aggregate demand of finitely many agents, the problem falls into the category of dynamic oligopolistic competition [9]. The system model studied in this paper belongs to this category.

In this work, we shall examine the *market dynamics* in a dynamical system framework. Here, "dynamics" refers to the property that the demand decisions of the agents depend on the "state" of the system, whose evolution over time is in turn, determined by the agents' decisions. The notion of "state" refers to the market configuration. This will become more clear when precise definitions are given later in the sequel. From the perspective of system operation, we study the impact of the aggregate dynamic demand response on electricity markets, measured by "efficiency" and "risk". We observe that on one hand, dynamic demand response may largely smooth the aggregate demand process, thereby improving the system efficiency. On the other hand, agent interactions can generate aggregate demand spikes, which we define as endogenous risk in the electricity markets. Such exceedingly large demand / price spikes, introduce a level of volatility that can not only cause serious economic damage, but also undermine viability of electricity markets as a whole [1], [10]. Note that, our notion of "endogenous risk" of demand or price spikes is not the same as the financial risks associated with the real-time price uncertainties [11], [12], [13]. On the contrary, we examine how price volatility is shaped by the demand response. Furthermore, these spikes are endogenously created, and are therefore, distinctively different from the predicted peak hour loading.

In a more general setting, load scheduling, i.e., agents with coupled interests distributing their demand for a resource over multiple periods in order to minimize the expected cost of consumption, plays a crucial role in a wide array of applications. Examples include load scheduling in cloud computing under quality of service constraints [14], and multi-period rebalancing of multiple portfolio accounts in the presence of transaction costs [15]. In addition, many engineering disciplines actually do share compatible, if not identical, notions of the terms like "efficiency", "robustness", and "risk". So it is plausible that the messages in this study can be carried over to other engineering applications. In particular, for multi-agent systems, the mechanisms that can channel or even amplify uncertainties in the environment, i.e., exogenous shocks, into endogenous risks within the systems are still not well understood. Previous research efforts have explored various possible origins of endogenous risk, for example, failure of endogenizing other agents' action [16] and heterogeneous beliefs [17]. In our work, we provide an alternative explanation through a comparative study, and posit that they can arise from the nature of the system dynamics and agent interactions even at a complete information rational expectation equilibrium.

We consider a market where there are a large number of competitive electricity suppliers, aggregated into one representative supplier, and price anticipating self-interested agents, which enter and exit the market following a random process. The agents can model both consumers and the distributed renewable generators with potential load scheduling and storage technologies. Each agent activates a job upon arrival in the market which needs to be completed before his deadline. The agents dynamically schedule their electricity demand in reaction to real-time prices, which are set to be the marginal cost of the supplier at each time slot. When the agents schedule their load in a non-cooperative way, each agent optimizes his own expected cost; when the agents cooperate in the decision making process, the loads are scheduled in a way to optimize aggregate expected cost of all agents.

We observe that, when the agents schedule their loads in a cooperative way, they are more aggressive in absorbing exogenous uncertainties. As a result, they produce an aggregate demand process that is smoother on average and associated with a higher market efficiency. However, the tradeoff is a higher endogenous risk in terms of more aggregate demand spikes. This observation is consistent with the "robust yet fragile" property of complex systems, studied in [18], in the sense that the dynamic demand response can make the market place more robust to the relatively small and highly probable disturbances, but fragile for a different property (aggregate demand spikes) to certain less probable disturbances. The market architecture determines the locus of the system performance with respect to the tradeoff curve.

In [19], [20], [21], agent direct participation in the electricity market is modeled in a similar game theory based approach; and in [22], [23], the authors examined the application to electrical vehicle charging in a game theoretic approach with rational agents. However, in all the above works, the heterogeneous deadline constraints of individuals are not explicitly considered. We identify that the heterogeneous deadline constraints, together with the dynamic decision making process in response to the uncertainties, play a critical role in generating the endogenous risk of aggregate demand spikes. Therefore, the deadline constraints deserve a closer examination as proposed in this paper.

The remainder of the paper unfolds as follows. In Section II, we introduce the system model. In Section III, we focus on a specific case for which analytical solutions are obtained, and examine how various architectural properties affect the efficiency-risk tradeoffs. The analysis of this special case gives us insight on how to approximate the system by a more tractable model. The approximation leads to a Linear Time Invariant (LTI) system which yields qualitative messages that are consistent with the exact analysis of the special case. Proceeding with the approximated model, in Section IV, we discuss how the system operator's decision on the pricing rule will affect agent load scheduling behavior in a non-cooperative setup. In Section V, we provide a convex characterization of the Pareto front of performance measures, which dictates the fundamental tradeoff of the system with load scheduling dynamics. We conclude the paper with discussions on assumptions of the abstracted model and future work in Section VII. Due to space limitations, we only provide a sketch of the proofs in the appendix. Please see [24] for the details and supplementary materials.

II. System Model

In this section we introduce the system model consisting of heterogeneous agents with deadline constraints, competitive electricity suppliers, and the pricing scheme. We also define the non-cooperative and cooperative market architectures which define the nature of the interactions among the agents. We model consumer participation in the future smart grid, and focus on how the demand side of the wholesale spot market react to the uncertainties, which are in the form of unpredictable disturbances which force the system to deviate from the planned ahead generation and demand schedule. For the sake of analytical tractability, we chose an abstracted modeling approach, while retaining the most essential elements of electricity consumers' response to prices, namely the dynamic decision making process, heterogeneous deadline constraints, and the load arrival uncertainties. This abstract model fascilitates the main purpose of our analysis, which is to provide high level intuition and help understand important tradeoffs induced by the architectural properties of the system. More discussion on the modeling approach is included in Section VI.

II-A Agent Arrival Process

Agents in our model represent both consumers and the distributed renewable generations with potential load scheduling and storage techniques, and they enter and exit the market over time following a random process. The agent arrival process is a discrete time random process with time intervals indexed by $t = 0, 1, 2, \cdots$. When an agent arrives at the market, he activates a job that requires a certain amount of electricity in order to be completed. The agent needs to complete the job within a finite number of time intervals, and leave the market at his deadline. We define the length of time that an agent stays in the market to be his *type*, denoted by $l \in \mathcal{L} = \{1, \cdots, L\}$. Note that type 1 agents include the group of consumers with no real-time metering information who do not respond to the real-time price information. We assume that agents of type l arrive according to a Bernoulli process $\{h_l(t) \in \{0,1\} : t \in \mathbb{Z}\}$, with rate q_l. Namely, at each time slot, with probability q_l there is a type l agent arrival, which is independent of all other events. Upon arrival at time t, the type l agent activates a job of workload $d_l(t)$, which requires $d_l(t)$ units of electricity. We assume that $\{d_l(t) \in \mathbb{R} : t \in \mathbb{Z}\}$ is a white process with distribution D_l of mean $\mu_l = \mathbb{E}[D_l]$ and variance $\sigma_l^2 = \text{Var}[D_l]$. Let the L-dimensional column vectors $\mathbf{h}(t) = [h_l(t)] \in \{0,1\}^L$, and $\mathbf{d}(t) = [d_l(t)] \in \mathbb{R}^L$ denote the vector forms of arrival events and the corresponding workloads. Moreover, we assume that the agents have complete information about the market, and form rational expectations [25] about the real-time prices and the behavior of other agents.

II-B Real-time Electricity Pricing

We consider the case where the supply side consists of a large number of competitive and homogeneous electricity suppliers, which respond to the market prices as price takers in the same way. Therefore, we adopt the representative agent paradigm [26]. It is thus, sufficient to consider a single representative supplier, as the number of suppliers will be irrelevant for characterizing the equilibrium. Note that this is consistent with the paradigm of marginal cost pricing by the ISOs if the producers submit their true cost function to the ISO and follow the dispatch instructions. We assume that the representative supplier incurs a quadratic production cost of $\frac{1}{2}cV(t)^2$ at time t, where $V(t)$ is the total electricity generation, and c is a constant. Thus the marginal cost of production is $cV(t)$ and is increasing in the total supply. Let $U(t)$ denote the instantaneous aggregate electricity demand from all agents in the market. At the equilibrium where total supply equates total demand, i.e., $V(t) = U(t)$, the instantaneous electricity price $p(t)$ equals the marginal cost of production, namely $p(t) = cU(t)$. To simplify the notation, we let $c = 1$ throughout the rest of this paper.

We adopt a quadratic cost function for two reasons: firstly it constitutes a second-order approximation to other types of nonlinear cost functions, and secondly it captures the important fact of increasing marginal cost. Finally, in electricity markets, marginal cost pricing is a widely used mechanism [27]. It is well known that in a static setup, when both the suppliers and the consumers are competitive, marginal cost pricing leads to social optimality.

II-C System State Evolution

At each time t, there can exist more than one agent of each type, corresponding to agents that enter the market at different times before t, and with different deadlines in time. For example, there are at most $(L+1-\tau)$ agents who will stay in the market for τ periods for all $\tau \in \mathcal{L}$. They correpond to the type τ arrival at time t, the type $(\tau+1)$ arrival at time $(t-1)$, etc. We index the active agents in the market at time t in the following way: a type l agent, which has τ time slots before his deadline, is denoted by a pair $(l,\tau)_t$. According the arrival process described in Section II-A, the pair $(l,\tau)_t$ may take value in the following ordered set:

$$\begin{aligned}\mathcal{C} = \{(1,1),(2,1),(3,1)\cdots,&(L,1),\\ (2,2),(3,2),\cdots,&(L,2),\\ \cdots,&(L,L)\}.\end{aligned}$$

Also, $D_c = L(L+1)/2$ denotes the cardinality of $\mathcal{C}$.

Let $u_{(l,\tau)}(t) \in \mathbb{R}$ denote the *instantaneous demand*[1] from agent $(l,\tau)_t$, with the vector form denoted by:

$$\mathbf{u}(t) = [u_{(l,\tau)}(t) : (l,\tau) \in \mathcal{C}] \in \mathbb{R}^{D_c}.$$

By default $u_{(l,\tau)}(t) = 0$ if there is no agent $(l,\tau)_t$ at time t, in other words $h_l(t+\tau-l) = 0$. The instantaneous aggregate demand $U(t) = \sum_{(l,\tau)\in\mathcal{C}} u_{(l,\tau)}(t) = \mathbf{1}'\mathbf{u}(t)$, where $\mathbf{1}$ is a D_c-dimensional column vector of all ones.

The aggregate demand process is determined by agents scheduling their load in response to the real-time prices as well as to the system state information. In our setup, agents represent significant market participants including the utility companies, load aggregators, or micro grids, each of which may consist of many small end consumers. Therefore we assume that they are price anticipating and schedule their loads strategically.

Similarly, we define the *backlog state* $\mathbf{x}(t)$ and the *existence state* $\mathbf{o}(t)$ as follows:

$$\mathbf{x}(t) = [x_{(l,\tau)}(t) : (l,\tau) \in \mathcal{C}] \in \mathbb{R}^{D_c}, \tag{1}$$

$$\mathbf{o}(t) = [o_{(l,\tau)}(t) : (l,\tau) \in \mathcal{C}] \in \{0,1\}^{D_c}, \tag{2}$$

where element $x_{(l,\tau)}(t)$ denotes agent $(l,\tau)_t$'s unsatisfied load at time t, and element $o_{(l,\tau)}(t) = 1$ if and only if there is an arrival of type l agent at time $(t+\tau-l)$.

In summary, the system state at time t is defined to be $\mathbf{s}(t) = (\mathbf{x}(t), \mathbf{o}(t)) \in \mathcal{S}$, and the state space $\mathcal{S}$ is specified as:

$$\mathcal{S} = \mathbb{R}^{D_c} \times \{0,1\}^{D_c}. \tag{3}$$

The system state $\mathbf{s}(t)$ is updated after the realization of $\mathbf{h}(t)$ and $\mathbf{d}(t)$ at the beginning of each time slot t as follows:

$$\mathbf{x}(t+1) = \mathbf{R}_1(\mathbf{x}(t) - \mathbf{u}(t)) + \mathbf{R}_2\mathbf{d}(t), \tag{4}$$

$$\mathbf{o}(t+1) = \mathbf{R}_1\mathbf{o}(t) + \mathbf{R}_2\mathbf{h}(t), \tag{5}$$

where $\mathbf{R}_1$ and $\mathbf{R}_2$ are some constant matrices[2].

We assume that the state information is available to all agents in the market. We acknowledge that the assumption of complete and perfect information is very strong, especially when the number of agents is large. Information structure, though an important issue in dynamic games, is not the focus of this paper, as the identified mechanism that produces endogenous risk of spikes also exists in incomplete information models. This simplification can serve as a benchmark for incomplete information models.

In summary, Fig. 1 illustrates the model of the two-sided electricity market between the competitive supplier and the agents which dynamically schedule their loads over time.

II-D Non-cooperative Market Architecture

With full information, agent $(l,\tau)_t$'s decision about his instantaneous demand $u_{(l,\tau)}(t)$ is a function of the system state $\mathbf{s}(t)$, and the function form is determined by the nature of the interactions among the agents.

[1]Note that we allow the load realizations as well as the instantaneous demand from the agents to be negative, which can model the situation where distributed agents are equipped with renewable generations or storage units, and are able to sell energy back to the power grid. We also ran extensive numerical simulations for the scenario where there is a lower bound on instantaneous demand or supply, and in all of our the simulations the main results hold qualitatively.

[2]$\mathbf{R}_1$ has dimension $D_c \times D_c$ and $\mathbf{R}_2$ $D_c \times L$. These constant matrices are specified as follows:

$$\mathbf{R}_1\Big((k-1)(L+\frac{2-k}{2})+i+1,\quad k(L+\frac{1-k}{2})+i\Big) = 1,$$

for all $1 \le i \le L-k$ and $1 \le k \le L-1$;

$$\mathbf{R}_2\Big((l-1)(L+\frac{2-l}{2})+1,\quad l\Big) = 1,$$

for all $1 \le l \le L$.

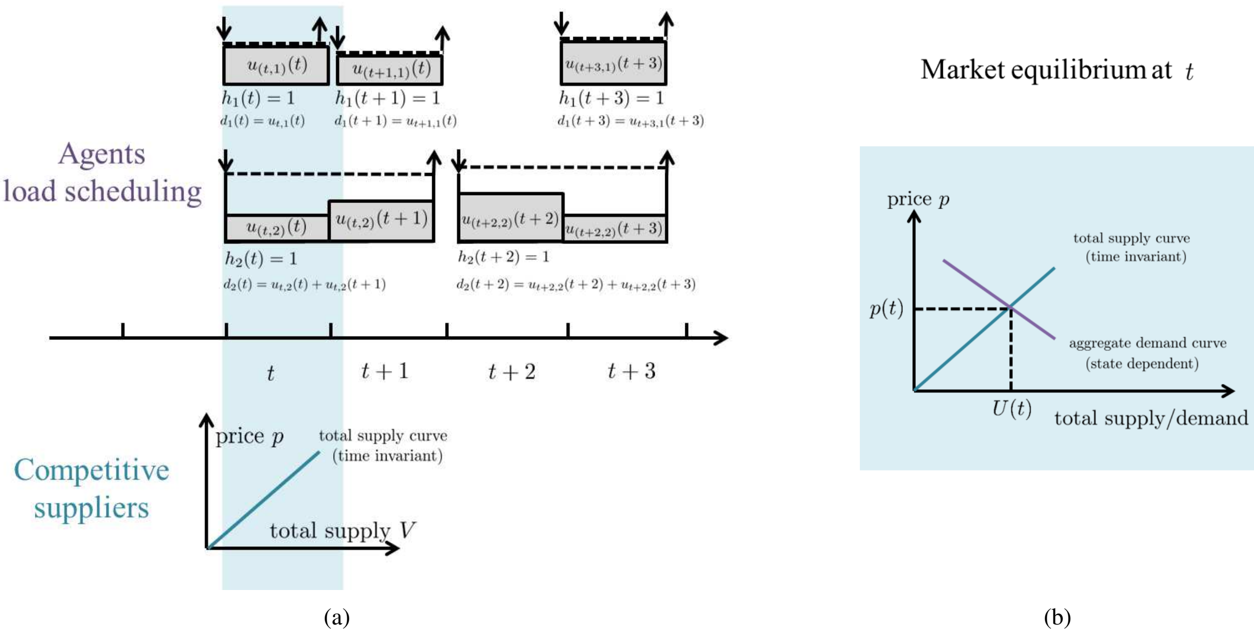

Fig. 1: Electricity market place with agents and the competitive supplier.

We define the *non-cooperative market architecture* to be the setup where there is no coordination among the agents. Under the assumption that the agents do not directly derive utility from electricity usage, their only objective is to minimize the expected procurement cost of electricity in order to complete the workload before the deadlines. Note that our framework can also be extended to cases where agents value their consumptions. In Section V we shall relax the deadline constraints and model the mismatch between the required workload and the completed load before deadline as disutility of the agents.

More specifically, a type l agent who arrives at time t dynamically optimizes his demand schedule

$$\{u_{(l,l-i)}(t+i):\ i=0,1,\cdots,l-1\}$$

to minimize his expected payment

$$\mathbb{E}[\sum_{i=0}^{l-1} p(t+i)u_{(l,l-i)}(t+i)].$$

Considering the cost coupling through endogenous pricing as well as the deadline constraints over time, we formulate the agent decision making as a stochastic dynamic game, specified as follows:

- **Players:** Over infinite time horizon, the players are indexed by $\{(l,\tau)_t : t\in\mathbb{Z},(l,\tau)\in\mathcal{C}\}$ according to their type and arrival time in the market.
- **State Space:** The state space is given by $\mathcal{S}$ defined in (3).
- **Action Set:** The action set is denoted by $\mathcal{A}$. The action set of player $(l,\tau)_t$ at time t in state $\mathbf{s}$ is given by:

$$A_{(l,\tau)}(\mathbf{s}) = \begin{cases} 0, & \text{if } o_{(l,\tau)}=0 \\ x_{(l,\tau)}, & \text{if } o_{(l,\tau)}=1 \text{ and } \tau=1 \\ \mathbb{R}, & \text{otherwise} \end{cases}. \quad (6)$$

- **Transition Probability:** For each state $\mathbf{s}$ and action vector $\mathbf{u}\in\prod_{(l,\tau)} A_{(l,\tau)}(\mathbf{s})$, the transition probability $\mathbb{P}(\mathbf{s}'|\mathbf{s},\mathbf{u})$ follows the system state evolution rule specified in (4), (5) and the agent arrival process described in II-A.

We shall focus on *Markov Perfect Equilibrium* (MPE) [28], [9] throughout our discussion. This equilibrium concept refers to a subgame perfect equilibrium of the stochastic dynamic game where the strategies only depend on the current system state. Compared to other standard equilibrium concepts such as Nash equilibrium (NE) and subgame perfect equilibrium (SPE), MPE is a more suitable equilibrium type to consider here. First, subgame perfection refines NE in the dynamic decision making setup. Second, when an agent makes a decision of instantaneous electricity demand, the load consumption trajectory in the past does not matter to him, rather he only focuses on the total amount of unfinished load he needs to fulfill before his departure from the market. This justifies the Markov property of MPE. The *Markov strategy* is defined as a function $\mathbf{u}:\mathcal{S}\to\mathcal{A}$. Note that as a result of the i.i.d. arrival process, all agents of the same type have the same cost structure. We shall focus on the symmetric equilibria where for every $(l,\tau)\in\mathcal{C}$, the demand $u_{(l,\tau)}(t)$ is a function of the system state $\mathbf{s}(t)$, and that function is time invariant. Namely, if $\mathbf{s}(t)=\mathbf{s}(t')$, agent $(l,\tau)_t$ and $(l,\tau)_{t'}$ will reach the same decision $u_{(l,\tau)}(t)=u_{(l,\tau)}(t')$. The symmetric MPE is defined below.

Definition 1 (Symmetric MPE): A strategy profile

$$\mathbf{u}^{nc} = \{u^{nc}_{(l,\tau)}(\mathbf{s}) : (l,\tau) \in \mathcal{C}, \mathbf{s} \in \mathcal{S}\}$$

is defined to be a Symmetric MPE strategy, if the following fixed point equations are satisfied for agent $(l,\tau)_t$, for any $(l,\tau) \in \mathcal{C}$, for any time t, and for any system states $\mathbf{s}(t) \in \mathcal{S}$:

$$\begin{aligned} u^{nc}_{(l,\tau)}(\mathbf{s}(t)) = \arg\min_{u} \\ \mathbb{E}\Big[p(t)u + \sum_{i=1}^{\tau-1} p(t+i)u^{nc}_{(l,\tau-i)}(\mathbf{s}(t+i))\Big|\mathbf{s}(t)\Big] \quad (7) \\ \text{subject to: } \sum_{i=0}^{l-1} u^{nc}_{(l,l-i)}(\mathbf{s}(t+i)) = d_l(t), \quad \forall t, l, \\ p(t) = u + \sum_{(l',\tau') \in \mathcal{C}, (l',\tau') \neq (l,\tau)} u^{nc}_{(l',\tau')}(\mathbf{s}(t)), \\ p(t+i) = \sum_{(l',\tau') \in \mathcal{C}} u^{nc}_{(l',\tau')}(\mathbf{s}(t+i)), \quad \forall i \geq 1, \end{aligned}$$

where $\mathbf{s}(t)$ evolves according to (4), (5).

II-E Cooperative Market Architecture

We define the *cooperative market architecture* to be the setup where the agents coordinate their decisions of load scheduling $\mathbf{u}(\mathbf{s}(t))$ to minimize the expected total cost $\mathbf{E}[\sum_t p(t)U(t)]$. This setup corresponds to the scenario where the agents delegate the task of load scheduling to a single load aggregator that schedules the loads in a centralized way on behalf of all the agents.

Definition 2 (Optimal Stationary Cooperative Strategy): A strategy profile $\mathbf{u}^c = \{\mathbf{u}^c_{(l,\tau)})(\mathbf{s}) : (l,\tau) \in \mathcal{C}, \mathbf{s} \in \mathcal{S}\}$ is defined to be an optimal stationary cooperative strategy if $\mathbf{u}^c(\mathbf{s}) = \big[u^c_{(l,\tau)}(\mathbf{s}) : (l,\tau) \in \mathcal{C}\big]$ solves the following fixed point equations for any system states $\mathbf{s}(t) \in \mathcal{S}$:

$$\begin{aligned} \mathbf{u}^c(\mathbf{s}(t)) = \arg \min_{\mathbf{u}^c = [u_{(l,\tau)}:(l,\tau)\in\mathcal{C}]} \mathbb{E}\Big[\sum_{(l,\tau)\in\mathcal{C}} p(t)u_{(l,\tau)} \\ + \sum_{t'=t+1}^{T} \sum_{(l,\tau)\in\mathcal{C}} p(t+i)u^c_{(l,\tau)}(\mathbf{s}(t'))\Big|\mathbf{s}(t)\Big], \quad (8) \\ \text{subject to: } \sum_{i=0}^{l-1} u^c_{(l,l-i)}(\mathbf{s}(t+i)) = d_l(t), \quad \forall t, l, \\ p(t) = \sum_{(l,\tau)\in\mathcal{C}} u_{(l,\tau)}, \\ p(t+i) = \sum_{(l,\tau)\in\mathcal{C}} u^c_{(l,\tau)}(\mathbf{s}(t+i)), \quad \forall i \geq 1, \end{aligned}$$

where $\mathbf{s}(t)$ evolves according to (4), (5).
The above fixed point equation is derived from an infinite horizon average cost MDP, and the associated Bellman equation can be solved via standard value iteration or policy iteration.

II-F Welfare Metrics

Different market architectures induce different agent decisions of load scheduling, and different aggregate demand processes $\{U(t) : t \in \mathbb{Z}\}$. We focus on two measures of the system performance: market efficiency and the risk of aggregate demand spikes. *Market efficiency* is defined to be the expected sum of the electricity provider's surplus W_p and the agents' surplus W_a:

$$W = \underbrace{\mathbb{E}[p(t)U(t) - \frac{1}{2}U(t)^2]}_{W_p} + \underbrace{\mathbb{E}[-p(t)U(t)]}_{W_a} = -\frac{1}{2}\mathbb{E}[U(t)^2], \tag{9}$$

where the second inequality is the result of marginal cost pricing $p(t) = U(t)$. Moreover, we have that $W = W_a/2$. Therefore, market efficiency is decreasing in $\mathbb{E}[U(t)^2]$ as well as $Var[U(t)^2]$, the variance of the aggregate demand process. By definition, the optimal stationary cooperative strategy maximizes the agents' surplus W^c_a, and it also leads to the highest system efficiency. Let W^{nc}_a denote the agents' surplus achieved by the non-cooperative equilibrium strategy. The agents' surplus loss $(W^c_a - W^{nc}_a)$ is commonly known as the "price of anarchy" due to the strategic behavior of non-cooperative agents when payoff externalities exist. In this setup, it also leads to a proportional system efficiency gap $(W^c - W^{nc})$.

We define *risk* to be the tail probability of the stationary aggregate demand process:

$$R = \Pr(U(t) > M), \tag{10}$$

for some positive large constant M. As a result of marginal cost pricing and increasing marginal cost, risk also captures the tendency for the spot price or aggregate demand to spike[3].

III. Tradeoff Analysis for $L = 2$ Case

III-A Non-cooperative Equilibrium and Optimal Cooperative Strategy

In general, there are no closed form solutions to either (7) or (8), and numerical solutions involve exponential complexity. In this section, we analyze the case with the number of types L is equal to 2. Type 1 agents have inflexible loads that must be satisfied upon arrival, and type 2 agents have the flexibility to split their workload between two consecutive time periods. In this case, both the equilibrium strategy and the optimal cooperative strategy can be found explicitly, and the analysis can shed light on understanding the agent behavior and the resulting system dynamics induced by different market architectures in the general setup.

At any time t, there are at most 3 agents in the market, which are indexed as: $(1,1)_t$, $(2,1)_t$, and $(2,2)_t$. Among them, the instantaneous demands of agents $(1,1)_t$ and $(2,1)_t$ are given by $u_{(1,1)}(\mathbf{s}) = x_{(1,1)}$ and $u_{(2,1)}(\mathbf{s}) = x_{(2,1)}$ due to the deadline constraints. Only agent $(2,2)_t$ needs to make a nontrivial decision to schedule his workload $(u_{(2,2)}(t), u_{(2,1)}(t))$, and for him the sufficient statistics of the system state is $(x(t), d_2(t))$, where $x(t) = x_{(1,1)}(t) +$

[3]Note that since the system dynamics is stationary, the expected long-term averages converge to the corresponding expectations for each time slot.

$x_{(2,1)}(t)$ is defined as the *aggregate backlog* state, and $d_2(t)$ is his total workload. We define a *linear strategy* as a strategy profile $\mathbf{u}(\mathbf{s})$ with $u_{(1,1)}(\mathbf{s}) = x_{(1,1)}$, $u_{(2,1)}(\mathbf{s}) = x_{(2,1)}$, and $u_{(2,2)}(\mathbf{s}) = u(x, d_2) = -ax + bd_2 + g$ for some constants a, b and g.

Proposition 1 (Existence of linear symmetric MPE): Under the non-cooperative market architecture, for $L = 2$, there exists a linear symmetric MPE with the linear strategy $u^{nc}(x, d_2)$ given by:

$$u^{nc}(x,d_2) = -\underbrace{\frac{1}{2(1+\sqrt{1-q_2/2})}}_{a^{nc}}x + \underbrace{\frac{1}{1+1/\sqrt{1-q_2/2}}}_{b^{nc}}d_2 + \underbrace{\frac{q_1\mu_1 + q_2\mu_2\frac{1}{1+\sqrt{1-q_2/2}}}{2(1+\sqrt{1-q_2/2})}}_{g^{nc}}. \tag{11}$$

Proposition 2 (Existence of optimal cooperative strategy): Under the cooperative market architecture, for $L = 2$, there exists a linear optimal stationary cooperative load scheduling strategy $u^c(x, d_2)$ given by:

$$u^{c}(x,d_2) = -\underbrace{\frac{1}{1+\sqrt{1-q_2}}}_{a^{c}}x + \underbrace{\frac{1}{1+1/\sqrt{1-q_2}}}_{b^{c}}d_2 + \underbrace{\frac{q_1\mu_1 + q_2\mu_2\frac{1}{1+\sqrt{1-q_2}}}{1+\sqrt{1-q_2}}}_{g^{c}}. \tag{12}$$

Remark 1 (Interpretation of the coefficients): For a linear strategy adopted by type 2 agents, the coefficient a can be interpreted as the sensitivity to the aggregate backlog $x(t)$. A larger a means that the strategy is more aggresive in absorbing the fluctuation of inflexible loads. Note that both a^{nc} and a^c are increasing in q_2. Intuitively, with a higher type 2 arrival rate q_2, each type 2 agent is more aggresive in responding to $x(t)$ at their first period, anticipating that in the next time slot another type 2 agent will arrive and respond to $x(t+1)$ in a similar aggresive way. Also note that for any arrival rate q_2, $a^{nc} < a^c$ always holds, and $a^c \in [0.5, 1]$, $a^{nc} \in [0.25, 0.2929]$, which means that type 2 agents always respond less aggresively to the aggregate backlog $x(t)$ under the non-cooperative market architecuture. This can be understood as a result of their strategic behavior at equilibrium.

III-B System Performance

Given a linear strategy with $u(x, d_2) = -ax + bd_2 + g$, we have the state evolution dynamics:

$$x(t+1) = o_{(1,1)}(t+1)d_1(t+1) + o_{(2,2)}(t)\big(d_2(t) - u(x(t), d_2(t))\big),$$

which determines the stationary distribution of the aggraegate backlog / aggregate demand process.

Assume that all type 2 agents adopt the linear strategy $u(x, d_2) = -ax+bd_2+g$, market efficiency, as defined in (9) can be evaluated explicitly. In particular, with linear strategies $u^{nc}(\cdot,\cdot)$ and $u^c(\cdot,\cdot)$, the market efficiency W^{nc} and W^c, as well as the difference $\Delta = W^c - W^{nc}$ can be obtained in closed form (not included here for brevity). Furthermore, Δ is positive and increasing in q_2, namely, the higher q_2 is, the larger efficiency loss of non-cooperative scheme will be. This observation suggests that the efficiency gap between cooperative and non-cooperative load scheduling scheme even widens as the arrival rate of flexible loads increases.

However, the stationary distribution of the aggregate demand process under the cooperative market architecture has a larger right tail, which corresponds to more aggregate demand spikes. An upper bound of this higher risk is quantified in the next proposition:

Proposition 3 (Upper bound on the risk R): Suppose that the workloads D_i have Normal distributions $\mathcal{N}(\mu_i, \sigma_i^2)$ for $i = 1, 2$. For the stationary aggregate backlog distribution $\mathcal{X}$ induced by a linear strategy $u(x, d_2) = -ax + bd_2 + g$ ($a \in (0, 1)$), the probability of aggregate backlog exceeding M is upper bounded as follows:

$$\Pr(x(t) > M) \le \frac{1}{\sqrt{2\pi}m_1}e^{-\frac{m_1^2}{2}},$$

where

$$m_1 = \frac{M - \frac{\mu_1+(1-b)\mu_2-g}{1-a}}{\sqrt{\frac{\sigma_1^2+(1-b)^2\sigma_2^2}{1-a^2}}}.$$

Moreover, if the following condition is satisfied:

$$\frac{1-(1-a)^2}{1-a^2} > \frac{b^2}{\sigma_1^2/\sigma_2^2 + (1-b)^2}, \tag{13}$$

for sufficiently large constant M, the risk of aggregate backlog exceeding M is upper bounded as follows:

$$R = \Pr(U(t) > M) \le q\Pr(x(t) \ge M) + o(e^{-M}). \tag{14}$$

The interpretation of condition (13) is that when the variance of flexible load realizations is sufficiently lower than that of the inflexible load realizations, and when the coefficient a is sufficiently larger than the coefficient b, the aggregate demand spikes are mostly contributed by the high aggregate backlogs. It is easy to verify that the stationary distribution of $x(t)$ induced by $u^c(\cdot,\cdot)$ has a larger mean and a larger variance than that induced by $u^{nc}(\cdot,\cdot)$. In other words, the state of the aggregate backlog is more volatile in the cooperative scheme, which can be associated with a higher upper bound of the risk under certain condition.

III-C Numerical Results

In this part, we will demonstrate the efficiency-risk tradeoffs via numerical experiments. First we compare the stationary distribution of the aggregate demand process induced by four different linear strategies. In addition to $u^c(\cdot,\cdot)$ and $u^{nc}(\cdot,\cdot)$, we include the "naive load scheduling" scheme $u^{naive}(x, d_2) = d_2/2$, and the "no load scheduling" scheme $u^{no}(x, d_2) = d_2$ for comparison.

Fig. 2a shows the efficiency performance, which is negatively proportional to the second order moment of the aggregate demand process, under the four strategies. We observe that for all arrival rate q_2, cooperative load scheduling is the most efficient, and the efficiency loss of the non-cooperative

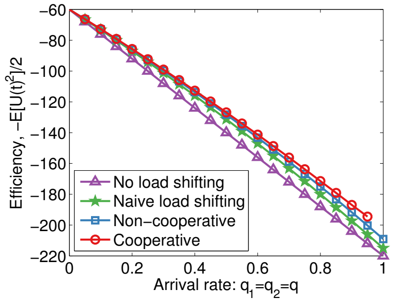

(a) The cooperative load scheduling scheme achieves the highest efficiency, and the efficiency loss $\Delta = W^c - W^{nc}$ increases in the flexible load arrival rate q_2.

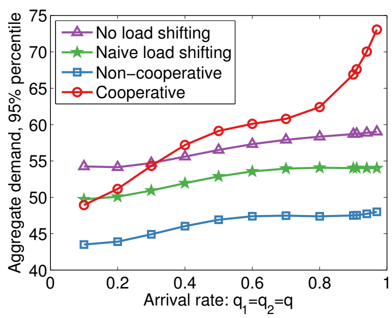

(b) The risk of aggregate demand spikes is the lowest for the non-cooperative load scheduling, and increases the fastest for the coopeartive load scheduling.

Fig. 2: Efficiency-risk tradeoff for $L = 2$

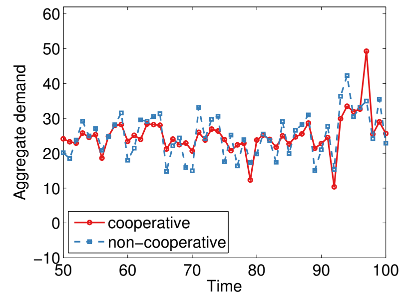

(a) At a short time scale, the cooperative load scheduling can better smooth out the aggregate demand process.

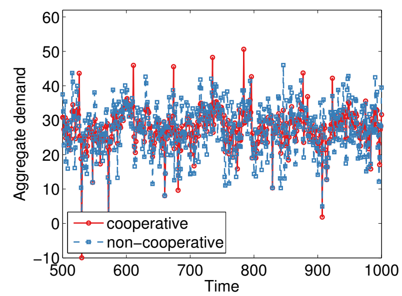

(b) At a large time scale, the cooperative load scheduling produces more severe aggregate demand spikes.

Fig. 3: Sample paths of the aggregate demand process

scheme when compared to the cooperative scheme increases in q_2. This suggests that the cooperative load scheduling becomes more effective in terms of attenuating the aggregate demand variance when the arrival rate of flexible loads increases.

Fig. 2b compares the risk of spikes across the four strategies. The 0.95-quantile of the stationary distribution of the aggregate demand process is plotted for each strategy. A higher 0.95-quantile is associated with a higher risk. We observe that for all arrival rates, the non-cooperative scheme achieves the lowest risk. Also, as q_2 increases, the risk increases most rapidly with the cooperative scheme, while the other three load scheduling schemes only increase slowly.

Fig. 3 shows the sample paths of the aggregate demand processes. In Fig. 3a, we observe that at a smaller time scale, the cooperative scheme can better smooth the aggregate demand process, which is consistent with the lower aggregate demand variance. However in Fig. 3b, at a larger time scale, we can identify more demand spikes produced endogenously by the cooperative load scheduling scheme, corresponding to the higher risk of aggregate demand spikes in the cooperative scheme.

Remark 2 (Understanding when spikes occur): On one hand, the aggregate demand spikes can happen when the workload realization $\mathbf{d}(t)$ is extremely high. We attribute this type of spikes to exogenous shocks as they correspond to rare events associated with the uncertainties in the environment. If the support of $\mathbf{D}_l$ is bounded, for sufficiently large constant M the spikes from exogenous shocks are negligible. On the other hand, an aggregate spike can be produced endogenously when a high aggregate backlog is built up during load scheduling and the type 2 agent does not arrive at that time slot. The lack of flexibility to further postpone the backlogged loads translates the aggregate backlog at the deadline into an aggregate demand spike. This intuition is confirmed by observing the conditional distributions of aggregate demand process in Fig. 4. We can see that the tail of the aggregate demand distribution is significantly larger, conditioned on the non-arrival of type 2 agent, and is also significantly larger, conditioned on the event that the aggregate backlog is high. This provides intuition to the high efficiency - high risk tradeoff of the cooperative load scheduling scheme: the more efficient a load scheduling strategy is, the more intense the backlog usage will be, and the resulting high backlog volatility leads to aggregate demand spikes.

IV. General L Analysis: Pricing

In general, as illustrated in Figure 5a, the agent load scheduling can be viewed as a full state feedback controller, the feedback control signal $\mathbf{u}(t)$ affects the system state evolution according to (4) and (5), and the system output is the aggregate demand process. As a result of the Bernoulli

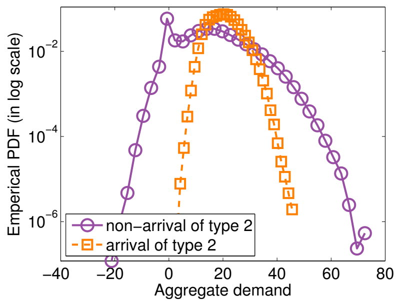

(a) The stationary distribution of the aggregate demand process has a heavier tail conditional on the event of non-arrrival of flexible loads.

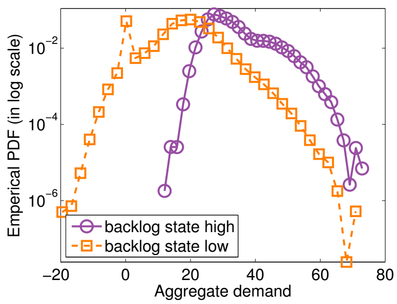

(b) The stationary distribution of the aggregate demand process has a heavier tail conditional on a high aggregate backlog state.

Fig. 4: Conditional distributions (in log scale) of the aggregate demand process reveal when aggregate demand spikes happen, for the case $L = 2$.

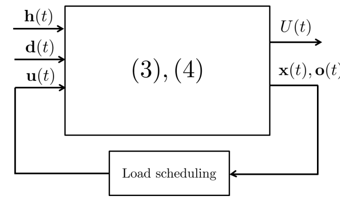

(a) Original system with nonlinear dynamics

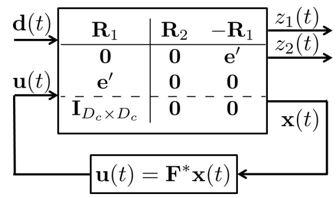

(b) LTI formulation with deadline constraints

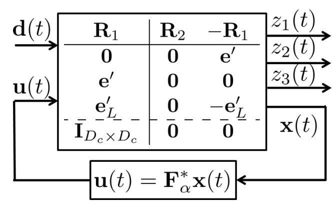

(c) LTI formulation with disutility of load mismatch upon deadline

Fig. 5: System diagrams

arrival processes, the agent load scheduling forms a time-invariant feedback controller that is nonlinear. Instead of analyzing the nonlinear dynamics directly, we introduce two key approximations to the system dynamics along with surrogate performance measures. The approximated system resembles the original system in the most essential ways and facilitates the analysis.

Approximation 1: The agent arrival rate $\mathbf{q} = \mathbf{1}$, namely agents of all types arrive at every period, so that $\mathbf{h}(t) = \mathbf{1}$ and $\mathbf{o}(t) = \mathbf{1}$ for all t.

Approximation 2: The second moment $\mathbb{E}[z_2(t)^2]$ of the aggregate backlog process $z_2(t) = \mathbf{e}'\mathbf{x}(t)$ is used as a substitute measure of risk.

The intuition behind the two approximations is as follows. We focus on the system dynamics with high arrival rate, i.e., $\mathbf{q}$ is close to $\mathbf{1}$. According to our observations in Remark 2, if flexible loads continue to arrive, spikes are unlikely to happen, however, with small probability $(1 - \prod_{i=1}^{L} q_i)$ when there is a non-arrival of flexible loads, a high backlog will be turned into an aggregate demand spike. We also normalize the load arrival process so that the average load realization $\boldsymbol{\mu}$, the average backlog state $\mathbb{E}[\mathbf{x}(t)]$, and the average demand $\mathbb{E}[\mathbf{u}(t)]$, are all zero vectors, and we assume the load arrival process $\{\mathbf{d}(t) : t \in \mathbb{Z}\}$ is a white Gaussian process.

In summary, the system diagram of the approximated system with linear dynamics is shown in Fig. 5b. Ideally, in order to achieve a high efficiency and a low risk, agent load scheduling should serve to suppress the volatility of both aggregate demand state and backlog state. The performance measures of efficiency and risk are therefore approximated by the second moments of the two outputs $z_1(t) = \mathbf{e}'\mathbf{u}(t)$ and $z_2(t) = \mathbf{e}'\mathbf{x}(t)$.

Many of the market architectural properties, for example the degree of cooperation and risk sensitivity of the agents, are usually fixed, while the system operator has the freedom to design the pricing rule to elicit the desired agent behavior and tradeoffs. Within the LTI framework, we focus on the non-cooperative setup and consider the system operator's problem of choosing a static linear pricing rule parameterized by coefficients $\mathbf{q}_1$ and $\mathbf{q}_2$ as follows:

$$p(t) = \mathbf{q}_1'\mathbf{x}(t) + \mathbf{q}_2'\mathbf{u}(t). \qquad (15)$$

The instantaneous demand decisions are made by individual agents in a non-cooperative way. In the linear symmetric MPE, assuming that the load scheduling strategy exists, and it is in the following form:

$$\mathbf{u}^*(t) = \mathbf{F}^*\mathbf{x}(t), \qquad (16)$$

we denote the (l,τ)-th row of $\mathbf{F}^*$ by $\mathbf{F}^*_{(l,\tau)} \in \mathbb{R}^{D_c}$. Applying the one-shot deviation principle [29] at the equilibrium, we obtain the optimal load scheduling decision $u^*_{(l,\tau)}(t)$ for any $(l,\tau) \in \mathcal{C}$ as follows:

if $\tau = 1$, $u^*_{(l,\tau)}(t) = \mathbf{e}_{(l,\tau)}\mathbf{x}(t)$;
if $\tau > 1$,

$$u^*_{(l,\tau)}(t) = \arg\min_{u\in\mathbb{R}} \{p(t)u + \mathbb{E}[\sum_{k=1}^{\tau-1} p(t+k)\mathbf{F}^*_{(l,\tau-k)}\mathbf{x}(t+k)]\} \tag{17}$$

$$\begin{aligned} \text{subject to: } & \mathbf{u}(t) = \mathbf{F}^*\mathbf{x}(t) + \mathbf{e}_{(l,\tau)}(u - \mathbf{F}^*_{(l,\tau)}\mathbf{x}(t)), \\ & \mathbf{u}(t+k) = \mathbf{F}^*\mathbf{x}(t+k), \quad \forall k > 0, \\ & p(i) = \mathbf{q}_1'\mathbf{x}(i) + \mathbf{q}_2'\mathbf{u}(i), \quad \forall i, \\ & \mathbf{x}(i+1) = \mathbf{R}_1\mathbf{x}(i) + \mathbf{R}_2\mathbf{d}(i) - \mathbf{R}_1\mathbf{u}(i), \quad \forall i, \end{aligned}$$

where $\mathbf{e}_{(l,\tau)}$ is a D_c dimensional vector with the only non-zero element being 1 at the (l,τ)-th position. Moreover, at the symmetric equilibrium the expected strategy adopted by all other agents should be consistent with the best response strategy of a particular agent. A direct application of the principle of optimality to (17) leads to:

$$\mathbf{F}^* = f_{(\mathbf{q}_1,\mathbf{q}_2)}(\mathbf{F}^*). \tag{18}$$

For given coefficients $\mathbf{q}_1, \mathbf{q}_2$, the (l,τ)-th row of the mapping $f_{(\mathbf{q}_1,\mathbf{q}_2)} : \mathbb{R}^{D_c\times D_c} \to \mathbb{R}^{D_c\times D_c}$ is specified in (19). Obtaining the conditions on the parameters which guarantee the existence of a fixed point solution to the nonlinear equation (19) is a challenging task. However, the equation provides a set of necessary conditions for the equilibrium strategies to satisfy.

Proposition 4 (System operator's problem): Assume that the social welfare can be modeled as a weighted sum of efficiency and risk as follows:

$$J(\mathbb{E}[z_1(t)^2], \mathbb{E}[z_2(t)^2]) = -(\alpha_1\mathbb{E}[z_1(t)^2] + \alpha_2\mathbb{E}[z_2(t)^2]).$$

The system operator optimizes the parameters of the pricing rule as defined in (15) to maximizes the social welfare. The optimal solution $(\mathbf{q}_1^*, \mathbf{q}_2^*)$ is a solution to the following optimization problem:

$$\min_{\substack{\mathbf{q}_1,\mathbf{q}_2\in\mathbb{R}^{D_c} \\ \mathbf{Q},\mathbf{F}\in\mathbb{R}^{D_c\times D_c}}} \alpha_1\mathbf{e}'\mathbf{F}_{(\mathbf{q}_1,\mathbf{q}_2)}\mathbf{Q}\mathbf{F}_{(\mathbf{q}_1,\mathbf{q}_2)}\mathbf{e} + \alpha_2\mathbf{e}'\mathbf{Q}\mathbf{e} \tag{19}$$

$$\begin{aligned} \text{subject to: } & \mathbf{R}_1(\mathbf{I}-\mathbf{F})\mathbf{Q}(\mathbf{I}-\mathbf{F}')\mathbf{R}_1' - \mathbf{Q} + \mathbf{R}_2\mathbf{R}_2' = \mathbf{0}, \\ & \mathbf{F} = f_{(\mathbf{q}_1,\mathbf{q}_2)}(\mathbf{F}), \end{aligned}$$

where $f_{(\mathbf{q}_1,\mathbf{q}_2)}$ is the mapping defined in (19).

V. FUNDAMENTAL TRADEOFF

A more fundamental question that we attempt to address is that in face of exogenous uncertainties, to what extent is system volatility inevitable and to what extent it can be controlled. More specifically, is there a limit of the feedback control, in the form of load scheduling, to achieve the dual goals of increasing market efficiency and reducing endogenous risk? We show that for a broad class of load scheduling strategies, the exogenous randomness cannot be completely eliminated, and the dual goals of suppressing the second moments of the two processes $z_1(t)$ and $z_2(t)$ cannot be achieved simultaneously.

A load scheduling strategy is defined to be Pareto optimal if there does not exist any other strategy that makes the variance of $z_1(t)$ smaller without making the variance of $z_2(t)$ larger, and a pair $(\mathbb{E}[z_1(t)^2], \mathbb{E}[z_2(t)^2])$ locates on the Pareto front if it is achieved by a Pareto optimal strategy. Unless it trivially includes the point $(0,0)$, the Pareto front dictates the limit of the system performances with a downward sloping tradeoff curve between efficiency and risk. Also, note that the concept of Pareto optimal load scheduling strategy does not rely on market architecture specifications, in the sense that the system performance achievable under any specific market architecture cannot outperform the Pareto front. The Pareto front thus serves as a benchmark to measure the welfare loss between the optimal strategies and a load scheduling strategy induced by a specific market architecture.

We introduce the third approximation, with which the Pareto front can be characterized by solving a standard Linear Quadratic (LQ) optimization problem [30], [31].

Approximation 3: The deadline constraints are relaxed. Instead of requiring that all agents fulfill their backlogged load when they exit the market, we track the total load mismatch upon their deadlines: $z_3(t) = \mathbf{e}_L'(\mathbf{x}(t) - \mathbf{u}(t))$, where $\mathbf{e}_L$ is a D_c-dimensional column vector with the first L elements being ones and all others zero. Agent disutility of the load mismatch upon deadline is modeled by the variance of the load mismatch $\mathbb{E}[z_3(t)^2]$, and is defined to be the third performance measure.

With Approximation 1, 2 and 3, the system diagram with the outputs $\mathbf{z}(t) = [z_1(t), z_2(t), z_3(t)]$ is shown in Fig. 5c. We generalize the tradeoff between efficiency and risk to a three-way tradeoff among efficiency, risk, and load mismatch upon deadline, with the three-way Pareto optimal strategies and the three-way Pareto front similarly defined. The problem of characterizing the Pareto front can be cast into an LQ optimization problem with an unconstrained feedback controller. We follow the standard multi-objective optimization technique to scalarize the objective. Consider the weighted output process:

$$\mathbf{z}_\alpha(t) = [\alpha_1 z_1(t), \alpha_2 z_2(t), \alpha_3 z_3(t)],$$

where $\alpha_i > 0$, for $i = 1,2,3$, and $\alpha_1^2+\alpha_2^2+\alpha_3^2 = 1$. A Pareto optimal load scheduling strategy minimizes the $\mathcal{H}_2$ system norm [30] for a given weight $\alpha = (\alpha_1,\alpha_2,\alpha_3)$:

$$\min_{\{\mathbf{u}(t):t\in\mathbb{Z}\}} \|\mathbf{z}_\alpha\|_2^2 \tag{20}$$

$$\text{subject to: } \mathbf{x}(t+1) = \mathbf{R}_1(\mathbf{x}(t) - \mathbf{u}(t)) + \mathbf{R}_2\mathbf{d}(t)$$

Proposition 5 (Three-way Pareto front): 1) For given non-negative weight $\alpha = (\alpha_1,\alpha_2,\alpha_3)$, the corresponding Pareto optimal load scheduling strategy is static and linear in the system state $\mathbf{x}(t)$ as follows:

$$\mathbf{u}(t) = \mathbf{F}^*_\alpha\mathbf{x}(t).$$

where $\mathbf{F}^*_\alpha = \mathbf{Q}^*\mathbf{P}^{*-1}$, and $(\mathbf{Q}^*, \mathbf{P}^*)$ is the unique

$$
f_{(\mathbf{q}_1,\mathbf{q}_2)}(\mathbf{F})_{(l,\tau)} = \begin{cases} \mathbf{e}'_{(l,\tau)}, \text{ if } \tau = 1, \\ \dfrac{\mathbf{e}'_{(l,\tau)}\mathbf{R}_1'\mathbf{A}_{(l,\tau)}\Big(\mathbf{R}_1(\mathbf{I}-\mathbf{F})+\mathbf{R}_1\mathbf{e}_{(l,\tau)}\mathbf{F}_{(l,\tau)}\Big)-\Big(\mathbf{q}_1'+\mathbf{q}_2'(\mathbf{F}-\mathbf{e}_{(l,\tau)}\mathbf{F}_{(l,\tau)})\Big)}{\mathbf{e}'_{(l,\tau)}\mathbf{R}_1'\mathbf{A}_{(l,\tau)}\mathbf{R}_1\mathbf{e}_{(l,\tau)}+2\mathbf{e}'_{(l,\tau)}\mathbf{q}_2}, \text{ if } \tau > 1, \end{cases} \tag{19}
$$

$$
\text{where } \mathbf{A}_{(l,\tau)} = \sum_{k=1}^{\tau-1}\Big(\big(\mathbf{R}_1(\mathbf{I}-\mathbf{F})\big)^{k-1}\Big)'\Big((\mathbf{q}_1+\mathbf{F}'\mathbf{q}_2)\mathbf{F}_{(l,\tau-k)}+\mathbf{F}'_{(l,\tau-k)}(\mathbf{q}_1'+\mathbf{q}_2'\mathbf{F})\Big)\Big(\big(\mathbf{R}_1(\mathbf{I}-\mathbf{F})\big)^{k-1}\Big).
$$

solution to the following convex optimization problem:

$$
\min_{\mathbf{Q},\mathbf{P}\in\mathbb{R}^{D_c\times D_c},\mathbf{M}\in\mathbb{R}^{3\times 3}} \quad \rho
$$

$$
\begin{aligned}
\text{subject to: } \; & \mathbf{Q} > \mathbf{0}, \\
& \text{Trace}(\mathbf{M}) \leq \rho, \\
& \begin{bmatrix} \mathbf{Q} & (\mathbf{R}_1\mathbf{Q}-\mathbf{R}_1\mathbf{P})' \\ (\mathbf{R}_1\mathbf{Q}-\mathbf{R}_1\mathbf{P}) & \mathbf{Q}-\mathbf{R}_2\mathbf{R}_2' \end{bmatrix} > \mathbf{0}, \\
& \begin{bmatrix} \mathbf{Q} & (\mathbf{C}_1\mathbf{Q}+\mathbf{D}_{12}\mathbf{P})' \\ (\mathbf{C}_1\mathbf{Q}+\mathbf{D}_{12}\mathbf{P}) & \mathbf{M} \end{bmatrix} > \mathbf{0}.
\end{aligned}
$$

where

$$
\mathbf{C}_1 = [0 \quad \alpha_2\mathbf{e} \quad \alpha_3\mathbf{e}_L]', \quad \mathbf{D}_{12} = [\alpha_1\mathbf{e} \quad \mathbf{0} \quad -\alpha_3\mathbf{e}_L]'.
$$

2) Given a matrix $\mathbf{F}$ such that the feedback rule $\mathbf{u}(t) = \mathbf{F}\mathbf{x}(t)$ stabilizes the system, the $\mathcal{H}_2$ system norm of the three performance measures is given by:

$$
\begin{aligned}
&\|z_1\|_2^2 = \mathbf{e}'\mathbf{F}\mathbf{Q}_F\mathbf{F}'\mathbf{e}, \quad \|z_2\|_2^2 = \mathbf{e}'\mathbf{Q}_F\mathbf{e}, \\
&\|z_3\|_2^2 = (\mathbf{e}'-\mathbf{e}_L'\mathbf{F})\mathbf{Q}_F(\mathbf{e}'-\mathbf{e}_L'\mathbf{F})',
\end{aligned}
$$

where $\mathbf{Q}_F$ is the controllability Gramian given by solving the following equation:

$$
\mathbf{R}_1(\mathbf{I}-\mathbf{F})\mathbf{Q}_F(\mathbf{I}-\mathbf{F}')\mathbf{R}_1' - \mathbf{Q}_F + \mathbf{R}_2\mathbf{R}_2' = \mathbf{0}.
$$

With different parameters of $\alpha = (\alpha_1, \alpha_2, \alpha_3)$, different Pareto optimal solutions are produced, and we can trace out the Pareto front. In particular, the curve when restricting the three-way Pareto front to the plane of $\|z_3(t)\|_2^2 = \epsilon$ for $\epsilon \ll 1$ ($\alpha_3/\alpha_1 \gg 1$ and $\alpha_3/\alpha_2 \gg 1$) approaches the efficiency-risk tradeoff curve with enforced deadline constraints. The second part of Proposition 5 provides a way to evaluate the system performance for any linear load scheduling strategy.

As an example, in Fig. 6a, we plot the Pareto front for the case with $L = 5$ to visualize the three-way tradeoff among the three system performance measures. In Fig. 6b, we observe that as we loosen the constraint on load mismatch upon deadline, namely with a larger β_3 in the constraint $\|z_3(t)\|_2^2 \leq \beta_3$, the two-way Pareto front of efficiency and risk shifts towards the origin. This inward shifting corresponds to a Pareto improvement as the volatility of both aggregate demand and aggregate backlog are reduced.

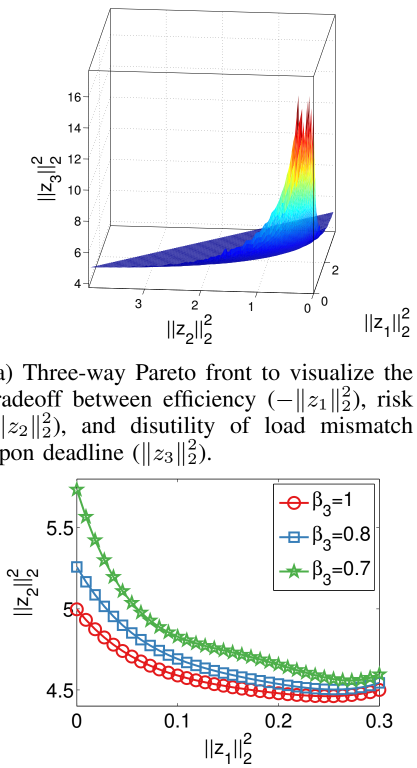

(a) Three-way Pareto front to visualize the tradeoff between efficiency ($-\|z_1\|_2^2$), risk ($\|z_2\|_2^2$), and disutility of load mismatch upon deadline ($\|z_3\|_2^2$).

(b) Efficiency-risk tradeoff can be Pareto improved at a cost of higher disutility of load mismatch upon deadline.

Fig. 6: Three way tradeoff, for the case $L = 5$.

VI. Discussion

VI-A Modeling

In this section we discuss the rationale behind the choice of the abstracted model adopted in this paper and justify the assumptions.

First, we focus on the demand side of the wholesale spot market, where the participants are the utility companies, load aggregators, or micro grids, each of which may consist of many small end consumers. Therefore, it is valid to assume that the agents are price anticipating and make strategic decisions. Second, we aim to capture the uncertainties in the spot market, in the form of unpredictable disturbances which force the system to deviate from the planned ahead generation and demand schedule. Therefore, even though at a much slower time scale, the daily base load profile is highly time dependent, it is still reasonable to impose the time invariant assumption on the agent arrival random process and the load realizations, which correspond to the spot market time scale disturbances.

Finally, it is not the purpose of our paper to propose a design framework or a pricing scheme that can be immediately matched to the demand response models in the field. Instead, our major contribution is twofold: first, we propose the abstracted model that captures the most important elements of dynamic demand response; secondly, we investigate the model and deliver the message of efficiency-risk tradeoff, which is an ignored tradeoff that has not been discussed in the literature.

In the model setup, we only include the elements which are essential to illustrate the corresponding efficiency-risk tradeoffs, namely the agent **dynamic decision making process**, **the heterogeneous deadline constraints**, and **the load arrival uncertainties**. These elements fundamentally shape the statistical characteristics of the aggregate demand process. With this modeling approach and the semi-analytic solutions, one is able to see the impact of the most important parameters more transparently. For example how the load arrival frequency $\mathbf{q}$ and the load realization $\mathbf{d}$ affect demand decisions in (11) - (12), and how the resulting agent behavior affect the system performance in Section III-B. Incorporating more details of the power system and electricity market in the modeling increases the dimension of the model parameters, and will likely leave us with only numerical analysis and simulations. Although it is possible to observe the tradeoff through extensive simulation, one might also easily miss the phenomenon. Moreover, as long as the three essential modeling elements are preserved qualitatively, we believe that the main messages will still hold.

VI-B Future work

We propose three main directions for future research. First, as discussed above, it is impossible to validate the assumptions of the abstracted model with actual load and metering data; however, it is possible to validate the findings and the implications of our analysis against actual data. Systematic and rigorous validation of the results with empirical data is an important direction for further research.

Second, in our analysis we have been maintaining the complete and perfect information assumption. However, the state information may be difficult to obtain in practice. One important research direction is to relax the complete information assumption, examine how agent behavior is shaped by the information structure, and study the corresponding impact on the performance of dynamic demand response, as well as on the efficiency-risk tradeoffs.

Third, in the electricity market, there exist hierarchical organizations and finer structures of cooperation, ranging from large load aggregators to the many end consumers. Adapting the proposed generic framework to accommodate such hierarchical structures, and examining the impact of the hierarchical aspect of market architecture on the tradeoffs is another interesting research direction.

VII. Conclusion

In this paper, we proposed a framework to examine the welfare impacts of dynamic demand response under different market architectures. We examined the efficiency-risk tradeoff rooted at the market architectural properties, and characterized the Pareto optimal front. Aside from its academic value, we believe that a better and deeper understanding of such tradeoffs will be helpful to system operators and regulatory agencies in designing the system architecture and operational policies.

Appendix

Due to space limitations, we only provide a sketch of the proofs in the appendix. Interested readers please see [24] for the details and supplementary materials.

Proof 1 (Proposition 1): For $L = 2$, under the non-cooperative market architecture, the equilibrium strategy $u^{nc}(x, d_2)$ is characterized by the solution to the following fixed point equation:

$$u^{nc}(x(t), d_2(t)) = \arg\min_u \Big\{ u(u + x(t)) + \mathbb{E}\Big[(d_2(t) - u)\Big(x(t+1) + h_2(t+1) u^{nc}(x(t+1), d_2(t+1)) \Big) \Big| x(t), d_2(t) \Big] \Big\}, \tag{21}$$

where $x(t+1) = d_1(t+1) + (d_2(t) - u)$. The best response strategy of any agent $(2, 2)_t$ to the conjectured linear strategy of all other agents is also linear, and the fixed point equation pins down the MPE.

Proof 2 (Proposition 2): For $L = 2$, under the cooperative market architecture, the optimal stationary cooperative strategy can also be obtained as a closed form solution of the following Belllman equation with value function $V^c(x)$ and average cost per period λ^c:

$$\lambda^c + V^c(x) = (1-q)\Big(x^2 + \mathbb{E}_{d_1}\Big[V^c(d_1) \Big] \Big) + q\mathbb{E}_{d_1, d_2}\Big[\min_u \big\{ (x+u)^2 + V^c(d_2 - u + d_1) \big\} \Big] \tag{22}$$

The value function of the corresponding Bellman equation is of quadratic form, and first order condition gives the optimal policy in linear form.

Proof 3 (Proposition 3): With Bernoulli arrival processes, the stationary distribution of $x(t)$ is a weighted sum of many Gaussian variables, and the sum of weights is one. The upper bound of $\Pr(x(t) > M)$ can be obtained by the Gaussian variable with the highest mean and variance in the limit.

Proof 4 (Proposition 4): Since $\mathbf{d}(t)$ is white Gaussian process, solving the problem of the utility optimization of the system operator is equivalent to minimizing the system $\mathcal{H}_2$ norm, under the agents rationality constraint, which is specified by (18).

Proof 5 (Proposition 5): Following an approach of changing of variables similar to that in [30], the optimization problem in (20) can be convexified into a standard linear matrix inequality (LMI) [32] problem.

References

[1] D. Kirschen, "Demand-side view of electricity markets," *Power Systems, IEEE Transactions on*, vol. 18, no. 2, pp. 520–527, May 2003.

[2] A. J. Conejo, J. M. Morales, and L. Baringo, “Real-time demand response model,” *Smart Grid, IEEE Transactions on*, vol. 1, no. 3, pp. 236–242, 2010.
[3] M. Roozbehani, M. Dahleh, and S. Mitter, “Volatility of power grids under real-time pricing,” *Power Systems, IEEE Transactions on*, vol. 27, no. 4, pp. 1926–1940, Nov 2012.
[4] Z. Fan, “A distributed demand response algorithm and its application to PHEV charging in smart grids,” *Smart Grid, IEEE Transactions on*, vol. 3, no. 3, pp. 1280–1290, 2012.
[5] J. Foster and M. Caramanis, “Energy reserves and clearing in stochastic power markets: The case of plug-in-hybrid electric vehicle battery charging,” in *Decision and Control (CDC), 49th IEEE Conference on*, 2010, pp. 1037–1044.
[6] G. Koutitas, “Control of flexible smart devices in the smart grid,” *Smart Grid, IEEE Transactions on*, vol. 3, no. 3, pp. 1333–1343, 2012.
[7] M. Roozbehani, D. Materassi, M. Ohannessian, and M. A. Dahleh, “Robust and optimal consumption policies for deadline-constrained deferrable loads,” *IEEE Transactions on Smart Grids, To Appear*, 2014.
[8] T. Kim and H. V. Poor, “Scheduling power consumption with price uncertainty,” *Smart Grid, IEEE Transactions on*, vol. 2, no. 3, pp. 519–527, 2011.
[9] E. Maskin and J. Tirole, “A theory of dynamic oligopoly, iii: Cournot competition,” *European Economic Review*, vol. 31, no. 4, pp. 947–968, 1987.
[10] S. Borenstein, *The trouble with electricity markets (and some solutions)*. University of California Energy Institute, 2001.
[11] C.-K. Woo, I. Horowitz, J. Moore, and A. Pacheco, “The impact of wind generation on the electricity spot-market price level and variance: The texas experience,” *Energy Policy*, vol. 39, no. 7, pp. 3939–3944, 2011.
[12] E. Baeyens, E. Bitar, P. Khargonekar, and K. Poolla, “Coalitional aggregation of wind power,” *Power Systems, IEEE Transactions on*, vol. 28, no. 4, pp. 3774–3784, Nov 2013.
[13] P. P. Varaiya, F. F. Wu, and J. W. Bialek, “Smart operation of smart grid: Risk-limiting dispatch,” *Proceedings of the IEEE*, vol. 99, no. 1, pp. 40–57, 2011.
[14] R. Buyya and M. Murshed, “A deadline and budget constrained cost-time optimisation algorithm for scheduling task farming applications on global grids,” *Arxiv preprint cs/0203020*, 2002.
[15] R. A. Stubbs and D. Vandenbussche, “Multi-portfolio optimization and fairness in allocation of trades,” *White paper., January*, 2009.
[16] J. Danielsson, H. Shin, and J. Zigrand, “Endogenous and systemic risk,” in *Quantifying Systemic Risk*. Univ. of Chicago Press, 2011.
[17] J. Geanakoplos, *The leverage cycle*. Yale University, Cowles Foundation for Research in Economics, 2009.
[18] D. L. Alderson and J. C. Doyle, “Contrasting views of complexity and their implications for network-centric infrastructures,” *Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions on*, vol. 40, no. 4, pp. 839–852, 2010.
[19] M. C. Caramanis and J. M. Foster, “Uniform and complex bids for demand response and wind generation scheduling in multi-period linked transmission and distribution markets,” in *Decision and Control and European Control Conference (CDC-ECC), 50th IEEE Conference*, 2011, pp. 4340–4347.
[20] A.-H. Mohsenian-Rad, V. W. Wong, J. Jatskevich, R. Schober, and A. Leon-Garcia, “Autonomous demand-side management based on game-theoretic energy consumption scheduling for the future smart grid,” *Smart Grid, IEEE Transactions on*, vol. 1, no. 3, pp. 320–331, 2010.
[21] L. Chen, N. Li, S. H. Low, and J. C. Doyle, “Two market models for demand response in power networks,” *IEEE SmartGridComm*, vol. 10, pp. 397–402, 2010.
[22] Z. Ma, D. Callaway, and I. Hiskens, “Decentralized charging control for large populations of plug-in electric vehicles,” in *Decision and Control (CDC), 49th IEEE Conference on*, 2010, pp. 206–212.
[23] W. Tushar, W. Saad, H. V. Poor, and D. B. Smith, “Economics of electric vehicle charging: A game theoretic approach,” *Smart Grid, IEEE Transactions on*, vol. 3, no. 4, pp. 1767–1778, 2012.
[24] Q. Huang, M. Roozbehani, and M. Dahleh, “Efficiency-robustness tradeoffs in dynamic oligopoly markets,” *Technical Report [online]. Available at http://arxiv.org/abs/1209.0229*, 2012.
[25] T. J. Sargent, “rational expectations,” in *The New Palgrave Dictionary of Economics*, S. N. Durlauf and L. E. Blume, Eds. Basingstoke: Palgrave Macmillan, 2008.
[26] J. E. Hartley, *The representative agent in macroeconomics*. Psychology Press, 1997, vol. 10.
[27] F. Schweppe, *Spot pricing of electricity*. Springer, 1988.
[28] L. Shapley, “Stochastic games,” *Proc. of the National Academy of Sciences of the U.S.A.*, vol. 39, no. 10, p. 1095, 1953.
[29] D. Blackwell, “Discounted dynamic programming,” *The Annals of Mathematical Statistics*, pp. 226–235, 1965.
[30] G. E. Dullerud and F. Paganini, *A course in robust control theory*. Springer New York, 2000, vol. 6.
[31] V. Sima, *Algorithms for linear-quadratic optimization*. CRC Press, 1996, vol. 200.
[32] L. Vandenberghe and S. Boyd, “Semidefinite programming,” *SIAM review*, vol. 38, no. 1, pp. 49–95, 1996.

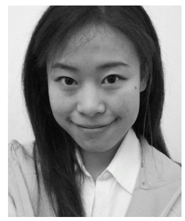

Qingqing Huang received the BEng and BBA degree from the Hong Kong University of Science & Technology, in 2011, and the SM degree from the Massachusetts Institute of Technology (MIT), in 2013. She is currently a graduate student in the EECS department at MIT.

Mardavij Roozebehani (M’03) received the B.Sc. degree from Sharif University of Technology, Tehran, Iran, in 2000, the M.Sc. degree in mechanical and aerospace engineering from the University of Virginia, Charlottesville, VA, in 2003, and the Ph.D. degree in aeronautics and astronautics from the Massachusetts Institute of Technology (MIT), Cambridge, MA, in 2008. Since 2012, he has held a Principal Research Scientist position at the Laboratory for Information and Decision Systems (LIDS) at MIT, where he previously held instructor, postdoctoral, and research scientist positions between 2008 and 2011. His main research interests include distributed and networked control systems, dynamics and economics of power systems with an emphasis on robustness and risk, analysis and design of software control systems, and finite-state systems. Dr. Roozbehani is a recipient of the 2007 AIAA graduate award for safety verification of real-time software systems.

Munther A. Dahleh Munther A. Dahleh (F’01) received the B.S. degree from Texas A & M university, College Station, TX, in 1983, and the Ph.D. degree from Rice University, Houston, TX, in 1987, both in electrical engineering. Since then, he has been with the Department of Electrical Engineering and Computer Science, the Massachusetts Institute of Technology (MIT), Cambridge, MA, where he is now the acting director of ESD. Previously, he was the Associate Department Head of EECS. He has been a visiting Professor at the Department of Electrical Engineering, California Institute of Technology, Pasadena, CA, for the Spring of 1993. He has held consulting positions with several companies in the U.S. and abroad. He is interested in problems at the interface of robust control, filtering, information theory, and computation which include control problems with communication constraints and distributed agents with local decision capabilities. In addition to methodology development, he has been interested in the application of distributed control in the future electric grid and the future transportation system with particular emphasis in the management of systemic risk. He is also interested in various problems in network science including information propagation over complex engineering and social networks. He is the co-author (with I. Diaz-Bobillo) of the book Control of Uncertain Systems: A Linear Programming Approach (Englewood Cliffs, NJ: Prentice-Hall, 1995), and the co-author (with N. Elia) of the book Computational Methods for Controller Design (New York: Springer, 1998).