

Contract Design for Energy Demand Response

Reshef Meir

Technion – Israel Institute of Technology
Haifa, Israel
reshefm@ie.technion.ac.il

Hongyao Ma

Harvard University
Cambridge, MA, US
hma@seas.harvard.edu

Valentin Robu

Heriot-Watt University
Edinburgh, UK
V.Robu@hw.ac.uk

Abstract

Power companies such as Southern California Edison (SCE) uses Demand Response (DR) contracts to incentivize consumers to reduce their power consumption during periods when demand forecast exceeds supply. Current mechanisms in use offer contracts to consumers independent of one another, do not take into consideration consumers' heterogeneity in consumption profile or reliability, and fail to achieve high participation.

We introduce DR-VCG, a new DR mechanism that offers a flexible set of contracts (which may include the standard SCE contracts) and uses VCG pricing. We prove that DR-VCG elicits truthful bids, incentivizes honest preparation efforts, and enables efficient computation of allocation and prices. With simple fixed-penalty contracts, the optimization goal of the mechanism is an upper bound on probability that the reduction target is missed. Extensive simulations show that compared to the current mechanism deployed by SCE, the DR-VCG mechanism achieves higher participation, increased reliability, and significantly reduced total expenses.

1 Introduction

Power system operation involves many challenges, driven by the requirement that supply equals demand at all times. Too much supply may lead to overload on the grid, whereas excessive demand may lead to shortages and blackouts. The problem is aggravated by the fact that consumption tends to vary sharply due to certain events (for example, surges in consumption during heatwaves), whereas increasing the supply is typically slow and costly. Even if the power company wants to shift some of the demand to a different time, it cannot coerce the consumers to do so, and may only affect their behavior by using monetary incentives, such as increasing electricity price during peak-demand times. As in other markets, consumers may respond to incentives in different ways based on their own preferences. Unlike some other markets, the serious consequences of failure to meet demand, and the large uncertainty about how consumers may react to incentives, requires the power company to guarantee there is enough slack on the supply side.

The DR-SCE mechanism Demand Response (DR) programs are used by power companies to handle surges in demand by reducing consumption rather than increasing production. Typically, when a surge is predicted one day ahead of time, the company lets consumers bid on how much consumption they can reduce. Each consumer is being paid \$0.5 per kWh, but only if the reduced consumption is between 50% and 150% of her bid. We call this system the *DR-SCE mechanism*, as it is used by Southern California Edison and PG&E [Patterson *et al.*, 2014; Hansen *et al.*, 2014].

DR-SCE has several shortcomings: incentives for participation are often insufficient (only 12% of registered participants in 2012-2013 submitted any bid), the system does not capture the very different consumption profiles of consumers, and does not filter out unreliable bidders. Yet, being a widely deployed DR system, we treat DR-SCE both as a starting point and as a benchmark for new mechanisms.

Contribution We propose a novel DR-VCG mechanism for selecting and incentivizing a subset of consumers to reduce consumption. The grid offers a set of contracts defined by some desired reduction target and a penalty scheme, and agents may bid how much they want to get paid on for accepting each contract. The mechanism then selects a subset of contracts that minimizes the *sum of bids*, and applies VCG prices to pay the agents. As a result, it is a dominant strategy for all agents to bid their true costs.

We show that for natural penalty schemes, the sum of bids is a good proxy for the reliability of the joint contract, as high bids are indicative of low individual reliability. We show that the current contracts used by SCE and PG&E can still be offered under DR-VCG (to allow for easy transition and backward-compatibility). We demonstrate via examples and simulations that even when restricted to offering SCE-like contracts, DR-VCG dominates DR-SCE in terms of reliability and grid expense. All omitted proofs are available in the full version of this paper [Meir *et al.*, 2017].

Related Work A number of recent works have discussed how groups of agents can be coordinated and incentivized to shift power demand [Haring *et al.*, 2016; Zhang *et al.*, 2015; Su *et al.*, 2014]. Considering strategic agents, [Rose *et al.*, 2012] and [Akasiadis and Chalkiadakis, 2013] propose the

use of *scoring rules* to incentivize truthful reports about expected future generation or consumption. However, scoring rule approaches are not concerned with selection of agents to satisfy a system-wide reliability constraint. [Li *et al.*, 2015] consider agents bidding using supply curves, and study the market equilibria for this setting. They do not, however take a mechanism design perspective or try to guarantee truthfulness. Mechanism design approaches for aggregating load apply either variations of VCG [Samadi *et al.*, 2012; Chapman and Verbic, 2017], offline optimization of dynamic pricing [Höning and La Poutré, 2013], or a new “staggered clock-proxy” auction [Nekouei *et al.*, 2015]. Neither work considers the crucial issue of *reliability*, i.e. in practice not all agents who intend to respond will do so. A different variation of VCG pricing was used by [Porter *et al.*, 2008] to align the incentives of agents in face of possible failures in general mechanisms. However their version requires full revelation (which is problematic in practical DR programs), and is aimed at maximizing social welfare rather than reliability.

The work closest to ours is [Ma *et al.*, 2016], who propose a mechanism that allows agents to bid on the maximal penalty for failing the DR contract, while [Ma *et al.*, 2017] extend this work to include uncertainty about costs. Our work takes a different approach, which generalizes currently used contracts and, in our view, is more geared to practical applications.

2 Model

We consider a single power utility or system operator (henceforth, the *grid*), and a set N of consumers (or *agents*) who are registered to a demand response program.

The most important task of the grid is to cut down consumption by at least M energy units (say, kWh) during the DR event. Given that this target is met, the grid would like to minimize payments to the agents. The baseline consumption profile of each agent is assumed known from past consumption data, so the grid can measure how much each agent reduced in practice.

2.1 Contracts

A contract in our model is defined by a pair (ℓ, F) , where ℓ is a *commitment goal* in energy units and $F : \mathbb{N} \rightarrow \mathbb{R}_+$ is a penalty function, mapping the realized energy reduction X to a monetary penalty $F(X)$. A-priori, F is unconstrained, and ℓ is merely a non-binding declaration of the agent’s intentions. It makes sense to consider more specific classes of penalties that attach the penalty to the commitment goal.

Fixed contracts A *Fixed penalty contract* is defined by a pair (ℓ, f_ℓ) , and the penalty is set to $F(X) = f_\ell$ if $X < \ell$ and 0 otherwise. In other words, the agent commits to reduce ℓ , or otherwise pay a penalty of f_ℓ . See Fig. 1(a).

Cliff contracts A *Cliff penalty contract* is defined by a tuple $(\ell, f_\ell, \alpha, \beta)$ where $\alpha < 1, \beta > 0, f_\ell \geq \ell(1 - \alpha)\beta$. It has the following form:

$$F(X) = \begin{cases} f_\ell, & X < \alpha \cdot \ell \\ (\ell - X)\beta, & \alpha \cdot \ell \leq X < \ell \\ 0, & \ell \leq X \end{cases}$$

We can think of a Cliff penalty function as a plateau where the penalty is 0 whenever the commitment ℓ is met. Failure




Figure 1: The penalty $-F$ as a function of the realized reduction X under a Fixed contract (a), a Cliff contract (b) and a general contract (c).

to meet the goal results in a linear penalty, where beyond a certain point the penalty becomes a constant, and the utility drops sharply (hence a “cliff”). See Fig. 1(b).

Clearly any Fixed contract is a Cliff contract. As we will later see, the SCE payment scheme can be implemented as a particular Cliff penalty scheme, so even restricting our mechanism to using Cliff contracts is sufficient to generalize the SCE system.

Optimal contract sets For what follows, we assume no structural restriction on contracts or penalty schemes. We simply assume that a set of k contracts J are offered, and $F(j, X)$ is the penalty for an agent who signs up for contract $j \in J$ and reduces consumption by X . Since the value of a contract to an agent is always non-positive, denote by $B_{ij} \geq 0$ the bid of agent i on contract j .¹ Then for a subset of contracts $S \subseteq N \times J$, we denote the *sum of bids* by $SB(S) = \sum_{(i,j) \in S} B_{ij}$.

In addition, the grid may pose a restriction on which sets of contracts are valid: we denote by \mathcal{S} all valid sets of contracts. In this work, we use this constraint to impose a lower bound on (declared) reduction in consumption, thus

$$\mathcal{S}(M) = \{S : \sum_{(i,j) \in S} \ell_j \geq M \text{ and } \forall i \{ (i,j) \in S \} \leq 1\}.$$

In other words, $\mathcal{S}(M)$ includes all sets of contracts that *claim* to reduce at least M units of consumption, and each agent has at most one contract. For an agent $i \in N$, we denote $SB_{-i}(S) = \sum_{(i',j) \in S: i' \neq i} B_{i'j}$, that is, the sum of bids over all agents in S except i . We sometimes denote $i \in S$ as a selected agent (meaning there is some j s.t. $(i,j) \in S$).

An *optimal contract set* is a set of individual contracts that minimizes the sum of bids, i.e. $\operatorname{argmin}_{S \in \mathcal{S}} SB(S)$.

2.2 The DR-VCG Mechanism

We define the *DR-VCG mechanism*, for assigning demand response contracts using Vickrey-Clarke-Groves (VCG) payments. The grid publishes a finite set of contracts J . Each agent i submits a single bid B_{ij} on each contract j . For now, we can think of B_{ij} as some proxy of the cost required from i when taking on contract j . The mechanism finds the optimal valid set of contracts by solving $\min_{S \in \mathcal{S}} SB(S)$.

¹The bid is supposed to reflect the various costs involved in taking the contract: preparation cost, online adjustment costs, the expected penalty and so on. We elaborate on this in the next section.

For a set of agents N , a set of contracts J , and a reduction goal M , we can plug in our more specific optimization goal and constraints. We get the optimal subset of contracts as

$$S^*(N, J, M) = \operatorname{argmin}_{S \in \mathcal{S}(M)} \left\{ \sum_{(i,j) \in S} B_{ij} \right\}.$$

We denote $SB^*(N, J, M) = SB(S^*(N, J, M))$, and omit some of the parameters when they are clear from the context. Then, for each selected agent $i \in S^*$, the individual rewards are computed as the VCG payments with the Clarke pivot rule [Clarke, 1971]. Informally, the sum of bids is analog of the *social cost*, and the VCG payment is the positive externality the agent's presence has on the rest of the agents. Formally, for each $i \in N$,

$$r_i = SB^*(N_{-i}, J) - SB_{-i}^*(N, J).$$

The reward is paid to the agent up front, regardless of how much reduction it eventually achieves in practice.

Finally, for each $(i, j) \in S^*(N, J)$, the selected agent i pays $F(j, X_i)$ to the grid, where each X_i is the realized reduction of agent i . The utility of agent i depends on the reward, the penalty, and the investment costs required to meet the contract. We analyze agents' incentives and utilities in the next section.

Example 1 (Running example). Suppose we apply the DR-VCG mechanism with a single fixed contract ($\ell = 100, f_\ell = 50$). There are three agents that submit bids of $B_1 = 0, B_2 = 5, B_3 = 15$, and the goal for the grid is set to $M = 200$.

The optimal set of contracts that meets $M = 200$ is $S^* = \{1, 2\}$ with $SB^* = 5$. The rewards are: $r_1 = SB^*(\{2, 3\}) - SB_{-1}^*(\{1, 2\}) = 20 - 5 = 15$; $r_2 = 15 - 0 = 15$, so in total the grid pays 30 (some of which it might get back as penalties).

3 Analysis

Complexity In order to run the DR-VCG mechanism, we should be able to efficiently compute the optimal contract set and prices. Suppose that energy units (including the reduction goal M) are integers.

Theorem 1. *For any sets of agents N and Cliff contracts J , both of $S^*(N, J)$ and $SB^*(N, J)$ (and thus also VCG prices) can be computed in time polynomial in n, k, M .*

In the general case, finding an optimal set of contracts is NP-hard even for fixed contracts, by a reduction from the Knapsack problem [Karp, 1972] (details omitted).

However, the knapsack problem is solvable by dynamic programming when the units are *bounded integers*, and a similar algorithm can be applied to our problem (intuitively, compute dynamically the optimal contract sets for agents $\{1, \dots, i\}$ for $i = 1, 2, \dots, n$). Hence we get Theorem 1.

3.1 Incentives

To make things concrete we will describe a particular probabilistic model from which we can derive agents' costs, and will show how under this model the incentives of the agents align nicely with those of the grid. However the claim the agents' dominant strategy is to reveal their true costs does not depend on this interpretation, and holds whenever agents can attribute a well-defined cost to each contract.

Effort and types In general, agents do not know with certainty the amount of energy they will be able to reduce, as this depends on some unknown factors such as urgent service orders, last minute clients and so on. Moreover, preparation may have some cost (e.g. due to changing the work schedule, or turning down orders). By investing a higher effort/cost, an agent might be able to commit to saving more energy.

The *type* of each agent i is given by a distribution p_i . In detail, $p_i(c, X)$ is the probability that by investing c , agent i will reduce exactly X units of consumption (thus $\sum_{X \geq 0} p_i(c, X) = 1$ for all c).² We assume agents are always trying to maximize their utility.

A straightforward approach would be to ask agents to report their types (and then apply some version of VCG). However, a language to report an arbitrary distribution may be very complicated. Further, unsophisticated agents like small households may not know their own distribution (or even *what is* a distribution). Fortunately, our DR-VCG does not require the agents to report any such distribution.

Fix an agent i who accepted a contract j (with penalty scheme F) and gets paid reward r_i . If she decides to invest c , she will pay an expected penalty of $EF_i(j, c) = E_{X \sim p_i(c)}[F(j, X)]$, and her expected utility would be:

$$u_i(j, c) = r_i - c - EF_i(j, c) = r_i - c - E_{X \sim p_i(c)}[F(j, X)].$$

Therefore, the *optimal investment* an utility maximizing agent should make for contract j should be

$$c_i^*(j) = \operatorname{argmax}_c u_i(j, c) = \operatorname{argmin}_{c \geq 0} (c + EF_i(j, c)).$$

In words, when agent i is signed up for contract j , investing $c_i^*(j)$ will minimize her total cost (investment + penalty). We denote this cost by $C_i^*(j) = c_i^*(j) + EF_i(j, c_i^*(j)) = \min_{c \geq 0} (c + EF_i(j, c))$. We refer to $C_i^* : J \rightarrow \mathbb{R}^+$ as the *cost type* of agent i , which is derived from her type p_i and F .

The *total expense* (TE) of the mechanism can be computed as the sum of rewards paid to the agents minus expected penalties: $TE(S) = \sum_{(i,j) \in S} (r_i - EF_i(j, c_i^*(j)))$ (assuming agents invest optimally).

A mechanism is *truthful* if it is a dominant strategy for any agent i to report her true cost $C_i^*(j)$ on any contract j .

A mechanism is *individually rational* (IR) if for every pair (i, j) selected by the mechanism, $u_i(j, c_i^*(j)) \geq 0$ (that is, by participating each agent does not lose in expectation).

Theorem 2. *Consider DR-VCG with arbitrary N, J, M .*

1. *For every contract $j \in J$, it is a dominant strategy for agent i to bid $B_{ij} = C_i^*(j)$;*
2. *If contract (i, j) is selected, it is a dominant strategy for i to invest $c_i^*(j)$;*
3. *The mechanism is IR.*

Proof sketch. Intuitively, we show that VCG payments are market clearing (following similar proofs in other domains, see [Nisan, 2007]), i.e. that no agent prefers a different contract (or no contract) under the given prices. Since the prices that agent i faces are independent of her bids, it is a dominant

²Investment c may include preparation costs, on-line actions required to produce the energy cut X , opportunity cost, and so on.

strategy to report truthfully. Once contract (i, j) is selected, then $u_i(j, c) = r_i - c - EF_i(j, c)$. By definition, this is maximized by investing $c_i^*(j)$. \square

Example 2. Consider three agents, where each one can reduce consumption by 100 kWh without effort. However, agents have different reliability and only manage to hold their commitment with a probability p_i of 1, 0.9, and 0.7, respectively (and otherwise reduce 0). Suppose that the goal of the grid is $M = 200$ kWh.

Consider the DR-VCG mechanism with a single fixed contract ($\ell = 100, f_\ell = 50$). Since agent 1 always meets her commitment, $C_1^* = 0$. For the others, $C_2^* = 0.1 \cdot 50 = 5$ (as agent 2 fails w.p. 0.1), and $C_3^* = 0.3 \cdot 50 = 15$, i.e. $B_i = C_i^*$ are exactly the bids in Ex. 1. The selected set is $S^* = \{1, 2\}$ (the DR-VCG mechanism filters out the least reliable agent). The total expense of the grid is $TE^{VCG} = 30 - 5 = 25$ (expected penalty of 5 from agent 2).

3.2 Fallback Options and Reserve Costs

In general, the grid may not find enough agents to meet the reduction goal M , and may thus need to use some fallback option like a standby generator or emergency blackouts. For every amount m we denote by R_m the cost for the grid of using its fallback option to reduce consumption/increase production by m units. E.g. if the total cost of demand response contracts exceeds R_M , then the power company is better off without assigning any contracts, or the fallback options can be used to fill up some gap between $\sum_{(i,j) \in S^*} \ell_j$ and M . A fallback option can be simply added to the mechanism as a ‘virtual agent’ that bids R_m on the fixed contract ($\ell = m, F \equiv R_m$). This is similar to the role of reserve prices in auctions. The reserve costs guarantee that: (I) The grid finds a cheap set of solutions, whether these solutions are contracts with agents or external options; and (II) For any (i, j) , the reward is bounded: $r_i \leq R_{M-L} - R_{M-(L+\ell_j)}$ where $L = \sum_{j' \in S^* \setminus \{j\}} \ell_{j'}$.

4 Reliability and Expenses

The incentive analysis we presented goes through for any set of contracts. Yet, the goal of the grid is to match demand and supply, preferably at low total expenses, which is a-priori not the same as minimizing the sum of bids. By restricting DR-VCG to use structured constructs we can relate this goal.

We next analyze how the sum of bids relates to *reliability*, i.e. the probability that the reduction target is met.

4.1 Fixed Penalty Contracts

Denote by $P(S, m)$ the probability that a quantity of at least m is reduced under contracts S . Then $P^*(N, J, M) = P(S^*(N, J, M), M)$ is the *reliability* of the DR-VCG mechanism. We would like to measure or bound $\bar{P}^*(N, J, M) = 1 - P^*(N, J, M)$, which is the probability that the mechanism *fails* to meet the lower bound reduction M (we may omit some of the parameters).

Proposition 3. Let $J = \{(\ell_j, f)\}_{j=1,2,\dots}$ for some fixed f , then $\bar{P}^*(M) \leq \frac{1}{f} SB^*(M)$. This bound is tight.

Proof. For any agent i that is assigned a contract j : The optimal investment is $c_i^*(j)$. The expected penalty is

$$EF_i(j, c_i^*(j)) = Pr_{X \sim p_i(c_i^*(j))}[X < \ell_j] \cdot f,$$

i.e., proportional to the probability it will undershoot the commitment ℓ_j . The DR-VCG mechanism minimizes

$$SB(S) = \sum_{(i,j) \in S} C_i^*(j) = \sum_{(i,j) \in S} c_i^*(j) + f \cdot \sum_{(i,j) \in S} Pr[X_i < \ell_j],$$

that is, a combination of the total investment and the sum of individual failure probabilities. Note that

$$\begin{aligned} \sum_{(i,j) \in S} Pr[X_i < \ell_j] &\stackrel{(a)}{\geq} Pr[\exists (i,j) \in S \text{ s.t. } X_i < \ell_j] \\ &\stackrel{(b)}{\geq} Pr[\sum_{(i,j) \in S} X_i < \sum_{(i,j) \in S} \ell_j] \stackrel{(c)}{\geq} Pr[\sum_{(i,j) \in S} X_i < M]. \end{aligned}$$

Thus,

$$\begin{aligned} SB^*(M) &\geq f \cdot Pr[\sum_{(i,j) \in S^*(M)} X_i < M] + \sum_{(i,j) \in S^*(M)} c_i^*(j) \\ &= f \cdot \bar{P}^*(M) + \sum_{(i,j) \in S} c_i^*(j) \stackrel{(d)}{\geq} f \cdot \bar{P}^*(M). \end{aligned}$$

To see why the bound is tight, observe that the inequalities in the proof are tight if (respectively): (a) failure events are disjoint (i.e. maximally negatively correlated); (b) agents never reduce more than ℓ_j ; (c) the reduction goal is met exactly ($\sum \ell_j = M$); and (d) investments are 0. \square

Thus a fixed penalty lets us bound the probability that the grid fails (reduction goal is not met). As we increase f , the (bound on) failure probability becomes smaller, at higher expense (due to higher bids). Of course, this is a worst-case bound. If, for example, some agents exceed their commitment then this would compensate for failures of others, and will increase the probability that the reduction goal M is met.

In general the grid may set a higher goal $M' = \gamma M$ than the expected surge M as a safety margin. Another result ties this safety margin with the sum of bids.

Proposition 4. Suppose that there is a single Fixed contract $(1, f)$. Then $M' - E[\sum_{i \in S^*(M')} X_i] \leq \frac{1}{f} SB^*(M')$.

Thus if the grid sets M' s.t. $M' - \frac{1}{f} SB^*(M') \geq M$, then actual reduction is at least M in expectation.

4.2 Cliff Penalty Contracts

We saw that having a constant penalty for a violation allows us to bound the failure probability. A Cliff penalty is more ‘forgiving,’ yet it provides similar guarantees.

Proposition 5. Suppose that the set of possible contracts J is composed of Cliff penalty contracts of the form $(\ell_j, f, \alpha, \beta_j)$ for some fixed f and α (same for all contracts), then $\bar{P}(S, \alpha \cdot M) \leq \frac{1}{f} SB(S)$. This bound is tight.

That is, we get a guarantee on the probability that we miss the reduction goal by a factor of α (tightness is achieved if either $\alpha = 1$ or $\beta_j = 0$ for all j). Note that this does not require any assumption on the types of the agents. The grid can then sign contracts that sum up to $M' = \frac{M}{\alpha}$ so as to bound the probability of missing its actual goal \bar{M} .

5 DR-SCE vs. DR-VCG

We argue that the DR-SCE mechanism can be simulated exactly using a Cliff payment contract. Formally, in DR-SCE and each agent submits a bid b_i , and the grid selects agents at random until $\sum b_i \geq M$. After the reduction X_i is realized, each selected agent gets a reward of $r_i = 0$ if $X_i < b_i/2$, $r_i = \frac{X_i}{2}$ if $X_i \in [b_i/2, 3b_i/2]$, and $r_i = 3b_i/2$ otherwise.

We argue that same SCE contracts used today can be offered via the DR-VCG mechanism. For any $\ell > 0$, we define a Cliff contract $j_\ell = (\ell, f_\ell = \frac{\ell}{2}, \alpha = \frac{1}{3}, \beta = \frac{1}{2})$.

Proposition 6. *For any agent i of type p_i , submitting optimal bid b_i to the DR-SCE mechanism is ex-post equivalent to being the only bidder in DR-VCG with $M \geq b_i$, $J^{SCE} = \{j_\ell\}_{\ell \geq 0}$, and reserve prices $R_m = m/2$ for all m .*

Proof sketch. We show that contract b_i in DR-SCE is completely equivalent to contract j_ℓ where $\ell = 3b_i/2$ (i.e. same behavior and same ex-post utility). This is by writing the penalty function $F(j_\ell, X)$, and considering the realization of X_i when: $X_i < b_i/2$, $X_i \in [b_i/2, 3b_i/2]$, and $X_i > 3b_i/2$. Then bidding b_i is optimal in DR-SCE if and only if DR-VCG assigns j_ℓ to i . \square

Proposition 6 has two important implications. First, transition from the currently used DR-SCE mechanism to DR-VCG can be gradual and backward-compatible: we can still allow bids on quantity (b_i) and internally convert them to the appropriate Cliff contract j_ℓ with reserve price 0.5ℓ . Second, it becomes obvious that DR-SCE is just a very restricted version of the more general DR-VCG, where parameter values are arbitrary and most likely suboptimal. By setting the proper reduction target and reserve prices, we expect DR-VCG to outperform DR-SCE. In particular:

1. DR-SCE makes no informed selection. With an explicit reduction goal M (based on the actual surge prediction), DR-VCG **selects agents** who are more reliable. Thus we expect that $P^{VCG} > P^{SCE}$ in most cases.
2. DR-VCG pays rewards based on **competition**. In fact, for J^{SCE} and any M and S , $TE^{VCG} \leq TE^{SCE}$.
3. DR-VCG allows agents to bid on **multiple contracts**, so they can reveal more information on their type.
4. DR-SCE uses arbitrary price of $0.5m$ for contracts of size m , whereas DR-VCG is **flexible**. In particular we may use the actual costs of generating m kWh, which are highly non-linear due to the cost of adding another generator.

Example 3. Consider the same 3 agents from Example 2. In the DR-SCE mechanism, all agents will submit a bid of $b_i = 100$, and the grid cannot distinguish between them. If it selects two of them at random, it pays $TE^{SCE} = (50 + 45 + 35) \cdot \frac{2}{3} = 83.33$ —much more than $TE^{VCG} = 25$.

If we compare failure probabilities, then $\bar{P}^{SCE} = \frac{1}{3}0.1 + \frac{1}{3}0.3 + \frac{1}{3}(1 - 0.9 \cdot 0.7) = 0.224$, whereas $\bar{P}^{VCG} = 0.1$, which is again an improvement over SCE.

On the other hand, if agents have to invest high costs c_i^* then they might not participate in DR-SCE at all, as their reward is bounded by $\$0.5$. Thus DR-SCE pays *too much* to agents with low c_i^* , and *too little* to agents with high c_i^* .

The parameters of the Cliff contract j_ℓ are also arbitrary (e.g. why $\alpha = \frac{1}{3}$?), however there is no obvious way to set them a-priori (see Discussion).

5.1 Simulations

Settings Each agent i has T “effort levels,” where each level is a triple (c_{it}, q_{it}, p_i) , meaning that with investment c_{it} agent i can reduce q_{it} kWh. The reduction succeeds w.p. p_i , and w.p. $1 - p_i$ reduction is 0 due to an unexpected event. The expected demand surge is M , and the grid uses a safety margin $\gamma \geq 1$. For each mechanism we denote by $TE = TE(S^*(N, J, \gamma M))$ the total expense, and by $P = P(S^*(N, J, \gamma M), M)$ the reliability, i.e., the probability that $\sum X_i > M$ when the mechanism collects contracts for γM . We run both DR-SCE and DR-VCG mechanisms, and measure TE^{SCE} , TE^{VCG} , P^{SCE} , and P^{VCG} .

To set up a realistic scenario of a typical demand response event, we used [Patterson *et al.*, 2014] that summarize previous DR programs. We fix the expected demand surge to $M = 10\text{MWh}$. In each economy we sample n agents i.i.d., where each agent has $T \in \{1, 3, 5\}$ effort levels. For each agent $i \in N$ and effort level $t \leq T$: the capacity (in kWh) is $q_{it} \sim \text{Zipf}(1, 500) \cdot 10$; individual reliability is $p_i \sim U[0.7, 1]$; and agents’ investment costs (in \$) are $c_{it} \sim U[0.2, 1]$, multiplied by q_i .³ Note that only agents with $\max_t \frac{c_{it}}{q_{it}} \leq 0.5$ will submit bids in DR-SCE. We generated populations of 3 sizes: $n = 100$ (small), $n = 200$ (medium) and $n = 400$ (large), and for each population varied the safety margin between $\gamma \in [1, 2]$.

We run simulations that demonstrate the four advantages of the DR-VCG mechanism mentioned above. Every datapoint in our simulations is an average over 100 instances.

Selection and Competition In our first simulation, agents each have a single effort level ($T = 1$). We use the set of “SCE-like” contracts $J^{SCE} = \{j_\ell\}_{\ell=10,20,\dots}$, and set linear reserve prices $R_m = 0.5m$ for all $m \geq 0$. Thus for a single bidder, DR-SCE and DR-VCG are equivalent by Prop. 6.

Fig. 2 shows the expense-vs.-reliability frontier under both mechanisms. Larger safety margin γ results in more recruited agents and higher reliability, but also higher costs. We can see that in both populations DR-VCG dominates DR-SCE by guaranteeing *any* reliability level at a much lower cost. We found that even if we pay the reserve prices to all selected agents, DR-VCG does somewhat better than DR-SCE, meaning that it does indeed select better agents.

Fig. 2 also shows that the advantage of DR-VCG becomes larger in large populations (or when the expected surge is small), as competition drives prices down. In contrast, in

³This roughly mimics the aggregate statistics in the data, where agents’ bids are highly skewed, with a minimum of 10 kWh up to several MWh, overall reliability is ~ 0.85 , and participation is low (about 100 bidders out of 1000 registered users). We also tried different distributions and got similar results.




Figure 4: Outcomes under SCE and VCG mechanism for small population ($n = 100$) and flexible reserve prices.




Figure 2: Reliability vs. expense for large and medium populations and single effort level.




Figure 3: Reliability vs. expense with for medium population $n = 200$, with multiple effort levels.

small populations DR-VCG and DR-SCE are the same, as both exhaust all agents with low investment costs. Our next simulations show how the other two advantages of DR-VCG overcome this problem.

Multiple levels Fig. 3 shows the reliability frontier for a medium population ($n = 200$) of agents with multiple effort levels. We can see from Figs. 2 and 3 that DR-VCG performance becomes better as population gets larger and/or agents’ types are more complex, whereas the performance of DR-SCE remains almost the same. Intuitively, selecting from n agents each submitting T independent bids is similar to selecting from Tn agents, i.e. there is more competition.

Flexible reserve prices The SCE contracts with their linear reserve price do not reflect correctly the outside options available to the grid. In reality, the grid cannot generate additional power at small quantities to fill the gaps between agents’ bids and the reduction goal. Failure to reach the reduction goal γM means that the grid cannot rely on the current DR contracts, and must increase supply by operating another generator at a large cost. The operating cost with modern gas turbines is $\$0.04 - 0.1$ per kWh, but each turbine generates at least 100 MWh. We thus set the reserve prices to $R_m = 4000 + 0.1m$, which creates a dichotomy between “success” (where the DR mechanism collected enough contracts to forgo the additional turbine) and “failure.” Increasing the reserve prices also requires higher penalties. Otherwise agents may bid for contracts they do not plan to keep, with reward higher than the maximal penalty. We did not optimize the contract (see Discussion) and instead just set the penalties to $f_\ell = \ell$ (double from j_ℓ).

Fig. 4 shows how flexible prices benefit DR-VCG in *small populations*. As the target capacity γM increases, this requires high reliability using a small population, which DR-SCE very often fails to achieve. This is because it may not find enough reliable agents willing to bid for a payment of $\$0.5$, and it must use the extra turbine for a high cost. In contrast, DR-VCG can increase the reward to agents, thereby attracting also bidders with investment costs higher than $\$0.5$.

6 Discussion

We suggested in this paper the DR-VCG mechanism for demand-response contracts that is based on individual “soft” commitments and flexible penalty schemes. While the details of the contracts and the analysis were specific to demand response programs, the general idea of offering flexible penalty contracts to multiple agents may be useful in other domains that require joint effort under uncertainty [Porter *et al.*, 2008].

We considered three natural parametric classes of penalties that allow for efficient computation of VCG prices, and showed how they generalize the currently deployed SCE contract. Power companies can adopt the new DR-VCG mechanism with the SCE contracts for painless migration at first. Then, they can gradually add more contracts and optimize their parameters based on distributional assumptions, data on consumption profiles, trial-and-error, and so on. Another benefit of DR-VCG is that the grid can focus on optimizing the set of contracts without worrying about agents’ strategic behavior, whereas agents can focus on accurately estimating their costs for each contract. Based on initial simulations, it seems that penalties should be higher than $\$0.5$ per kWh, and perhaps superlinear in the size of the contract (as reliability of large consumers is more important). Finding optimal parameters under various assumptions is a major topic for future research.

Other future directions include extending the mechanism to consider temporal shifts of consumption (as in [Höning and La Poutré, 2013]), and obtaining a better understanding of the connections between reliability, penalties, and expenses.

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.
- [Chapman and Verbic, 2017] A. C. Chapman and G. Verbic. An iterative on-line auction mechanism for aggregated demand-side participation. *IEEE Transactions on Smart Grid*, 8(1):158–168, Jan 2017.
- [Clarke, 1971] E. H. Clarke. Multipart pricing of public goods. *Public choice*, 11(1):17–33, 1971.
- [Hansen *et al.*, 2014] Daniel G. Hansen, Steven D. Braithwait, David A. Armstrong, and Marlies C. Hilbrink. 2013 Load Impact Evaluation of California Statewide Demand Bidding Programs (DBP) for Non-Residential Customers: Ex Post and Ex Ante Report. http://www.calmac.org/publications/PY13_DBP_Ex_Ante_Report_20140401_PUBLIC_Revised.pdf, 2014.
- [Haring *et al.*, 2016] T. W. Haring, J. L. Mathieu, and G. Andersson. Comparing centralized and decentralized contract design enabling direct load control for reserves. *IEEE Transactions on Power Systems*, 31(3):2044–2054, May 2016.
- [Höning and La Poutré, 2013] Nicolas Höning and Han La Poutré. Reducing electricity consumption peaks with parametrised dynamic pricing strategies given maximal unit prices. In *Database and Expert Systems Applications (DEXA), 2013 24th International Workshop on*, pages 171–175. IEEE, 2013.
- [Karp, 1972] Richard M Karp. Reducibility among combinatorial problems. In *Complexity of computer computations*, pages 85–103. Springer, 1972.
- [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.*, 2016] Hongyao Ma, Valentin Robu, Na Li, and David C. Parkes. Incentivizing Reliability in Demand-Side Response. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI’16)*, 2016.
- [Ma *et al.*, 2017] Hongyao Ma, David C. Parkes, and Valentin Robu. Generalizing Demand Response Through Reward Bidding. In *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems, (AAMAS’17)*, 2017.
- [Meir *et al.*, 2017] Reshef Meir, Hongyao Ma, and Valentin Robu. Contract design for energy demand response. *arXiv preprint arXiv:1705.07300*, 2017.
- [Nekouei *et al.*, 2015] E. Nekouei, T. Alpcan, and D. Chattopadhyay. Game-theoretic frameworks for demand response in electricity markets. *IEEE Transactions on Smart Grid*, 6(2):748–758, March 2015.
- [Nisan, 2007] Noam Nisan. Introduction to mechanism design (for computer scientists). In Nisan *et al.*, editor, *Algorithmic game theory*. Cambridge University Press, 2007.
- [Patterson *et al.*, 2014] Olivia Patterson, Mary Sutter, and Alan Elliott. 2012-2013 PG&E and SCE Demand Bidding Program Process Evaluation. http://www.calmac.org/publications/California_2012-2013_Demand_Bidding_Program_Process_Evaluation_FINAL.pdf, 2014.
- [Porter *et al.*, 2008] Ryan Porter, Amir Ronen, Yoav Shoham, and Moshe Tennenholtz. Fault tolerant mechanism design. *Artificial Intelligence*, 172(15):1783–1799, 2008.
- [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.
- [Samadi *et al.*, 2012] Pedram Samadi, Hamed Mohsenian-Rad, Robert Schober, and Vincent WS Wong. Advanced demand side management for the future smart grid using mechanism design. *IEEE Transactions on Smart Grid*, 3(3):1170–1180, 2012.
- [Su *et al.*, 2014] Wencong Su, Jianhui Wang, and Jaehyung Roh. Stochastic energy scheduling in microgrids with intermittent renewable energy resources. *IEEE Transactions on Smart Grid*, 5(4):1876–1883, 2014.
- [Zhang *et al.*, 2015] Baosen Zhang, R. Johari, and R. Rajagopal. Competition and coalition formation of renewable power producers. *IEEE Transactions on Power Systems*, 30(3):1624–1632, 2015.