

Incentivizing Reliability in Demand-Side Response

Hongyao Ma

Harvard University
hma@seas.harvard.edu

Valentin Robu

Heriot Watt University
V.Robu@hw.ac.uk

Na Li

Harvard University
nali@seas.harvard.edu

David C. Parkes

Harvard University
parkes@eecs.harvard.edu

Abstract

We study the problem of incentivizing reliable demand-response in modern electricity grids. Each agent is uncertain about her future ability to reduce demand and unreliable. Agents who choose to participate in a demand-response scheme may be paid when they respond and penalized otherwise. The goal is to reliably achieve a demand reduction target while selecting a minimal set of agents from those willing to participate. We design incentive-aligned, direct and indirect mechanisms. The direct mechanism elicits both response probabilities and costs, while the indirect mechanism elicits willingness to accept a penalty in the case of non-response. We benchmark against a spot auction, in which demand reduction is purchased from agents when needed. Both the direct and indirect mechanisms achieve the reliability target in a dominant-strategy equilibrium, select a small number of agents to prepare, and do so at low cost and with much lower variance in payments than the spot auction.

1 Introduction

A crucial aspect of operating an electric power system is that an exact balance between supply and demand must be maintained at all times. Electricity grids are facing a number of new challenges in this regard, due to both the increasing number of intermittent renewable generation resources [Su *et al.*, 2014; Zhang *et al.*, 2015], and the presence of more volatile types of loads, such as those from electric vehicle charging [Robu *et al.*, 2013].

These challenges have led to an increased interest in *demand-side response* (DR), in which consumers (industrial, commercial, or domestic) commit to temporarily reduce or shift consumption away from periods with low generation capacity [Palensky and Dietrich, 2011]. A number of DR aggregation services exist to facilitate this process, and intermediate between distribution network operators (DNOs) and consumers. These range from those operated directly by DNOs, to commercial services.¹

We model demand response as a two-step process. First, consumers opt-in to a DR scheme, possibly sharing information about their cost to respond if asked, and their probability of being able to respond. Some, perhaps all, of these consumers are selected and asked to prepare for the possibility of demand reduction. Second, in the event that a demand reduction is required, then some, perhaps all of the selected consumers are asked to reduce demand, and may receive a payment or pay a penalty depend on whether or not they follow-through and respond.

DR aggregators are not interested only in those consumers with lowest cost, but also in consumers who are reliable, and most likely to respond if needed. Selecting reliable consumers allows a target to be met with high confidence while asking fewer consumers to prepare, leading to less economic disruption. It is natural that consumers will not always be able to respond. Consider an industrial factory for example, which uses electricity for the production line, transporting raw materials, and cooling. Its ability to respond in a DR event is highly uncertain; e.g., it may depend on the production procedure, time of day, customer requests, and weather conditions.

From a mechanism design perspective, the goal is to reliably meet a target reduction while minimizing the number of agents who are selected. There are a number of interrelated challenges: (1) incentivize consumers to opt-in to the scheme, (2) truthfully elicit information about reliability and cost, (3) select a small, reliable set of agents to ask to prepare, and (4) set up payments and penalties so that those who are selected will choose to follow-through if asked and if able to reduce demand. Simple approaches fail. For example, setting a high fixed penalty for not reducing demand when asked would ensure follow-through, but not provide incentives for opt-in.

We advance two related designs: a direct mechanism that elicits both costs and probabilities directly, and an indirect mechanism that elicits only willingness to accept a penalty in the case of non-response, from which response probabilities are inferred. The mechanisms fix a payment that will be made to agents for demand response, and select agents in decreasing order of the maximum acceptable non-response penalties until the reliability target is met. Both mechanisms have a simple dominant-strategy equilibrium, meaning truthful reporting in the direct mechanism and truthfully reporting the maximum acceptable penalty in the indirect mechanism.

In an experimental evaluation with a wide range of param-

¹Examples include Enernoc, Kiwi Power or Upside Energy.

eter values, we show that the mechanisms achieve close to the *first best* (i.e. assuming the mechanism knows agent types and can choose the most reliable ones) with regard to the number of agents who are selected and asked to prepare. We also benchmark against a spot auction, in which demand reduction is purchased from agents when needed. The spot auction has a higher total cost as well as a very large variance in payments, making the scheme risky for both the agents and the grid, as well as susceptible to collusion.

1.1 Related Work

To our knowledge, this is the first application of mechanism design for eliciting cost and probability information from consumers in order to achieve reliable demand response. A number of papers in the power systems literature discuss the use of DR aggregation [Zhang *et al.*, 2015; Su *et al.*, 2014]. Although not an approach of mechanism design, there is also some prior work on market design for demand response (e.g. [Li *et al.*, 2015; Johari and Tsitsiklis, 2011; PJM, 2015]), proposing to allow for bids of supply or cost curves, and discussing market equilibria. Other papers [De Vries and Heijnen, 2008; Kwag and Kim, 2014] have considered reliability in DR but from a policy (non mechanism-design) perspective. There are also works that design contracts for load control, aiming at for example maximizing the system payoff while satisfying other operational constraints [Balandat *et al.*, 2014; Haring *et al.*, 2013; Yang *et al.*, 2015], but again without taking a mechanism design viewpoint. Others have proposed to use *scoring rules* to incentivize truthful reports about expected future consumption in power grids [Rose *et al.*, 2012; Robu *et al.*, 2012; Akasiadis and Chalkiadakis, 2013], which do not apply to our setting because they do not set penalties correctly in order to provide incentives for follow-through if needed.

On the mechanism design side, the inspiration for our approach is the work to promote capacity utilization [Ma *et al.*, 2015]. In a sense, ours' is the inverse problem where we want to promote capacity reduction rather than utilization. An additional, technical aspect of the present problem is to consider the question of how to select a number of agents in order to probabilistically meet a system reliability target.

2 The Demand Side Response

In the demand side response problem, the planner is the electricity grid (or DR aggregator) and the agents are consumers of electricity interested in offering DR services. Let $N = \{1, 2, \dots, n\}$ denote the set of agents. Each agent can prepare for demand response ahead at a cost of $c_i \geq 0$. We consider the simple setting in which each agent demands the same quantity (a single unit, w.l.o.g.). Our results easily extend to agents with heterogeneous demand sizes. If an agent prepares to reduce demand, then she is able to reduce demand with probability $p_i \in (0, 1)$, at the cost of $v_i \geq 0$. The amount v_i represents the opportunity cost for the loss of electricity. The triple $\theta_i = (v_i, p_i, c_i)$ defines an agent's *type*, and is agent i 's private information. Let $\theta = (\theta_1, \dots, \theta_n)$ denote a type profile. We assume that the agents' abilities to respond are independent to each other. In our model, an agent can




Figure 1: The timeline of a two-period mechanism.

only respond if she first prepares.²

Reliability Target. Denote $M \in \mathbb{N}_+$ as the *target capacity reduction* that needs to be achieved. The designer's objective is to select a minimal set of agents to prepare for DR ahead of time such that the target reduction is met with probability at least τ , where $\tau \in (0, 1)$ is a system-wide *reliability target*. We make a *deep market assumption*, which is that there are enough agents in the economy that it is possible to meet the reliability target. This holds for most real DR markets.

Example 1. Suppose target $M = 1$ and probability $\tau = 0.9$, and there are three agents with $p_1 = 0.8$, $p_2 = 0.6$, and $p_3 = 0.4$. If only agent 1 prepares, 1 unit of power is provided with probability $0.8 < \tau$ thus the reliability target is not met. If both agents 1 and 2 prepare, the probability with which no one is able to respond is $(1 - p_1)(1 - p_2) = 0.08 < 1 - \tau = 0.1$, and the reliability requirement is met.

Two-Period Mechanism. We consider mechanisms that run over two time periods which use a fixed reward $R > 0$ and a variable penalty per agent. The timeline is as follows (see Figure 1):

Period 0:

- Agents can choose to report information to the mechanism, with knowledge of their type and the demand-response reward R .
- The mechanism determines for each selected agent i the period-one penalty $t_i \geq 0$ in case of non-response.
- With the knowledge of t_i and R , each selected agent decides whether to prepare for demand response.

Period 1:

- If agent i can respond (with probability p_i) she decides whether or not to do so with the knowledge of t_i and R .
- For selected agent i , the mechanism pays R upon response, and charges $t_i \geq 0$, otherwise.

Note that neither the selected agents' choice on preparation or their ability to respond are observable.

The structure of the design is well motivated, representing a small change from the current practice by DNOs and aggregators. In current practice, as in our model, selected agents do not receive a payment when asked to prepare, and the amount of payment made to an agent in the event of demand response is fixed. The new ingredients are that we elicit type information in period zero, and use this both to decide who to select, and also to set a penalty in the event of non-response.

A demand-response mechanism is *dominant strategy incentive compatible* (DSIC) if truthful reporting maximizes

²In practice, advance notification is required to perform demand response at a reasonable cost [Borenstein *et al.*, 2002].




Figure 2: Expected utility as a function of the penalty t_i .

each agent's expected utility regardless of the reports of other agents, and conditioned on the agent making rational decisions in period one (see below). A demand-response mechanism is *individually rational* (IR) if each agents' expected utility for (truthful) participation is non-negative. Informally, a DSIC mechanism is "truthful" and an IR mechanism ensures that agents will choose to participate.

Rational Decisions and Expected Utilities. How should an agent i with type $\theta_i = (v_i, p_i, c_i)$, who reports truthfully and is selected, and now faces a reward R and penalty t_i , behave in the mechanism? Consider the following cases:

1. If the agent does not prepare, she is unable to respond and her utility will be $-t_i$.
2. If the agent does prepare, but is not able to or decides not to respond in period one, her utility will be $-t_i - c_i$.
3. If the agent does prepare, and is able to respond in period one and decides to do so, her utility is $R - v_i - c_i$.

For any agent with $v_i > R$, her utility is negative for all $t_i \geq 0$. As we will see later such agents will not be selected by the mechanisms for DR. Assume now that $R > v_i$. If the agent prepares and is able to respond, then it is rational to respond because $R - v_i - c_i > -t_i - c_i$. If the agent decides to prepare, her expected utility is,

$$u_i(R, t_i) = p_i(R - v_i) - (1 - p_i)t_i - c_i. \quad (1)$$

For $(R - v_i)p_i - c_i < 0$, the expected utility, whether preparing or not, would be negative as long as the penalty $t_i \geq 0$, and she will not be selected by the mechanisms. Assume for now that $(R - v_i)p_i - c_i \geq 0$. Such an agent will choose to prepare (getting $u(R, t_i)$ rather than $-t_i$) and also choose to respond if possible. The expected utility $u_i(R, t_i)$ decreases linearly with slope $1 - p_i$ as the penalty t_i increases. A simple example with two agents is shown in Figure 2.

Let $z \geq 0$ denote the penalty that represents the unique zero crossing point of $u(R, t)$, such that $u(R, z) = 0$. This is the maximum penalty that the agent is willing to pay:

$$z_i = \frac{(R - v_i)p_i - c_i}{1 - p_i}. \quad (2)$$

Fixing v_i and c_i , the higher the probability p_i and the more reliable the agent, the slower her expected utility $u_i(R, t_i)$ decreases with t_i , the shallower the utility curve, and the larger the crossing point, z_i . For example, agent 1 in Figure 2 has a larger willingness to pay, although the expected reward she gets from the grid minus her cost (the y-intercept) is lower.

3 The Direct Mechanism

We now design a truthful, direct mechanism for demand response. The mechanism selects agents in decreasing order of willingness to pay until the reliability target is met.

Definition 1 (The Direct Mechanism with Reward R). *The direct mechanism $M_{\text{dir}}(R)$ collects reported type profile $\hat{\theta} = (\hat{\theta}_1, \dots, \hat{\theta}_n)$, and computes the willingness to pay \hat{z}_i according to (2). Assume w.l.o.g. that $\hat{z}_1 \geq \hat{z}_2 \geq \dots \geq \hat{z}_n$ (breaking ties arbitrarily). Let X_i be a Bernoulli random variable with parameter \hat{p}_i if $\hat{z}_i \geq 0$ and let $X_i \equiv 0$ if $\hat{z}_i < 0$.*

- *Selection rule (period zero): $x_i(\hat{\theta}) = 1$ for $i \leq m$ and $\hat{z}_i \geq 0$, and $x_i(\hat{\theta}) = 0$ for $i > m$, where the last agent selected is,*

$$m = \min_{\ell \in N} \ell \text{ s.t. } \mathbb{P} \left[\sum_{i=1}^{\ell} X_i \geq M \right] \geq \tau. \quad (3)$$

- *Payment rule (evaluated in period zero, payments made in period one): make payment R to any selected agent who reduces demand. If selected, charge agent i a penalty amount $t_i(\hat{\theta}) = \hat{z}_{m-i}$ for no response, where*

$$m_{-i} = \min_{\ell \neq i} \ell \text{ s.t. } \mathbb{P} \left[\sum_{i'=1, \dots, \ell, i' \neq i} X_{i'} \geq M \right] \geq \tau,$$

is the agent with the smallest willingness to pay that would be selected in an economy without agent i . Demand reduction is not accepted from any unselected agent, and her payment is zero.

Note: We've focused for clarity on describing the typical case (given our assumption about a deep market) that (3) has a solution, i.e. there are enough agents participating at current reward level R . If not, the mechanism simply selects all agents and charges no penalty.

For intuition for payments, the quantity $t_i(\hat{\theta}) = \hat{z}_{m-i}$ is the minimum payment that agent i needs to be willing to accept as penalty in order to be selected, and this is independent of agent i 's own report.

Example 2. Suppose $M = 1$ needs to be reduced with probability at least $\tau = 0.75$, and the grid pays a reward $R = 6$ for demand reduction. There are three agents, with types

- Agent 1: $v_1 = 1, p_1 = 0.9, c_1 = 1$
- Agent 2: $v_2 = 1, p_2 = 0.7, c_2 = 1$
- Agent 3: $v_3 = 1, p_3 = 0.6, c_3 = 1$

For truthful reports, the zero-crossings are $z_1 = 35, z_2 = 8.3$ and $z_3 = 5$. We first allocate to agent 1. Since $\mathbb{P}[X_1 \geq M] = p_1 > \tau$, agent 1 is the only agent that is selected.

In the economy without agent 1, we would first allocate to agent 2. Since $\mathbb{P}[X_2 \geq M] = p_2 < \tau$, allocating to only agent 2 is not enough to satisfy the reliability requirement. Thus we also allocate to agent 3 and we can check $\mathbb{P}[X_2 + X_3 \geq M] > \tau$. Therefore, $m_{-1} = 3$, and agent 1's penalty is $t_1(\hat{\theta}) = z_{m-1} = 5$ if she does not respond.

The rational decision of agent 1 is to prepare and respond if possible, and her expected utility is $u_1(R, t_1) = 3 > 0$. It's easy to see that for all report $\hat{\theta}_1$ of agent 1 s.t. $\hat{z}_1 \geq z_3$, agent 1 would be selected and face the same payments since agent 2 herself cannot meet the reliability target. Making reports s.t. $\hat{z}_1 < z_3$, agent 1 would not be selected thus gets utility zero. Truthful reporting is therefore a dominant strategy. \square

For each agent, the two possible alternatives under the direct mechanism are that she is selected or not. Now we show that the mechanism satisfies *agent-independence* and *agent maximizing*, thus is DSIC [Nisan, 2007]. Agent-independence requires that each agent faces prices for each alternative that are independent of their own report, and agent maximizing means that the mechanism selects the alternative that maximizes her utility under such prices.

Theorem 1. *The direct mechanism is DSIC, IR and always meets the reliability target for a sufficiently large R .*

Proof. First, note that for a selected agent the penalty $t_i = z_{m-i}$ and reward R are independent of $\hat{\theta}_i$. Also, there is no payment to or from unselected agents thus payments satisfy agent-independence.

Fix an agent i . If $x_i(\hat{\theta}) = 1$, we know $i \leq m$ and $\mathbb{P}\left[\sum_{i'=1}^{i-1} X_{i'} \geq M\right] < \tau$. This implies that $m-i \geq i+1$, and thus $t_i = z_{m-i} \leq z_i$. Agent i 's expected utility for preparing $u_i(R, t_i) \geq u_i(R, z_i) = 0$, i.e. selecting agent i is agent-maximizing for her, comparing with not selecting which results in utility zero. If agent i is not selected, then $i > m$ (as computed in (3)) or $z_i < 0$. If $z_i < 0$, $u_i(R, t_i) < 0$ for all $t_i \geq 0$ thus if selected agent i gets negative utility. If $i > m$, $m-i = m$ and $t_i = z_{m-i} \geq z_i$ thus expected utility from preparing $u_i(R, t_i) \leq 0$. Not selecting her is agent-maximizing, which gives her utility zero.

Combining the two cases we see that the mechanism is DSIC. From the above argument, we also see that the expected utility for all agents are non-negative thus IR follows.

With R high enough, $z_i \geq 0$ for all agents i thus (3) can be met due to the deep market assumption. It follows from rational decisions in period one that all selected agents will choose to prepare and then reduce demand when possible, and that the mechanism will achieve its reliability target. \square

The mechanism does not necessarily select the most reliable agents, since the zero-crossings z_i 's are not always aligned with the p_i 's, and we cannot allocate in decreasing order of the reported p_i 's while retaining incentive alignment.

3.1 Reliability Evaluation

The total quantity of demand reduction by a set S of selected agents, $X = \sum_{i \in S} X_i$, is a Poisson-binomial distributed random variable [Chen and Liu, 1997]. The CDF of X is

$$\mathbb{P}[X \leq k] = \sum_{\ell=0}^k \sum_{A \in F_\ell} \prod_{i \in A} p_i \prod_{j \in A^c} (1 - p_j), \quad (4)$$

where F_ℓ is the set of all subsets of S that are of cardinality ℓ , and $A^c = S \setminus A$. A naive way of evaluating the CDF requires computing the sum of an exponential number of terms. However, exact polynomial time algorithms based on recursive methods [Radke Jr and Evanoff, 1994] or Fourier Transform [Fernández and Williams, 2010] exist, and return the CDF of a summation of a thousand random variables within seconds. Hence, we adopt the Fourier Transform approach in the experiments reported here.³

³The Chernoff bound [Chernoff, 1952] and other large deviation bounds can also be used to approximate this expression, and would

4 The Indirect Mechanism

The direct mechanism asks for the full type of each agent and computes the willingness to pay using the reported costs and reliability. For simplicity, and also to better protect the private cost information of participants, we introduce an *indirect mechanism*, which elicits a single bid from each agent, estimates the p_i 's of each agent from the bid, and then evaluates the reliability target using the estimated p_i 's.

We can compute a bound on reliability p_i from the willingness to pay. We know from (2) that

$$p_i = \frac{z_i + c_i}{z_i + R - v_i} \geq \frac{z_i}{z_i + R}, \quad (5)$$

where the inequality holds since $c_i, v_i \geq 0$. Note that with fixed R , the expression $z_i/(z_i + R)$ is monotone in z_i . Let b_i be agent i 's bid, and $b = (b_1, \dots, b_n)$ denote a bid profile.

Definition 2 (The Indirect Mechanism with Reward R).

1. *Reports:* The indirect mechanism $M_{\text{ind}}(R)$ collects a single bid b_i from each agent, representing the maximum willingness to accept a penalty in the case of non-response. Assume w.l.o.g. that $b_1 \geq \dots \geq b_n$ (breaking ties arbitrarily).

2. *Inference:* Let $\tilde{p}_i = b_i/(b_i + R)$ and \tilde{X}_i be a Bernoulli random variable with parameter \tilde{p}_i .

3. *Outcome:* The indirect mechanism determines selection and payments as in the direct mechanism, simply replacing X_i with \tilde{X}_i and z_i with b_i .

Theorem 2. *It is a dominant strategy for each agent to bid $b_i^* = z_i$ under the indirect mechanism. The indirect mechanism is IR and meets the reliability target for R large enough.*

Proof. Consider an agent with zero-crossing z_i . If she bids $b_i = z_i$ and does not get selected, we know that $m-i = m$ as computed in (3) and $i > m$. Bidding lower than b_{m-i} would not change her utility. Bidding higher than b_{m-i} means that agent i would be selected and get charged $b_{m-i} > z_i$, thus get negative utility in expectation.

If an agent bids $b_i = z_i$ and gets selected, we know $m-i > i$ must hold and $t_i(b) \leq z_i$, thus agent i gets non-negative utility in expectation. Bidding weakly above b_{m-i} would not change her utility. Bidding below b_{m-i} , the agent i would not get selected thus would get utility zero.

Combining the above two cases, we know that it's a dominant strategy to bid $b_i^* = z_i$. Therefore under the DSE, $\tilde{p}_i = b_i^*/(b_i^* + R) = z_i/(z_i + R)$ is a lower bound on the reliability p_i of each agent. With similar argument as in the direct mechanism we know that selected agents always prepare and choose to respond when possible, thus the reliability requirement is always met with large enough R . \square

Example 2. (continued) Continuing the earlier example, in the indirect mechanism, agents bid their willingness to pay, $b_1^* = 35$, $b_2^* = 8.3$ and $b_3^* = 5$, and the estimated probabilities are $\tilde{p}_1 = 0.85$, $\tilde{p}_2 = 0.58$, and $\tilde{p}_3 = 0.45$. With $M = 1$ and $\tau = 0.75$, it suffices to select only agent 1 since she will respond with probability at least $\tilde{p}_1 = 0.85$. We can check that in the economy without agent 1, both agents 2 and 3 need

select slightly more agents than necessary but provide an easy way of evaluating reliability while retaining truthfulness.

to be selected, thus the penalty agent 1 would be charged in case of non-response is $b_3 = 5$. In this case, the outcome is the same as the outcome of the direct mechanism.

Suppose instead that $\tau = 0.9$. Now, from the estimated reliability, agent 1 does not appear to be sufficient to meet the target in the indirect mechanism, even though she is actually able to meet the reliability target. In this case, the two mechanisms would have different outcomes. \square

Observe that the gap $p_i - \tilde{p}_i$ is small while v_i and c_i are small, since the (conservative) inference approach approximates them by adopting zero. Also, the gap is small when p_i is large, since with large p_i the zero-crossing z_i is high thus v_i and c_i are less significant. Since we select agents in decreasing order of z_i , we are selecting the set of agents for which the bounds are relatively tight.

5 A Comparison with a Spot Auction

For a simple comparison, we adopt a spot auction, where we run an $M + 1$ st-price auction with a reserve price r , i.e. the reserve sets an upper bound on the reward payment. This auction is run in period one in the event that DR is required, and without pre-selection in period zero of which agents should invest effort and prepare. Rather, agents need to reason about possible payments from the auction in period one in order to decide whether to prepare in period zero.

For simplicity, we will study a complete information Nash equilibrium of the period-zero preparation decisions under the spot auction mechanism (i.e., assuming bidders know each others' values.)

Definition 3 (Spot Auction with Reserve r). *The Spot Auction Spot(r) collects a single bid b_i from each agent at period one, and chooses the M lowest bidders to reduce their consumption, making payment to each agent equal to the minimum of the $M + 1$ st lowest bid and the reserve price r . Denote $b = (b_1, \dots, b_n)$ as a bid profile and assume w.l.o.g. that $b_1 \leq \dots \leq b_n$ (breaking ties arbitrarily.)*

- *Allocation rule:* $x_i(b) = 1$, if $i \leq M$ and $b_i \leq r$.
- *Payment rule:* pay allocated agents $\min(r, b_{M+1})$. No payment to or from unallocated agents.

The period one strategic problem for an agent in the spot auction is straightforward. An agent who has prepared, and is now able to respond at a cost of v_i has a dominant strategy to bid $b_i^* = v_i$. What is challenging for an agent is to decide whether or not to invest effort and prepare in period zero.

5.1 Period-Zero Preparation Decision

For simplicity, assume that every agent has the same cost of preparation $c_i = c$, and the same probability of being able to respond $p_i = p$. As earlier in the paper, we assume $np \gg M$, so that if all agents prepare the capacity requirement will be met with high probability.

Assume w.l.o.g. $v_1 \leq \dots \leq v_n$ so that the agents with lower index find it less costly to reduce demand, and denote π_i as the probability that agent i prepares for demand response. Let X_i be a Bernoulli random variable with probability $p_i (=p)$ and let $X^{(m)} = \sum_{i=1}^m X_i$.

Proposition 1 (Pure Strategy Nash Equilibrium in preparation). *Under the spot auction with capacity M and reserve price r , there is a pure Nash equilibrium where each agent i prepares according to:*

$$\pi_i^* = \begin{cases} 1, & \text{for } i \leq m_r \\ 0, & \text{for } i > m_r \end{cases} \quad (6)$$

The threshold m_r is the largest index $m_r = \max_{m \in N} m$ s.t.

$$(r - v_m) \cdot p \cdot \mathbb{P}[X^{(m-1)} \leq M - 1] - c \geq 0. \quad (7)$$

If there is no such agent, then $m_r = 0$ and $\pi_i^* = 0, \forall i$.

We give the intuition for this result. First, observe that if an agent prepares and is able to respond, the highest amount that she can get paid is bounded by r . When $m_r = 0$, we have $(r - v_1)p - c < 0$, which implies that even the lowest-cost agent always loses in expectation by preparing and thus $\pi_i^* = 0$ is the unique equilibrium.

Assume now that $m_r > 0$. We can check that the left hand side of (7) is the expected utility of agent m assuming agents $i \leq m$ all prepare with probability 1 and agents $i > m$ all prepare with probability 0. Therefore agent m_r gets non-negative utility under (6), and we can check agents $i < m$ get non-negative utility from preparation, and that for an agent $i > m$, deviate from π_i^* and prepare would result in an expected utility smaller than zero.

5.2 Reserve r and Reliability Target

Under the same setup as in Section 5.1, in order to meet the reliability target τ , we need to make sure that a minimum number of agents m^* such that $\mathbb{P}[X^{(m^*)} \geq M] \geq \tau$ are preparing in equilibrium. We know from (7) that this requires $(r - v_{m^*}) \cdot p \cdot \mathbb{P}[X^{(m^*-1)} \leq M - 1] - c \geq 0$ which can be rewritten as a lower bound on the reserve price $r \geq v_{m^*} + c / (p \cdot \mathbb{P}[X^{(m^*-1)} \leq M - 1])$. When M is large, $\mathbb{P}[X^{(m^*-1)} \leq M - 1]$ and $\mathbb{P}[X^{(m^*)} \leq M]$ are close thus the reserve price needs to be $r \approx v_{m^*} + c / (p(1 - \tau))$. As τ becomes close to 1, the minimum reserve required to meet the reliability target is a lot larger than both c and v_i .

Remark: The only way that the spot auction can meet the reliability target is to set a huge reserve price, which is paid with a very small probability close to $1 - \tau$. Although an equilibrium strategy exists, participation in the spot auction is essentially a lottery for both the agent and mechanism designer. Most of the time the capacity is met and the $M + 1$ st bid is paid, but once in $1/(1 - \tau)$ events the mechanism makes a very large payment due to the huge reserve price. Hence, agents must be willing to gamble by preparing for DR. Moreover, the potential high payoff (much higher than in the direct mechanism) makes the spot auction susceptible to collusion.

6 Simulation Results

We show in this section via simulation that the direct and indirect mechanisms have good performance, comparing with the best possible outcome (in a world without private information) as well as the spot auction.




Figure 3: Comparison between the number of selected agents in the direct and indirect mechanisms.

Direct, Indirect Mechanisms v.s. First Best We compare the number of agents selected by the direct and indirect mechanisms with what we call the “first best,” which assumes that the mechanism knows the types of agents and can select agents in decreasing order of p_i 's (this would not be truthful). Let the total number of agents be $n = 500$ and the types be iid from the distributions: $v_i \sim U[0, 2]$, $c_i \sim U[0, 2]$, $p_i \sim U[0, 1]$. We first assume that the grid pays reward $R = 10$, and would like to achieve a capacity of $M = 100$.

With τ varying from 0.9 to 0.999, the average number of selected agents over 1000 economies are as shown in Figure 3(a). The horizontal axis $-\log_{10}(1 - \tau)$ translates $\tau = 0.9$ to 1 and $\tau = 0.999$ to 3. We can see that under both the direct and indirect mechanisms, more agents are selected when τ increases. Both mechanisms are doing well comparing with the first best, and the indirect mechanism selects roughly 10 more agents due to the estimation of p_i 's.

Fixing $\tau = 0.98$, the effect of varying reward R is as shown in Figure 3(b). The numbers of selected agents decrease in both mechanisms as R increases, since with higher R , the willingness to pay z_i 's are more aligned with the reliability p_i so the mechanisms are effectively selecting agents that are more reliable. For the indirect mechanism, increasing R also improves the inference on p_i and we see an additional significant drop in the number of selected agents.

As the reward R increases, both the reward and the penalty increase, however the increase in total paid rewards outweigh the additional penalty collected. The effect is that the total (expected) cost increases as the reward R increases (figure not included due to space limit). The cost under the indirect mechanism is typically 5% to 10% (depending on R) larger since more agents are selected.

Direct Mechanism v.s. Spot Auction. Consider a fixed economy with $n = 1000$ agents, each of whom has the same $p_i = 0.8$ and $c_i = 2$. Assume that $v_i = i/100$, so that the agents with lower index find it less costly to reduce demand. For comparison between the direct mechanism and the spot auction, we set the reward R and reserve r (in the direct and spot, respectively) to be the minimum values such that a minimum number of agents (denoted m^*) prepares thus $M = 100$ units are guaranteed with probability τ . For the direct mechanism, R needs to be large enough such that m^* agents have




Figure 4: Comparison between the average and std. dev. of total cost to the mechanism under the direct mechanism and spot auction.

non-negative zero crossings, and for the spot auction r needs to incentivize m^* agents to prepare in equilibrium.

The mean and standard deviation (std) of the total costs (reward payments minus collected penalties for the direct mechanism) computed over 1 million economies are shown in Figure 4. The total cost under the direct mechanism is lower than that of spot auction, moreover, the std of the total costs under the spot auction is extremely high.

This is because under the spot auction, the total cost are low most of the times (when the $M+1$ th bid is paid), however, with probability close to $1 - \tau$ the huge reserve price is paid, and this results in a huge variance in the total payments. Given the high variability of the payment (and the high payoff from colluding), we conclude a spot auction is less practical than the two-period mechanisms.

7 Conclusions

We studied the problem of incentivizing truthfulness when selecting from a number of unreliable providers in demand response. We introduce two new, dominant-strategy equilibrium mechanisms. The mechanisms are almost first-best in their ability to select a small number of reliable agents, and achieve lower payments and lower variance in payments than the spot auction. In future work, we plan to understand whether it is possible to meet the reduction target with high probability without reducing beyond the target and while retaining dominant-strategy equilibrium. Another interesting direction is to consider agents for whom the probability of responding is a function of the effort invested in preparation.

Acknowledgments

Parkes is supported by the SEAS TomKat fund, Li by NSF CAREER 1553407, and Robu by EPSRC EP/P001173/1: Centre for Energy Systems Integration.

References

- [Akasiadis and Chalkiadakis, 2013] Charilaos Akasiadis and Georgios Chalkiadakis. Agent cooperatives for effective power consumption shifting. In *Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence*, 2013.
- [Balandat *et al.*, 2014] Maximilian Balandat, Frauke Oldewurtel, Mo Chen, and Claire Tomlin. Contract design for frequency regulation by aggregations of commercial buildings. In *Communication, Control, and Computing (Allerton)*, 2014 52nd Annual Allerton Conference on, pages 38–45. IEEE, 2014.
- [Borenstein *et al.*, 2002] Severin Borenstein, Michael Jaske, and Arthurenfeld Ros. Dynamic Pricing, Advanced Metering, and Demand Response in Electricity Markets. *Journal of the American Chemical Society*, 128(12):4136–45, 2002.
- [Chen and Liu, 1997] Sean X Chen and Jun S Liu. Statistical applications of the Poisson-Binomial and Conditional Bernoulli distributions. *Statistica Sinica*, 7:875–892, 1997.
- [Chernoff, 1952] Herman Chernoff. A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations. *The Annals of Mathematical Statistics*, pages 493–507, 1952.
- [De Vries and Heijnen, 2008] Laurens De Vries and Petra Heijnen. The impact of electricity market design upon investment under uncertainty: The effectiveness of capacity mechanisms. *Utilities Policy*, 16(3):215–227, 2008.
- [Fernández and Williams, 2010] Manuel Fernández and Stuart Williams. Closed-form expression for the poisson-binomial probability density function. *IEEE Transactions on Aerospace and Electronic Systems*, 46(2):803–817, 2010.
- [Haring *et al.*, 2013] Tobias Haring, Johanna L Mathieu, and Goran Andersson. Decentralized contract design for demand response. In *European Energy Market (EEM)*, 2013 10th International Conference on the, pages 1–8. IEEE, 2013.
- [Johari and Tsitsiklis, 2011] R. Johari and J. N. Tsitsiklis. Parameterized supply function bidding: Equilibrium and efficiency. *Operations Research*, 59(5):1079–1089, 2011.
- [Kwag and Kim, 2014] Hyung-Geun Kwag and Jin-O Kim. Reliability modeling of demand response considering uncertainty of customer behavior. *Applied Energy*, 122:24–33, 2014.
- [Li *et al.*, 2015] Na Li, Lijun Chen, and Munther Dahleh. Demand response using linear supply function bidding. *IEEE Transactions on Smart Grid*, 6(4):1827–1838, 2015.
- [Ma *et al.*, 2015] Hongyao Ma, Reshef Meir, David C. Parkes, and James Zou. Are you going to do that? contingent-payment mechanisms to improve coordination. The Third Conference on Auctions, Market Mechanisms and Their Applications (AMMA’15), 2015.
- [Nisan, 2007] Noam Nisan. Introduction to mechanism design (for computer scientists). In Noam Nisan *et al.*, editor, *Algorithmic game theory*. Cambridge University Press, 2007.
- [Palensky and Dietrich, 2011] P. Palensky and D. Dietrich. Demand side management: Demand response, intelligent energy systems, and smart loads. *Industrial Informatics, IEEE Transactions on*, 7(3):381–388, 2011.
- [PJM, 2015] PJM. Energy & ancillary services market operations. Technical report, PJM, AUGUST 2015.
- [Radke Jr and Evanoff, 1994] George E Radke Jr and Javon Evanoff. A fast recursive algorithm to compute the probability of m-out-of-n events. In *Reliability and Maintainability Symposium, 1994. Proceedings., Annual*, pages 114–117. IEEE, 1994.
- [Robu *et al.*, 2012] Valentin Robu, Ramachandra Kota, Georgios Chalkiadakis, Alex Rogers, and Nicholas R. Jennings. Cooperative virtual power plant formation using scoring rules. In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence*, 2012.
- [Robu *et al.*, 2013] Valentin Robu, Enrico H. Gerding, Sebastian Stein, David C. Parkes, Alex Rogers, and Nick R. Jennings. An online mechanism for multi-unit demand and its application to plug-in hybrid electric vehicle charging. *J. Artif. Intell. Res. (JAIR)*, 48:175–230, 2013.
- [Rose *et al.*, 2012] Harry Thomas Rose, Alex Rogers, and Enrico H. Gerding. A scoring rule-based mechanism for aggregate demand prediction in the smart grid. In *International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2012*, pages 661–668, 2012.
- [Su *et al.*, 2014] Wencong Su, Jianhui Wang, and Jaehyung Roh. Stochastic energy scheduling in microgrids with intermittent renewable energy resources. *Smart Grid, IEEE Transactions on*, 5(4):1876–1883, 2014.
- [Yang *et al.*, 2015] Insoon Yang, Duncan S Callaway, and Claire J Tomlin. Indirect load control for electricity market risk management via risk-limiting dynamic contracts. In *American Control Conference (ACC)*, 2015, pages 3025–3031. IEEE, 2015.
- [Zhang *et al.*, 2015] Baosen Zhang, R. Johari, and R. Rajagopal. Competition and coalition formation of renewable power producers. *Power Systems, IEEE Transactions on*, 30(3):1624–1632, 2015.