

Approximately Stable Pricing for Coordinated Purchasing of Electricity

Andrew Perrault and Craig Boutilier*

Department of Computer Science
University of Toronto
{perrault, cebly}@cs.toronto.edu

Abstract

Matching markets are often used in exchange settings (e.g., supply chain) to increase economic efficiency while respecting certain global constraints on outcomes. We investigate their application to pricing and cost sharing in *group purchasing* of electricity in smart grid settings. The task is complicated by the complexities of producer cost functions due to constraints on generation from different sources (they are sufficiently complex that welfare-optimal matchings are not usually in equilibrium). We develop two novel cost sharing schemes: one based on Shapley values that is “fair,” but computationally intensive; and one that captures many of the essential properties of Shapley pricing, but scales to large numbers of consumers. Empirical results show these schemes achieve a high degree of stability in practice and can be made more stable by sacrificing small amounts ($< 2\%$) of social welfare.

1 Introduction

Coordinating the group purchase and consumption of electricity can offer significant benefits in terms of economic efficiency, predictability, and fairness. These benefits emerge for several reasons: i) Consumers who are able to shift their loads from periods of high consumption can be compensated by others for the inconvenience or discomfort of doing so, with the resulting flatter demand profiles reducing overall cost of generation. ii) The formation of groups of consumers can increase competition by transferring large numbers of consumers between providers.¹ iii) Group purchasing can result in prices that are more responsive to market conditions, decreasing the ability of generators to exercise market power [Rassenti *et al.*, 2003]. iv) Groups can predict aggregate consumption levels more reliably and incentivize their members to consume at predicted levels [Robu *et al.*, 2014]. To facilitate group-level coordination, the tools of cooperative game

*Currently on leave at Google, Inc., Mountain View.

¹Group purchasing has been explored as a method of increasing electricity market competition, as in The Big Switch is a UK program where 30,000 households agreed to have their demand auctioned collectively to the lowest-bidding provider.

theory can be used to determine mutually acceptable cost sharing among users in a group.

The *smart grid*, particularly smart meters, makes it possible to analyze electricity consumption on an individual level, which is important for coordination. However, electricity markets are difficult to optimize and to price due to their size and complexity. Furthermore, reasonable properties of pricing functions in this market setting are not well understood. Producers have complex cost functions that depend on: minimum and maximum production levels; multiple *layers* of generation with different costs; and ramp constraints that constrain production adjustments over time. In addition, consumers have variable preferences, and may be willing to trade off cost for comfort/convenience and shift their loads.

These attributes pose challenges from a game theoretic perspective. Efficient outcomes under realistic modeling assumptions, as we will show, do not support stability from a coalitional perspective (e.g., core) or from a purely strategic perspective (e.g., Nash equilibrium); indeed, we will show that Nash equilibria can have arbitrarily lower quality than the best solution out of equilibrium. Thus, there is a natural tradeoff between stability and social welfare. We explore this tradeoff using two different cost sharing schemes, one based on Shapley values, and the other based on a new notion of *similarity-based envy freeness*. We show that small sacrifices in social welfare can provide large gains in stability. Furthermore, our similarity-based envy-free cost sharing scheme, while not as conceptually simple as Shapley, achieves greater stability and has significantly better computational properties.

Our contributions are: First, we develop a tractable market model for matching consumers to producers while reflecting many of the complexities of electricity production and consumption. Second, we explore the stability properties of this model under various cost sharing schemes. Finally, we develop two payment algorithms that exhibit high stability and fairness, while allowing tradeoffs between social welfare and stability to be made. In Sec. 2, we describe our basic market model and related work. Sec. 3 addresses stability in our model and describes our two payment models. Experimental results in Sec. 4 demonstrate their efficacy.

2 Setting

Let N be a set of n consumers and let M be a set of m electricity producers that each control a set of generation facili-

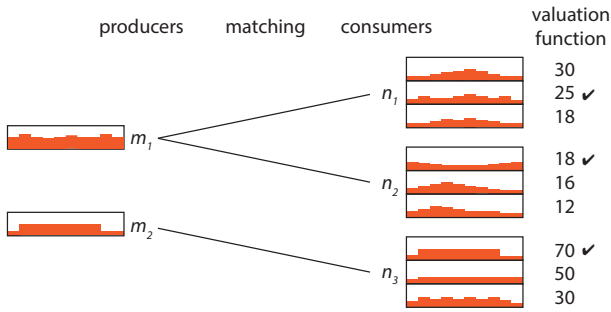


Figure 1: Our market model. The demand profiles chosen by the consumers depend on the prices charged by the producers. Each producer has a posted price function, which is not shown here.

ties. We assume T time periods, representing, for example, the hours in a day or week. Each consumer i has a non-empty set of *demand profiles* Π_i , where each profile $\pi \in \Pi_i \subset \mathbb{R}^T$ reflects an “acceptable” consumption pattern (electricity use per period in kilowatt-hours (kwh)) for i . Each consumer has a *valuation function* $V_i : \Pi_i \rightarrow \mathbb{R}$ indicating her value (in dollars) for each of her demand profiles (this captures preferences for and potential flexibility in consumption). Such profiles may be explicitly elicited or estimated using past consumption data. This is an abstraction of reality, of course, because consumer’s valuations derive from the actions that use the available electricity, not the electricity itself.

A *matching* μ maps each consumer i to a producer $\mu(i)$, from whom she purchases electricity, and demand profile $\mu^p(i) \in \Pi_i$, indicating her consumption. In this matching, i pays the price per unit *posted* by producer $\mu(i)$, which may depend on the aggregate demand of all agents matched to $\mu(i)$. (We discuss pricing and cost sharing further below). We include a *null producer* which represents any consumer’s best outside option (i.e., not participating in the market); the value i being matched to null is θ_i (we assume i ’s profile is chosen based on external factors). Fig. 1 shows a diagram of an instance of the model with two producers and three consumers, each with three demand profiles. The total demand that is served by each producer is the sum of the demands across the profiles chosen by the consumers that are matched to that producer.

Each producer j has a *price function* $P_j : \mathbb{R}^T \rightarrow \mathbb{R}$ that maps total demand in each time period to a price. We treat P_j as exogenous—it represents j ’s *posted prices*. We assume that P_j captures the fundamental features of generation costs (see next subsection), and that producers are not strategic, instead they simply recover their costs.² The null producer has a fixed zero price (net values θ_i includes any price charged).

The *social welfare* (SW) of matching μ is the net utility realized by all consumers acting on the demand profiles selected by μ , i.e., the sum of the consumers’ valuation minus the sum of the producers’ prices:³

$$SW(\mu) = \sum_{i \in N} V_i(\mu^p(i)) - \sum_{j \in M} P_j(\sum_{i' \in \mu^{-1}(j)} \mu^p(i'))$$

²This reflects the reality in many regulated markets, where producers must make only a fixed percentage return on their costs.

³Note that we consider the producer’s profit to be part of the cost of generating electricity.

Naturally, we’d like to maximize SW given user valuations and producer price functions (as elaborated below). Given the SW -optimal matching, our aim is to find a cost-sharing scheme that ensures producer costs are recovered. In addition, we want the matching to be *stable* in the sense that no consumer has an incentive to defect either by changing her matched profile or by switching to a different producer (or both). We turn to this in Sec. 3.

Optimization. Even for trivial producer price functions, maximization of SW is NP-hard by reduction to the model of Lu and Boutilier [2012], henceforth LB.⁴ Given the price functions we consider below, SW maximization can be formulated as a mixed integer program (MIP) in a straightforward way. The MIP can be solved by relaxing the binary matching variables, leaving a number of binary variables proportional to MT , and independent of the number of consumers and profiles. As in many matching problems, most relaxed variables are integral at the optimal solution in practice (just a few consumers may have their demand split across several generators). It may be acceptable for such agents to have contracts split across generators; otherwise, LP rounding may be used. SW maximization can be solved for large instances, since the number of producers (requiring integer variables) is generally very small compared to the number of consumers. We are able to solve instances with 5000 consumers, 2 producers, 4 profiles and 24 time periods in less than 15 minutes on a 12x2.6GHz, 32GB machine using CPLEX 12.51.

Producer Price Functions (PPFs). Deciding the optimal output levels of a group of generation facilities in order to meet system demand is complex and has been studied extensively [Kirschen and Strbac, 2004]. We focus on three of its most important features. i) Generation facilities have limited *ramp rate*—the amount by which they can change their output from one time period to the next. Ramp rates of different generation facilities vary radically (e.g., demand tracking plants such as natural gas can ramp up or down in half an hour, whereas nuclear plants take days. ii) Different kinds of generation have different variable costs (e.g., most renewables have low variable cost, while natural gas has high variable cost). iii) Finally, certain kinds of plants (e.g., coal) have high costs when run below a certain level—it imposes considerable wear on the components. Shutting down these plants also incurs costs. To capture these features, we model each producer as follows. It has a *base layer* that has low generation costs, but has a low ramp rate, and is expensive to take below a certain level of generation in any time period (the *minimum economic generation level (MEGL)*). It also has a *tracking layer* that can be adjusted rapidly or shut off entirely, but has high generation costs and limited capacity.

We provide a brief overview of the form of PPFs that we use below. The online appendix contains a complete specification as well as its formulation within a MIP. For producer

⁴LB can be simulated in our model by assuming a different time period for each agent and having each agent demand one unit of power in that period. While we do not allow consumer preferences over producers, these can be represented in the producer price functions. The LB model is NP-hard via reduction to Knapsack.

j 's base layer, let $c_j^{(l)}$ be the price per kwh, $d_j^{(l)+}$ be the capacity (kwh), $d_j^{(l)-}$ be the MEGL (kwh), and r_j be the maximum ramp rate between periods (kwh). Let s_j be the *shutdown cost* (in dollars) that is incurred when demand is reduced below the MEGL. Let $c_j^{(h)}$ be the price of the tracking layer per kwh and $d_j^{(h)+}$ be its capacity (kwh).

- If demand is smooth and does not exceed the maximum base layer capacity or fall short of the MEGL, only base layer costs are incurred. Formally, if demand in every period is in the interval $[d_j^{(l)-}, d_j^{(l)+}]$, and the largest period-to-period change in demand does not exceed r_j , the unit price is $c_j^{(l)}$. Demand that exceeds the base layer capacity will be met using the tracking layer, if capacity is available, at a price $c_j^{(h)}$. In Fig. 2a, the total cost to meet demand in the first time period is $c_j^{(l)} d_j^{(h)+}$ plus $c_j^{(h)}$ times the amount of demand that exceeds $d_j^{(h)+}$.
- A shutdown cost is charged if demand in a period is less than the MEGL and the demand in the previous period was greater than the MEGL. If demand in the previous period is greater than $d_j^{(l)-}$ and demand in the current period is less than $d_j^{(l)-}$, the shutdown cost of s_j is charged. In Fig. 2a, a shutdown occurs in the second period.
- If there is a large increase in demand between two periods, the first r_j units of the increase are met using the base layer at price $c_j^{(l)}$, and the remaining units of the increase are met using the tracking layer at price $c_j^{(h)}$. Fig. 2b shows the ramp costs that are incurred at time $t + 1$ given a moderate demand at time t . Note that the base or tracking layer may have insufficient capacity, which would result in a demand profile that cannot be served, i.e., it has infinite cost.
- If there is a large decrease in demand from period to period, an additional fee of $c_j^{(h)} - c_j^{(l)}$ per unit of decrease exceeding r_j is charged, which represents the cost of meeting the necessary amount of the previous period's demand using the tracking layer. Fig. 2c shows the ramp costs incurred at time $t + 1$ given a high demand at time t .

Our PPFs have the Markov property: the price paid in any period depends only on demand in that period and in the previous one, which makes them easy to compute. It is also gives a lower bound on the cost of meeting the demand by optimizing base and tracking layer production levels in each time period. Every cost incurred in the price function must be also be incurred by any solution that satisfies the generation constraints, but the price function may underestimate costs (e.g., it assumes that base layer ramping can be performed within two time periods). Our general approach for optimizing social welfare and stability can be applied to a variety of PPFs—the form of the PPF may be application-dependent.

Related Work. Assignment games and matching markets have been extensively studied using different stability

concepts and pricing models [Shapley and Shubik, 1971; Gale and Shapley, 1962; Demange *et al.*, 1986]. Research in real-world markets has largely focused on revenue maximization for monopolistic sellers, though strategic aspects are sometimes considered. The literature on group buying, summarized in [Anand and Aron, 2003; Chen and Roma, 2010], considers the value of offering discounts to groups of buyers who purchase items in bulk. Several group buying models are similar to ours.

Our work extends that of Lu and Boutilier [2012], who focus on a more restrictive model of buyer preferences (unit demand, only the supplier affects utility) and seller price functions (volume discounts). Similarly to them, we focus on the strategic behavior of buyers and treat seller prices as exogenous (strategic behavior of sellers was later investigated by Meir *et al.* [2014]). Two similar group buying models are those of Anand and Aron [2003] and Chen *et al.* [2007]. Both have seller prices that are affected by the amount purchased, but neither allow for sufficiently complex price functions to model electricity generation. Anand and Aron focuses on a single vendor and does not consider buyer coordination, while Chen *et al.* uses a multi-stage auction mechanism.

In the AI literature, the process of finding a group of fully cooperative buyers an optimal seller has been studied [Sarne and Kraus, 2005; Manisterski *et al.*, 2008]. In the context of electricity, group purchasing has been suggested as a way of reducing seller uncertainty about stochastic buyer demands [Robu *et al.*, 2014], an aspect which we do not consider here.

3 Cost Sharing and Stability Concepts

Finding a SW-optimal matching is relatively straightforward, although somewhat involved due to the complexities of the producer price functions. More difficult is the question of appropriate cost sharing among the group of consumers. By coordinating demand to maximize social welfare, some consumers sacrifice their own utility for the benefit of the group and thus should be compensated. Various notions of *stability* can be used for this purpose. Given some cost sharing scheme, stability measures the incentive for any consumer to defect, i.e., change their profile or producer. Of course, defining the stability of a cost sharing scheme, requires that defections themselves be priced, i.e., what does a consumer pay if they change their matching. We approach the issue of defection pricing from two perspectives: a *marginal cost defection model*, where a producer accepts any defector who pays the marginal cost they impose by defecting; and an *envy-free defection model* where a defector pays the same as any other consumer original matched to that producer with a similar profile.

Other than stability and budget-balance (i.e., all producers' costs are paid), there are several other desiderata of a cost sharing scheme. A matching is *envy-free* if no consumer would prefer the matched pair of any other consumer. This notion requires some generalization in our model (as we discuss below). A scheme should be transparent: it should be clear why a consumer is paying what they are, and what they can do to change what they pay. It should be deterministic and easy to describe so that outcomes do not appear arbitrary.




Figure 2: Diagrams showing how the various components of the producer price function are calculated.

We also desire computational scalability. Finally, we would like the ability to sacrifice SW to achieve these properties, especially stability, in a controllable way. History shows that stability in particular is an essential component of successful matching mechanisms [Roth, 2002].

First, we analyze cost sharing under the marginal-cost defection model and find that stability is difficult to achieve. Second, we present two cost-sharing schemes under the envy-free defection model that achieve many of our desiderata.

Cost Sharing under the Marginal-Cost Defection Model.

One difficulty in defining the cost/price of a defection is that the cost imposed by a consumer on a producer depends on those already matched to that producer. We address this with two cost sharing schemes below. Here we begin by considering *marginal cost payments*—where a consumer defecting to a new producer (or changing profile) pays the marginal cost imposed on the new producer—and its stability properties.

Ideally, we would like payments to be *core stable* [Shapley and Shubik, 1971], wherein no group of consumers can benefit by defecting. Core stability is not always achievable in this game: since LB can be reduced to our model, their proof that the core may be empty (under much simpler producer cost functions) applies to our model as well. Core stability is a very strong notion that is hard to attain without sacrificing SW. Fortunately, achieving it is not critical in practice if large groups of consumers cannot effectively discover, communicate and coordinate their actions.

A weaker, but more practical concept is *Nash stability* [Nash, 1950], which accounts for defections by single consumers. It is informative to examine the pure Nash equilibria (NE) with the worst and best SW. The ratio of optimal SW to the SW of the worst NE is the (*pure*) *price of anarchy* (PoA) [Koutsoupias and Papadimitriou, 1999], and to that of the best NE is the (*pure*) *price of stability* (PoS) [Schulz and Moses, 2003]. PoS is usually more appropriate when a centralized mechanism constructs the matching.

Tab. 1 shows our results for the PoS and PoA under producer price functions with various combinations of model features. We see that, under marginal-cost defection pricing, capacity constraints plus either ramp constraints or shutdown constraints are sufficient to ensure the non-existence of (pure)

NE. We sketch two proofs here—full versions are in the appendix.

Theorem 1. *There is no cost sharing scheme that achieves a price of stability better than ∞ when producer price functions may have capacity constraints and ramp constraints.*

Proof. Consider an instance with two consumers, two producers and a single time period. Each consumer has a single demand profile of 1 with value 2 and outside option value 0. Producer m_1 has an MEGL of 1.5, base layer cost 1, shutdown cost 1 and large base layer capacity. Producer m_2 has no MEGL, base layer cost 0.75 and base layer capacity 1. This instance has no pure NE. There are three feasible matchings. First, we can match both consumers to m_1 : total cost is 2, so one of the consumers pays at least 1. This consumer has incentive to defect to m_2 to pay 0.75. The other matchings have one consumer matched to m_1 and the other to m_2 . Because the shutdown cost is removed when the m_2 -consumer moves to m_1 , the net cost she imposes by defecting is 0. Thus, the consumer who is matched to m_1 must pay the entire cost imposed on both agents for the assignment to be stable, which is $3 > 2$, making defection to the null producer attractive. Thus, no matching is stable. \square

Theorem 2. *There is a cost sharing scheme that achieves a price of stability of 1 when producer price functions have only tracking layers and capacity constraints.*

Proof. Consider the SW-optimal matching μ and suppose each producer charges each matched consumer the average price per unit times the number of units she consumes. By way of contradiction, suppose consumer n_1 benefits by defecting to m using profile π . Let the matching after defection be μ' . Let $p_{\mu(n_1)}$ and $p'_{\mu(n_1)}$ be the average price per unit for $\mu(n_1)$ before and after n_1 's departure, respectively. When n_1 defects, the average unit cost for demand on $\mu(n_1)$ decreases because some demand that was previously met with the tracking layer may be met with the base layer. Thus $p'_{\mu(n_1)} \leq p_{\mu(n_1)}$. Since n_1 defected, $V(\mu^p(n_1)) - |\mu^p(n_1)|p_{\mu(n_1)} \leq V(\pi) - C_m(\mu'^{-1}(m)) + C_m(\mu^{-1}(m))$. These inequalities can be rearranged to show that social welfare before the defection is less than the social welfare after, which is a contradiction. Note: the argument can be extended to group defec-

Feature	w/ capacity constraints	w/o capacity constraints
Shutdown costs	PoS = ∞	?
Ramp constraints	PoS = ∞	PoS=?, PoA = ∞
Tracking layer	PoS = 1, PoA = ∞	N/A
Base layer only	PoS = 1, PoA = ∞	PoA = 1

Table 1: Table of stability results for combinations of producer price function features under the marginal cost defection model.

tions (to show strong Nash stability), and still applies in the absence of a tracking layer. \square

While the best NE may be arbitrarily worse than the SW-optimal matching, one might hope that the optimal matching is close to a NE in practice. We have found find this is generally *not* the case, but we do not focus on that question in this paper. Since marginal cost defection pricing fails to induce stability, we consider two cost sharing schemes that assume “envy-free” defection pricing, in which a defector is treated no differently than a consumer who was originally matched to that producer. Envy-free defection pricing assumes that while producers are free to make offers to any consumer, they cannot offer a deal to a potential defector that they do not offer to other consumers. This assumption is realistic in a setting with many small consumers.

Shapley-Like Payments.

Since the underlying problem is a cooperative game, one natural approach to cost sharing is to use the *Shapley value* [Shapley, 1953]. We consider the group of consumers matched to a single producer to be a coalition. The Shapley value charges each agent the average marginal cost (or benefit) they contribute to their coalition over all possible *join orders*. Formally, the Shapley value of consumer n_0 matched to producer m_0 under μ is:

$$s(n_0) = \alpha \sum_{S \in \mu^{-1}(m_0) \setminus \{n_0\}} P_{m_0}(\text{dem}_\mu(S \cup \{n_0\})) - P_{m_0}(\text{dem}_\mu(S))$$

where α is a normalization constant (number of permutations) and $\text{dem}_\mu(x)$ is the total demand of the set x of consumers when using the profiles assigned under μ . Our setting is atypical because some join orders induce demand profiles that cannot be feasibly served (e.g., due to ramp or capacity constraints), which is not accounted for in the standard definition of the Shapley value; to deal with this, we impose a large cost when joining a coalition causes infeasibility. Since all costs must be recovered, we normalize the payments so that the total paid by consumers matched to a producer j equals the total charged by j .

Shapley values provide a conceptually simple approach to cost sharing that captures price functions well and is “fair.” However, it is computationally intractable: #P-complete in general [Deng and Papadimitriou, 1994] and difficult to approximate [Fatima *et al.*, 2008]. To overcome this, we sample permutations to approximate Shapley costs. In addition, Shapley payments do not explicitly aim for stability, and indeed, we’ll see they are not inherently stable. Hence, we allow Shapley values to be adjusted $\pm 10\%$ within each generator to increase stability, though even this modification does not admit stability. Ideally, we desire to sacrifice some SW to

improve stability (indeed, find matchings on the Pareto frontier of SW and degree of stability). This is difficult, however, because we can’t efficiently maximize stability: producer price functions are far from concave and do not admit good concave upper bounds.

However, we find that sampling the matchings with high SW allows us to gain a significant amount of stability without losing much SW.⁵ We sample matchings in two ways: through *exclusions* and *cuts*. Both use the well-known linear constraint that precludes a particular assignment of binary variables $\{X_0 = x_0, X_1 = x_1, \dots, X_n = x_n\}$ from being selected by an optimization:

$$\sum_{i:x_i=1} X_i - \sum_{i:x_i=0} X_i \leq \sum_{i \in [n]} x_i - 1$$

The MIP formulation of the SW-maximization contains two types of binary variables: the matching variables $y_{i,j,k}$ that indicate that consumer i is matched to producer j and is using profile k and the support variables of the PPFs, such as $I_{j,t}^{(SC)}$ which indicates whether producer j incurred a shutdown cost at time period t . The exclusions method focuses on matching variables only: each iteration is the standard SW maximization plus constraints that exclude the settings of the matching variables corresponding to the matchings found in previous iterations. The cuts method requires that *both* the matching variables and the support variables are different than the values used in previously found matchings. Note that the behavior of the cuts method is highly dependent on the form of the PPFs and may not be applicable to all PPFs whereas exclusions can be applied in any matching setting. Cuts require a more drastic change to the matching, decreasing SW by a larger amount, but sampling more diverse areas of the matching space. These two methods are compared below.

Similarity-Based Envy-Free Payments.

The standard notion of envy-freeness, that no agent would prefer to receive the outcome that any other agent received, is too weak in our setting. Since demand profiles are real-valued vectors, they are generally unique in that no two consumers share an identical profile. To handle this, we consider a generalization, *similarity-based envy-freeness (SBEF)*, where vectors that are “close” (we use L_2 -distance) are priced identically (on a per-unit basis). Specifically, we use a clustering algorithm to partition the *demand profiles*, and constrain the unit price for any profile in a given partition to be equal. Our experimental model uses 24 1hr. time periods. While we could use demand in each period as the feature-vector for clustering, we instead use higher level features: the average and standard deviation of the demand across all periods; the global maximum and minimum demands, and the gap between them; and the average and standard deviation of demand in 6-hour windows. These high-level features blur the boundaries between partitions, which could be misleadingly granular if demand profiles were used directly. For instance, with demand profiles, we might distinguish two partitions based on consuming more or less than x units from

⁵The same techniques can be used to enumerate matchings when searching for approximate NEs with high social welfare.

1–2PM, which might lead consumers to respond to these specific features (e.g., by shifting some tasks from 1–2PM to 12–1PM). Such specific responses are unlikely to have a large effect on generation cost, especially if other consumers behave likewise. By using abstract features, consumer responses tend to have a greater effect on generation cost.

We consider Ward clustering [Ward Jr, 1963], which groups similar profiles quite well and induces stable payments, while maintaining a high degree of scalability. But it constructs partition boundaries that are difficult to communicate. To address this, we approximate the resulting clusters by building a bounded-depth decision tree (using CART [Breiman *et al.*, 1984]) with easily understandable partition boundaries based on a small number of features, without sacrificing much stability. In general, the choice of partitioning scheme should support the desiderata for price functions. In summary, the procedure we use for calculating SBEF payments is as follows: (1) partition demand profiles using Ward clustering on high-level feature vectors of the demand profiles; (2) approximate the resulting clusters with a decision tree; (3) find a matching that maximizes stability, subject to (i) recovering all costs, and (ii) requiring a fixed unit price for profiles in any given partition. The procedure takes polynomial time and is very fast in practice.

SBEF ensures that consumers are indifferent as to which producer they are matched.

Observation 1. *Assume two or more generators. Suppose we have a matching μ and a set of SBEF payments p . If at least one profile in each partition is assigned to each generator and the maximum incentive to defect is 0, all generators must offer the same unit price in each partition.*

If this were not the case, some consumer would be matched to a more expensive generator given her profile. This consumer would have an incentive to defect that is at least equal to the (positive) difference between the costs of two generators for that profile.

The price differences between “adjacent” partitions will be “reasonable” if there are enough consumers with demand profiles in multiple partitions. When a consumer has demand profiles in two partitions, stability puts pressure on the difference in price between two partitions to be small w.r.t. the difference in their valuations. The SBEF price procedure is somewhat more conceptually complex than Shapley-like payments, but it is much more computationally efficient and it addresses envy-freeness more directly than Shapley. While the Shapley payments within a coalition may be intuitively fair, the assignment of similar profiles to particular producers by SW optimization may be somewhat arbitrary and result in payments that are far from envy-free. We see below that SBEF payments achieve much better stability in practice than Shapley.

4 Experiments

To test our algorithms, we use a model of the US residential energy market. Building characteristics are based on the 2011 Buildings Energy Data Book [D&R International, Ltd., 2012]. The building thermal model, which includes temperature, solar radiation and a miscellaneous factor, is derived




Figure 3: Max and average incentive to defect for maximally-stable matchings using Shapley-like payments. The corresponding social welfare is shown in Fig. 4.




Figure 4: Social welfare for maximally stable matchings using Shapley-like payments.

from [Huang *et al.*, 1999]. Roughly, we independently sample square footage and insulation level from known US distributions. Using appliance surveys, we randomly generate appliances and load events for each appliance. We then calculate air conditioner loads for a variety of target interior temperatures. External conditions are those of July 10, 2010 in San Antonio, Texas: since most home electricity use is due to air conditioning, hot summer days stress generation heavily and induce larger incentives for consumers who are willing to alter their behavior.

In all experiments, we use 50 consumers, 2 producers, 4 profiles per consumer, 24 time periods, and run 50 trials for each experiment. We use a small number of consumers because Shapley values are expensive to compute.⁶ In this setting, the SW-optimal matching has a mean SW of \$837.5 with std. dev. of \$138.8, or around \$16.75 per consumer.

Shapley-Like Payments. Figs. 3 and 4 show results using Shapley payments. Each trial requires about 1hr. of computation, almost all of which is to approximate Shapley values. We sample 30 random join orders, a number which was de-

⁶We are able to optimize SW and find SBEF payments for instances with 2500 consumers, 2 producers, 4 profiles per consumer, 24 time periods in 30 minutes on average. Scalability could be increased further by: i) using a simpler clustering algorithm and ii) compressing the optimization by grouping similar consumers/demand profiles together.

	Avg. 12-6pm	Std. dev.	Std. dev. 12-6pm	Max
Gini importance	0.46	0.32	0.14	0.034

Table 2: Table of stability results for combinations of producer price function features under the marginal cost defection model.

terminated empirically to induce convergence. Initial *maximum (over consumers) incentive to defect (MitD)* is \$17.6 on average (std. dev. \$16.4). After 14 iterations, mean MitD was \$7.8 with cuts (44% of the original) and \$13.9 with exclusions (79% of the original). Cuts decreased MitD faster than exclusions—after two cuts, MitD decreased to \$13.8, on average, a greater reduction than 14 exclusions. (Standard deviations are large because they include the variation among instances.) Using a paired t-test, the difference between MitD using cuts vs. exclusions is statistically significant after 2 iterations ($p < 0.05$).

Since the MitD is primarily influenced by large agents, the average incentive to defect is also shown on Fig. 3. For consumers with a positive incentive, the mean decreased from \$2.81 to \$2.57 with exclusions (91.3% of the original) and to \$1.52 with cuts (54% of the original), correlating with the decreases in MitD, but showing a less dramatic drop when cuts are used. The percentage of agents with positive incentive increased slightly after 14 iterations, from 50% to 52.6% with exclusions and 50.3% with cuts. This appears to be a spurious consequence of the enumeration process.

Exclusions reduced SW by less on average than cuts. The % of max SW under exclusions had a mean of 99.9% after 14 iterations, while cuts had a mean of 98.6%. Since exclusions enumerate every matching, exclusions enumerate those returned by cuts, but they are much slower—the first cut has greater effect than 14 exclusions.

Similarity-Based Envy-Free Payments. Fig. 5 shows the effect of using different numbers of partitions and decision tree depth within SBEF payments on mean MitD. Each trial takes only a few seconds (in stark contrast to Shapley). It is important to note that these results all use the SW optimal matching—since stability is so high, we do not explore the trade-off between stability and SW (though cuts and exclusions could be used here, as above). We see that SBEF payments are highly stable under all tested conditions. Stability increases with the number of partitions: mean MitD is \$1.76 with two partitions (std. dev. \$2.03) and \$0.71 with nine (std. dev. \$1.21). Having more partitions tends to reduce potential “envy-freeness” as fewer consumers are in each. The figure also suggests that a minimum tree depth is needed for the decision-tree approximation to be as stable as the original partitioning: from one level for two partitions, up to four levels for six or more partitions. Mean incentive for customers with a positive defection incentive increases slightly with the number of partitions: \$0.21 with 2 partitions and \$0.24 with 9 partitions. The number of customers with positive defection incentive decreased from 34% with 2 partitions to 24% with 9.

We use Gini importance to assess the feature importance in the partition-approximating decision trees. Tab. 2 shows the




Figure 5: Stability of different numbers of partitions and decision tree depths approximating those partitions.

four most important features across all instances with four partitions and a decision tree of depth 3. Because temperatures peak in the afternoon, the fact that afternoon consumption is the most important determinant of production price makes sense. While this cost sharing scheme resembles time-of-use pricing, the features that affect overall cost change on the fly, dynamically reflecting their impact on generation cost.

5 Conclusion

We have presented a market model for matching electricity producers and consumers, which can be tractably optimized for a large number of consumers. The model allows for consumers to present multiple demand profiles, which allows the matching mechanism to offer discounts to consumers if they are willing to shift demand in a way that reduces production costs. We showed that Nash-stable matchings may not exist in settings with realistic producer price functions and presented two alternate cost sharing schemes, which we tested on synthetic residential energy preference data.

One major question that our system does not address is how to elicit demand profiles from consumers. Direct surveying may be subject to strategic manipulation. While historical data can be used to learn utilities through revealed preference, the static nature of most pricing systems means that it is difficult to learn about behavior outside of standard conditions.

The efficient computational properties of SBEF payments could make it useful in other mechanism design domains. Beyond allowing for approximate envy-freeness in domains that lack that a natural extension of that concept, type-space compression may be useful when there are too many different profiles to reason about efficiently. In this domain, the question of whether compression schemes can be made sensitive to the goals of the mechanism is an interesting one.

Acknowledgments. Perrault was supported by OGS. We gratefully acknowledge the support of NSERC. We thank the anonymous reviewers for their suggestions.

References

- [Anand and Aron, 2003] Krishnan S. Anand and Ravi Aron. Group buying on the web: A comparison of price-discovery mechanisms. *Management Science*, 49(11):1546–1562, 2003.
- [Breiman *et al.*, 1984] Leo Breiman, Jerome Friedman, Charles J. Stone, and Richard A. Olshen. *Classification and Regression Trees*. Chapman and Hall/CRC, 1984.
- [Chen and Roma, 2010] Rachel R. Chen and Paolo Roma. Group buying of competing retailers. *Production and Operations Management*, 20(2):181–197, 2010.
- [Chen *et al.*, 2007] Jian Chen, Xilong Chen, and Xiping Song. Comparison of the group-buying auction and the fixed pricing mechanism. *Decision Support Systems*, 43(2):445–459, 2007.
- [Demange *et al.*, 1986] Gabrielle Demange, David Gale, and Marilda Sotomayer. Multi-item auctions. *Journal of Political Economy*, 94:863–872, 1986.
- [Deng and Papadimitriou, 1994] Xiaotie Deng and Christos H. Papadimitriou. On the complexity of cooperative solution concepts. *Math. Oper. Res.*, 19(2):257–266, 1994.
- [D&R International, Ltd., 2012] D&R International, Ltd. *2011 Buildings Energy Data Book*. Department of Energy, Office of Energy Efficiency and Renewable Energy, 2012.
- [Fatima *et al.*, 2008] Shaheen S. Fatima, Michael Wooldridge, and Nicholas R. Jennings. A linear approximation method for the Shapley value. *Artificial Intelligence*, 172(14):1673–1699, 2008.
- [Gale and Shapley, 1962] D. Gale and L. S. Shapley. College admissions and the stability of marriage. *American Mathematical Monthly*, 69(1):9–15, 1962.
- [Huang *et al.*, 1999] Joe Huang, James Hanford, and Fuqiang Yang. Residential heating and cooling loads component analysis. Technical report, Building Technologies Department, Environmental Energy Technologies Division, Lawrence Berkeley National Laboratory, University of California, 1999.
- [Kirschen and Strbac, 2004] Daniel S Kirschen and Goran Strbac. *Fundamentals of Power System Economics*. John Wiley & Sons, Chichester, UK, 2004.
- [Koutsoupias and Papadimitriou, 1999] Elias Koutsoupias and Christos Papadimitriou. Worst-case equilibria. In *Proceedings of the Sixteenth Symposium on Theoretical Aspects of Computer Science (STACS-99)*, pages 404–413, Trier, Germany, 1999.
- [Lu and Boutilier, 2012] Tyler Lu and Craig Boutilier. Matching models for preference-sensitive group purchasing. In *Proc. of EC’12*, pages 723–740, Valencia, Spain, 2012.
- [Manisterski *et al.*, 2008] Efrat Manisterski, David Sarne, and Sarit Kraus. Enhancing cooperative search with concurrent interactions. *Journal of Artificial Intelligence Research*, 32(1):1–36, 2008.
- [Meir *et al.*, 2014] Reshef Meir, Tyler Lu, Moshe Tennenholtz, and Craig Boutilier. On the value of using group discounts under price competition. *Artificial Intelligence*, 216:163–178, 2014.
- [Nash, 1950] John F. Nash. Equilibrium points in n-person games. *Proc. Natl. Acad. Sci. U.S.A.*, 36(1):48–49, 1950.
- [Rassenti *et al.*, 2003] Stephen J. Rassenti, Vernon L. Smith, and Bart J. Wilson. Controlling market power and price spikes in electricity networks: Demand-side bidding. *Proc. Natl. Acad. Sci. U.S.A.*, 100(5):2998–3003, 2003.
- [Robu *et al.*, 2014] Valentin Robu, Meritxell Vinyals, Alex Rogers, and Nicholas R Jennings. Efficient buyer groups for prediction-of-use electricity tariffs. In *Proc. of AAAI-14*, pages 451–457, Quebec City, 2014.
- [Roth, 2002] Alvin E. Roth. The economist as engineer: Game theory, experimentation, and computation as tools for design economics. *Econometrica*, 70(4):1341–1378, 2002.
- [Sarne and Kraus, 2005] David Sarne and Sarit Kraus. Cooperative exploration in the electronic marketplace. In *Proc. of AAAI-05*, pages 158–163, Pittsburgh, 2005.
- [Schulz and Moses, 2003] Andreas S. Schulz and Nicolás E. Stier Moses. On the performance of user equilibria in traffic networks. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA-03)*, pages 86–87, San Francisco, 2003.
- [Shapley and Shubik, 1971] L. S. Shapley and M. Shubik. The assignment game I: The core. *International Journal of Game Theory*, 1(1):111–130, 1971.
- [Shapley, 1953] Lloyd S Shapley. A value for n-person games. In Harold W. Kuhn and Albert W. Tucker, editors, *Contributions to the Theory of Games II*, pages 307–317. Princeton University Press, Princeton, 1953.
- [Ward Jr, 1963] Joe H. Ward Jr. Hierarchical grouping to optimize an objective function. *Journal of the American Statistical Association*, 58(301):236–244, 1963.