Symbol	Meaning
P_j	Price function. $P_j = BC_j + RC_j + SC_j$
$d_j^{(l)-}$	Minimum economic generation level for the base layer
$d_j^{(l)+}$	Maximum capacity of the base layer
$c_j^{(l)}$	Per unit price for electricity provided through the base layer
r_j	Maximum ramp rate of the base layer
s_j	Shutdown cost that is incurred when base layer generation is below the minimum economic layer
$d_j^{(h)+}$	Maximum capacity of the tracking layer
$c_j^{(h)}$	Per unit price for electricity provided through the tracking layer
BC_j	Base cost function
RC_j	Ramp cost function
SC_j	Shutdown cost function

Figure 1: Table of notation for producer cost function for producer j.

1 Specification of Producer Posted Prices

The problem of deciding the optimal output levels of a group of power generation facilities in order to meet system demand has been studied extensively [1]. If the set of active generation facilities is fixed (i.e., we cannot choose which facilities are active or inactive), the problem is referred to as *economic dispatch*, whereas it is referred to as *unit commitment* if the choice of active generation facilities is part of the optimization. The problem of generation scheduling is typically decoupled of the problem of distribution to some extent; however, distribution constraints may be added to the problem, forcing some generators to be active or require extra power to be generated to account for losses. The generation scheduling problem has several internal constraints:

1. Generation facilities have limited *ramp rate*—the amount of power output change that can be sustained from time period to time period. Ramp rates of different generation facilities vary radically: demand tracking plants such as natural gas can ramp up or down in half an hour, whereas nuclear plants take days.

2. Different kinds of generation have different variable costs—renewables have near zero variable cost and natural gas has quite a high variable cost.

3. For certain kinds of plants (e.g., coal), reducing production below a certain level is expensive—it imposes additional wear on the components. Shutting down these plants also incurs costs.

4. The number of available maintenance crews may be limited, which may limit how many plants can be adjusted at the same time.

5. Plants may have minimum uptime and downtime requirements in order to reduce maintenance costs.

6. Due to the fact that the generation scheduling problem is solved ahead of the realization of demand, it is necessary to keep a reserve of active generation capacity as well as generation capacity that can be quickly taken offline to account for uncertainty in prediction.

We will present a model that captures the first three constraints on this list and abstracts away the others. We have the following model of an electricity producer in mind: they have a *base layer* that has low generation costs, but is slow to adjust and expensive to go below a certain level of generation in any time period and a *tracking layer* that can be adjusted rapidly or shut off entirely, but has high generation costs and limited capacity. For producer j, let $c_j^{(l)}$ be the generation cost of the base layer, let $d_j^{(l)+}$ be the maximum capacity, let $d_j^{(l)-}$ be the minimum generation before additional costs are incurred (the *minimum economic generation level*), and r_j be the maximum ramp rate of the base layer in a single period. Let s_j be the cost of reducing

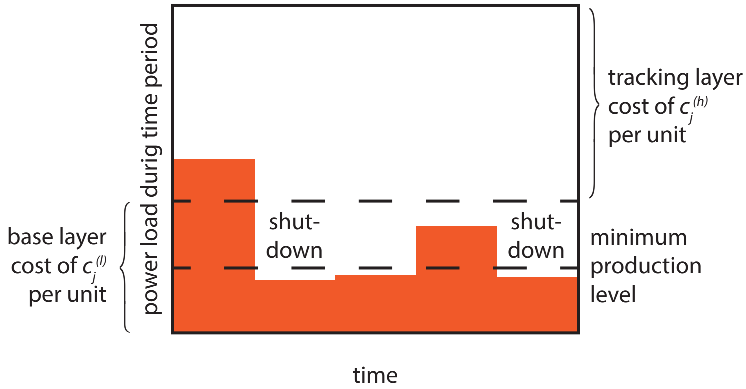

Figure 2: Diagram of the base cost component of the producer price function.

generation on the base layer below the minimum economic generation level (the *shutdown cost*). Let $c_j^{(h)}$ be the generation cost of the tracking layer and $d_j^{(h)+}$ be the maximum capacity of the tracking layer. We will assume that the maximum ramp rate is small compared to the capacity levels, i.e., $r_j < d_j^{(l)-}$, $r_j < d_j^{(l)+}$ and $r_j < d_j^{(h)+}$.

Our form of producer price functions will have the Markov property and will have no additional state beyond the demand in the previous period—the price paid in a particular period only depends on the demand in the current period and the demand in the previous period. We will first describe the form at a high level:

- There is a preferred rate for demand that is easy to serve. If demand is smooth and does not exceed the maximum base layer capacity or fall short of the minimum economic generation level, the price is the base rate only. 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 price is $c_j^{(l)}$ per unit demanded.
- 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 (if there is remaining capacity available on the base layer) at price $c_j^{(l)}$, and the remaining units of the increase are met using the tracking layer (supposing that the size of the increase does not exceed the capacity of the tracking layer) at price $c_j^{(h)}$.
- If there is a large decrease in demand from period to period, we will charge an additional fee $c_j^{(h)} - c_j^{(l)}$ per unit of decrease exceeding $-r_j$, representing the cost of meeting the necessary amount of the previous period's demand using the tracking layer.
- Demand that exceeds the base layer capacity will be met using the tracking layer, if capacity is available, at a price $c_j^{(h)}$.
- If demand in the previous period exceeds the minimum economic generation level of the base layer $d_j^{(l)-}$ and demand in the current period is less than $d_j^{(l)-}$, the demand in the current period is charged a price of $c_j^{(l)}$ plus a fixed additional fee s_j.

Next, we will formalize it. Let $BC_j(x)$ represent the base cost of meeting demand at level x for producer j, assuming no additional ramping costs. Demand is preferentially met using the base layer, and the tracking layer is used only if necessary.

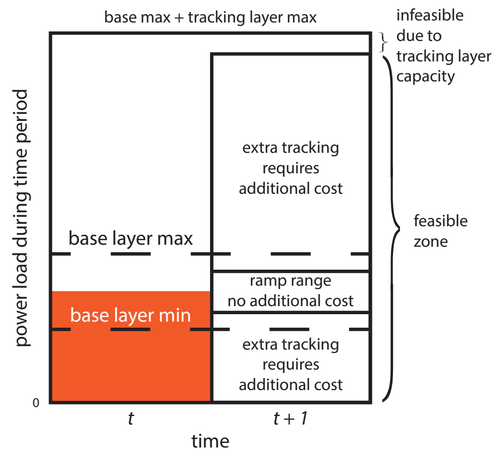

Figure 3: Diagram showing the ramp cost component of the producer price function in the case of an increasing ramp.

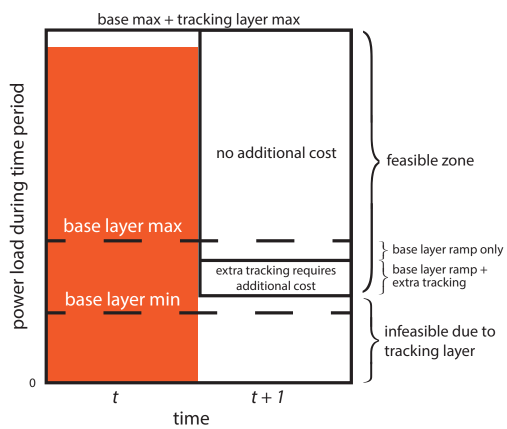

Figure 4: Diagram showing the ramp cost component of the producer price function in the case of a decreasing ramp.

$$BC_j(x) = \begin{cases} c_j^{(l)} x & \text{if } 0 \leq x \leq d_j^{(l)+} \\ c_j^{(l)} d_j^{(l)+} + c_j^{(h)} \left(x - d_j^{(l)+} \right) & \text{if } d_j^{(l)+} < x \leq d_j^{(l)+} + d_j^{(h)+} \\ \infty & \text{if } x > d_j^{(l)+} + d_j^{(h)+} \end{cases} \tag{1}$$

Let $RC_j(x_1, x_2)$ represent the cost of ramping from generation level x_1 to generation level x_2. The ramp cost is zero if both demand levels are on the tracking layer or they differ by no more than the ramp rate of the base layer. Otherwise, the ramp cost is the cost of using the tracking layer to cover the amount of change exceeding the ramp rate by substituting it for generation that would ordinarily be served using the base layer. This substitution may occur in the generation schedule of the current period or that of the previous period. The amount of additional ramp capacity required may not exceed the total unused capacity of the tracking layer in the relevant (current or previous) period.

$$RC_j(x_1,x_2) = \begin{cases} 0 & \text{if } x_1 \in \left(d_j^{(l)+}, d_j^{(l)+} + d_j^{(h)+}\right] \wedge x_2 \in \left(d_j^{(l)+}, d_j^{(l)+} + d_j^{(h)+}\right] \\ 0 & \text{if } x_1 \in \left[0, d_j^{(l)+} + d_j^{(h)+}\right] \wedge x_2 \in \left[0, d_j^{(l)+} + d_j^{(h)+}\right] \wedge \\ & \quad |x_2 - x_1| \leq r_j \\ (c_j^{(h)} - c_j^{(l)})(x_2 - x_1 - r_j) & \text{if } x_1 \in \left[0, d_j^{(l)+}\right] \wedge x_2 \in \left[0, d_j^{(l)+}\right] \wedge x_2 - x_1 - r_j \in \left(0, d_j^{(h)+}\right] \\ (c_j^{(h)} - c_j^{(l)})(x_1 - x_2 - r_j) & \text{if } x_1 \in \left[0, d_j^{(l)+}\right] \wedge x_2 \in \left[0, d_j^{(l)+}\right] \wedge x_1 - x_2 - r_j \in \left(0, d_j^{(h)+}\right] \\ (c_j^{(h)} - c_j^{(l)}) \max\left(d_j^{(l)+} - x_1 - r_j, 0\right) & \text{if } x_1 \in \left[0, d_j^{(l)+}\right] \wedge x_2 \in \left(d_j^{(l)+}, d_j^{(l)+} + d_j^{(h)+}\right] \wedge \\ & \quad x_2 - x_1 - r_j \in \left(0, d_j^{(h)+}\right] \\ (c_j^{(h)} - c_j^{(l)}) \max\left(d_j^{(l)+} - x_2 - r_j, 0\right) & \text{if } x_1 \in \left(d_j^{(l)+}, d_j^{(l)+} + d_j^{(h)+}\right] \wedge x_2 \in \left[0, d_j^{(l)+}\right] \wedge \\ & \quad x_1 - x_2 - r_j \in \left(0, d_j^{(h)+}\right] \\ \infty & \text{otherwise} \end{cases} \tag{2}$$

The last component component required is the shutdown cost $SC_j(x_1, x_2)$. The shutdown cost is charged if the previous demand level is above the minimum economic generation level and the current demand is below it.

$$SC_j(x_1,x_2) = \begin{cases} s_j & \text{if } x_1 \geq d_j^{(l)-} \wedge x_2 < d_j^{(l)-} \\ 0 & \text{otherwise} \end{cases} \tag{3}$$

Note that there are several cases where our approximation will underestimate the cost of generation significantly that we ignore because accounting for them would require introducing additional state into the pricing algorithm. One issue is that in the case of a large ramp downward, the amount of power that is required to be transferred from the base layer to the tracking layer in the previous period could be large enough that it forces the base layer generation to decrease below the minimum economic generation level. Another is that adjusting the generation mix between the layers in the previous period may have an effect on the cost of ramping to that generation mix in the previous period. A third issue is that infeasible ramps may be performed because we are essentially assuming that the base layer can ramp to any level below its maximum capacity within two time periods.

If one were to write an optimization problem to set base and tracking layer generation at each time period subject to ramp and capacity constraints, the value of the optimal solution to this problem would be greater than or equal to the estimation we have defined. This is because every cost incurred in the estimation must be also be incurred by any solution that satisfies the generation constraints.

There are several reasons why we restrict the form of the price function. Allowing arbitrary price functions would require potential consumers to query the producer for a quote, which can cause solving the matching optimization directly to become intractable because the cost function cannot be efficiently represented in the constraints of the optimization. If the generation scheduling problem is itself tractable, it is possible to approximate the optimal matching by finding a fixed point of the generation scheduling and matching optimizations, but this complicates the problem of finding an optimal matching significantly. However, it may be possible in practice to allow the matching problem to be solved initially using our approximate representation of the producer's prices and the price function could be queried directly for final prices. The resulting matching would be approximately stable and approximately optimal depending on how closely our model aligns with the true costs incurred by a producer.

2 Formulation of the Social Welfare Optimization

In this section, we will formulate the social welfare optimization as a mixed integer program (MIP). For clarity of exposition, we will present a MIP form of the approximate producer price function, which we will

Symbol	Meaning
$y_{i,j,k}$	Binary variable indicating whether consumer i is matched to producer j and using profile k
$w_{j,t}^{BC}$	Continuous variable indicating the base layer cost for producer j in time period t
$I_{j,t}^{RC}$	Auxiliary binary variable used in the calculation of the ramp cost
$w_{j,t}^{RC}$	Continuous variable indicating the amount of ramp beyond the ramp rate required
$I_{j,t}^{SC}$	Binary variable indicating whether production is above or below minimum economic generation level
$J_{j,t}^{SC}$	Auxiliary binary variable used in the calculation of the shutdown cost
$\Phi(\cdot)$	Abbreviation used to represent the constraints of a particular cost component

Figure 5: Table of variables for MIP optimization for producer j in time period t.

then integrate into the matching problem. Since the price function is the sum of three non-negative functions on the demand of each consecutive pair of time intervals, the social welfare optimization will attempt to minimize the posted prices—we will use this "pressure" to simplify the required constraints.

The first component function is the minimum cost function $BC_j(x_{j,t})$. This function can be written as:

$$BC_j(x_{j,t}) = c_j^{(l)} \min\left(x_{j,t}, d_j^{(l)+}\right) + c_j^{(h)} \max\left(0, x_{j,t} - d_j^{(l)+}\right) \tag{4}$$

plus the condition that $x_{j,t}$ is less than the capacity of the generation system, $x_{j,t} \leq d_j^{(l)+} + d_j^{(h)+}$. This will be represented in the social welfare optimization as the following linear constraints $\Phi(BC_j, t)$, introducing one continuous variables (we will use I and J for auxiliary variables that are intended to take on binary values, even if they are relaxed to continuous, and w for continuous variables):

$$\begin{aligned} BC_j(x_{j,t}) &= w_{j,t}^{(BC)} \\ w_{j,t}^{(BC)} &\geq c_j^{(l)} x_{j,t} \\ w_{j,t}^{(BC)} &\geq c_j^{(h)} x_{j,t} - d_j^{(l)+}\left(c_j^{(h)} - c_j^{(l)}\right) \\ x_{j,t} &\leq d_j^{(l)+} + d_j^{(h)+} \end{aligned} \tag{5}$$

The second component function is the ramping cost between two periods, $RC_j(x_{j,t}, x_{j,t+1})$. The ramping function can be written in the following way:

$$RC_j(x_{j,t}, x_{j,t+1}) = \left(c_j^{(h)} - c_j^{(l)}\right) \max\left(\min\left(x_{j,t+1}, d_j^{(l)+}\right) - x_{j,t} - r_j, \min\left(x_{j,t}, d_j^{(l)+}\right) - x_{j,t+1} - r_j, 0\right) \tag{6}$$

plus the constraint that $|x_{j,t} - x_{j,t+1}| - r_j \leq d_j^{(h)+}$. To write these in the MIP, we introduce a binary variable $I_{j,t}^{(RC)}$ which is 1 if and only if the base layer of producer j is saturated at time t, i.e., $x_{j,t} \geq d_j^{(l)+}$. We will denote the follow constraints as $\Phi(RC_j, t, t+1)$:

$$\begin{aligned}
RC_j(x_{j,t}, x_{j,t+1}) &= \left(c_j^{(h)} - c_j^{(l)}\right) w_{j,t}^{(RC)} \\
I_{j,t}^{(RC)} &\in \{0,1\} \\
I_{j,t+1}^{(RC)} &\in \{0,1\} \\
w_{j,t}^{(RC)} &\geq x_{j,t} - x_{j,t+1} - r_j - U I_{j,t+1}^{(RC)} \\
w_{j,t}^{(RC)} &\geq d_j^{(l)+} - x_{j,t+1} - r_j - U\left(1 - I_{j,t}^{(RC)}\right) \\
w_{j,t}^{(RC)} &\geq x_{j,t+1} - x_{j,t} - r_j - U I_{j,t+1}^{(RC)} \\
w_{j,t}^{(RC)} &\geq d_j^{(l)+} - x_{j,t} - r_j - U\left(1 - I_{j,t+1}^{(RC)}\right) \\
w_{j,t}^{(RC)} &\geq 0 \\
x_{j,t} &\leq d_j^{(l)+} + U\left(1 - I_{j,t}^{(RC)}\right) \\
d_j^{(l)+} &\leq x_{j,t} + U I_{j,t}^{(RC)} \\
x_{j,t} - x_{j,t+1} - r_j &\leq d_j^{(h)+} \\
x_{j,t+1} - x_{j,t} - r_j &\leq d_j^{(h)+}
\end{aligned} \tag{7}$$

where U is a large constant. The fourth-to-last and third-to-last constraints confirm that the minimum selected by the optimization is the true minimum. This will be needed later, but it is not necessary for the simple social welfare optimization.

The third and final component is the shutdown cost $SC_j(x_{j,t}, x_{j,t+1})$, which can be written as:

$$SC_j(x_{j,t}, x_{j,t+1}) = s_j I\left[x_{j,t} \geq d_j^{(l)-}\right] I\left[x_{j,t+1} < d_j^{(l)-}\right] \tag{8}$$

We charge the shutdown cost only when a producer goes from being above the minimum economic level to below it; after incurring the shutdown cost once, the producer can remain below the economic level for any number of consecutive periods without incurring additional penalties. This will be represented by the following constraints $\Phi(SC_j, t, t+1)$ in the MIP, introducing one binary variable per time period $I_{j,t}^{(SC)}$ which represents whether the load in period t on producer j is below the minimum economic level, i.e., $x_{j,t} \leq d_j^{(l)-}$:

$$\begin{aligned}
SC_j(x_{j,t}, x_{j,t+1}) &= s_j J_{j,t}^{(SC)} \\
I_{j,t}^{(SC)} &\in \{0,1\} \\
I_{j,t+1}^{(SC)} &\in \{0,1\} \\
x_{j,t} - d_j^{(l)-} &\geq -U I_{j,t}^{(SC)} \\
d_j^{(l)-} - x_{j,t} &\geq -U(1 - I_{j,t}^{(SC)}) \\
x_{j,t+1} - d_j^{(l)-} &\geq -U I_{j,t+1}^{(SC)} \\
d_j^{(l)-} - x_{j,t+1} &\geq -U(1 - I_{j,t+1}^{(SC)}) \\
J_{j,t}^{(SC)} &\geq I_{j,t+1}^{(SC)} - I_{j,t}^{(SC)} \\
J_{j,t}^{(SC)} &\geq 0
\end{aligned} \tag{9}$$

With the components assembled, we will combine them into a single optimization problem. Our decision variables will be the binary matching variables $y_{i,j,k}$, which is true if and only if household i is matched to producer j and household i is using demand profile $\boldsymbol{\pi}_i^{(k)}$.

$$\max_{\boldsymbol{y},\boldsymbol{w},\boldsymbol{I}} \sum_{i\in N}\sum_{k\in\Pi_i}\sum_{j\in M} y_{i,j,k} V_i\left(\boldsymbol{\pi}_i^{(k)}\right) - \sum_{j\in M}\sum_{t\in[T]} BC_j(x_{j,t}) - \sum_{j\in M}\sum_{t\in[T-1]} (RC_j(x_{j,t},x_{j,t+1}) + SC_j(x_{j,t},x_{j,t+1})) - \sum_{j\in M} s_j(1-I_{j,0,2}) \tag{10}$$

subject to the following constraints. Note that we define several abbreviations: $x_{j,t}$ represents the total demand assigned to producer j in period t and the function symbols BC_j, RC_j and SC_j, which are defined as before. These abbreviations can be eliminated by substituting the expression for each, which is specified as a constraint, in each place where they occur. The Φ symbols represent sets of constraints as defined earlier in this section.

$$\begin{aligned}
& y_{i,j,k} \in \{0,1\} && \forall i\in N, \forall j \in M, \forall k\in \Pi_i \\
& \sum_{j\in M}\sum_{k\in\Pi_i} y_{i,j,k} = 1 && \forall i \in N \\
& x_{j,t} = \sum_{i\in N}\sum_{k\in\Pi_i} l_{i,j}\pi_{i,t}^{(k)} I_{i,j,k} && \forall j\in M, \forall t\in[T] \\
& \Phi(BC_j, t) && \forall j\in M, \forall t\in[T] \\
& \Phi(RC_j, t, t+1) && \forall j\in M, \forall t\in[T-1] \\
& \Phi(SC_j, t, t+1) && \forall j\in M, \forall t\in[T-1]
\end{aligned} \tag{11}$$

where $l_{i,j}$ is an arbitrary scaling factor that represents producers preferences for consumers. In our case, we will interpret it as represent the amount of power loss incurred when power is delivered from producer j to consumer i, which we will assume is proportional to the distance between them.

3 Cost Sharing and Stability Concepts

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

Proof. Consider an instance with two consumers, two producers and a single time period. Suppose each consumer has a demand of 1, value of 2 for their demand and outside option of value 0. Suppose that producer m_1 has a minimum economic generation level of 1.5, base layer cost of 1, shutdown cost of 1 and a large base layer capacity. Producer m_2 has no minimum economic generation level, base layer cost of 0.75 and base layer capacity of 1. This instance has no Nash equilibria. There are three feasible matchings. First, we can match both consumers to producer m_1. In this case, since the total cost is 2, one of the consumers is paying ≥ 1. This consumer has incentive to defect to m_2 to pay 0.75. The other feasible matchings have one consumer matched to m_1 and the other to m_2. Because the shutdown cost is removed when the consumer that is matched to m_2 moves to m_1, the net cost that customer imposes by defecting is 0. Thus, the consumer who is matched to m_1 would have to pay the entire cost imposed on both agents for the assignment to be stable, which is $3 > 2$, which makes defecting to the null producer attractive. Thus, none of the matchings are stable. □

Theorem 2. *There is no cost sharing scheme that achieves a price of stability better than ∞ when producer price functions may have ramp constraints and capacity constraints.*

Proof. Consider a problem instance with two consumers, two producers and two time periods. Suppose consumer n_1 has demand vector $(1, \epsilon)$ and consumer n_2 has demand vector $(\epsilon, 1)$. Producer m_1 has ramp rate 0 and base layer cost 1. Producer m_2 has ramp rate 1, base layer cost 0.75 and base layer capacity 1. The only feasible configuration is to match both agents to m_1, but this is not stable because the consumer that is paying ≥ 1 will defect to m_2. Thus, no NE exists. □

Theorem 3. *There is no cost sharing scheme that achieves a price of anarchy better than* ∞ *when producer price functions may have ramp constraints.*

Proof. Suppose we have a two period model with two producers with no ramp capacity, i.e., the only feasible matchings have an identical amount of load on a particular producer in both periods. Suppose that there exists a feasible matching with SW > 0. If there does not exist a consumer with identical loads in both periods, the matching where every consumer is matched to the null producer is also stable. In the case where the outside options are 0, the price of anarchy is ∞. □

Theorem 4. *There is a cost sharing scheme such that achieves a price of stability of 1 when producer price functions may tracking layers and capacity constraints.*

Proof. Consider the social welfare-optimal matching and suppose that each producer charges each consumer that is matched to them the average price per unit times the number of units they consume. For purposes of contradiction, suppose consumer n_1 can benefit by defecting from the producer and profile it is matched to under μ to producer m using profile π. Let the matching after defection be μ'. Let p_1 and p_1' be the average price per unit for producer 1 before and after n_1's departure, respectively. When n_1 defects, the average unit cost for demand on $\mu(n_1)$ will decrease because some demand that was previously met with the tracking layer may be met with the base layer. Thus $p_1' \leq p_1$. Since n_1 defected, $V(\mu^p(n_1)) - |c|p_1 \leq V(\pi') - C_m(\mu'^{-1}(m)) + C_m(\mu^{-1}(m))$. Using these inequalities, the social welfare before the defection can be shown to be less than the social welfare after, which is a contradiction.

$$\sum_{n\in\mu^{-1}(m_1)} (V(\pi_n) - |\pi_n|p_1) + V(\pi) - |c|p_1 \leq \sum_{n\in\mu^{-1}(m_1)} (V(\pi_n) - |\pi_n|p_1') + V(\pi') - C_{m_2}(\mu'^{-1}(m_2)) + C_{m_2}(\mu^{-1}(m_2))$$

Since the social welfare of consumers matched to m_2 is unaffected, this inequality shows that social welfare increases as a result of n_1's defection, which contradicts the fact that we assumed μ was social welfare optimal. Note that the same logic can be extended to defections of groups of consumers (Strong Nash stability). The same logic still holds without the presence of a tracking layer, where n_1's departures and arrivals have no effect on the average price. □

Theorem 5. *There is no cost sharing scheme that achieves a price of anarchy better than* ∞ *when producer price functions may have capacity constraints.*

Proof. Consider an instance with two producers, two consumers and a single time period. Producer m_1 has base layer cost 1.5 and capacity 1 and producer m_2 has base layer cost 0 and capacity 1. Suppose consumer n_1 has demand $\epsilon << 1$, valuation ϵ and outside option 0 and consumer n_2 has demand 1, valuation 1 and outside option 0. The social welfare optimum matching μ^* is $\mu^*(n_1) = m_1$ and $\mu^*(n_2) = m_2$ which has $SW(\mu^*) = 1 + \epsilon - 1.5\epsilon = 1 - 0.5\epsilon$. Consider the matching $\mu(n_1) = m_2$ and n_2 unmatched, and where n_1 pays the entire cost of ϵ^2. $SW(\mu) = \epsilon - \epsilon^2$. n_1 does not want to defect to m_1 because she would incur a cost of $1.5\epsilon > \epsilon^2$, nor does she want to be unmatched as $\epsilon - \epsilon^2 \geq 0$. n_2 does not want to defect to m_1 because her net change in utility would be $1 - 1.5 < 0$. n_2 cannot defect to m_2 because m_2 does not have enough capacity. Thus μ is stable. $SW(\mu^*)/SW(\mu) = \frac{1-0.5\epsilon}{\epsilon-\epsilon^2}$. As ϵ approaches 0, this ratio approaches ∞. □

Theorem 6. *There exists a cost sharing scheme that achieves a price of anarchy of* 1 *when producer price functions have a base layer only with no capacity constraints.*

Proof. Consider the cost sharing scheme where each consumer pays the number of units consumed times the price of the producer they are matched to. By using this scheme in this setting, the prices paid are not affected by the behavior of other consumers. Thus, the SW optimum is obtained by maximizing the utility of each consumer individually. Note that by using this scheme in this setting, the worst-case defection model and the envy-free defection models are equivalent. Also, the prices paid under the defection model are the same as those that would have been paid if the consumer had originally been matched to the producer they were defecting to. The SW optimum is stable because switching profiles and producers cannot increase the net utility of any consumer since this quantity is maximized in the SW optimum. Any other matching is not stable because this quantity can be increased. Thus, the price of anarchy is 1. □

4 Model of Consumer Demand

The general framework for the model is a simulation with T periods and a population of agents that make a decision for each period. Agents differ primarily in how much power they demand and how flexible those demands are. They have many common features that have to do with the building they inhabit. Each agent type has the following features that operate similarly between the types, although the distribution for each may depend on the agent type, i.e., factories are typically larger than houses.

1. *Building floorspace* (square feet). The floorspace of a building is the main metric of size that is used in surveys. We will use it to determine the general level of electricity consumption of a building as well as to extrapolate the volume and surface area of a building for heating and cooling purposes.

2. *Building volume* (cubic feet). Volume is one important component of determining heating and cooling costs for buildings. We will calculate the building volume by assuming that the ceilings are 8 ft high.

3. *Heater and air conditioner properties.*

   (a) *Heater type* (gas or electric). If heater is electric,

      i. *Heater efficiency* (Btu per Watt-hour). Typical efficiency is 3.413 Btu per Watt-hour since we are only concerned with electric heaters. We assume that heaters have infinite output power.

   (b) *Air conditioner efficiency* (Btu per Watt-hour). Air condition efficiency at a given moment depends on outside temperature, but for simplicity, we assume that it is fixed. Typical air conditioner efficiency ratings are averages across a variety of conditions. The typical range for these average efficiency ratings is 10–14 Btu per Watt-hour. We assume that air conditioners have infinite output power.

4. *Window surface area* (square feet) SA_{window}.

5. *Wall surface area* (square feet) SA_{wall}.

6. *Roof surface area* (square feet) SA_{roof}.

7. *Window shading coefficient* (dimensionless) WSC. This number represents the fraction of solar radiative heat transmitted by a window; range 0.57–0.74.

8. Insulation levels ($\text{ft}^2 \cdot {}^\circ\text{F} \cdot \text{hr/Btu}$). These represent the amount of heat per unit surface area carried through a particular medium per degree of difference between interior and exterior temperatures.

   (a) *Window insulation level* R_{window}. The range of insulation values will be roughly 1.3–1.7.

   (b) *Wall insulation level* R_{wall}. The range of insulation values will be roughly 2.5–7.

   (c) *Roof insulation level* R_{roof}. The range of insulation values will be roughly 9–14.

9. *Temperature preference surface* (input units are time and degrees Farenheit, and output units are dollars). The temperature preference curve represents the cost (in dollars) of maintaining a particular temperature at a particular time; this cost may represent discomfort and/or lost revenue depending on the type of agents involved.

10. *Lighting power density* (Watts per square foot). Heating generated by lighting is an important component of heating and cooling costs. Typical power density for lights is between 1–2.5 Watts per square foot. We approximate that 75% of this power is emitted as heat.

External Conditions. The population of agents shares a common external context consisting of several factors.

1. *Temperature* (units time to degrees Farenheit).

2. *Solar radiation* (units time to Btu per square foot per second) $SR(t)$.

Heat Transfer Calculation. Given these numbers, we can calculate the overall heat transfer. For conductive transfer, the instantaneous heat transfer is given by:

$$1.75\Delta T\left(\frac{SA_{window}}{R_{window}}+\frac{SA_{wall}}{R_{wall}}+\frac{SA_{roof}}{R_{roof}}\right) \tag{12}$$

where ΔT is difference between the interior and exterior building temperatures. Note that 1.75 is a corrective factor to make up for heat transfer not modeled (i.e., ventilation, foundation, and infiltration).

For radiative transfer, it is:

$$0.33 * SR(t) * SA_{window} * WSC \tag{13}$$

where 0.33 is a guess about the amount of time a particular window is exposed to the sun each day.

References

[1] D. S. Kirschen and G. Strbac. *Fundamentals of Power System Economics.* John Wiley & Sons, Chichester, UK, 2004.