On Coordinating Electricity Markets: Smart Power Scheduling for Demand Side Management and Economic Dispatch

A thesis presented

by

Ashin D. Shah

to

Applied Mathematics

in partial fulfillment of the honors requirements
for the degree of
Bachelor of Arts
Harvard University
Cambridge, Massachusetts

March 30, 2012

On Coordinating Electricity Markets: Smart Power Scheduling for Demand Side Management and Economic Dispatch

Abstract

Information asymmetry in retail electricity markets is one of the largest sources of inefficiency in electricity production. In this thesis, we analyze the structure of and interaction between wholesale and retail electricity markets. We model these markets as optimization programs and propose a two-way system of communicating real-time prices and hourly demand between electricity producers and households to improve production outcomes. We formulate the interactions between retail market participants as a game using formal theory and prove the existence of Nash equilibrium in the market under complete information. Under a system of minimal information transfer between the producer and households, we structure a mechanism to approximate equilibrium prices and quantities for use as a real-time pricing scheme. Our goal is to maximize social welfare through implementing scheduling algorithms for producers and households that coordinate optimal production and household appliance use in the market with hidden information.

Our work incorporates ideas from market design, linear optimization, congestion game theory, and mechanism design to design a scheduling mechanism to improve production efficiency in the retail electricity market.

Acknowledgments

I would first like to thank my advisor Yiling Chen for her time, support, and guidance throughout this project. Under your instruction, a full year ago, I first learned the techniques of linear programming that have brought to life the models in this thesis, and I hope my work a year later has made you proud.

I also would like to extend my gratitude to my readers, Doug Fearing and Ozlem Ergun, for their help in crafting my early ideas and for their time spent in review of my final product. I specifically would like to thank Ozlem for introducing me to network flow optimization this past fall in her class, without which much of my thesis would not have been possible. I also extend my thanks to Michael Brenner, Margo Levine, and Mauricio Santillana for each of their time and help over the many years for advising and in support of the senior thesis writers.

This thesis would also not have been possible without the invaluable help of a very good friend, Rick Goldstein. I thank you for your advice and time in the last several months and over four years as we sat together through the many classes encompassing the ideas utilized in this thesis. I also thank Priya for keeping me calm and focused through the many months spanning my relationship with this project. A great deal of gratitude also goes out to my friends for their encouragement, empathy, and motivation.

And at last, but never least, I thank my parents, Avani and Dhiren, for their continued love and support in all of my endeavors. I am here, in every sense of the phrase, because of you.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

In 2003, the National Academy of Engineering (NAE) compiled a ranked list of the greatest engineering achievements of the past century [4]. Among the engineering marvels on this list, nuclear technologies ranked nineteenth, with the Internet ranked at thirteenth. Ahead of even "highways" at eleventh and the automobile at second place, "electrification" as made possible by the power grid was ranked by the NAE as "the most significant engineering achievement of the 20th century." In the United States, today it is an interconnected system of over 18,000 electric utility generators spanning three distinct grids. The combined U.S. power grid serves an online capacity of 1,000,000 megawatts to over one hundred million residential, commercial, and industrial customers, with a net generation of 4,125,060 thousand megawatt-hours in 2010. A staggering feat of engineering, many have called our power grid, "the largest interconnected machine on Earth."

Nonetheless for the past three decades, the power grid has largely seen advancement only in its size. Since the 1980s, spending on research and development among electric utilities has been among the lowest of any industry, at a mere 2% of total revenue. And so, as the power grid has sprawled in scale, struggling to keep up with the rate of growth in demand, what this has produced today is a rapidly aging, overburdened system tasked with enduring the demand challenges of tomorrow. As even the Department of Energy remarks, "the irony is profound: In a society where technology reigns supreme, America is relying on a centrally planned and controlled infrastructure created largely before the age of microprocessors" [5]. Today, it has become a system fraught with inefficiency and overproduction.

In recent years, climate change concerns, rising fuel prices, and renewed political support have since revitalized attention and investment in "cleantech" and renewable energies. The success of these technologies nonetheless depends on the implementation of more fundamental innovations governing the power grid as a whole. Reengineering the aged power grid of the 20th century to face the challenges of the 21st and beyond may very well constitute the greatest engineering achievement of tomorrow.

1.1 Background

As a market, electricity generation is the epitome of the "just-in-time" production strategy, a concept codified into management science by Sakichi Toyoda and the Toyota Production System. In the late 1940s, Toyota Motor Corporation began orienting production to incorporate the idea that inventory was waste, instead making parts only in the needed amounts at the right times. Until large-scale energy storage becomes an economically-sound and scalable reality, electricity can only be consumed at the very moment it is produced. Excess production in one time period cannot be transferred to meet demand in another. Mismatch between supply and demand in the market constitutes one of the worst structural inefficiencies to the electric power system. Unlike an auto original equipment manufacturer (OEM), for an electric utility the carrying cost of excess electricity held as inventory is effectively infinite.

The production side of the electricity market is organized into utilities responsible for the generation, transmission, and distribution of electricity across a physically-connected subregion of the power grid. Dispatch centers coordinate the flow of electricity through these various stages in production in order to match supply with demand. The interconnectedness of the power system allows generating utilities to themselves function as a wholesale market, separate from the end-use retail market for electricity. This wholesale market for electricity generation executes instantaneous auctions on the hour (or some fixed time interval standardized across the specific submarket) that sets a wholesale price of electricity, determines which generators will produce in the coming hour, and at what levels.

However, under most conventional retail pricing contracts, electricity distributors typically charge end-use customers a fixed retail block pricing (called **time-of-use (TOU) pricing**) that may not be adjusted for many months. Different rates are set ahead of time for each time *block*: peak, shoulder, and off-peak, which are roughly analogous to evening, day-time, and night-time hours, and prices are adjusted relatively infrequently. Because distributors fix retail pricing ahead of time, though these prices incorporate the general costs of production, they are independent of the instantaneous wholesale market price set at each hour.

TOU pricing, for its relatively simple design, has been the most common industry practice for decades, though it has received much criticism from economists because it fails to reflect the hour-by-hour variation in the true underlying wholesale cost of electricity. Generators face both price risk, largely from fuel cost fluctuations, and volume risk in production, and additionally, the two risks are heavily correlated. Once again, because electricity markets depend on "just-in-time" production, supply must always orient itself to meet demand. Nonetheless, whereas production may lag demand orders in the automobile market, delivery of electricity must occur simultaneously with the realization of demand. And so, this demand realization occurs ex-post to earlier production commitment decisions like which generators to run. Because there is virtually zero short-term price-responsiveness in demand under TOU pricing, the supply side has to rely on forecasting models based on historical consumption in order to predict future demand. The nature of electricity generation technology forces suppliers to fix production decisions ahead of time as fixed costs and supply frictions prevent generator capacity from being matched to demand perfectly.

Structurally, information asymmetry between supply and demand in the market leads to overproduction and a deviation from the socially efficient level of electricity generation. As will be discussed in Chapter 2, backup generators or *peaking power plants* are the most marginally expensive plants to operate. Peaking power plants are named as so because they are designed to produce only at periods with peak demand, but must be keep at standby to provide slack production capacity to the grid. Peaker plants produce electricity at typically lower efficiencies than *base load power plants*, the latter of which are designed to be always on, continuously providing base load supply at the cheapest marginal rates. Given uncertainty in demand, the supply side must maintain sufficient production capacity in backup or reserve power to hedge the risk of blackout or grid failure. It is worth distinguishing between the system costs of generating electricity for transmission and the costs imposed in order to maintain a requisite level of online, but non-producing capacity. The latter stems from fixed costs of operation and maintenance plus depreciation and the opportunity cost of capital investment in expanding marginal capacity [2]. Inefficiency arises because the value generated from additional power produced at this incremental marginal capacity is less than the cost of the capacity given that it is utilized only during peak demand periods. This off-peak marginal value, or shadow price, of additional capacity is zero.

Beyond the inherent negative externalities to the environment, over-investment in generation capacity is economically inefficient and a misallocation of capital. As electricity use continues to grow, the U.S. power grid will eventually struggle in scaling its current level of production to continue to meet demand unless these inefficiencies are removed from the system and generation streamlined.

1.2 Motivation

Inefficiencies in the electricity market stem primarily from the supply side's inability to reliably forecast demand. In recent years, researchers have almost unanimously agreed that the demand side of the market must play a more active role [2]. Already, a number of demand-side participation programs have been implemented, all of which aim to provide economic incentives to end-users in order to help balance supply and demand, but they have had varying degrees of success. Borenstein *et al.* [2] analyzes a few of these programs, including interruptible contracts and paying large industrial customers to reduce their demand as counterweight during peak demand periods.

Interruptible demand contracts are agreements made by the system to cut power to large end-use customers at times when the grid approaches collapse. Likewise, rolling blackouts are a crude and last resort effort to curtail power demand to prevent a total blackout. Whether voluntary or involuntary, outages are estimated by the Department of Energy to cost American businesses more than $100 billion each year [5]. In voluntary contracts, subscribers usually receive some fixed payment for compliance when a program is in effect, though subscribers often reserve the right to continue use at a much higher marginal rate.

From an economic perspective, demand reduction programs first face a problem of *adverse selection*: if such a program is voluntary, it will be joined by a disproportionate number of customers that already know they intend to lower consumption. As [2] argues, if a program uses the past year's consumption as a baseline, companies that have shrunk in the past year will be the first to sign up. Much of the reduction in demand from this program would have thus occurred anyway and payments from the program will have a much smaller effect than expected. Furthermore, demand reduction programs also introduce a *moral hazard* problem: if baseline consumption is used to measure conservation, the program discourages demand reduction during non-peak periods in which payments are not in effect.

The most agreed upon solutions in the literature advocate **real-time pricing (RTP)** schemes for the retail market. Under RTP, these end-use prices would reflect the variations incurred in wholesale prices, which by nature of their auction design, in turn reflect the true costs of producing electricity. With time-varying retail prices, customers are incentivized to consume less at peak times due to higher prices, reducing the need for expenditure on excess capacity.

As motivation for using prices to incentivize consumption behavior, we consider the implication of changes in price on demand for electricity: higher prices at peak times lead customers to either shift peak consumption to off-peak hours (where prices are relatively cheaper) or reduce consumption more generally. These two behavioral changes correspond to a *substitution* and an *income effect*, respectively, and can be modeled in terms of the uncompensated Marshallian, $x_t(p, Y)$, and compensated Hicksian, $h_t(p, u)$, demands at wealth Y and utility u for electricity consumption in period t:

$$\underbrace{\frac{\partial x_t(p,Y)}{\partial p_t}}_{\text{net price-responsiveness} < 0} = \underbrace{\frac{\partial h_t(p,u)}{\partial p_t}}_{\text{substitution effect} < 0} - \underbrace{\frac{\partial x_t(p,Y)}{\partial Y} \cdot x_t(p,Y)}_{\text{income effect} > 0}$$

Because the price p_t has increased, the individual is poorer in real terms of wealth and his budget constraint, and therefore experiences an income effect. Electricity is likely a normal good as individuals consume more as they become wealthier. Regardless of the actual magnitudes of either effect, the net effect of imposing real-time prices is to unambiguously reduce consumption in the peak period.

The move in the literature and practice toward market-based mechanisms of price-responsiveness (*dynamic* pricing) and demand curtailment is a significant paradigm shift from almost 75 years of electricity planning that has assumed peak demand to be price-invariant. As [2] argues, devising a system that permits adjustments on both the supply and demand side will improve efficiency, reduce costs, and reduce environmental harm. The increasing availability of smart meter technology at prices affordable for grid-wide installation proves promising for the implementation of a dynamic demand-response mechanism of this type. However, our work presented in this thesis is motivated by recent literature on real-time pricing that investigates the impact of RTP on grid stability and volatility. Researchers at MIT in 2011 showed that exposing retail customers to real-time wholesale market prices creates a closed-loop feedback system that is likely to increase volatility and even destabilize the system [11]. Roozbehani *et al.* proposes that the design of any real-time pricing mechanism must take system stability issues into consideration.

Beyond merely providing retail customers with pricing information, we instead propose a mechanism that exploits the feedback loop between wholesale supply and retail demand discussed in [11]. We consider the dynamical system formed by these two groups of market participants across the wholesale and retail market partial equilibria and propose a method for communicating expected demand schedules from consumers to the producer. As discussed, the problem of information asymmetry is at the core of electricity market inefficiency. Through this supply-demand feedback, we structure a mechanism in which each side of the market "sells" a relevant piece of its hidden information to the other. As in dynamic pricing, the supply side continues to provide spot prices for the day ahead to retail customers, but in order to balance out the possible destabilizing effect from an increase in price-elasticity of demand in the market, we require the demand side to communicate its day-ahead predicted load profile to the producer as payment for these prices.

Before the start of each day, the two sides of the market iteratively communicate load and price information back and forth until our mechanism converges to a set of prices for the day ahead that correspond to a stable market equilibrium. At each computational iteration, our demand component updates its predicted day-ahead load profile given the updated set of day-ahead prices it has received, before communicating this updated load profile back to the producer. Likewise, the producer updates prices based on the demand information received at each iteration in this computation. For every iteration in this mechanism until a stable solution is found, the producer and consumers continue to update and propose hourly schedules for generation and household appliance use, respectively, both seeking to minimize their respective costs. If an equilibrium is found, the producer can set prices accordingly such that households will choose to consume according to the corresponding equilibrium quantities.

Our model draws inspiration for scheduling household appliances from the promised increase in electricity demand due in part to the electrification of the transportation market. Plug-in hybrid electric vehicles (PHEVs) require 0.2-0.3 kWh of charging power per single mile of driving. Researchers measure that during charging time, PHEVs approximately double average household load [10], and if they increasingly become popular, will impose a significant strain on the current system. As overall household load increases, a more systematic approach to load management will become increasingly necessary and so we turn to scheduling appliances to reduce peak demand. To further motivate this approach, we consider the likely increase over time in *schedulable* appliances as opposed to *real-time* appliances, ones that consume power as they desire and cannot be rescheduled [16]. Whereas it is possible overall household appliance use will continue to increase into the future as it has trended in the past, it is unlikely this increase will continue from the addition of appliances that require real-time operation, like a hairdryer or stovetop. More realistically, sizable increases in household electricity load in the future will result from the introduction of appliances and household jobs that can be scheduled ahead of time with minimal disruption to consumers, like recharging PHEVs. We consider the growing popularity of mobile or "on-the-go" devices that run on energy stored from the grid but may not actively rely on grid energy at time of use. The development of cost-effective and scalable electricity storage furthermore provides motivation for our work. This changing nature of household electricity demand implies that there will be a greater fraction of jobs that can be shifted or scheduled from peak to off-peak hours in the future, in order to reduce *peak-to-average ratio* (PAR) in consumption and overall system costs during peak hours. Additionally, as determined earlier, by reducing the variance in consumption and smoothing demand, we reduce the need on the production side to invest in costly excess slack capacity for generation that is otherwise underutilized for most of the day.

The objective to our mechanism is to achieve a stable and efficient market equilibrium in which neither side has an incentive to deviate from its proposed prices or load profile. We seek to prove that our iterative mechanism converges to a Nash equilibrium for production and households, first proving whether such an equilibrium exists. If the mechanism converges to a Nash equilibrium, then given others' participation in the mechanism, no agent or the production side will strictly benefit from deviating from the mechanism's result. Furthermore, we seek to converge to this result with minimal information transfer, so as not to impose any undue burden on either the producer or, in particular, consumers. We also desire the property that participation in the program leaves consumers at least as well off as by not participating, though individuals reserve the right to opt-out (*individual rationality*).

Our mechanism provides a framework for the transfer of hidden information in order to facilitate convergence to market equilibrium. Our primary goal then is to provide a quantification of the "value" of information in closed-loop electricity markets. By enabling a transfer of hidden information between supply and demand, are we able to compute equilibrium and construct prices that will allow for more efficient production?

1.3 Summary of Contributions

In this thesis, we introduce and develop a pricing mechanism to compute and set retail prices at equilibrium values in order to coordinate supply and demand in the market for electricity generation. Our model proposes a program in which the production side sells its day-ahead prices to households in exchange for each household's proposed load schedule as implicit payment. The transfer of this information allows each side of the market, the producer and consumers, to better optimize scheduling of electricity generation or appliance use in the home, respectively. We construct a dynamic mechanism for computing equilibrium pricing in which the supply-side problem of economic dispatch and a centralized program coordinating the demand-side household problems of minimum cost appliance scheduling are iteratively solved for an endogenous set of prices and production quantities. By considering the substitutability of household energy consumption across time periods in a day, we intend to prove that a market equilibrium exists and that we may converge to this result through the computation outlined in our proposed mechanism.

The first contribution made here is a modeling contribution that departs from the current literature by formulating both the supply-side and demand-side optimization problems using network flows and with endogenous quantities for price and hourly demand. We present the complete model in Chapter 4 and motivate our need to formulate these programs to be as computationally efficient as possible. Our formulation emphasizes the need to be able to scale these problems in practice, in order to solve for prices and demand schedules for thousands of households. The formulation proposed is a departure from the current literature that relies on integer programs and allows our proposed programs to be efficiently solved thousands of times iteratively as needed in our computations of equilibrium. Network flow formulations are proposed because there exist well-known polynomial time algorithms to find the optimal solution to the circulation network flow problem. We empirically test our models with a commercial solver for mathematical programs, IBM ILOG CPLEX Optimization Studio (CPLEX), and simulate our proposed iterative computations to determine whether they converge. Our formulation of both sides of the problem using network flows is a strong contribution to modeling the market for electricity generation in a way that is scalable and still computationally efficient. It is furthermore a departure from the current literature to formulate both sides of the market as problems with endogenous variables for price and quantity. While in Chapter 2 we construct formulations for both optimization problems as standalone programs with exogenous parameters for demand and prices, in order to coordinate the wholesale and retail markets for electricity, we require a complete model for the market to jointly solve for market equilibrium.

In our complete model, we solve the unit commitment problem on the part of suppliers in order to determine which generators to run and at which capacities, a production scheduling problem. Our program coordinates the scheduling of household appliances and jobs for consumers, optimizing a minimum-cost solution constrained by iso-utility bounds set by household users for job release and deadline times. Our formulation of both sides of the market is then unique in the current literature in which models have typically emphasized either the production aspect of the market or the household problem, but no work has been done as extensively to coordinate the two problems as we do. Additionally, beyond merely modeling aggregate production or consumption levels, our model solves for optimal hourly schedules for both the producer and households given prices and demand at each iteration. In practice, our mechanism and programs, as design contributions, are intended to be implementable by producers and households as currently formulated.

As another original contribution in modeling this market, we have proposed a centralized program for solving the demand-side problem by approaching it from the perspective of a benevolent social planner who seeks to minimize the aggregate social costs of production. Our formulation again requires us to solve only two programs, one for the demand side and one for supply regardless of the number of households and market participants. Our centralized optimization program allows us to feasibly determine optimal solutions for every household simultaneously by solving just this system of programs. In a subsequent chapter, we also argue that any optimal solution found through this centralized optimization program must also be individually optimal for each household given that each household in the program is modeled as its own independent subnetwork. There exist no coupling constraints between households that would otherwise constrain a solution to the centralized program. This property allows us to break larger problems of thousands of households into smaller decomposed subsets of households in order to reduce computational latencies. This method of scaling our model is another original contribution to the literature on electricity markets and on modeling large-scale optimization problems.

In this thesis, in order to motivate coordination between supply and demand, we propose a two-way communication protocol for relaying information between these two sides of the market to compute equilibrium prices. We argue that the feedback loop established by this iterative mechanism will, in the limit, converge to a stable equilibrium balancing supply and demand. By modeling the energy consumption game played between households and production, we first confirm that a stable equilibrium, defined as a vector of hourly prices balancing production and consumption, exists for the market and define under which assumptions. We furthermore characterize the nature of this equilibrium in terms of the cost structure of producers, providing an intuition relevant toward computing equilibrium. In subsequent chapters, we propose a number of algorithms enabling us to solve for equilibrium prices. We demonstrate that by myopically allowing households to all best respond simultaneously to changes in price induces a congestion game in which prices fail to reflect the dynamic impact of other households' consumption changes. We discover that the mechanism fails to converge when every household is able to respond simultaneously. We provide a formal game theory proof for why this mechanism fails and confirm our result empirically. We show that under myopic best response, we produce oscillations in price, quantity supplied, and quantity demanded without convergence.

The result we obtain here then motivates an application of **congestion games**, a class of games used to model situations where player payoffs depend on the number of other players who select the same action, to the market for electricity generation, another unique contribution to the current literature. Given the complexity in solving for nonlinear congestion prices that fully incorporate the externality imposed by shifting consumption to a given hour, we instead propose a tractable method to best approximate congestion pricing. We develop a stochastic algorithm derived from the concept of Nash to solve for the market equilibrium by asynchronously updating households in response to price changes. We demonstrate that this method, as the number of total households increases and the proportion of households able to update their demand in any single iteration grows smaller, allows prices to virtually incorporate the congestion effects of household updates and in the limit enables the market to converge toward the desired equilibrium. We conclude with empirical results comparing our proposed algorithm against allowing households and the supply side to myopically best respond to one another, as a means to computing equilibrium.

Our research, in summary, provides a unique formulation of the problem and complete model of the market that is distinct from the current literature. We determine that an equilibrium in this market exists under complete information. Furthermore, we construct a mechanism enabling us to best approximate and in the limit converge to the equilibrium set of prices in the market, and demonstrate that our method is tractable for finding an approximation of equilibrium prices within some error bound of the actual values.

1.4 Related Work

Though the widespread installation of the physical technology necessary for our mechanism still lags our work, research on the "Grid of the Future" provides examples of technologies nearly ready for implementation and important to our solution. Ipakchi and Albuyeh [7] discuss the implications of PHEVs becoming online on the grid in the coming years and confirm the role these appliances may play in load shaping. Furthermore, they discuss current policy initiatives supporting demand response as the grid begins a transition from a load-following to load-shaping strategy, in which demand-side resources will be managed to meet the available generation and grid power delivery capacity at any time. Given their favorable outlook on demand response technology, we affirm the motivation for our research. In fact, given the central role smart meter technology will play in our proposed system, we refer to our mechanism as **smart power scheduling** for production and households.

From a theory perspective, foundational work on dynamic pricing by researchers at the University of California motivates the underlying approach to supply-demand coordination presented in this thesis. Borenstein *et al.* [2] provide an overview of demand response programs currently in effect or that have been proposed in the literature, and evaluates the economics behind these incentive programs and their effect on market efficiency and competitiveness. Though they consider programs as wide-ranging as interruptible contracts to positive reinforcement programs like paying customers to reduce demand, [2] provides a strong economic argument for dynamic pricing. They conclude that RTP captures a far larger share of wholesale price variation and therefore is among the most effective means to improving market outcomes.

Roozbehani *et al.* [11], referenced earlier, caution the dangers of real-time pricing in electricity grids. Their result points to an increased volatility in demand and the raised potential for grid instability when real-time prices are communicated to customers without a dampening constraint on the rate of demand response permissible. In the development of our model, we take this result into consideration and incorporate the supply side into our feedback loop in order to stabilize the effect of demand responses. For example, if demand is shifted away from a given hour, from $t = 1$ to $t = 2$, at the next iteration in which the supply problem is updated, prices for period $t = 1$ will decrease in response, which will at the next iteration incentivize some of the original customers to shift their consumption back to $t = 1$. This feedback loop in which the supply side provides a response to each change in demand, in effect stabilizes some of the volatility of the demand responses. This mechanism, as described later, additionally guarantees the existence of an equilibrium price in each hour.

Other similar papers to ours that apply game theory in analyzing the market for electricity generation include Chen *et al.* [3] and Mohsenian-Rad, Wong *et al.* [10]. Like our model, [3] proposes an RTP-based scheduling scheme to provide demand response for residential power use. Their model utilizes a Stackelberg game to model the responses between the service provider and individual households. The paper derives a sequential equilibrium for this game from the information exchange between the two sides of the market. Our work shares this feature of incorporating information exchange between the supply side and households. We, however, depart from the formulation proposed in [3], which optimizes over individual appliance requests to minimize the cost of inconvenience from delaying appliances summed with the cost of actual electricity consumption. Our model instead proposes a complete demand schedule for the coming day at each iteration in our computation, providing a schedule for each household's entire set of jobs to be run, and utilizes prices as a means to incentivize households into committing to these proposed load profiles during the day. [3] provides a set of optimal delays times for appliances as its solution, whereas our model schedules tasks feasibly into user-provided time constraints for each task. Computationally, our problem maintains a linear programming formulation for both sides of the market as we emphasize the practical need to set schedules for thousands of households simultaneously and to be able to solve for equilibrium prices.

The paper by Chen *et al.* also alludes to the problem of congestion pricing that we experience in our model when applying a myopic best response strategy by households in solving for equilibrium. They refer to the appearance of "rebound" peaks in hours that were supposed to be low loads, which occur because households make similar usage decisions without prices concurrently incorporating the congestion impact of these changes [3]. We approach solving this problem by randomizing and constraining which households are permitted to respond before the supply side updates prices again. We also provide a more rigorous treatment of this effect and develop an understanding of the congestion game theory underlying it in the process of proposing an algorithm that converges to a stable equilibrium.

A 2010 paper by Mohsenian-Rad and collaborators at the University of British Columbia

[10] takes another game theoretic approach to this problem, designing a consumption scheduling program based on nonlinear congestion-based costs. This paper treats the total electricity demand for each household as fixed and distributes the power required for each appliance across a feasible time range between parameterized release and deadline times. We borrow release and deadline times from this formulation for our own model, but constrain appliances to discrete power modes specific to each task, as is more realistic in practice. Whereas [10] also proves a demand smoothing result, the solution proposed by this model in practice would be feasible only for appliances with continuous gradation in power modes. Though the model in this paper can be extended to limit appliance power only to discrete levels, augmenting the model as is would require the introduction of integer decision variables, whereas our formulation for the demand side is restrictedly a linear program without any integrality constraints. Furthermore, the objective function in [10] is nonlinear, whereas our formulation preserves linearity in cost and still models congestion effects by iteratively updating prices as small perturbative responses in demand are made. This paper also provides a proof of the existence of a Nash equilibrium in this market, which proves important toward proving the existence of an equilibrium to our own model.

Another paper by Xiong, Chen *et al.* [16] focuses on the development of algorithms for appliance selection to set hourly household power budgets. In our model, we instead hard code the set of appliances for each household and define a household's utility as fixed, given that each task is scheduled within its feasible time range. From [16], we borrow and adapt the power request message format proposed for signaling appliance requests, defining our jobs as 5-tuples composed of the appliance, power level, duration, release time, and deadline time. [16] forgoes modeling the supply side and instead only considers the problem of demand smoothing through optimizing delays to appliance start times, a simplification in comparison to our model.

Other papers by Mohsenian-Rad and his co-authors focus on different aspects of residential load control relevant to our work. A paper co-authored with Leon-Garcia proposes regression-based techniques for price prediction in wholesale markets [9]. The authors argue that any residential load control strategy in real-time electricity pricing environments requires price prediction capabilities if utility companies are to provide price information only one or two hours ahead of time. Given the computational constraints of households, in practicality it is infeasible for individual households to run a complex price prediction model. As a result, in our communication mechanism, we compel the supply side to release a full set of prices for the day-ahead market forgoing the need for households to invest in any costly price prediction technology. Like the other papers, the model presented in [9] again emphasizes the determination of optimal delay times for appliances. By assuming households to be indifferent between schedules within the set time ranges for each task, we reduce our objective function to ignore the inconvenience costs of delayed start times. We approach the problem as though households are only minimizing their daily electricity bill by determining when to run tasks. [9] also proposes a stochastic method for overcoming load synchronization that is similar to our approach. They introduce short random starting delays in addition to the actual scheduled delay to diversify the start times of different appliances further and prevent load synchronization. Furthermore, whereas [9] solves a non-linear program with price prediction, which is transformed with a few modeling tricks into a mixed-integer program for tractability, we construct our formulation as a linear program from the start.

As we noted earlier, while the related literature also proposes a two-way feedback loop between the supply side and households, no formulation in the current literature is as complete as ours in modeling the full market as two interconnected linear programs. Our work, however, builds on the papers discussed here, taking their treatment of appliance job scheduling as inspiration but attempts to construct a more tractable model utilizing network flows, a formulation unique among the recent literature on this topic.

1.5 Outline

In Chapter 2, we introduce formal notation to describe the supply-side problem of economic dispatch and the demand-side problem of minimum-cost appliance scheduling as standalone optimization programs. In this chapter, we construct programs that allow us to solve for optimal production and appliance schedules given exogenous demand and prices. We also formally introduce our complete model of the market in which we seek to determine endogenous solutions to these quantities, and discuss the assumptions facilitating further analysis. In Chapter 3, we determine the existence of a Nash equilibrium for a market under complete information and describe the nature of the optimal solution. We motivate our mechanism design for prices by confirming convergence to Nash equilibrium if prices are set at equilibrium values. In the chapters that follow this, we direct our discussion toward the goal of computing equilibrium for our market, in which there is restricted transfer of hidden information as proposed for our retail market dynamics.

In Chapter 4, we propose a network flow formulation of the complete market and discuss the computational advantages of this model. In this chapter, we develop our centralized optimization program coordinating the demand-side problem across all households and determine that the solution found from this program will also be individually optimal for each household's individual optimization program. Chapter 4 also presents our initial attempts at solving for equilibrium by applying `MYOPIC_BEST_RESPONSE` as the demand response updating rule for each household in our computation. We provide a theoretical proof of why this algorithm fails to converge based on arguments from congestion game theory. Utilizing an understanding of congestion pricing, we conclude Chapter 4 by proposing a stochastic method to approximate market equilibrium. In Chapter 5, we provide empirical confirmation of our results from the previous chapter via simulation and compare the performance of our two updating algorithms from Chapter 4 in solving for equilibrium. Chapter 6 concludes with a recapitulation of our work and proposed extensions for further research.

Chapter 2

Model Preliminaries

In this chapter, we provide a more mathematical treatment of the wholesale and retail markets for electricity generation, in order to motivate our discussion of market equilibrium in the subsequent chapter.

2.1 Wholesale Market Structure

As discussed in the previous chapter, the wholesale electricity market is cleared in hourly auctions held by a centralized auctioneer. Generating plants available for production submit silent day-ahead bids to the dispatch center for each auction period based on their known production costs. In contrast to traditional goods auctions, the wholesale electricity auction has only one buyer and many sellers. This market allocates contracts for production to bidding generating plants at minimum cost to the dispatch center, which is responsible for the transfer of payment from retail customers to generating electric utilities.

In price-quantity space, bid i is submitted by a generating plant to the dispatch center as a promised quantity q_i^S for production in period t at a constant marginal cost mc_i. When arranged by marginal cost in order from lowest to highest, and summed cumulatively, these bids can be used to sketch the supply curve for wholesale electricity, forming what is also known as the **bid stack** for each time period. In electricity generation, the term *merit order* is used to denote a ranking of available sources of energy in increasing order of short-run marginal costs of production. Technologies with the lowest marginal costs are the first to be brought online to meet demand by merit order.

More formally, the bid stack is equivalent to the quantile function or inverse of the cumulative distribution function (CDF) of the distribution of marginal cost bids P. The quantile function over bids returns the price p below which q percent of the aggregated bids fall:

$$F(p) = \Pr(P \leq p) = \frac{Q^S}{Q^{max}} \equiv q \times 100\%.$$

Q^S ("quantity supplied") is the aggregate quantity available for production below the given price p, and Q^{max} is the maximum generation capacity across all power plants. The dispatch center in effect sets a daily Q^{max} as a result of deciding which plants to turn on for the day, constraining which plants are available and able to submit bids. The wholesale spot price P_t^{w*} of electricity at time t, or the **locational marginal price (LMP)**, is then the highest bid price from the submitted bids needed to match demand Q_t^D ("quantity demanded"), and it becomes the market-clearing wholesale price from each auction: $P_t^{w*} = \inf \{p \in \mathbb{R} : Q_t^D / Q^{max} \leq F(p)\}$. This auction mechanism creates a built-in incentive for generating plants not to overbid as plants that offer too high of a price are not selected to run. Furthermore, we will assume generating plants bid truthfully in this model given regulatory pressures that prevent market collusion.

For our model, it is worth noting that demand in the wholesale market is modeled as being perfectly inelastic because it is a *realized* demand set from the retail market equilibrium: retail demand connects the two markets. Furthermore, if end-user prices are set using time-of-use block pricing, in which retail prices are published possibly months ahead of time before the wholesale price P_t^w is known, we see that Q_t^D will not vary with P_t^w. Demand in the wholesale market, where the corresponding retail market observes block pricing, should be taken to be inelastic and is treated as a constant.

The bid stack, on the other hand, describes how Q^S varies as a function of the wholesale price p: $Q^S(p, Q^{max}) = Q^{max} \times F(p)$. The supply side must anticipate expected demand before it is realized in each production period in order to maintain optimal generation capacity for the day. Generation capacity corresponds to the number of power plants turned on for the day, setting the daily value for Q^{max}.

Because demand in each time period, until it is realized, will be unknown, we may treat it as a stochastic variable that follows some unknown distribution, $Q_t^D \sim F(\cdot)$. Under the current market structure, the supply side depends on load prediction programs based on historical load patterns to estimate $\mathbb{E}[Q_t^D]$, but fails to capture the associated error term ϵ_t, where $Q_t^D = \mathbb{E}[Q_t^D] + \epsilon_t$. If the conditional distribution of demand against its underlying explanatory variables like weather patterns or seasonal fluctuations can be empirically estimated, an electric utility can further condition its expectation of demand by these factors: $\mathbb{E}[Q_t^D | X_1 = x_1, X_2 = x_2, \ldots, X_n = x_n]$. Nonetheless, it is still infeasible to capture the effect of all of the explanatory variables underlying demand.

In this thesis, we propose a pricing mechanism by which the producer can compute and set retail prices to equilibrium values such that hourly demand converges to the respective equilibrium quantities that minimize the total costs of production. By designing prices, we allow production to actively induce a socially optimal profile for consumption instead of merely passively predicting hourly demand.

2.2 Economic Dispatch Problem

In this section, we briefly introduce the relevant notation and set up the problem faced by the supply side. We also provide a discussion of generation cost.

Economic dispatch is a class of optimization problems in which a central supply planner must make short-term capacity decisions about the optimal number of generating plants to turn on in a given production period and subsequent production decisions about where to produce electricity (i.e. at which generating plants and at what levels), in order to meet system load as it is realized, at minimum cost and while maintaining grid stability. The problem is an important one because generation capacity decisions must be made ex ante to production decisions made later in every hour. Setting too high of a generation capacity imposes high fixed costs that are then spread out over the entire day; if capacity falls short of demand in any hour, the system may face blackouts or grid failure.

In our simplified model of the bid stack, setting Q^{max} and deciding from which generating plants to accept bids represents the first step of economic dispatch. Subsequently, given the set of plants turned on, in each hour, the dispatch center must determine production schedules for every online plant under its purview. The key decisions in this problem are:

- which generating units to turn on to provide day-long system capacity, and
- the generation level of each online plant, across each time period.

The dispatch center is tasked with finding a minimum cost solution for hourly production that is subject to a series of technical constraints, including:

1. maximum and minimum stable operating levels at each plant,
2. maximum generation capacity at each plant online for the day, and
3. minimum system-wide generation in each period of production to meet demand.

We assume in this model that the dispatch center determines daily capacity and hourly production for a set $\mathcal{G}$ of available generating plants for every hour $h \in \mathcal{H}$, the set of time periods in the day. For each generating plant, the dispatch center faces increasing total costs of production, as well as non-decreasing marginal costs. In practice, the marginal production cost curve for each plant can be modeled as piecewise, as shown in Figure 2.1. In our model, we refer to each piece of this curve with constant marginal cost as a discrete "production mode" for the plant. A production mode $\tau \in \mathcal{T}(g)$, the complete set of production modes for plant $g \in \mathcal{G}$, is a fixed unit of capacity for production available at constant variable cost. Once the plant has produced up to $q^{max}_{g,\tau}$ units of electricity, the width of each piece of the curve in Figure 2.1, at variable cost $c^v_{g,\tau}$, the plant must carry out the next unit of production at a higher marginal cost, corresponding to the next production mode after τ, $c^v_{g,(\tau+1)}$. The highest production mode at which a plant is currently operating determines the marginal cost for the next unit of electricity to be produced.

To introduce more notation, production variable $f^h_{g,\tau}$ is the generation in hour $h \in \mathcal{H}$ at production mode τ with constant variable cost $c^v_{g,\tau}$ at plant g. Variable costs at every

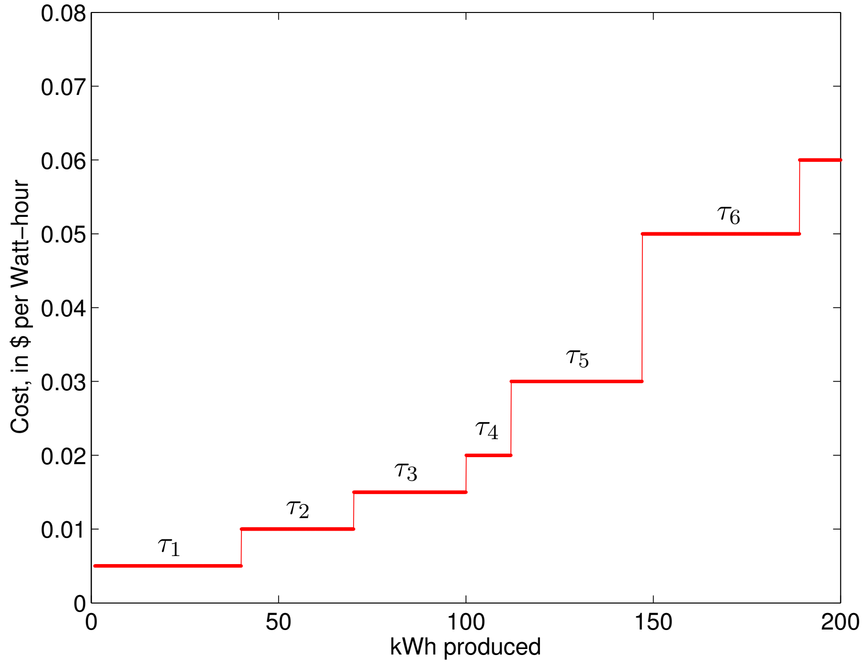

Figure 2.1: Production Modes $\mathcal{T}(g)$ for Plant $g \in \mathcal{G}$ and Marginal Cost $c^v_{g,\tau}$

generating plant increase in τ, thus $0 < c^v_{g,\tau} < c^v_{g,\tau'}$, $\forall \tau, \tau' \in \mathcal{T}(g) : \tau < \tau'$. Additionally, if $c^v_{g,\tau} = c^v_{g,\tilde{\tau}}$ then implicitly $\tau = \tilde{\tau}$, by our definition of a *production mode*. Given this, the dispatch center, in minimizing cost, will always allocate generation to lower and less costly production modes until the capacity constraints at the lower modes bind: $f^h_{g,\tau} \leq q^{max}_{g,\tau}$, forcing marginal production to be scheduled for the lowest production mode for which this capacity constraint is not yet binding. Thus, when we say a plant is operating at a given mode τ, we mean marginal production occurs at mode τ, but the plant is still operating at capacity at all modes $\tau' < \tau$, $\forall \tau' \in \mathcal{T}(g)$, $g \in \mathcal{G}$. Additionally, c^f_g is the daily fixed cost of running capacity online at plant g for the day. The fixed cost of turning on the plant is the same regardless of which mode(s) for production are later utilized.

Given these assumptions, the hourly production cost curve $C^h_g(\cdot)$, $h \in \mathcal{H}$ for every plant $g \in \mathcal{G}$ is then increasing and convex in the quantity produced in hour h at that plant. In our model, given that variables costs are assumed to be the same throughout the day, we know that: $C^0_g(\cdot) = C^1_g(\cdot) = \ldots = C^h_g(\cdot) = C_g(\cdot)$, $\forall h \in \mathcal{H}$, $g \in \mathcal{G}$. The daily fixed cost to turning on a plant as online generation capacity is then shared uniformly across each hour. Given the set of production curves for each generating plant $g \in \mathcal{G}$, we may derive the aggregate hourly production cost curve $C(\cdot)$, if each $q^S_{g,h}$ is the hourly production at plant g:

$$C\big(q^S_{1,h}, q^S_{2,h}, \ldots, q^S_{|\mathcal{G}|,h}\big) = \sum_{g \in \mathcal{G}} C_g\big(q^S_{g,h}\big), \quad \forall h \in \mathcal{H}.$$

However, in our model, aggregate production $Q^S_h = \sum_{g \in \mathcal{G}} q^S_{g,h}$ is coordinated by the dispatch center in order to minimize cost. Given this, we may apply a merit order to production as we have done before to generation capacities in the bid stack. Under a composition of all production modes across all plants, a cost minimizing dispatch center will always allocate production to a mode with lower marginal cost until its capacity constraint binds before scheduling production at a higher cost mode across plants. As a result, the aggregate marginal production cost curve will always be non-decreasing in the aggregate quantity of hourly production Q^S_h. Finally, because marginal production costs will always be non-decreasing, we know that the aggregate production curve is also convex in the aggregate quantity of hourly production.

We conclude this section by introducing the remaining notation for our model, as well as the mixed-integer programming formulation of the economic dispatch problem.

Table 2.1 defines the remaining sets relevant to the problem. $\mathcal{G}$ is the set of generating plants and $\mathcal{T}(g)$, the set of production modes at plant g. $\mathcal{H}$ is, once again, the set of time periods for production, corresponding to hours in the day.

Tables 2.2 and 2.3 define the decision variables and parameters relevant to this problem, in notation consistent with the rest of the thesis. It is important to note for this problem that Q^D_h, hourly demand, is taken to be exogenous. The program we define below simply schedules production according to a given demand for each hour. In our complete market model and at the end of this chapter, we will begin to treat this quantity as an endogenous variable. Before we do so, we must first introduce the household problem in the next section and provide the framework to structure these two problems together at the end of this chapter.

The economic dispatch problem is defined as follows:

$$\min_{\sigma,\mathbf{f}} \Big\{ \sum_{g\in\mathcal{G}} c_g^f \cdot \sigma_g + \sum_{h\in\mathcal{H}} \sum_{g\in\mathcal{G}} \sum_{\tau\in\mathcal{T}(g)} c_{g,\tau}^v \cdot f_{g,\tau}^h \Big\} \tag{2.1}$$

subject to:

$$\sum_{h\in\mathcal{H}} \sum_{\tau\in\mathcal{T}(g)} f_{g,\tau}^h \leq M_g \cdot \sigma_g, \, \forall g \in \mathcal{G} \tag{2.2}$$

$$q_g^{min} \cdot \sigma_g \leq \sum_{\tau\in\mathcal{T}(g)} f_{g,\tau}^h, \, \forall g \in \mathcal{G}, \, h \in \mathcal{H} \tag{2.3}$$

$$\sum_{g\in\mathcal{G}} \sum_{\tau\in\mathcal{T}(g)} f_{g,\tau}^h \geq Q_h^D, \, \forall h \in \mathcal{H} \tag{2.4}$$

$$0 \leq f_{g,\tau}^h \leq q_{g,\tau}^{max}, \forall \tau \in \mathcal{T}(g), \, \forall g \in \mathcal{G}, \, h \in \mathcal{H} \tag{2.5}$$

$$\sigma_g \in \{0,1\}, \, \forall g \in \mathcal{G} \tag{2.6}$$

Table 2.4 provides descriptions of the constraints formulated above. Constraint (2.2) restricts total daily production, across all modes and time periods, at plant g to its online daily generation capacity, which is restricted to 0 if $\sigma_g = 0$. Constraint (2.3) sets the minimum stable operating level across all production modes at plant g for each hour if plant g is online. Given the next constraint, (2.4), hourly production over all plants must meet demand in each time period. Constraint (2.5) limits marginal generation capacity at each production mode $\tau \in \mathcal{T}(g)$, $g \in \mathcal{G}$ in each hour. Finally, the generation binaries σ_g, $g \in \mathcal{G}$ indicate whether a plant has been turned on for production or not. As per the structure of electricity markets, the supply side is able to set generation binaries σ only once per day, at midnight, defining the partition of generating plants $\mathcal{G}$ from which it may accept hourly bids, and then is permitted to readjust production flows $\mathbf{f}$ at each hour once the aggregate hourly load Q_h^D is realized.

The objective function (2.1) is a minimization over the total aggregate daily production cost, the sum of the fixed costs of all generating plants turned on with the hourly variable costs of production at every generating plant.

Table 2.1: Sets for the Economic Dispatch Problem

Set	Description
$g \in \mathcal{G}$	set of generating plants
$\mathcal{T}(g) \in \mathcal{T}$, $\forall g \in \mathcal{G}$	set of sets of production modes, indexed by plant g
$\tau \in \mathcal{T}(g)$, $\forall g \in \mathcal{G}$	set of production modes at plant g
$h \in \mathcal{H}$	time periods

Table 2.2: Decision Variables and Cost Parameters for the Economic Dispatch Problem

Variable	Description	Cost
σ_g, $\forall g \in \mathcal{G}$	turning on power plant g for the day	c_g^f
$f_{g,\tau}^h$, $\forall \tau \in \mathcal{T}(g), g \in \mathcal{G}, h \in \mathcal{H}$	hourly generation at plant g at production mode τ	$c_{g,\tau}^v$

Table 2.3: Capacity Parameters for the Economic Dispatch Problem

Parameter	Description
$q_g^{min}, \forall g \in \mathcal{G}$	minimum hourly operating level at plant g
$q_{g,\tau}^{max}, \forall \tau \in \mathcal{T}$, $g \in \mathcal{G}$	marginal hourly capacity at plant g for production mode τ
$M_g, \forall g \in \mathcal{G}$	maximum daily capacity of plant g
$Q_h^D, \forall h \in \mathcal{H}$	hourly demand load in time period h

Table 2.4: Constraints for the Economic Dispatch Problem

Constraint	Description
(2.2)	generation binary restricting production to plants with online capacity
(2.3)	minimum stable hourly operating level at plant g if online
(2.4)	hourly production over all plants must meet demand in each time period
(2.5)	min/max marginal generation capacity at production mode τ
(2.6)	binary decision constraint for turning on power plant g

2.3 Demand Side Management

In this section, we introduce formal notation for the household problem and define the constraints relevant to our program for appliance scheduling. We also propose necessary and sufficient conditions for scheduling feasibility.

We first introduce the sets relevant to our model for households. Let $\mathcal{N}$ be the complete set of household subscribers, with $N \triangleq |\mathcal{N}|$. For each household $n \in \mathcal{N}$, $\mathcal{A}_n$ denotes the set of appliances with tasks that require scheduling. Set $\mathcal{M}(a)$, $a \in \mathcal{A}_n$ is the set of discrete power modes at which appliance a can be run. This is an improvement to the model setup in [10], in which appliance power levels are taken to be continuous. Set $\mathcal{J}(a)$, $a \in \mathcal{A}_n$, $n \in \mathcal{N}$ is the set of jobs to be run in the given day on appliance a belonging to household n that have been submitted for scheduling. We define job $j \in \mathcal{J}(a)$ as the tuple $j = (a, m, d, \alpha, \beta)$, where m is a power mode in the set $\mathcal{M}(a)$, $a \in \mathcal{A}_n$, and duration d for task j is less than or equal to the the time range (inclusive of the deadline hour) set by the household to schedule the task: $\beta - \alpha + 1 \geq d > 0$. Release and deadline times α and β are specified by the subscriber when the task is submitted for scheduling such that $\alpha \leq \beta$. We make the strong assumption that in submitting a time range for scheduling, the subscriber indicates uniform preference across the interval for completing the task. If the time range [6:00 PM, 9:00 PM] is submitted for a task of duration $d = 1$ hour, the program assumes this to indicate that the subscriber has equal utility between the task being scheduled for 6:00-7:00 PM, 7:00-8:00 PM, 8:00-9:00 PM, or 9:00-10:00 PM, given that the hour of our time deadline is included. Set $\mathcal{H} \triangleq \{0, \ldots, H-1\}$ is the set of hours or time periods in the current day-ahead scheduling time frame. All deadlines are set to expire before the end of the day: $0 \leq \alpha \leq \beta \leq H - 1$. For notational convenience, we also define set $\mathcal{J}_n = \cup_{a \in \mathcal{A}_n} \mathcal{J}(a)$ as the full set of jobs to be scheduled in household n across all of its appliances. Table 2.5 summarizes.

Our setup further extends the model in [10] by allowing multiple tasks to be submitted for each appliance during the day. For example, a household may need to schedule the dishwasher once for an hour in the morning between [8:00 AM, 10:00 AM], inclusively, and

Set	Description
$n \in \mathcal{N}$	set of households
$a \in \mathcal{A}_n$	set of appliances owned by household n
$m \in \mathcal{M}(a)$, $\forall a \in \mathcal{A}_n$, $n \in \mathcal{N}$	set of power modes for appliance a
$j \in \mathcal{J}(a)$, $\forall a \in \mathcal{A}_n$, $n \in \mathcal{N}$	set of jobs for appliance a
$\mathcal{J}_n = \cup_{a \in \mathcal{A}_n} \mathcal{J}(a)$, $\forall n \in \mathcal{N}$	union of all jobs to be scheduled in household n
$h \in \mathcal{H}$	time periods

Table 2.5: Sets for the Household Appliance Scheduling Problem

again at night for an hour between [6:00 PM, 9:00 PM].

In order to compose an implementable schedule, we impose the constraint that an appliance can only be run at a single power mode at any given time and two tasks on a single appliance cannot be scheduled to run simultaneously. We impose checks to user inputs to evaluate whether a schedule can feasibly be set given the tasks submitted by the user for scheduling, and determine a criteria to characterize infeasible inputs. For two tasks on a single appliance, task 1 and task 2, with durations d_1 and d_2 and corresponding time ranges during which the tasks can be scheduled: closed intervals $[\alpha_1, \beta_1]$ and $[\alpha_2, \beta_2]$, we consider all possible cases by which the two tasks may overlap. Assuming without loss of generality that $\alpha_1 \leq \alpha_2$, we consider all three logical permutations of the time bounds below. To simplify notation, we define β_i' to be the *time* deadline by which the task must be completed, with hour β_i' unavailable for running the task: $\beta_i' = \beta_i + 1$. We also assume that each task by itself has been feasibly defined, in that the time range submitted is at least as long as the required time for the task: $\boxed{\beta_i' - \alpha_i \geq d_i,\ \forall i}$.

Case 1 (disjoint time ranges: $\alpha_1 < \beta_1' \leq \alpha_2 < \beta_2'$)**. If a feasible schedule exists, then** $\boxed{\beta_i' - \alpha_i \geq d_i > 0}$ **holds** $\forall i = \{1, 2\}$**.**

Proof. The proof in this case is trivial. We had imposed this condition for each task that the user submitted for scheduling. □

Case 2 (overlapping time ranges: $\alpha_1 \leq \alpha_2 \leq \beta_1' \leq \beta_2'$)**. If a feasible schedule exists, then** $\boxed{\beta_2' - \alpha_1 \geq d_1 + d_2}$ **must hold**.

Proof. The remaining time between $[\alpha_2, \beta_2')$ after the second job has been scheduled plus the allotted time in the disjoint interval $[\alpha_1, \alpha_2)$, in which only work toward the first task may be scheduled, must be greater than the duration required for the first job: $\alpha_2 - \alpha_1 + \underbrace{\beta_2' - \alpha_2 - d_2}_{\text{time remaining}} \geq d_1 \implies \beta_2' - \alpha_1 \geq d_1 + d_2$, which we sought to prove. □

Case 3 (inscribed time ranges: $\alpha_1 \leq \alpha_2 \leq \beta_2' \leq \beta_1'$). **If a feasible schedule exists, then** $\boxed{\beta_1' - \alpha_1 \geq d_1 + d_2}$ **must hold**.

Proof. The time remaining in the range $(\alpha_2, \beta_2']$ after task 2 has been scheduled plus the time in the segments $(\alpha_1, \alpha_2]$ and $(\beta_2', \beta_1']$, during which only task 1 may be scheduled, must be greater than the duration required for the first job: $\underbrace{\beta_2' - \alpha_2 - d_2}_{\text{time remaining}} + \alpha_2 - \alpha_1 + \beta_1' - \beta_2' \geq d_1 \implies \beta_1' - \alpha_1 \geq d_1 + d_2$, which we sought to prove. □

The contrapositives of these statements imply that if any of the time range constraints do not hold, then a feasible solution does not exist and the inputted jobs cannot be feasibly scheduled:

- **All Cases**: If $\beta_i' - \alpha_i < d_i$ for either $i \in \{1, 2\}$, then a feasible solution does not exist.
- **Case 1**: Nothing else.
- **Case 2**: If $\beta_2' - \alpha_1 < d_1 + d_2$, then a feasible solution does not exist.
- **Case 3**: If $\beta_1' - \alpha_1 < d_1 + d_2$, then a feasible solution does not exist.

The statements here correspond to the cases defined on the previous page.

By generalizing the above conditions to apply to any number of k overlapping tasks, we establish that: *a feasible solution exists if and only if the total width of* **all subsets** *of overlapping time ranges for an appliance is greater than or equal to the sum of the durations of all tasks in the subset.* Moreover, for k overlapping tasks, we must check that the above conditions hold for all $2^k - 1$ nonempty subsets of tasks. The conditions are trivially true for the empty subset.

Proof. This requirement is a consequence of Hall's Marriage Theorem for a perfect matching [6]. We construct a bipartite graph with a partition of nodes corresponding to every hour in duration for each task ("task-hours"). For example, a three-hour long task j would have three nodes $\{j_1, j_2, j_3\}$. In the second partition, we define a node corresponding to each hour of the day that appliance a can be run ("appliance-hours"). To produce a perfect matching, we append dummy nodes in the first partition, for as many hours as an appliance will sit idle during the day. There are $H - \sum_{j \in \mathcal{J}(a)} d_j$ such dummy nodes for each $a \in \mathcal{A}_n$, $n \in \mathcal{N}$. Edges in the graph are drawn between task-hour nodes to the respective appliance-hour nodes in the range $[\alpha_j, \beta_j]$. A dummy node is adjacent to all appliance-hour nodes for the corresponding appliance. A perfect matching of all task-hours to appliance-hours, implying a feasible schedule, is guaranteed if and only if the condition in Hall's Marriage Theorem, stated above, holds true. $\square$

We have so far directed our attention at the mechanics of household appliance scheduling at the local level. Given real-time prices for each hour, in our model, households seek to optimally schedule tasks to be completed in order to minimize their household electricity bill. **Demand side management** refers to this process by which household appliances are autonomously scheduled by a built-in scheduling program in order to obtain a minimum-cost solution.

Like our model for economic dispatch, the household appliance scheduling model presented here solves for only an optimal solution for each household's schedule given a vector of hourly prices as exogenous parameters. In the next section, we begin to introduce our problems in conjunction with one another, as a framework to compute prices and quantities endogenously.

We restate the parameters introduced in this section in Figure 2.6 for reference.

In order to construct a linear program for the household appliance scheduling problem, we consider the bipartite graph formulation discussed earlier; our problem corresponds to a minimum weight perfect matching problem. For this formulation, we first construct a bipartite graph $G = (X, Y, E)$ where nodes in set X corresponds to task-hours and nodes

Table 2.6: Parameters for the Household Appliance Scheduling Problem

Parameter	Description
$[\alpha_j, \beta_j]$, $\forall j \in \mathcal{J}_n$, $n \in \mathcal{N}$	feasible time ranges submitted for all jobs
d_j, $\forall j \in \mathcal{J}_n$, $n \in \mathcal{N}$	submitted durations for all jobs
m_j, $\forall j \in \mathcal{J}(a)$, $a \in \mathcal{A}_n, n \in \mathcal{N} : m_j \in \mathcal{M}(a)$	submitted appliance mode for all jobs
P_h, $h \in \mathcal{H}$	hourly price per unit of consumption

in set Y represent appliance-hours. Node $i \in X$ is defined as the tuple $(n, j, t) \in X$, $\forall n \in \mathcal{N}$, $j \in \mathcal{J}_n$, $t = \{1, 2, \ldots, d_j\}$, where d_j is the duration of job j. Node $k \in Y$ corresponds to the tuple $(n, a, h) \in \mathcal{N} \times \mathcal{A}_n \times \mathcal{H}$, such that we have constructed a node for every appliance-hour. Arc $\big((n, j, t), (n, a, h)\big) \in E$, the set of edges of graph G, if $\alpha_j \leq h \leq \beta_j$ and job $j \in \mathcal{J}(a)$, $\forall n \in \mathcal{N}$.

To solve the household scheduling problem using this combinatorial optimization approach, we must find a perfect matching of task-hours to appliance-hours. In any feasible solution, the total duration of tasks to be scheduled equals the total number of hours of appliance use. The mapping from X to Y is bijective or perfect when accounting for the appended "dummy nodes" in X for each appliance. For a matching to be feasible in this problem, no appliance must be scheduled for longer than an hour for each appliance-hour. Likewise, no task-hour should be scheduled for less than one hour to ensure that all tasks are completed, and nor should a task-hour be scheduled for more than one hour. We introduce variables x_{ik}, each corresponding to an arc $(i, k) \in E$. Each of these arcs has cost $c_{(n,j,t)(n,a,h)} \equiv P_h$, $\forall \big((n, j, t), (n, a, h)\big) \in E : h \in \mathcal{H}$. In the objective function, each of these costs will be weighted by m_j, the power setting at which each job is scheduled to run. All outgoing arcs from a dummy node are assumed to have a weighted cost of 0.

As notation, we introduce a few other terms and sets that will be useful in our formulation below. For any arc (i, k), we refer to $i \in X$ as the start node and $k \in Y$ as the end node. Given this, let $I(k)$ be the set of all start nodes in X of arcs incoming to node k, and let $O(i)$ be the set of end nodes in Y of arcs outgoing from node i. More formally, for any bipartite graph $G = (X, Y, E)$, we define $I(k) \equiv \{i \in X \mid (i, k) \in E\}$ and $O(i) \equiv \{k \in Y \mid (i, k) \in E\}$.

Our linear program can be formulated as follows:

$$\min_{\mathbf{x}} \sum_{\big((n,j,t),(n,a,h)\big)\in E} c_{(n,j,t)(n,a,h)} \cdot m_j \cdot x_{(n,j,t)(n,a,h)} \tag{2.7}$$

subject to:

$$\sum_{k\in O(i)} x_{ik} = 1, \quad \forall i \in X \tag{2.8}$$

$$\sum_{i\in I(k)} x_{ik} = 1, \quad \forall k \in Y \tag{2.9}$$

$$0 \le x_{ik} \le 1, \quad \forall (i,k) \in E \tag{2.10}$$

We leave this program as a linear program relaxation as we permit solutions in which an appliance has been scheduled for part of an hour as long as our constraints above are still satisfied. Given our formulation, though our decision variables are not constrained to $\{0,1\}$, the upper and lower bound constraint on x_{ik} still restrict an appliance from being scheduled for more than an hour for each time period.

We provide some intuition to the rest of the constraints from above. Constraint (2.8) guarantees that every task-hour (n,j,t) is scheduled for exactly one hour on the respective appliance in the time range $[\alpha,\beta]$. Likewise, constraint (2.9) guarantees that every appliance is scheduled for exactly one hour of use or idle time in period h, $\forall h \in \mathcal{H}$, across all jobs submitted for that appliance. The minimum cost perfect matching formulation captures our problem at hand. By Hall's Marriage Theorem, we are guaranteed a feasible solution to this problem if all tasks submitted for scheduling meet the conditions outlined earlier in this section.

Nonetheless, there is a big assumption implicit in this formulation, that we later improve in Chapter 4 with our network flow formulation of this problem: we may only submit tasks with integer durations to this program. Each task is represented by an integral number of nodes in the partition X, equal in number to the duration of the task. Thus, in this program, while the formulation allows for x_{ik} to take factional values, allowing for a task to be scheduled on an appliance for less than an hour during period h, it still requires that the duration of every task be integral.

Finally, as worth mentioning again, this model only sets the schedule for a household given exogenous parameters for prices in each hour. In the next section, we begin to introduce a framework for connecting the household appliance scheduling problem to the economic dispatch problem from the previous section as the first step toward solving for prices and quantities endogenously. By treating price and quantity as endogenous variables, we may begin to model the complete market. The framework introduced here enables us to begin to solve for the market equilibrium that optimizes both programs simultaneously, achieving a minimum cost appliance schedule for households that is also minimizes the total social costs of production.

2.4 Retail Market Model

In the previous two sections, we constructed linear programming formulations of the supply-side and demand-side problems with the important caveat that both require either demand or price parameters to be defined exogenously. The economic dispatch program we formulated takes as a parameter, the vector $\mathbf{Q}^D$, demand for each hour. Likewise, the household appliance scheduling program from the previous section requires pricing information for each hour, $\mathbf{P}$.

Whereas we began the chapter with a discussion of the wholesale market structure, we conclude by introducing the framework for the retail market as we envision it, one that takes these two standalone programs and determines prices and quantities endogenously. First, it is worth restating the current market structure, as described in the previous chapter, with more formal notation. Under the time-of-use block pricing schemes discussed in Chapter 1, the producer determines exogenous prices for consumers ahead of time, setting $\mathbf{P}$ typically by a three-tiered block pricing scheme:

$$P_h = \begin{cases} P_{\text{peak}} & \text{if } h \in \mathcal{H}_{\text{peak}} \\ P_{\text{shoulder}} & \text{if } h \in \mathcal{H}_{\text{shoulder}} \\ P_{\text{non-peak}} & \text{if } h \in \mathcal{H}_{\text{non-peak}} \end{cases}, \qquad \mathcal{H} = \mathcal{H}_{\text{peak}} \cup \mathcal{H}_{\text{shoulder}} \cup \mathcal{H}_{\text{non-peak}}.$$

In this pricing model, the electric utility determines partitions of the day into different blocks with fixed hourly prices for each time of day. The actual prices P_{peak}, P_{shoulder}, and $P_{\text{non-peak}}$

are readjusted only a few times a year to reflect general changes in the costs of production.

Wholesale prices, on the other hand, are still determined by hourly auctions that are carried out by the dispatch center, which also determine the production schedules for each period. This auction mechanism is roughly captured by the economic dispatch problem we have constructed. Producers depend on forecasting models for $\hat{Q}_h^D$, $\forall h \in \mathcal{H}$, in order to set day-long generation capacity and then production schedules for every hour. Under block pricing, the market for electricity generation operates as two disjoint programs as we have constructed in this chapter thus far.

We now introduce our endogenous model for the retail market in which these two problems are solved jointly in order to construct a socially efficient solution. In this model, we first communicate optimal consumption schedules from the household program to set parameters for hourly demand in the producer optimization program: $Q_h^D \triangleq Q_h^D(\mathbf{x}^*)$. In return, real-time prices are derived from the costs of production in each period and are communicated from the producer to households, where they are used as parameters in the household appliance scheduling program: $P_h \triangleq P_h(\mathbf{f}^*, \sigma^*; \mathbf{c}^f, \mathbf{c}^v)$. Each side of the market continues to optimize its respective objective function, but parameters are now determined by the optimal solution for the other side of the market. Demand Q_h^D in each hour depends on how households schedule production. Prices, in an ideal sense, should correlate with the actual costs of production, which are determined by the optimal production schedule set in the economic dispatch problem.

Figure 2.2 visualizes the structure of the retail market with two-way communication in place between the producer and households. In the figure, we reference the respective objective functions that each side of the market optimizes. As we will introduce in a later chapter, we intend for each side of the market to solve a single optimization program that will centrally coordinate production and consumption.

By designing prices in this market based on the true costs of production in each hour, we are able to align the individual cost of consumption with the social cost of production. The point at which these two valuations are equal reflects the socially optimal point of

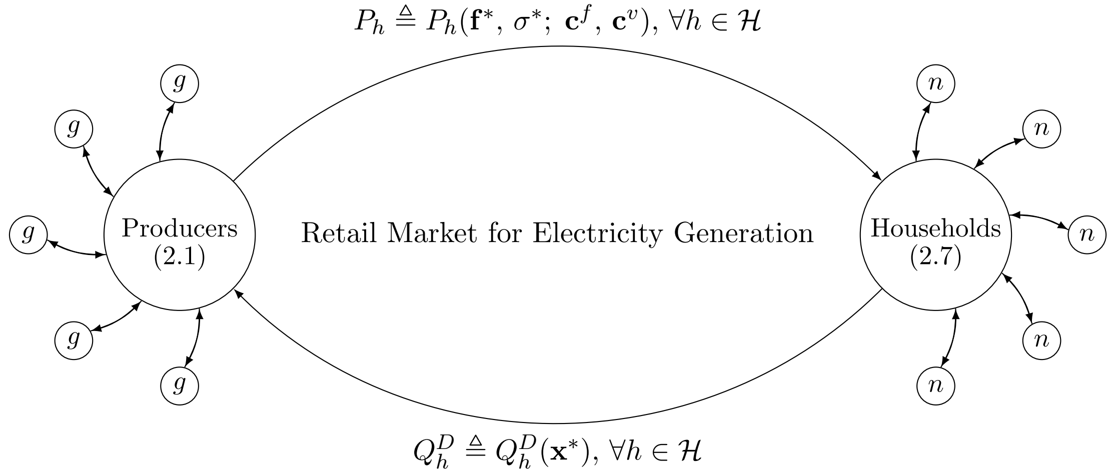

Figure 2.2: Complete Market Model, $\forall n \in \mathcal{N},\ g \in \mathcal{G}$

consumption and production in the market and defines an equilibrium where supply and demand are in balance.

In our model, we base the retail pricing function on the average total cost of production. Let $\sum_{h\in\mathcal{H}} C\big(\mathbf{f},\, \sigma;\ \mathbf{c}^f,\, \mathbf{c}^v\big)$ denote the objective value (2.1) of the economic dispatch problem at the solution $(\mathbf{f},\, \sigma)$. The per unit cost of production in each hour is then:

$$\frac{C\big(\mathbf{f}^h,\, \sigma;\ \mathbf{c}^f,\, \mathbf{c}^v\big)}{\mathbf{f}^h \cdot \mathbf{1}} = \frac{C\big(Q_h^S\big)}{Q_h^S}, \quad \forall h \in \mathcal{H}.$$

Given that we seek to coordinate supply and demand, we assume real-time retail prices to be some linear scaling of this per unit cost of production. In our notation, we define κ as the gross profit margin for the producer. The retail price charged to consumers per unit of consumption is then:

$$P_h \triangleq \kappa \times \frac{C\big(\mathbf{f}^h,\, \sigma;\ \mathbf{c}^f,\, \mathbf{c}^v\big)}{\mathbf{f}^h \cdot \mathbf{1}}, \quad \forall h \in \mathcal{H}. \tag{2.11}$$

Throughout our model and discussion in subsequent chapters, we implicitly assume $\kappa \equiv 1$, as it is really just a scaling to household currency. In our analysis, we are more concerned with the implications for total costs of production than the profit made by the producer. With any constant value κ, profit will always be equal to $(\kappa - 1) \times 100\%$ of the total costs of production.

We conclude this section by discussing the rationale for minimizing the total costs of production, particularly why the supply side agrees to minimize cost if profit is defined to be a fixed percent of total cost. In our retail market, we assume some benevolent social planner will implement this proposed communication design. Though the producer is no longer profit maximizing by our definition of retail prices, minimizing total costs of production corresponds to a maximization of social welfare: the energy consumption needs of households are met, given any feasible solution to the household problem, at a minimum cost to production. Our end goal in this system is to produce a more efficient outcome in the market for electricity generation by devising real-time prices for each hour that forces the market to converge to equilibrium.

In the next chapter, we extend our framework for understanding the retail market and prove results on the existence of market equilibria.

Chapter 3

Market Equilibrium

In this chapter, we introduce an energy consumption game played between the producer and households as a consequence of economic dispatch and demand side management in the retail market. After explaining the rules, we analyze strategies of this game, its implications for the producer, and analyze the conditions under which a Nash equilibrium exists for the game. We assume from this chapter onward that the producer and households interact in the retail market we constructed at the end of Chapter 2. Prices and quantities for consumption and production are taken to be endogenous variables in our complete market model.

3.1 Retail Market Energy Consumption Game

In this section, we utilize game theory to model the interaction between the producer and households in the retail market. Before introducing our formal model, we outline a few assumptions about household preference.

3.1.1 Household Utility and Payments

In our model for demand side management, any feasible solution that satisfies the constraints in (2.8), (2.9), and (2.10), falls on the same indifference curve for households; that is, household n has iso-utility for any solution that feasibly schedules all of its tasks for the day. This follows from the assumption that in submitting a time range $[\alpha_j, \beta_j]$ for each task, subscriber n indicates uniform preference across the interval for completing the task, $\forall j \in \mathcal{J}_n$. We define the utility function for household n formally in terms of its solution $\mathbf{x}_n$ to the household appliance scheduling program from the previous chapter:

$$u_n(\mathbf{x}_n) = \begin{cases} +1 & \text{if } \mathbf{x}_n \text{ satisfies (2.8), (2.9), and (2.10)} \\ -\infty & \text{otherwise} \end{cases} \tag{3.1}$$

Given that households are indifferent between any feasible solution to an appliance scheduling program, the only rational move for households is to then minimize their payment to the producer for this fixed amount of daily electricity consumption.

We define household n's daily bill b_n to the electric utility as the sum over its hourly consumption $q^D_{n,h}$, which is tabulated from $\mathbf{x}_n$, times the hourly price P_h:

$$b_n = \sum_{h \in \mathcal{H}} q^D_{n,h} \times P_h, \quad \forall n \in \mathcal{N}.$$

Because we are now endogenously solving for prices and quantity, we may substitute our expression for price from (2.11), as defined in terms of the economic dispatch decision variables:

$$b_n = \sum_{h \in \mathcal{H}} q^D_{n,h} \times \underbrace{\kappa \times \frac{C\big(\mathbf{f}^h,\, \sigma;\; \mathbf{c}^f,\, \mathbf{c}^v\big)}{\mathbf{f}^h \cdot \mathbf{1}}}_{P_h}, \quad \forall n \in \mathcal{N}.$$

Replacing Q^S_h for aggregate production in each hour and setting, without loss of generality, $\kappa \equiv 1$, our expression for household n's daily bill simplifies to:

$$b_n = \sum_{h \in \mathcal{H}} \frac{C\big(Q^S_h\big)}{Q^S_h} \cdot q^D_{n,h}, \quad \forall n \in \mathcal{N}. \tag{3.2}$$

This equation is an endogenous expression for household payments in terms of the decision variables from the economic dispatch problem.

3.1.2 Formal Model

The behavior and interaction between households and the producer in the market can be formally represented as a game in which households choose between consumption profiles, minimizing their energy bill, and the producer, in response to household consumption, determines his optimal production strategy for the day. Defining the **retail market energy**

consumption game as a model for households and the producer, we are able to prove our intuition as to the behavior of the retail market in response to real-time prices that lead households to shift consumption from one hour to another.

Game 1 (Retail Market Energy Consumption Game)**.** The energy consumption game is a tuple $(\mathcal{N} \cup P,\ X \times F,\ u)$, where:

- **Players**: $\mathcal{N}$ is a finite set of N households, corresponding to the set defined earlier and indexed by n, and P is the dispatch center coordinating production for the market;

- **Actions**: The cross $X \times F$ defines the action space for the game. $X = X_1 \times X_2 \times \ldots \times X_N$, where X_n is the finite set of energy consumption scheduling profiles $\mathbf{x}_n$ available to player n. Load profile $\mathbf{x}_n$ is an $|\mathcal{A}_n| \times |\mathcal{H}|$ matrix denoting the energy allocation to each appliance in each hour for household n, of the form:

$$\mathbf{x}_n = \underbrace{\begin{matrix} 1 \\ 2 \\ \ldots \\ |\mathcal{A}_n| \end{matrix}}_{a \in \mathcal{A}_n} \underbrace{\begin{bmatrix} x_{n,1}^0 & x_{n,1}^1 & x_{n,1}^2 & \ldots & x_{n,1}^{23} \\ x_{n,2}^0 & x_{n,2}^1 & x_{n,2}^2 & \ldots & x_{n,2}^{23} \\ \ldots & \ldots & \ldots & x_{n,a}^h & \ldots \\ \ldots & \ldots & \ldots & \ldots & x_{n,|\mathcal{A}_n|}^{23} \end{bmatrix}}_{\text{columns correspond to } h \in \mathcal{H}}$$

  Each element $x_{n,a}^h$ in this matrix can be tabulated from the decision variables defined for the household appliance scheduling program in the previous chapter, and is the energy consumption (in kilowatt-hours) by appliance a in hour h for household n. Set X_n is the set of all schedules that can be constructed for household n's consumption. Each vector of matrices $\mathbf{x} = (\mathbf{x}_1, \ldots, \mathbf{x}_n) \in X$ describes the joint action of all households. We use $\mathbf{x}_{-n}$ to denote the energy consumption scheduling profiles of all households other than household n. Set F is the set of all production schedules $(\mathbf{f}, \sigma)$ available to the dispatch center to play. Each vector $\big(\mathbf{x}, (\mathbf{f}, \sigma)\big) \in X \times F$ is called an *action profile*;

- **Utility Function**: $u = (u_1, u_2, \ldots, u_N, u_P)$, where $u_n : X_n \mapsto \mathbb{R}, \forall n \in \mathcal{N}$ is the real-valued utility function[1] for player n defined above in (3.1), and $u_P : F \mapsto \mathbb{R}$ for the producer P. We define the producer's utility as:

[1] An implicit assumption is made that the set of outcomes O in the game equals the set of actions $X \times F$.

$$u_P(\mathbf{f}, \sigma) = \begin{cases} -\sum_{h\in\mathcal{H}} C\big(\mathbf{f},\, \sigma;\; \mathbf{c}^f,\, \mathbf{c}^v\big) & \text{if } (\mathbf{f}, \sigma) \text{ satisfies (2.2), (2.3), (2.4),} \\ & \text{(2.5), and (2.6)} \\ -\infty & \text{otherwise} \end{cases}$$

That is, any element of set F that satisfies the constraints in (2.2), (2.3), (2.4), (2.5), and (2.6) gives the producer utility equal to the negative of the total daily production cost, which corresponds to the objective value in (2.1) given the action played by the producer. If the action played by producers does not satisfy these constraints, the producer has utility of $-\infty$.

In this game, the supply side (player P) sets a production schedule to *minimize* the total daily generation cost subject to the constraints defined for the economic dispatch problem, or, likewise, maximize his utility. Given that household n is indifferent between playing any action that is a feasible solution to the household appliance scheduling problem, we assume that any rational household would select an action in order to minimize its daily payment to the electric utility for its fixed amount of daily consumption. In this model, the actions of households impact the set of feasible actions available to the producer with bounded utility through the hourly demand constraint imposed on production.

In a later section, we will prove that in our complete market model, households and the producer play the energy consumption game outlined here.

3.2 Fair Billing Axioms

In this section, we explore two properties desired in any retail billing system, which, when these conditions are satisfied, the literature calls a **fair billing** model. The theory laid out here will be important later in analyzing equilibrium solutions for the market. Mohsenian-Rad, Wong *et al.* [10] introduce the following axiomatic requirements:

1. The aggregate daily charge to households must be greater than or equal to the aggre-

gate daily cost of production.

$$\underbrace{\sum_{n\in\mathcal{N}} b_n^{FB}}_{\text{total daily charge}} \geq \underbrace{\sum_{h\in\mathcal{H}} C\big(Q_h^S\big)}_{\text{total daily cost}} \tag{3.3}$$

2. The total daily bill for each household should be linearly proportional to its total daily load: $\sum_{h\in\mathcal{H}} q_{n,h}^D$, or equivalently, $\sum_{h\in\mathcal{H}}\sum_{a\in\mathcal{A}_n} x_{n,a}^h$.

$$b_n^{FB} \propto \sum_{h\in\mathcal{H}}\sum_{a\in\mathcal{A}_n} x_{n,a}^h, \quad \forall n\in\mathcal{N} \tag{3.4}$$

The constant of proportionality in this expression is common for all households.

In our model, we desire to set hourly prices in order to achieve a fair billing, that is, an outcome where each household pays for only its proportion of the aggregate social cost of production and does not subsidize any other subscriber's consumption. From an economic efficiency perspective, fair billing captures the true cost of consumption for each subscriber.

Utilizing the axioms above, we may derive the general expression for each household's total daily energy bill under any fair billing system. Manipulating (3.4) leads to the following equality:

$$\frac{b_m^{FB}}{b_n^{FB}} = \frac{\sum_{h\in\mathcal{H}}\sum_{a\in\mathcal{A}_m} x_{m,a}^h}{\sum_{h\in\mathcal{H}}\sum_{a\in\mathcal{A}_n} x_{n,a}^h}, \quad \forall n,m\in\mathcal{N}.$$

After rearranging, we sum over all households $m\in\mathcal{N}$:

$$b_m^{FB} = b_n^{FB}\frac{\sum_{h\in\mathcal{H}}\sum_{a\in\mathcal{A}_m} x_{m,a}^h}{\sum_{h\in\mathcal{H}}\sum_{a\in\mathcal{A}_n} x_{n,a}^h}, \quad \forall n,m\in\mathcal{N},$$

$$\sum_{m\in\mathcal{N}} b_m^{FB} = \sum_{m\in\mathcal{N}} b_n^{FB}\frac{\sum_{h\in\mathcal{H}}\sum_{a\in\mathcal{A}_m} x_{m,a}^h}{\sum_{h\in\mathcal{H}}\sum_{a\in\mathcal{A}_n} x_{n,a}^h}, \quad \forall n\in\mathcal{N},$$

$$\sum_{m\in\mathcal{N}} b_m^{FB} = b_n^{FB}\frac{\sum_{m\in\mathcal{N}}\sum_{h\in\mathcal{H}}\sum_{a\in\mathcal{A}_m} x_{m,a}^h}{\sum_{h\in\mathcal{H}}\sum_{a\in\mathcal{A}_n} x_{n,a}^h}, \quad \forall n\in\mathcal{N}.$$

We rearrange the expression in terms of b_n^{FB}:

$$b_n^{FB} = \frac{\sum_{h\in\mathcal{H}}\sum_{a\in\mathcal{A}_n} x_{n,a}^h}{\sum_{m\in\mathcal{N}}\sum_{h\in\mathcal{H}}\sum_{a\in\mathcal{A}_m} x_{m,a}^h}\sum_{m\in\mathcal{N}} b_m^{FB}, \quad \forall n\in\mathcal{N}.$$

Next, by rearranging inequality (3.3), we note that the ratio of the left-hand side to the right-hand side is the same as the gross profit margin κ, as defined in the previous chapter:

$$\kappa = \frac{\sum_{n\in\mathcal{N}} b_n^{FB}}{\sum_{h\in\mathcal{H}} C\big(Q_h^S\big)} \geq 1.$$

When we set $\kappa \equiv 1$, we may refer to the billing system as *budget balanced* [10], indicating that the revenue raised from households equals the aggregate daily cost of production. Finally, by rearranging and substituting this into the expression above, we find that:

$$b_n^{FB} = \kappa \times \frac{\sum_{h\in\mathcal{H}} \sum_{a\in\mathcal{A}_n} x_{n,a}^h}{\sum_{m\in\mathcal{N}} \sum_{h\in\mathcal{H}} \sum_{a\in\mathcal{A}_m} x_{m,a}^h} \times \sum_{h\in\mathcal{H}} C\big(Q_h^S\big), \quad \forall n \in \mathcal{N}. \tag{3.5}$$

In the billing expression above, each household pays for its own fraction of the aggregate daily cost imposed to the system, as is desired.

From our derivation, we see that any billing system that meets axioms (3.3) and (3.4) takes the form of equation (3.5). The first term in the right-hand side product in (3.5) is the producer's gross markup, κ, followed by the total daily load for household n as a fraction of the aggregate daily load, a ratio that is a constant term under our modeling assumptions.

The billing system derived in (3.5) is ideal because under such a system, each household must pay for its own fraction of the aggregate cost, $\sum_{h\in\mathcal{H}} C\big(Q_h^S\big)$. However, we still require the determination of hourly retail prices that can be published ahead of time to households by electric utilities. Fair billing is a distributional principle we desire from our prices, describing the incidence of cost to households, but does not imply any ex ante pricing mechanism in and of itself. The challenge in retail distribution is to devise prices that produce an ex post outcome of fair billing for all households.

Finally, we conclude this section by demonstrating that household payments to the electric utility in our retail market can be re-expressed as a fair billing system, and under conditions we derive later in this chapter, the payments under our proposed system will be equivalent to fair billing payments for all households.

Starting with each household's payment to the electric utility, defined earlier in (3.2):

$$b_n = \sum_{h\in\mathcal{H}} \frac{C\big(Q_h^S\big)}{Q_h^S} \cdot q_{n,h}^D,$$

we sum over all households $n \in \mathcal{N}$:

$$\sum_{n \in \mathcal{N}} b_n = \sum_{n \in \mathcal{N}} \sum_{h \in \mathcal{H}} \left[\frac{C\big(Q_h^S\big)}{Q_h^S} \cdot q_{n,h}^D \right].$$

We rearrange the double summation, which does not affect the sum, and reduce:

$$= \sum_{h \in \mathcal{H}} \left[\frac{C\big(Q_h^S\big)}{Q_h^S} \times \sum_{n \in \mathcal{N}} q_{n,h}^D \right],$$

$$= \sum_{h \in \mathcal{H}} \left[C\big(Q_h^S\big) \times \frac{Q_h^D}{Q_h^S} \right].$$

The ratio $\frac{Q_h^D}{Q_h^S}$ derived here, quantity demanded to quantity supplied, is a measure of over-production in the market for hour h. On the other hand, if this ratio is > 1, there is a shortage of production and demand will not be met. At the market equilibrium, $Q_h^D = Q_h^S$.

Without affecting the sum, we may multiply this by a factor of 1:

$$= \frac{\sum_{h \in \mathcal{H}} Q_h^D}{\sum_{h \in \mathcal{H}} Q_h^D} \times \sum_{h \in \mathcal{H}} \left[C\big(Q_h^S\big) \times \frac{Q_h^D}{Q_h^S} \right].$$

Substituting and rearranging,

$$= \frac{\sum_{h \in \mathcal{H}} \sum_{n \in \mathcal{N}} q_{n,h}^D}{\sum_{h \in \mathcal{H}} Q_h^D} \times \sum_{h \in \mathcal{H}} \left[C\big(Q_h^S\big) \times \frac{Q_h^D}{Q_h^S} \right],$$

$$= \frac{\sum_{n \in \mathcal{N}} \sum_{h \in \mathcal{H}} q_{n,h}^D}{\sum_{h \in \mathcal{H}} Q_h^D} \times \sum_{h \in \mathcal{H}} \left[C\big(Q_h^S\big) \times \frac{Q_h^D}{Q_h^S} \right],$$

we find that the total payment from all households is equal to:

$$= \sum_{n \in \mathcal{N}} \left(\frac{\sum_{h \in \mathcal{H}} q_{n,h}^D}{\sum_{h \in \mathcal{H}} Q_h^D} \times \sum_{h \in \mathcal{H}} \left[C\big(Q_h^S\big) \times \frac{Q_h^D}{Q_h^S} \right] \right).$$

Comparing this to (3.5), and if we assume for a moment that $Q_h^D = Q_h^S$, $\forall h \in \mathcal{H}$, we see that the total amount collected from all households by our pricing scheme is equivalent to the total collected under a fair billing scheme. The first factor in each term in the summation is household n's total daily consumption as a fraction of total aggregate daily consumption, which is by our model a constant. The second factor, the internal sum, is the total payment from all households. As this expression implies, for each household n, we desire to structure a billing system such that each household pays for only its fraction of total consumption.

The expression derived above corresponds to a fair billing system, and shows that under certain conditions, our proposed pricing scheme will match the outcome desired from a fair billing system. However, we have yet to determine the required conditions under which each term of the sum will be equal to the payment expected of the respective household under fair billing. Later in this chapter, we prove that those required conditions are satisfied by the structure of our model at the optimal equilibrium for the market.

3.3 Nash Equilibrium

In this section, we provide a proof of our main result, that under certain conditions, the market where we implemented communication between the producer and households converges to a Nash equilibrium for the producer as well as for households. We first prove the existence of Nash equilibrium under stated assumptions, then determine a set of minimal conditions under which an equilibrium always exists. In order to construct this proof, we must also first establish and prove several lemmas important to our main result.

3.3.1 Perfect Production Smoothing

We begin by showing a few theoretical results on the socially optimal joint consumption profile for households. We first construct a relaxation of the household problem in which appliances and tasks can be freely rescheduled even outside of the bounds $[\alpha_j, \beta_j]$ for each submitted task $j \in \mathcal{J}_n, \forall n \in \mathcal{N}$. By considering this unconstrained relaxation of the problem, we may begin to build intuition as to the socially optimal strategy for households and the producer, and establish optimality bounds on the constrained problem. Constraints limiting each appliance to an hour of use across all of its power settings for each time period still hold and so our relaxation still characterizes an implementable schedule for tasks assigned to appliances.

Consider a strategy we propose for production called **perfect production smoothing**.

Definition 1 (Perfect Production Smoothing)**.** Perfect production smoothing (PPS) implies that the producer sets a production schedule that minimizes the variance in hourly production. Thus,

$$\mathbf{Q}^{PPS} = \arg\min_{\mathbf{Q}^S} \sum_{h \in \mathcal{H}} \frac{1}{H} \Big(Q_h^S - \bar{Q}_h^S \Big)^2,$$

where $\bar{Q}_h^S = \frac{1}{H} \sum_{h \in \mathcal{H}} Q_h^S$. If we take aggregate hourly production to be equal to aggregate hourly demand, thereby assuming there is no overproduction in any time period, a desired property, this implies that total aggregate daily production is also a constant, and therefore so is $\bar{Q}_h^S$, the average aggregate production per hour. Thus, in perfect production smoothing, the producer sets schedules such that an equal portion of production is carried out in each hour. This assumes that by doing so, production can load shape demand into smoothing perfectly across each period. Perfect production smoothing depends on this relaxation we propose for the household problem.

For notational shorthand, we define the aggregate hourly production given perfect production smoothing in each hour as Q^{PPS}.

$$Q^{PPS} = Q_0^S = Q_1^S = \ldots = Q_h^S = \ldots = Q_{H-1}^S = \frac{\sum_{h \in \mathcal{H}} Q_h^S}{H}$$

Lemma 1. *If the hourly production cost function $C(\cdot)$ is strictly convex and the same for each hour and by assuming a relaxation of the household appliance scheduling problem, for the producer to perfectly smooth electricity production maximizes his utility in Game 1.*

Proof. From the previous chapter, we showed that our hourly production cost curve $C_h(\cdot)$ is the same in each time period:

$$C_h(\cdot) = C_0(\cdot) = \ldots = C_{H-1}(\cdot) = C(\cdot).$$

We also demonstrated that given our assumptions for fixed and variable costs, $C(\cdot)$ is strictly convex:

$$C(\theta \cdot \hat{Q}_h^S + (1-\theta) \cdot \tilde{Q}_h^S) \leq \theta \cdot C(\hat{Q}_h^S) + (1-\theta) \cdot C(\tilde{Q}_h^S), \quad \forall \hat{Q}_h^S, \tilde{Q}_h^S \geq 0,\, 0 \leq \theta \leq 1.$$

Producer utility is maximized when the total costs of production are minimized. If perfect production smoothing minimizes aggregate daily production cost, then the following must hold:

$$\sum_{h \in \mathcal{H}} C(Q^{PPS}) \leq \sum_{h \in \mathcal{H}} C(Q_h^S)$$

for all aggregate production profiles $\mathbf{Q}^S = [Q_0^S, Q_1^S, \ldots, Q_{H-1}^S] \in F$, where F is the producer's action space, such that:

$$\sum_{h \in \mathcal{H}} Q_h^S = \sum_{h \in \mathcal{H}} Q_h^D \quad \text{and} \quad \mathbf{Q}^S \neq \mathbf{Q}^{PPS}.$$

We begin by simplifying the left-hand side summation and rearranging terms. Given that $C(Q^{PPS})$ is the same for each hour:

$$C(Q^{PPS}) \sum_{h \in \mathcal{H}} 1 \leq \sum_{h \in \mathcal{H}} C(Q_h^S),$$

$$C(Q^{PPS}) \cdot H \leq \sum_{h \in \mathcal{H}} C(Q_h^S),$$

$$C(Q^{PPS}) \leq \sum_{h \in \mathcal{H}} \frac{1}{H} \cdot C(Q_h^S).$$

We know from its definition that $Q^{PPS} = \sum_{h \in \mathcal{H}} \frac{1}{H} \cdot Q_h^S$. Substituting:

$$C\left(\sum_{h \in \mathcal{H}} \frac{1}{H} \cdot Q_h^S\right) \leq \sum_{h \in \mathcal{H}} \frac{1}{H} \cdot C(Q_h^S).$$

Given that $\sum_{h \in \mathcal{H}} \frac{1}{H} = 1$ and taking $C(\cdot)$ to be strictly convex, the above expression is true by Jensen's Inequality, that for any convex function $\phi(\cdot)$ the following holds true:

$$\phi\left(\sum_{i=1}^{n} a_i \cdot x_i\right) \leq \sum_{i=1}^{n} a_i \cdot \phi(x_i), \text{ where } \sum_{i=1}^{n} a_i = 1.$$

Thus, we have proven that if the producer perfectly smoothes electricity generation in each period, this strategy maximizes producer utility. □

An important result to discuss is that this solution is also an absolute lower bound on the total aggregate cost of production for the constrained household problem. Given that the production cost curve is the same for each hour and convex, the producer can do no better than when demand can be perfectly smoothed and Q_h^S is the same for each hour. In treating these problems as endogenous, it is important to note that the time range constraints imposed in the household problem restrict feasible values to the vector

$\mathbf{Q}^D$. Hourly demand enters the economic dispatch problem via the production constraints, specifically through (2.4), and thus this will determine which production schedules in the set F are feasible solutions given the production constraints.

Under perfect substitutability between hours for demand, the production smoothing solution derived above is the absolute minimum for total production cost. We can then see why it is favorable for production and social welfare. Additionally, it is also desirable for the demand side, as we prove in the following lemma: perfect demand smoothing in the aggregate, which corresponds to perfect production smoothing, guarantees a fair billing outcome for every household.

Lemma 2. *If there exists a solution to the household problem that perfectly smoothes aggregate demand in each time period, the payment made by each household is guaranteed to be the payment made under a fair billing system.*

Proof. We will show that payments under our billing are equivalent to a fair billing for every household by showing that the two expressions are equal. At the solution proven optimal in the previous lemma, the components of $\mathbf{Q}^S$ take the following values:

$$Q^{PPS} = Q_0^S = Q_1^S = \ldots = Q_h^S = \ldots = Q_{H-1}^S = \frac{\sum_{h\in\mathcal{H}} Q_h^S}{H}.$$

As we had assumed in the previous lemma, $Q_h^S = Q_h^D$, $\forall h \in \mathcal{H}$, thus perfect production smoothing also implies demand smoothing.

We begin with expressions for b_n and b_n^{FB}, in terms of Q_h^S and $q_{n,h}^D$ and substitute in our optimal values for production:

$$\sum_{h\in\mathcal{H}} \frac{C(Q_h^S)}{Q_h^S} \cdot q_{n,h}^D = b_n \stackrel{?}{=} b_n^{FB} = \frac{\sum_{h\in\mathcal{H}} q_{n,h}^D}{\sum_{h\in\mathcal{H}} Q_h^S} \cdot \sum_{h\in\mathcal{H}} C(Q_h^S),\ \forall n \in \mathcal{N},$$

$$\sum_{h\in\mathcal{H}} \frac{C(Q^{PPS})}{Q^{PPS}} \cdot q_{n,h}^D \stackrel{?}{=} \frac{\sum_{h\in\mathcal{H}} q_{n,h}^D}{\sum_{h\in\mathcal{H}} Q^{PPS}} \cdot \sum_{h\in\mathcal{H}} C(Q^{PPS}),\ \forall n \in \mathcal{N}.$$

Because the quantities $C(Q^{PPS})$ and Q^{PPS} are all independent of h, this simplifies to the following:

$$\frac{C(Q^{PPS})}{Q^{PPS}} \cdot \sum_{h\in\mathcal{H}} q_{n,h}^D = \frac{\cancel{H} \cdot C(Q^{PPS})}{\cancel{H} \cdot Q^{PPS}} \cdot \sum_{h\in\mathcal{H}} q_{n,h}^D.$$

The sum $\sum_{h \in \mathcal{H}} q_{n,h}^D$ is the total daily consumption for household n, a constant. The two sides are equivalent and so, at the global optimum, the payment by each household will be equal to payment under a fair billing system. □

This result provides an additional motivation for why the social welfare maximizing solution found before is desirable for households as well. We have already demonstrated that the result corresponding to perfect production smoothing forms an absolute lower bound on the social cost of production. Additionally, as a result of this solution, each household pays for only its fraction of total aggregate daily consumption. Payments for households align with the fair billing axioms. At other solutions for $\mathbf{Q}^D$, a household may inadvertently subsidize another household's daily consumption because of discrepancies in hourly prices.

In order to prove that this solution is a Nash equilibrium in the market, it is sufficient to show that $\mathbf{Q}^D$ and the underlying consumption schedules for each household $\mathbf{x}_n$, $\forall n \in \mathcal{N}$, are still feasible in the household problem when time range constraints are reinstated. The solution we found is utility maximizing for the producer and if it is a feasible solution for the household appliance scheduling problem, then every household also maximizes its respective utility function when playing this solution. The solution of the form $\big(\mathbf{x}, (\mathbf{f}, \sigma)\big)$ that corresponds to perfect demand and production smoothing is then a Nash equilibrium for the producer and households.

Ordinarily, however, time range constraints prevent households from pursuing a perfect demand smoothing strategy: $\mathbf{Q}^{PPS}$ may not always be a feasible solution due to the time range constraints for demand. This solution to the relaxation problem still provides some important intuition as to the nature of the optimal solution for the market: a solution that is variance-minimizing in hourly consumption and production will draw the market closer to the perfect production smoothing solution we proposed. In general, the producer is able to minimize his total daily costs of production when fluctuations in hourly production, which correspond to fluctuations in hourly demand, are minimized.

3.3.2 Nash Equilibrium under Complete Information

Though it may not be possible for households to perfectly smooth demand, we can construct a solution that optimizes the economic dispatch problem and also impose the constraints corresponding to the household appliance scheduling problem. If a solution exists to this joint program, we know it to be feasible for the household problem and cost-minimizing for the producer. The optimal solution to this joint problem, as we will prove, corresponds to a pure-strategy Nash equilibrium for the consumption game played between the producer and households.

In order to find the minimum cost solution for the producer that is feasible for households, we introduce an endogenous program for economic dispatch that is subject to the household appliance scheduling constraints as well.

In short, this program can be expressed as follows:

$$\min_{\mathbf{x},\mathbf{f},\sigma} \Big\{ (2.1) \Big\}$$

subject to the following constraints from the economic dispatch problem: (2.2), (2.3), (2.5), and (2.6). We also impose the constraints from the household problem: (2.8), (2.9), and (2.10). In order to couple these two problems, we present an endogenous formulation for the economic dispatch constraint for meeting hourly demand, previously formulated as (2.4), utilizing the endogenous definition for demand we presented at the end of the last chapter:

$$\sum_{g \in \mathcal{G}} \sum_{\tau \in \mathcal{T}(g)} f^h_{g,\tau} \geq Q^D_h(\mathbf{x}), \, \forall h \in \mathcal{H}. \tag{3.6}$$

The solution to this mixed-integer program sets values for $\mathbf{x}, \mathbf{f},$ and σ that are optimal for the economic dispatch problem and also feasible solutions to the household problem.

The full program, which we refer to as the economic dispatch program with endogenously-defined demand and subject to the constraints from the household problem or more compactly as the *joint program*, is defined here:

$$\min_{\mathbf{x},\mathbf{f},\sigma} \left\{ \sum_{g \in \mathcal{G}} c_g^f \cdot \sigma_g + \sum_{h \in \mathcal{H}} \sum_{g \in \mathcal{G}} \sum_{\tau \in \mathcal{T}(g)} c_{g,\tau}^v \cdot f_{g,\tau}^h \right\} \tag{3.7}$$

subject to:

$$\sum_{h \in \mathcal{H}} \sum_{\tau \in \mathcal{T}(g)} f_{g,\tau}^h \leq M_g \cdot \sigma_g, \ \forall g \in \mathcal{G}$$

$$q_g^{min} \cdot \sigma_g \leq \sum_{\tau \in \mathcal{T}(g)} f_{g,\tau}^h, \ \forall g \in \mathcal{G}, \ h \in \mathcal{H}$$

$$0 \leq f_{g,\tau}^h \leq q_{g,\tau}^{max}, \forall \tau \in \mathcal{T}(g), \ \forall g \in \mathcal{G}, \ h \in \mathcal{H}$$

$$\sigma_g \in \{0,1\}, \ \forall g \in \mathcal{G}$$

$$\sum_{j \in O(i)} x_{ij} = 1, \quad \forall i \in X$$

$$\sum_{i \in I(j)} x_{ij} = 1, \quad \forall j \in Y$$

$$0 \leq x_{ij} \leq 1, \quad \forall (i,j) \in E$$

$$\sum_{g \in \mathcal{G}} \sum_{\tau \in \mathcal{T}(g)} f_{g,\tau}^h \geq Q_h^D(\mathbf{x}), \ \forall h \in \mathcal{H}$$

The constraints in this problem correspond to the constraints and formulation introduced in Chapter 2. Given this complete program, we may solve for the optimal solution for the decision variables in the original economic dispatch program as well as the variables introduced from the household problem, corresponding to demand.

Given a solution to this program, we may prove our main result:

Theorem 1. *An optimal solution* $\left(\mathbf{x}^*, (\mathbf{f}^*, \sigma^*)\right)$ *to the joint program defined above is a pure-strategy Nash equilibrium for Game 1.*

Proof. If the solution $\big(\mathbf{x}^*, (\mathbf{f}^*, \sigma^*)\big)$ is optimal for the joint program defined above, then $(\mathbf{f}^*, \sigma^*)$ must satisfy all of the constraints for the economic dispatch program and must be a feasible solution for that problem. Likewise, if $\mathbf{x}^*$ satisfies the constraints above, it is a feasible solution for the household appliance scheduling program.

Because we know the solution to be feasible and given that $(\mathbf{f}^*, \sigma^*)$ is found by minimizing the objective value in (3.7), the aggregate daily cost of production, then the solution maximizes the producer's utility:

$$u_P(\mathbf{f}^*, \sigma^*) \geq u_P(\mathbf{f}', \sigma'), \quad \forall (\mathbf{f}', \sigma') \in F : (\mathbf{f}', \sigma') \neq (\mathbf{f}^*, \sigma^*).$$

Also, given that $\mathbf{x}^*$ is feasible, all households have positive and maximum utility:

$$u_n(\mathbf{x}_n^*) \geq u_n(\mathbf{x}_n'), \, \forall \mathbf{x}_n' \in X_n : \mathbf{x}_n' \neq \mathbf{x}_n^*, \, \forall n \in \mathcal{N}.$$

Having shown that this solution, corresponding to an action profile for the game, maximizes utility for both the producer and households, we haven proven that the solution $\big(\mathbf{x}^*, (\mathbf{f}^*, \sigma^*)\big)$ corresponds to a pure-strategy Nash equilibrium for Game 1. □

Given our formulation for this joint problem, we can furthermore prove that a pure-strategy Nash equilibrium always exists for Game 1.

Theorem 2. *If there is sufficient generation capacity available to meet demand in each period, then a Nash equilibrium always exists for Game 1.*

Proof. To prove this, it is sufficient to show that the set of feasible solutions to our joint program is nonempty.

In Chapter 2, we established necessary and sufficient conditions for tasks submitted by users for the household problem that guarantees that a feasible solution to the household problem always exist: a feasible solution $\mathbf{x}$ always exists that satisfies the constraints for the household problem. For the economic dispatch problem, we assume that the producer always has infinite generation capacity available from which to choose generating plants to turn on for production. Thus, there exists at least one feasible solution to the joint problem that serves the demand set by $\mathbf{x}$ in each hour.

By Corollary 2.3 in Bertsimas, for any linear programming problem of minimizing an objective over a nonempty polyhedron, either the optimal cost is equal to $-\infty$ or there exists an optimal solution [1]. Given that the cost vectors $\mathbf{c}^f$ and $\mathbf{c}^v$ in (3.7) are strictly positive and that we have imposed non-negativity constraints on the variables in $\mathbf{f}$, there must exist a finite optimum value for the objective.

Applying Theorem 1, we know that an optimal solution to the joint problem corresponds to a Nash equilibrium for Game 1. □

We have proven here that given our assumptions made in the previous chapter on the set of jobs and respective time range constraints accepted for scheduling and the total generation capacity available to the producer, there always exists a pure-strategy Nash equilibrium for Game 1.

Finally, we prove formally that the interaction between households and the producer in Game 1 models the retail market we constructed at the end of the previous chapter, if there is complete information across the market.

Theorem 3. *If a mechanism is introduced to the retail market such that households adjust consumption in response to prices communicated from the supply side and the producer is able to optimize the production schedule given complete information about hourly demand from households, the behavior of players in this ideal market will be captured by Game 1.*

Proof. The set of players defined for Game 1 corresponds to the market participants in the retail market, households belonging to set $\mathcal{N}$ and the producer. In the retail market, given hourly prices, household n minimizes its daily payment for consumption, which is captured in the objective (2.7), $\forall n \in \mathcal{N}$. A joint optimal solution $\mathbf{x}^*$ to this program satisfies the constraints in the household problem, and so for households to optimize their objective in (2.7) corresponds to playing an action in Game 1 that maximizes household utility.

The producer, given demand in each period from households, finds an optimal solution to the economic dispatch problem. Given that an optimal solution to the economic dispatch problem satisfies its constraints, by optimizing the total production cost in (2.1), the producer maximizes his own utility function from Game 1.

Thus, we have shown that the consumption game introduced in this chapter describes the behavior of the market under complete information. □

By introducing corollaries to Theorem 1, we conclude the chapter by characterizing properties about the Nash equilibrium that is guaranteed for the energy consumption game by the results above.

Corollary 1 (to Theorem 1)**.** The optimal solution $\big(\mathbf{x}^*, (\mathbf{f}^*, \sigma^*)\big)$ corresponds to the Nash equilibrium for Game 1 that maximizes social welfare.

Proof. The proof almost directly follows from the statements established earlier. We recall that the solution is found by minimizing the total aggregate daily cost of production, which thereby maximizes social welfare by our definition. □

Corollary 2 (to Theorem 1)**.** The optimal solution $\big(\mathbf{x}^*, (\mathbf{f}^*, \sigma^*)\big)$ is a Pareto efficient outcome for Game 1. That is, no deviation from this solution is both a weakly better outcome for all players and a strictly better outcome for some player.

Proof. The optimal solution found to the joint problem minimizes cost in the economic dispatch problem and also satisfies the constraints in the household problem. Thus, it maximizes utility for both the producer and households. Given that households are indifferent between any two actions that both satisfy the constraints in the household problem, we cannot strictly improve the utility of any household. Likewise, the solution $\big(\mathbf{x}^*, (\mathbf{f}^*, \sigma^*)\big)$ minimizes the objective value in (3.7), and thus maximizes the utility of the producer in Game 1 subject to the constraints of the economic dispatch problem and household problem. To improve the objective value in the joint problem requires a relaxation of one of the constraints, which makes another player strictly worse off. Thus, the solution found is Pareto optimal. □

Corollaries 1 and 2 provide further insight into why we care about the Nash equilibrium found for Game 1. Though it is uncertain whether $\big(\mathbf{x}^*, (\mathbf{f}^*, \sigma^*)\big)$ is a unique Nash equilibrium to our game, this solution corresponds to the equilibrium that maximizes social welfare, a property desired in improving efficiency in the market for electricity generation.

We conclude this chapter with a brief discussion and summary of our result, as well as motivation for our work in the remaining chapters.

3.3.3 Discussion

In this chapter, we determined the existence of a Nash equilibrium for Game 1, describing the interaction of households and the producer in the retail electricity market under complete information. We show that given our assumptions made for the complete market model, a social welfare-maximizing Nash equilibrium is guaranteed to exist and can be found by solving the joint program we construct in this chapter, if the full set of constraints for households is known to the producer. However, in practice, it is generally not possible for the producer to solve this joint program because of hidden information in the market that prevents the producer from having full knowledge of every household's constraints.

In our proposed model for the complete market, we had structured a communication system by which the producer provides real-time prices to households, corresponding to the true costs of electricity generation. After best responding to hourly price changes, households communicate their aggregate consumption profile, the joint action profile for households $\mathbf{x} = (\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_N)$, back to the supply side where it is used to update cost and retail prices. In the next computational iteration of this mechanism to compute equilibrium, once the updated prices are sent back to consumers, households again resolve their optimization program. By communicating prices and the joint consumption profile of households, we may improve the outcome of the market to approximate the equilibrium found in the market under complete information. Our intent in structuring this mechanism is to provide a minimal transfer of information between the two sides of the market that we hope is sufficient for convergence or approximation of the market equilibrium.

In the subsequent chapter, we develop and explore algorithms for determining equilibrium prices and production quantities that utilize the communication of prices and hourly demand across the market. Our overarching motivation has been to increase efficiency in the market given minimal information transfer between participants and preserving the property of individual rationality for household participation. It is unrealistic in practice that households would truthfully communicate the complete set of constraints for their optimization problem to the producer.

Finally, in general it is computationally difficult to solve for Nash equilibria, particularly for a game with potentially tens of thousands of households acting simultaneously. Given the nature of real-time markets and production decision-making for electricity generation, we require the ability to converge to a stable equilibrium within a relatively quick running time of our model, such that it remains practical for real-world implementation for economic dispatch. While we prove the existence of a Nash equilibrium in this chapter under complete information, we now turn our attention to tractably approximating it in the next chapter.

Chapter 4

Computing Equilibria

The previous chapter provides a proof of the existence of a Nash equilibrium in the retail market under complete information. In this chapter, we now turn to computing equilibrium prices and production quantities given only the restricted information transfer between market players defined by our complete model for the retail market. In our model, we intend to utilize the information about aggregate hourly demand communicated to the producer to devise real-time prices for households to push household consumption toward the market equilibrium quantities. In the next chapter, we test this model empirically to compare performance against outcomes in the market under complete information.

We begin this chapter with a reformulation of our model from Chapter 2 that permits large-scale computability and provides a framework for simulation of our mechanism. At the end of this chapter, we utilize our network flow formulation for the complete market model to propose algorithms for computing equilibria.

4.1 Network Flows

Network flow problems are among the most frequently solved linear programming problems in practice, arising from the design of communication, transportation, and logistics networks. Their structure provides a substantial simplification in terms of modeling constraints, and can be solved efficiently using any general algorithms for linear programming.

We apply a separate network flow formulation that permits scalability to inputs and allows us to efficiently solve for each of the economic dispatch and household appliance scheduling problems. Our programs be easily solved for their optimal values at each iteration, permitting the rapid iterative communication between both programs, as outlined in Chapter 2, that is necessary for the process of converging to equilibrium.

For each of these programs, we apply a capacitated minimum cost network flow design.

4.2 Economic Dispatch Problem Revisited

We define the supply side economic dispatch problem, utilizing network flows, as follows: let $G^S = (V^S, E^S)$ be a directed graph on a set of nodes V^S and a set of edges E^S. G^S is 4-partite graph consisting of four partitions in the set of nodes V^S. The first partition, V_1^S, consists of H nodes denoting elements of the set $\mathcal{H}$, the set of time periods in the day. The second partition, V_2^S, consists of nodes for every unique element $(g, \tau, h) \in \mathcal{G} \times \mathcal{T}(g) \times \mathcal{H}$, the complete set of all generator modes crossed with hours in the day for production. The third partition V_3^S consists of nodes corresponding to every plant crossed with every hour, $(g, h) \in \mathcal{G} \times \mathcal{H}$. Finally, partition V_4^S consists of nodes for every element of set $\mathcal{G}$, the full set of generating plants available to the producer for scheduling production. Table 4.1 summarizes.

Edges, which connect nodes across partitions, enable the flow of produced electricity from the nodes denoting each hour of demand in V_1^S to nodes representing each generating plant that produces electricity to meet that demand in V_4^S. Paths through this network determine which plant and at which production mode electricity generation will be scheduled in order to meet hourly demand.

We formally define each partition in the set of edges E^S in set-builder notation as

$$E_{AB}^S \equiv \Big\{(i, j) \in V_A^S \times V_B^S : \rho(i, j)\Big\},$$

for pairs of adjacent partitions V_A^S and V_B^S in V^S. Furthermore, the statement above says there exists arc $(i, j) \in E_{AB}^S$, and in the graph, if and only if the function $\rho(i, j)$ evaluates as true at i, j.

Between node partitions V_1^S and V_2^S, directed arc $(v_1, v_2) \in V_1^S \times V_2^S$ is a part of graph G^S if nodes v_1 and v_2 are indexed by the same hour. This implies that production to satisfy demand in hour h can be scheduled at any production mode only in the same time period.

$$E_{12}^S \equiv \Big\{ \big(\underbrace{h}_{\text{node } v_1} \quad , \quad \underbrace{(g, \tau, h')}_{\text{node } v_2} \big) \in V_1^S \times V_2^S : h = h' \Big\}$$

Between elements of node partitions V_2^S and V_3^S, directed arc $(v_2, v_3) \in V_2^S \times V_3^S$ is a part of the graph if again the two nodes share the same element from $\mathcal{H}$ as well as the same element in $\mathcal{G}$. These edges collapse the production modes for every plant into a single node in V_3^S for that plant g, for every hour.

$$E_{23}^S \equiv \Big\{ \big((g, \tau, h), (g', h') \big) \in V_2^S \times V_3^S : g = g' \text{ and } h = h' \Big\}$$

Finally, for elements in node partitions V_3^S and V_4^S, the directed arc (v_3, v_4) is a part of the graph if these two nodes contain the same element of set $\mathcal{G}$:

$$E_{34}^S \equiv \Big\{ \big((g, h), g' \big) \in V_3^S \times V_4^S : g = g' \Big\}.$$

This set of edges collapses the nodes for each hour of production at a plant into a single node for that generating plant, $\forall g \in \mathcal{G}$.

In order to formulate this problem as a minimum cost flow problem on graph G^S, we add a source node s and a sink node t to G^S, as partitions V_0^S and V_5^S, respectively. To connect these nodes to the graph, we draw directed arcs from the source to every element in node partition V_1^S, labeling this edge partition as E_{01}^S, and a directed arc from every node in node partition V_4^S to the sink, labeling these arcs as E_{45}^S.

$$E_{01}^S \equiv \{(s, h) \in V_0^S \times V_1^S\}$$

$$E_{45}^S \equiv \{(g, t) \in V_4^S \times V_5^S\}$$

Additionally, we also add a directed arc (t, s), from the sink back to the source, with associated cost $c_{ts} = 0$ and upper bound on flow $u_{ts} = \infty$ in order to model this problem as a circulation. This arc belongs to the set E_{50}^S.

$$E_{50}^S \equiv \{(t, s) \in V_5^S \times V_0^S\}$$

To summarize, the full set of edges E^S in graph G^S is the union over the edges defined above: $E^S = E^S_{01} \cup E^S_{12} \cup E^S_{23} \cup E^S_{34} \cup E^S_{45} \cup E^S_{50}$. Likewise, $V^S = V^S_0 \cup V^S_1 \cup V^S_2 \cup V^S_3 \cup V^S_4 \cup V^S_5$.

As notation, we reintroduce a few other terms and sets we have used previously in Chapter 2 that will again be useful in our formulation below. For any directed arc (i, j), we refer to i as the start node and j as the end node. Given this, let $I(i)$ be the set of all start nodes of arcs incoming to node i, and let $O(i)$ be the set of end nodes of arcs outgoing from node i. More formally, for any directed graph $G = (V, E)$, we define $I(i) \equiv \{j \in V \mid (j, i) \in E\}$ and $O(i) \equiv \{j \in V \mid (i, j) \in E\}$.

Our *network*, in this case, is the directed graph $G^S = (V^S, E^S)$ we have set up above, along with the following additional numerical information:

- parameters b_i representing the external supply at each node $i \in V^S$,
- parameters u_{ij} representing the maximum capacity of each arc $(i, j) \in E^S$,
- parameters l_{ij} representing the minimum capacity of each arc $(i, j) \in E^S$, and
- parameters c_{ij} representing the cost per unit of flow along arc $(i, j) \in E^S$.

Setting these parameters to edges and nodes in our network allows us to impose flow balance and capacity constraints on flows over G^S, which we may use to reformulate the constraints from Chapter 2 to model the economic dispatch problem.

Our network flow formulation is constructed as a circulation problem, and so $b_i \equiv 0$, $\forall i \in V^S$. Each node is treated as a transshipment node, where the total inflow at the node equals the total outflow. The upper bounds on flow on all arcs are left as infinite, except for arcs belong to E^S_{23}, for which the upper bound on flow u_{ij} is defined as $q^{max}_{g,\tau}$ from our original formulation of the economic dispatch problem. For every arc $\big((g, \tau, h), (g, h)\big) \in E^S_{23} : g \in \mathcal{G}$, $\tau \in \mathcal{T}(g)$, we define $u_{(g,\tau,h)(g,h)} \equiv q^{max}_{g,\tau}$. This upper bound on flow constrains electricity generation at each production mode τ belonging to generating plant $g \in \mathcal{G}$ to its marginal production capacity.

For every arc in E^S_{45}, we also impose an upper bound on flow, for the maximum total daily generation capacity at each generating plant,

$$u_{gt} \equiv \begin{cases} M_g & \text{if } \sigma_g = 1 \\ 0 & \text{if } \sigma_g = 0 \end{cases}, \quad \forall (g,t) \in E_{45}^S : g \in \mathcal{G},$$

where $\sigma_g \in \{0,1\}$ denotes the decision variable determining whether generating plant g has been turned on or not for production. If the plant has not been turned on, the flow capacity at arc (g,t) is zero, as modeled above. The expression above can be written more compactly as a linear mixed-integer constraint: $u_{gt} \equiv M_g \cdot \sigma_g$, $\forall (g,t) \in E_{45}^S : g \in \mathcal{G}$.

Likewise, we impose lower bounds on these same set of arcs, for the minimum operating level at each plant that is on: $l_{gt} \equiv q_g^{min} \cdot \sigma_g$, $\forall (g,t) \in E_{45}^S : g \in \mathcal{G}$. We also impose lower bounds on the flows over arcs in E_{01}^S, which enable flow from the source node to each hour: $l_{sh} \equiv Q_h^D, \forall (s,h) \in E_{01}^S : h \in \mathcal{H}$. The capacity constraints formed by these lower bounds guarantee that demand is met in each hour.

Finally, we impose the cost values $c_{g,\tau}^v$, $\forall g \in \mathcal{G}$, $\tau \in \mathcal{T}(g)$, on the arcs in E_{23}^S. For arc $\big((g,\tau,h),(g,h)\big) \in E_{23}^S : g \in \mathcal{G}$, $\tau \in \mathcal{T}(g)$, let $c_{(g,\tau,h)(g,h)} \equiv c_{g,\tau}^v$. Fixed costs $c_g^f, \forall g \in \mathcal{G}$ are later incorporated into the objective function.

For reference, Tables 4.1 and 4.2 provide a summary of our notation for graph $G^S = (V^S, E^S)$, as well as the capacity bounds and costs associated with each arc in the constructed network. Additionally, Figure 4.1 provides a visualization of the structure of the graph. Each arc in this figure is labeled with an interval denoting feasible flow along that arc.

By this formulation, we may visualize the network as the flow of electricity on each arc from hourly demand nodes to nodes for each generating plant. We use f_{ij} to denote the amount of flow through arc (i,j). As a part of our formulation, we impose the following conditions on the flow variables f_{ij}, $(i,j) \in E^S$:

$$\sum_{j \in I(i)} f_{ji} = \sum_{j \in O(i)} f_{ij}, \quad \forall\, i \in V^S \tag{4.1}$$

$$l_{ij} \leq f_{ij} \leq u_{ij}, \quad \forall (i,j) \in E^S \tag{4.2}$$

Constraint (4.1) corresponds to a flow balance that ensures circulation in the system

Table 4.1: Nodes in the Network Flow Formulation of the Economic Dispatch Problem

Node Partition	Corresponding Sets from Table 2.1
V_0^S	source
V_1^S	$h \in \mathcal{H}$
V_2^S	$(g, \tau, h) \in \mathcal{G} \times \mathcal{T}(g) \times \mathcal{H}$
V_3^S	$(g, h) \in \mathcal{G} \times \mathcal{H}$
V_4^S	$g \in \mathcal{G}$
V_5^S	sink

Table 4.2: Edges in the Network Flow Formulation of the Economic Dispatch Problem

E_{AB}	$i \in V_A^S$	$j \in V_B^S$	$\rho(i,j) \Longrightarrow (i,j) \in E_{AB}$	u_{ij}	l_{ij}	c_{ij}
E_{01}^S	source	$h \in \mathcal{H}$	all	∞	Q_h^D	0
E_{12}^S	$h \in \mathcal{H}$	$(g, \tau, h') \in \mathcal{G} \times \mathcal{T}(g) \times \mathcal{H}$	if $h = h'$	∞	0	0
E_{23}^S	$(g, \tau, h) \in \mathcal{G} \times \mathcal{T}(g) \times \mathcal{H}$	$(g', h') \in \mathcal{G} \times \mathcal{H}$	if $g = g'$ and $h = h'$	τ	0	$c_{g,\tau}^v$
E_{34}^S	$(g, h) \in \mathcal{G} \times \mathcal{H}$	$g' \in \mathcal{G}$	if $g = g'$	∞	0	0
E_{45}^S	$g \in \mathcal{G}$	sink	all	$M_g \cdot \sigma_g$	$q_g^{min} \cdot \sigma_g$	0
E_{50}^S	sink	source	all	∞	0	0

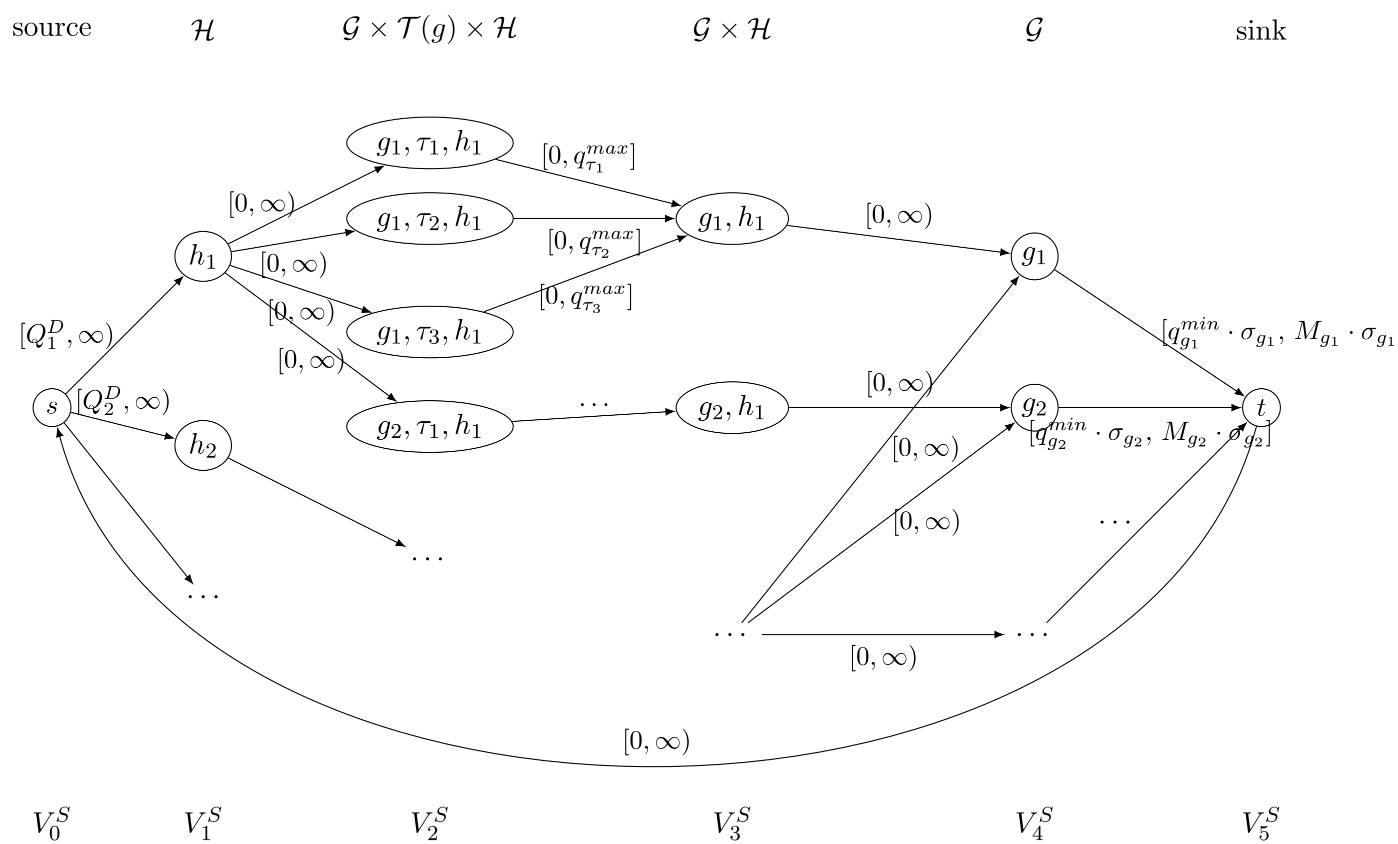

Figure 4.1: Network Flow Formulation for the Economic Dispatch Problem

from the hourly demand nodes in V_1^S, which correspond to hours in the day, to the nodes in V_4^S which denote each generating plant. Constraint (4.2) captures the minimum and maximum capacity for flow on each arc that restricts production scheduling to the marginal capacities for each plant and production mode.

The objective function for the economic dispatch problem is a modified version of (2.1) to utilize the new flow variables we have introduced:

$$\min_{\sigma,\mathbf{f}} \left\{ \sum_{g\in\mathcal{G}} c_g^f \cdot \sigma_g + \sum_{(i,j)\in E_{23}^S} c_{ij} \cdot f_{ij} \right\}. \tag{4.3}$$

The only other constraint that we require apart from this conventional network flow formulation are the unit commitment decision variables for each generating plant, the vector σ, which determines whether a plant is operating or not:

$$\sigma_g \in \{0,1\}, \quad \forall g \in \mathcal{G}. \tag{4.4}$$

Together, objective function (4.3) with constraints (4.1), (4.2), and (4.4) form the network flow formulation of the economic dispatch problem.

Utilizing this program, we may generate a production schedule for each available generator $g \in \mathcal{G}$ for the entire day. Flows along arcs in E_{34}^S contain the production schedule for each generating plant. The value of $f_{(g,h)h}$, $((g,h),h) \in E_{34}^S$ reflects the scheduled production at generating plant g for hour h.

We restate the complete network flow formulation for the economic dispatch program here:

$$\min_{\sigma,\mathbf{f}} \left\{ \sum_{g\in\mathcal{G}} c_g^f \cdot \sigma_g + \sum_{(i,j)\in E_{23}^S} c_{ij} \cdot f_{ij} \right\}$$

subject to:

$$\sum_{j\in I(i)} f_{ji} = \sum_{j\in O(i)} f_{ij}, \quad \forall\, i \in V^S$$
$$l_{ij} \le f_{ij} \le u_{ij}, \quad \forall (i,j) \in E^S$$
$$\sigma_g \in \{0,1\}, \quad \forall g \in \mathcal{G}$$

4.3 Household Appliance Scheduling Revisited

In this section, we outline our network flow formulation for the household appliance scheduling problem. We propose a single centralized optimization model to solve for every household's optimal consumption profile simultaneously, a unique aspect of our formulation. We refer to this program as the aggregate household appliance scheduling program.

As in the previous section, we again consider a capacitated minimum cost network flow design for this model. Let $G^D = (V^D, E^D)$ be the corresponding graph for the demand side problem. G^D is a 6-partite graph, which we again augment with two additional partitions for the source and sink nodes, s and t. The first partition in V^D, which we label as V_1^D, consists of N nodes corresponding to elements of set $\mathcal{N}$, the set of households. The second partition contains a node for every job that must be completed among all households. Nodes in V_2^D appear as tuples of the form $(n, \underbrace{a, m, d, \alpha, \beta}_{\text{job } j \in \mathcal{J}_n})$, with household n, appliance a belonging to household n, and then the 5-tuple defining job $j \in \mathcal{J}_n$ submitted for scheduling on that appliance. The parameter m is the power setting for appliance a, d is the duration of the task, and the interval $[\alpha, \beta]$ defines release and deadline time periods for the task. The cardinality of this partition, $|V_2^D|$ equals $\sum_{n \in \mathcal{N}} |\mathcal{J}_n|$, the total number of jobs submitted for scheduling summed over all households. Formally, we define V_2^D as the cross between $\mathcal{N} \times \mathcal{A}_n \times \mathcal{J}(a)$, if $n \in \mathcal{N}$ and $a \in \mathcal{A}_n$.

The third partition, V_3^D, contains nodes of the form (n, a, m, h), tuples containing household n, appliance a from $\mathcal{A}_n$, a power setting m for the appliance from the set $M(a)$, $a \in \mathcal{A}_n$, and an hour of the day $h \in \mathcal{H}$. Formally, it is the cross of $\mathcal{N} \times \mathcal{A}_n \times \mathcal{M}(a) \times \mathcal{H}$. In the context of the problem, this set corresponds to every possible mode that any appliance for every household can be run in any hour of the day. The next partition, V_4^D, begins to collapse nodes from the previous partition, and contains the cross of households by the set of appliances in each household by hours in the day: $\mathcal{N} \times \mathcal{A}_n \times \mathcal{H}$. There is a node in this partition for every distinct appliance owned by a household for every hour of the day that it can be run. The fifth partition, V_5^D, is the set of nodes representing the cross between households and hours in the day: $\mathcal{N} \times \mathcal{H}$. Partition V_6^D consists of nodes for every element of set $\mathcal{H}$. As before, V_0^D is the source node and V_7^D is the sink.

We apply analogous notation to denote sets of edges between nodes in these partitions. Every node in V_1^D is the end node to an arc in E_{01}^D, which connect the source to the node for each household:

$$E_{01}^D \equiv \Big\{(\text{source}, n) \in V_0^D \times V_1^D\Big\}.$$

Every node $n \in V_1^D$ is adjacent to the nodes in V_2^D which correspond to the jobs submitted by household n by an arc in the following set:

$$E_{12}^D \equiv \Big\{\big(n, (n', a, m, d, \alpha, \beta)\big) \in V_1^D \times V_2^D : n = n'\Big\}.$$

Every node of the form $(n, a, m, d, \alpha, \beta) \in V_2^D$ is adjacent to those nodes in V_3^D that correspond to household n, appliance a, power setting m, and hours $h \in \mathcal{H}$ in the interval $\alpha \leq h \leq \beta$:

$$E_{23}^D \equiv \Big\{\big((n, a, m, d, \alpha, \beta), (n', a', m', h)\big) \in V_2^D \times V_3^D : n = n',\ a = a',\ m = m', \text{ and } \alpha \leq h \leq \beta\Big\}.$$

By interpretation, arcs $(i, j) \in E_{23}^D$ connect jobs for scheduling for each household to feasible hours and appliance modes at which they may be completed. Flow along these arcs permits the distribution of the job duration d across the different hours within the time range $[\alpha, \beta]$ in which the job may be scheduled for completion, at the particular appliance power setting $m \in \mathcal{M}(a)$ at which the job must be run. The next set of edges, E_{34}^D, collapse the nodes representing different power modes of the same appliance for each hour into single nodes for each appliance crossed with $\mathcal{H}$, belonging to V_4^D. *Flow capacity constraints along arcs in E_{34}^D will restrict multiple jobs that are scheduled on a single appliance at a given power mode to run for at most one hour in each time period.*

$$E_{34}^D \equiv \Big\{\big((n, a, m, h), (n', a', h')\big) \in V_3^D \times V_4^D : n = n',\ a = a', \text{ and } h = h'\Big\}$$

Arcs in the next set of edges, E_{45}^D, restrict every appliance to be scheduled for at most one hour of use across all of its power settings in each hour $h \in \mathcal{H}$. Additionally, they collapse all the appliances owned in a single household to nodes in V_5^D which correspond to every household in each hour:

$$E_{45}^D \equiv \Big\{\big((n,a,h),(n',h')\big) \in V_4^D \times V_5^D : n = n' \text{ and } h = h'\Big\}.$$

Edges in E_{56}^D collapse subgraphs for every distinct households into single nodes for aggregate demand for every hour:

$$E_{56}^D \equiv \Big\{\big((n,h),h'\big) \in V_5^D \times V_6^D : h = h'\Big\}.$$

The next partition of edges are the auxiliary edges added to the graph to connect each of the nodes in V_6^D to the sink node in V_7^D:

$$E_{67}^D \equiv \Big\{(h,\ \text{sink}) \in V_6^D \times V_7^D\Big\}.$$

Finally, the last edge in the graph connects the sink back to the source, permitting circulation:

$$E_{70}^D \equiv \Big\{(\text{sink, source}) \in V_7^D \times V_0^D\Big\}.$$

As before, V^D and E^D are the unions over the partitions enumerated above. This concludes the description of the demand side graph $G^D = (V^D, E^D)$ for the combined household appliance scheduling problem. As before, we provide reference tables to summarize the structure of $G^D = (V^D, E^D)$, as well as the corresponding parameters u_{ij}, l_{ij}, and c_{ij}, $\forall (i,j) \in E^D$, that we outline below. Figure 4.2 also provides a visual reference for the structure of our formulation. The shading in this figure denotes scheduling across different time periods from set $\mathcal{H}$. Labels over edges in the graph are the feasible intervals for flow along that arc, denoted by l_{ij} and u_{ij}, respectively.

Once again, we define parameters for bounds and costs over our graph, in order to formulate flow balance and capacity constraints for the network flow. Because our problem is constructed as a circulation, we omit the supply value at each node, setting $b_i \equiv 0$, $\forall i \in V^D$. In this formulation, we consider the flow of time across our network, as the duration for each job is allocated to the appropriate appliances at permissible power modes in feasible hours. As such, the upper bound capacity on a majority of the arcs in graph G^D is 1, the

Table 4.3: Nodes in the Network Flow Formulation of the Aggregate Household Appliance Scheduling Problem

Node Partition	Corresponding Sets
V_0^D	source
V_1^D	$n \in \mathcal{N}$
V_2^D	$(n, a, m, d, \alpha, \beta) \in \mathcal{N} \times \mathcal{A}_n \times \mathcal{J}(a)$
V_3^D	$(n, a, m, h) \in \mathcal{N} \times \mathcal{A}_n \times \mathcal{M}(a) \times \mathcal{H}$
V_4^D	$(n, a, h) \in \mathcal{N} \times \mathcal{A}_n \times \mathcal{H}$
V_5^D	$(n, h) \in \mathcal{N} \times \mathcal{H}$
V_6^D	$h \in \mathcal{H}$
V_7^D	sink

Table 4.4: Edges in the Network Flow Formulation of the Aggregate Household Appliance Scheduling Problem

E_{AB}	$i \in V_A^D$	$j \in V_B^D$	$\rho(i,j) \Longrightarrow (i,j) \in E_{AB}$	u_{ij}	l_{ij}	c_{ij}
E_{01}^D	source	V_1^D	all	∞	0	0
E_{12}^D	V_1^D	V_2^D	if $n = n'$	∞	d	0
E_{23}^D	V_2^D	V_3^D	if $n = n'$, $a = a'$, $m = m'$, and $\alpha \leq h \leq \beta$	1	0	0
E_{34}^D	V_3^D	V_4^D	if $n = n'$, $a = a'$, and $h = h'$	1	0	P_h
E_{45}^D	V_4^D	V_5^D	if $n = n'$ and $h = h'$	1	0	0
E_{56}^D	V_5^D	V_6^D	if $h = h'$	∞	0	0
E_{67}^D	V_6^D	sink	all	∞	0	0
E_{70}^D	sink	source	all	∞	0	0

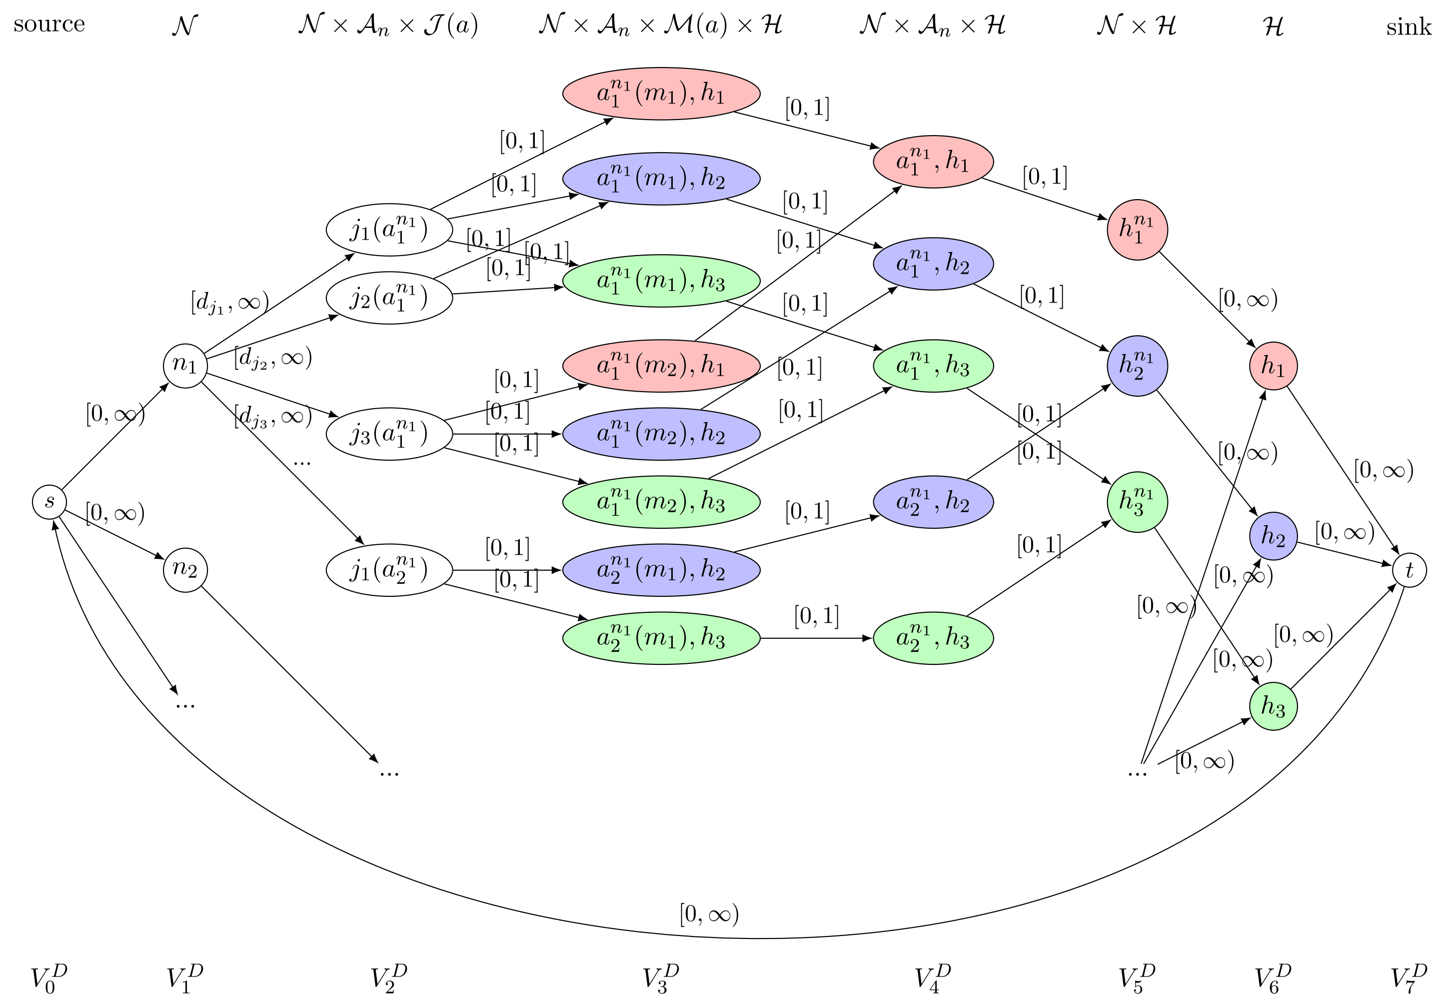

Figure 4.2: Network Flow Formulation for the Aggregate Household Appliance Scheduling Problem

interval at which elements of the set $\mathcal{H}$ are defined, hours in the day. We introduce the following exceptions:

- $u_{ij} \equiv \infty$, $\forall (i,j) \in V_{01}^D$. The flow over each of these arcs is the total duration of all jobs for the household denoted by the end node of the arc.
- $u_{ij} \equiv \infty$, $\forall (i,j) \in V_{12}^D$. Flow over these arcs is the number of appliance-hours allocated for each job, with the lower bound being the duration of the job denoted by the end node of the arc.
- $u_{ij} \equiv \infty$, $\forall (i,j) \in V_{56}^D$. Flow over these arcs corresponds to the number of appliance-hours being consumed by household n in hour h. Using two distinct appliances for thirty-minutes each sums to one appliance-hour. The upper bound on this flow depends on the number of appliances in each household.
- $u_{ij} \equiv \infty$, $\forall (i,j) \in V_{67}^D$. Flow on arcs in V_{67}^D sum flow over arcs in V_{56}^D across all households for each hour. The upper bound on this flow depends on the total number of appliances across all households available for use in hour $h \in \mathcal{H}$.
- $u_{ij} \equiv \infty$, $\forall (i,j) \in V_{70}^D$. This arc circulates total network flow from sink to source.

The lower bound on all arcs in this network are 0, the minimum amount of time that can be allocated for work on a job, with the exception of arcs in E_{12}^D:

$$l_{n(n,a,m,d,\alpha,\beta)} \equiv d,\ \forall \big(n,(n,a,m,d,\alpha,\beta)\big) \in E_{12}^D : (n,a,m,d,\alpha,\beta) \in \mathcal{J}_n,\ n \in \mathcal{N}.$$

The lower bound on these arcs is the duration of the respective job. This ensures that each job will be scheduled for at least d hours in order to permit completion. The structure of the network and problem as a cost minimization also prevents a job from being scheduled for more hours over the network than required by this lower bound. Finally, pricing is done only over arcs in E_{34}^D, after which arcs begin to collapse the schedule but do not impact the cost of the solution. We define:

$$c_{(n,a,m,h)(n,a,h)} \equiv P_h,\ \forall \big((n,a,m,h)(n,a,h)\big) \in E_{34}^D : h \in \mathcal{H}.$$

The cost of each arc (i, j) in E^D_{34} is the retail price for the hour of the day represented by that arc. The cost for all other arcs is 0.

The beauty of network flow formulations permits us to keep the same structure as the constraints from the supply side problem, and so we provide the flow balance and capacity constraints analogous to (4.1) and (4.2). For notational clarity, we instead define the flow variables for the demand side problem as x_{ij}, $\forall (i, j) \in E^D$. The constraints on these flow variables follow:

$$\sum_{j \in I(i)} x_{ji} = \sum_{j \in O(i)} x_{ij}, \quad \forall\, i \in V^D \tag{4.5}$$

$$l_{ij} \leq x_{ij} \leq u_{ij}, \quad \forall (i, j) \in E^D \tag{4.6}$$

Constraint (4.5) corresponds to a flow balance that ensures circulation in the system from each job requiring scheduling in V^D_2, into the network allocating the schedule across appliances, power settings, and hours in the day. Constraint (4.6) captures the minimum and maximum capacities defined above.

The objective function for the aggregate household scheduling problem is as follows:

$$\min_{\mathbf{x}} \left\{ \sum_{h \in \mathcal{H}} \sum_{n \in \mathcal{N}} \sum_{a \in \mathcal{A}_n} \sum_{m \in \mathcal{M}(a)} c_{(n,a,m,h)(n,a,h)} \cdot x_{(n,a,m,h)(n,a,h)} \cdot m \right\}. \tag{4.7}$$

To provide intuition into this expression, we approach each factor in the product under the summation. The term $c_{(n,a,m,h)(n,a,h)}$ is the cost per unit of electricity consumed, in \$/kWh. The term $x_{(n,a,m,h)(n,a,h)}$ is the flow in units of time, hours. Finally, we amplify the time flow by the energy requirement of each appliance at the mode in which it is operating, m, which is in units of kW. Together, the objective provides the desired sum, the total cost in dollars:

$$\frac{[\$]}{[\text{kilowatt-hours}]} \times [\text{hours}] \times [\text{kilowatts}] = [\$].$$

Utilizing the aggregate household appliance scheduling program, we may generate the consumption schedule for household $n \in \mathcal{N}$ for each day. Flows along the arcs in E^D_{23} correspond to whether task j is scheduled for appliance a at mode m in hour h, for each household n. The value of $x_{(n,a,m,d,\alpha,\beta)(n,a,m,h)} \cdot m$, $\big((n, a, m, d, \alpha, \beta), (n, a, m, h)\big) \in E^D_{23}$

reflects the electricity consumption of task $j = (n, a, m, d, \alpha, \beta)$ on appliance a at mode m in hour h. An important distinction worth noting, flow variables x_{ij} by themselves correspond to work done on a job in units of time, whereas the energy consumption variables defined in Chapter 2 and 3, $x_{n,a}^h$ are measured in kilowatt-hours. By weighting flow variables by the power setting m at which the appliance is run, we may derive the values of entries in $\mathbf{x}_n$ from our network.

For nodes $j = (n, a, m, h) \in V_3^D$ and $i = (n, a, h) \in V_4^D$, we define the energy consumption of appliance a in hour h for household n as follows:

$$x_{n,a}^h = \sum_{j \in I(i)} x_{ji} \cdot m, \quad \forall\, i \in V_4^D,$$

where $I(i)$ once again is the set of start nodes of arcs incoming to node i. $I(i) \subset V_3^D$. As before, we utilize the following notation to denote summations over $x_{n,a}^h$:

$$q_{n,h}^D = \sum_{a \in \mathcal{A}_n} x_{n,a}^h, \; \forall n \in \mathcal{N}, h \in \mathcal{H},$$

$$Q_h^D = \sum_{n \in \mathcal{H}} q_n^h = \sum_{n \in \mathcal{H}} \sum_{a \in \mathcal{A}_n} x_{n,a}^h, \; \forall h \in \mathcal{H}.$$

As before, we conclude this section with a complete formulation for our aggregate household appliance scheduling program utilizing network flows:

$$\min_{\mathbf{x}} \left\{ \sum_{h \in \mathcal{H}} \sum_{n \in \mathcal{N}} \sum_{a \in \mathcal{A}_n} \sum_{m \in \mathcal{M}(a)} c_{(n,a,m,h)(n,a,h)} \cdot x_{(n,a,m,h)(n,a,h)} \cdot m \right\}$$

subject to:

$$\sum_{j \in I(i)} x_{ji} = \sum_{j \in O(i)} x_{ij}, \quad \forall\, i \in V^D$$

$$l_{ij} \le x_{ij} \le u_{ij}, \quad \forall (i, j) \in E^D$$

4.3.1 Computability and Decomposition

An important property of our aggregate household appliance scheduling model is the ability to solve for each household's schedule simultaneously with this single program. Because the joint objective function is linearly additive of each household's local minimization problem and no constraint in (4.5) and (4.6) couples households together, by solving the household problems jointly we do not prune any solution that would otherwise be feasible by solving for each household's optimal schedule separately.

Theorem 4. *Given a fixed price vector* $\mathbf{P}$*,* $\mathbf{x}^*$ *is the optimal solution to* (4.7) *if and only if every component* $\mathbf{x}_n^*$ *of* $\mathbf{x}^*$*, for* $n \in \mathcal{N}$*, is locally optimal for the decomposed appliance scheduling subproblem for household* n*.*

Proof. We define the subgraph $G^D(n)$ for each household $n \in \mathcal{N}$ as the subset of nodes $V^D(n) \subset V^D$, for $|\mathcal{N}| > 1$, that are indexed by household n and the subset of arcs from E^D incident to any node $i \in V^D(n)$. The parameter values from the complete network also apply to the decomposition. The subproblem for each household n is a minimum cost flow over the subgraph for $G^D(n)$.

($\Longrightarrow$) If $\mathbf{x}^*$ is the optimal solution to the aggregate household appliance scheduling problem that minimizes (4.7), then we intend to show that the component of flow $\mathbf{x}^*$ over subgraph $G^D(n)$, is also optimal for the subproblem of minimizing household n's flow cost.

The only edges with nonzero cost in $G^D(n)$ belong to the set E^D_{34}, and so the total cost of flow through $G^D(n)$ equals $\sum_{h\in\mathcal{H}} \sum_{a\in\mathcal{A}_n} \sum_{m\in\mathcal{M}(a)} c_{(n,a,m,h)(n,a,h)} \cdot x_{(n,a,m,h)(n,a,h)} \cdot m$, the total daily payment for household n. Given there are no capacity constraints that restrict flow over arcs in E^D_{01}, E^D_{56}, E^D_{67}, or E^D_{70} in the aggregate household problem, there are no coupling constraints in the full network to restrict flow across subgraphs of G^D. If the solution $\mathbf{x}^*$ is optimal for the aggregate network, the component $\mathbf{x}_n^*$ must also minimize the cost of flow through each subgraph $G^D(n)$, $n \in \mathcal{N}$.

($\Longleftarrow$) If $\mathbf{x}_n^*$ is the optimal flow over each subgraph $G^D(n)$, $\forall n \in \mathcal{N}$, then the joint consumption profile $\mathbf{x}^*$ must also be the optimal solution for the aggregate problem.

All arcs belonging to E^D_{01}, E^D_{56}, E^D_{67}, or E^D_{70} are uncapacitated and so there are no coupling constraints restricting flow across subgraphs. Thus, if a solution $\mathbf{x}_n$ is feasible for subgraph $G^D(n)$, it is a feasible flow over the complete network. Then, if $\mathbf{x}_n^*$ is feasible for G^D and minimizes $\sum_{h\in\mathcal{H}} \sum_{a\in\mathcal{A}_n} \sum_{m\in\mathcal{M}(a)} c_{(n,a,m,h)(n,a,h)} \cdot x_{(n,a,m,h)(n,a,h)} \cdot m$, $\forall n \in \mathcal{N}$, then the joint solution $\mathbf{x}^*$ minimizes the sum of the total payments:

$$\sum_{n\in\mathcal{N}}\sum_{h\in\mathcal{H}}\sum_{a\in\mathcal{A}_n}\sum_{m\in\mathcal{M}(a)} c_{(n,a,m,h)(n,a,h)} \cdot x_{(n,a,m,h)(n,a,h)} \cdot m.$$

Given that $\mathbf{x}^*$ is feasible and minimizes the objective in (4.7), it is an optimal solution to the aggregate household problem. □

Corollary (to Theorem 4)**.** The optimal solution $\mathbf{x}^{P*}$ to any decomposition of graph G^D according to a partition P of the set $\mathcal{N}$ is equivalent to the optimal solution $\mathbf{x}^*$, found by optimizing (4.7) over the complete graph.

Proof. By Theorem 4, any optimal solution to an aggregate household appliance scheduling problem over the set $S \subseteq \mathcal{N}$ must be locally optimal for every household $n \in S$. By definition of a partition, every household $n \in \mathcal{N}$ belongs to some element of the partition P of $\mathcal{N}$ (elements of the partition are collectively exhaustive). Therefore, if an optimal solution is found for every subset in the partition P of $\mathcal{N}$, the combined solution $\mathbf{x}^{P*}$ for the partition P must be equivalent to $\mathbf{x}^*$, the optimal solution for the entire set $\mathcal{N}$: $\mathbf{x}^{P*} = \mathbf{x}^*$, by Theorem 4. □

The corollary we prove above suggests a natural algorithm for solving large-scale instances of our household appliance scheduling problem computationally. Instead of solving for every household simultaneously, we may hide latency in setting up our network by instead decomposing the graph G^D for a selected partition of $\mathcal{N}$ into distinct subgraphs for every element of the partition. This decomposes the original problem into subproblems that can be solved for their respective optimal solutions, which by Theorem 4 still constitute an optimal solution $\mathbf{x}^*$ for the full problem. The network formulation for the household appliance scheduling problem is engineered to be scalable and can be implemented with data parallelization over multiple processes to computationally generate an optimal schedule for all households in the set $\mathcal{N}$.

In contrast to the matching program we constructed in Chapter 2 for the household problem, we no longer need to restrict task durations to be integers in our network flow formulation. Though we define set $\mathcal{H}$ as containing only discrete time values, if the duration of any job submitted for scheduling is non-integral, our network still permits work on appliances to be scheduled for fractional values as work done corresponds to flows over the network. There are no integer constraints imposed on the flow variables in our formulation.

Nonetheless, we may still exploit a useful feature of network flows when the problem data is integral. When the duration for all tasks submitted are integers, the simplex method on our network can be implemented using integer arithmetic, as opposed to floating point. This permits faster computation and the issues of finite precision and truncation errors disappear [1]. Additionally, given that all flow balance and capacity parameters are integral in the network, every basic solution also has integer flow. We attribute this result to Bertsimas and Tsitsiklis [1]. As a consequence of parameter integrality, though we do not constrain flow variables to integers, we are still able to generate optimal solutions with integer values utilizing only standard linear programming techniques.

Finally, we conclude this discussion with a mixed-integer programming extension for our model to address the property of *preemptive scheduling,* defined below. We also consider the implications of this extension in terms of computability.

Definition 2 (Preemption)**.** Preemption is a property in scheduling where the processing of job j on appliance a can be broken down and carried out non-consecutively in smaller segments within the time range allocated. That is, another job on the same appliance or appliance idling may preempt the processing of job j if such a solution is more optimal as per the objective function [1].

In our model as presented above, we permit task preemption. This formulation as well as the cases outlined for overlapping tasks in Chapter 2 allow scheduled tasks to be interrupted by other tasks given that all tasks are scheduled within their feasible time ranges $[\alpha, \beta]$ and are seen to completion. We propose an extension to our model that prevents this property of preemption, at the expense of model simplicity.

Whereas the network routes flow from nodes in V_2^D over capacity constrained arcs to nodes in V_3^D, which correspond to distinct hours available on the appliance for scheduling,

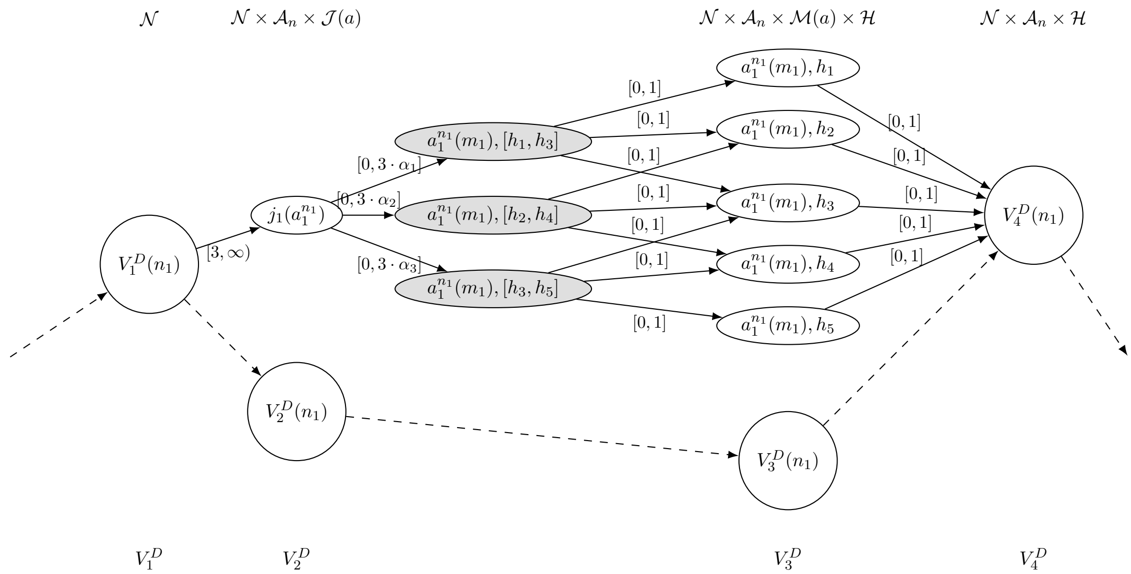

Figure 4.3: Network Flow Formulation for Nonpreemptive Scheduling

we may introduce an additional set of nodes in between these to denote uninterruptible time intervals available for scheduling. The intent is to define time intervals of width d within the time range parameters $[\alpha, \beta] : \beta - \alpha + 1 \geq d$ in which the task may be scheduled without interruption.

In Figure 4.3, we illustrate the addition of nodes between V_2^D and V_3^D, in shaded fill, for each time interval in which we may schedule task j_1 on appliance a_1: $\{[h_1, h_3], [h_2, h_4], [h_3, h_5]\}$. In this extended model, we restrict flow between these nodes by introducing binary variables α_1, α_2, and α_3, such that $\sum_{i=1}^{3} \alpha_i = 1, \forall \alpha_i \in \{0, 1\}$. Given that only one path will be unconstrained for flow in the network, our model will select the continuous time interval in which to schedule the task that minimizes the objective value. To implement this extension, we assume for each household's submitted schedule that at least one feasible solution does exist. The subsets of V^D that are indexed by n in Figure 4.3 are the node sets of $G^D(n)$, the subgraph corresponding to household n.

While this extension increases the complexity of the model by introducing a set of additional integer variables $\{\alpha_i^j\}_{i=1}^k$ for each task j to be scheduled, where k is the number of continuous time intervals over which decision j can be made, each constraint $\sum_{i=1}^k \alpha_i^j = 1, \forall j$ is restricted to only the set of α^j decision variables for job j. This grossly simplifies the integer program. While we facilitate nonpreemptive scheduling in our model, we are not concerned that it will unravel our model's practicality. We were primarily interested in household tasks, like recharging PHEVs, that are specifically interruptible and can be load shifted without consequence to the job being performed.

4.4 Algorithms for Computing Equilibria

The motivation for developing the network flow formulation in this chapter has been to compute equilibrium prices and quantities from our proposed two-way communication model for the retail market. In this section, we discuss the iterative process of communicating information between the two problems introduced before and proceed to describe algorithms for computing equilibrium values.

4.4.1 Tabulating Market Prices and Quantities

As outlined in our formulations of the economic dispatch problem and the aggregate household appliance scheduling problem, the respective programs each take communicated information from the other as parameters. We defined hourly demand $Q_h^D, \forall h \in \mathcal{H}$ in the economic dispatch problem and utilized it to assign lower bounds on flow in the network over arcs in E_{01}^S. In the household program, parameters $P_h, \forall h \in \mathcal{H}$ were exogenously defined as prices for each hour. We now begin to incorporate these quantities, which correspond to the system state vectors $(\mathbf{P}, \mathbf{Q}^D)$, as endogenous variables with respect to the complete market model. The derivations below redefine price and hourly demand parameters in terms of the flow variables for our network flow formulations.

Hourly demand corresponds to the aggregate flow over G^D, when summed over the arcs in E_{34}^D. As we defined before,

$$Q_h^D \triangleq \sum_{n \in \mathcal{N}} q_{n,h}^D = \sum_{n \in \mathcal{N}} \sum_{a \in \mathcal{A}_n} x_{n,a}^h, \; \forall h \in \mathcal{H}. \tag{4.8}$$

The hourly price P_h per kWh can be derived from the total aggregate cost of production, which is equivalent to the objective function in (4.3):

$$\sum_{h\in\mathcal{H}} C\big(Q_h^S\big) = \sum_{g\in\mathcal{G}} c_g^f \cdot \sigma_g + \sum_{(i,j)\in E_{23}^S} c_{ij} \cdot f_{ij}.$$

If we assume the fixed cost to be distributed uniformly over each hour and sum over only the arc flows in E_{23}^S that correspond to hour h, we produce the following expression for the aggregate hourly cost of production for hour $h \in \mathcal{H}$:

$$C\big(Q_h^S\big) = \underbrace{\frac{\sum_{g\in\mathcal{G}} c_g^f \cdot \sigma_g}{H}}_{\text{hourly fixed costs}} + \overbrace{\sum_{g\in\mathcal{G}} \sum_{\tau\in\mathcal{T}(g)} c_{(g,\tau,h)(g,h)} \cdot f_{(g,\tau,h)(g,h)}}^{\text{variable costs of production}}, \quad \forall h \in \mathcal{H}.$$

The per unit cost of production is the above quantity divided by hourly production, which is the flow over the arc f_{sh} in E_{01}^S, from the source node $s \in V_0^S$ to $h \in V_1^S$:

$$\frac{C\big(Q_h^S\big)}{Q_h^S} = \frac{\sum_{g\in\mathcal{G}} c_g^f \cdot \sigma_g}{H \cdot f_{sh}} + \frac{\sum_{g\in\mathcal{G}} \sum_{\tau\in\mathcal{T}(g)} c_{(g,\tau,h)(g,h)} \cdot f_{(g,\tau,h)(g,h)}}{f_{sh}}, \quad \forall h \in \mathcal{H}. \tag{4.9}$$

As before, we take the end-user retail prices to be a linear scaling of (4.9) by the factor κ, the gross profit margin for electric utilities, as discussed in previous sections: $P_h = \kappa \times \frac{C(Q_h^S)}{Q_h^S}$, $\forall h \in \mathcal{H}$. Once again, because the factor κ for prices simply scales the total daily household bill by this constant factor, there is no loss of generality by setting $\kappa \equiv 1$, as we may assume we have normalized currency in our retail market by factor $1/\kappa$.

$$P_h \triangleq \kappa \times \left[\frac{\sum_{g\in\mathcal{G}} c_g^f \cdot \sigma_g}{H \cdot f_{sh}} + \frac{\sum_{g\in\mathcal{G}} \sum_{\tau\in\mathcal{T}(g)} c_{(g,\tau,h)(g,h)} \cdot f_{(g,\tau,h)(g,h)}}{f_{sh}}\right], \quad \forall h \in \mathcal{H}. \tag{4.10}$$

4.4.2 `MYOPIC_BEST_RESPONSE` Algorithm

In this section, we propose an initial algorithm for computing the market equilibrium with our model for communication implemented. Our algorithm iterates between solving the supply side and demand side optimization programs we have outlined, simulating the best response of players at each computational iteration. Our goal is to iteratively simulate best responses to determine prices and aggregate hourly consumption. We are interested in whether the market converges to equilibrium with our proposed two-way communication of prices and hourly demand.

The process of communicating prices between the two sides of the market in computing for equilibrium reduces the payment by each household to a linear objective function of price times quantity consumed, which when summed together also corresponds to the objective value for the producer. Reducing payments to linear functions allows us to tractably solve for a Nash equilibrium, which is otherwise computationally difficult. Formulating the market interactions between households and the producer in terms of tractable linear programs is a central contribution we make to the literature. Because the payment made by each household is itself a function of the total aggregate cost of production, the market structure we devise wholly captures the cost imposed to society and aligns consumption decisions with the true cost of electricity generation, a property currently missing from retail electricity pricing. By structuring the household problem as a minimization over payments made to the producer, both sides of the market intend to maximize social welfare as desired.

In order to formulate termination conditions for our algorithms, we introduce a payoff function for the household players in Game 1 that corresponds to our proposed retail market design, for use in place of the utility function u_n defined earlier in Chapter 3. Payoff function $P_n : X_n \mapsto \mathbb{R}$, and is defined for each household as follows, $\forall n \in \mathcal{N}$:

$$P_n(\mathbf{x}_n) = \begin{cases} -b_n = -\sum_{h \in \mathcal{H}} \left[P_h \times \sum_{a \in \mathcal{A}_n} x^h_{n,a} \right] & \text{if } \mathbf{x}_n \text{ satisfies (4.5) and (4.6)} \\ -\infty & \text{otherwise} \end{cases} \tag{4.11}$$

The payoff for households is the additive inverse of their total daily payment to the electric utility if their schedule is feasible, and as before, $-\infty$ if the schedule $\mathbf{x}_n$ does not satisfy all of the respective constraints for household n. As we proved in our decomposition, given that no coupling constraints restrict flow across subgraphs of G^D, household n's component of total flow is feasible for G^D if and only if it is feasible for the subproblem of minimum cost flow over $G^D(n)$. Thus, checking that $\mathbf{x}_n$ satisfies the respective constraints in (4.5) and (4.6) is necessary and sufficient to prove that the schedule $\mathbf{x}_n$ is feasible for the appliance constraints for household n.

Defining the payoffs for households in Game 1 according to (4.11), reformulates the game as a model for the market with our two-way communication protocol in place. The solution to the aggregate household appliance scheduling program $\mathbf{x}^*$ minimizes (4.7), which implies that it maximizes the sum of the total payments made by all households. Once again, given our decomposition result, this means that at an optimal $\mathbf{x}^*$, the individual payoff for every household is also maximized: $-\sum_{n\in\mathcal{N}} P_n(\mathbf{x}_n)$ equals the value of the objective (4.7) for the aggregate household appliance scheduling program, and so the cost for flow over each subgraph must also be minimized.

The goal for our mechanism is to engineer prices such that the best responses of households and the producer also maximize social welfare and minimize the total aggregate cost of production, whichs correspond to a Nash equilibrium for the market as proven in the previous chapter. By simulating these best responses, we intend to see whether the market converges to equilibrium, and if so, determine the retail prices $\mathbf{P}$ corresponding to equilibrium.

Consider the `MYOPIC_BEST_RESPONSE` procedure below, which we base on the definition of Nash equilibrium. We take it that the algorithm is invoked for some initial action profile $\mathbf{x}^0$ that is feasible for the demand side network. Computing this initial feasible solution is simple and involves a single iteration of solving the demand side program given an initial set of prices.

Algorithm 1 `MYOPIC_BEST_RESPONSE`($\mathbf{x}^0$)

$i \leftarrow 0$
$\mathbf{Q}^{D,0} \leftarrow$ evaluate $(4.8)\big|_{\mathbf{x}^0}$
$(\mathbf{f}^{0*}, \sigma^{0*}) \leftarrow \arg\min_{\mathbf{f},\sigma}$ (4.3)
$\mathbf{P}^0 \leftarrow$ evaluate $(4.10)\big|_{(\mathbf{f}^*,\sigma^*)}$
while $\exists\, n \in \mathcal{N} : P_n(\mathbf{x}_n^*) \ngeq P_n(\mathbf{x}_n),\ \forall\, \mathbf{x}_n \in X_n, \mathbf{x}_n \neq \mathbf{x}_n^*$ **do**
  $i \leftarrow i + 1$
  $\mathbf{x}^{i*} \leftarrow \arg\min_{\mathbf{x}}$ (4.7)
  $\mathbf{Q}^{D,i} \leftarrow$ evaluate $(4.8)\big|_{\mathbf{x}^*}$
  $(\mathbf{f}^{i*}, \sigma^{i*}) \leftarrow \arg\min_{\mathbf{f},\sigma}$ (4.3)
  $\mathbf{P}^i \leftarrow$ evaluate $(4.10)\big|_{(\mathbf{f}^*,\sigma^*)}$
end while
$\mathbf{P}^* \leftarrow \mathbf{P}^i$
$\mathbf{Q}^{D*} \leftarrow \mathbf{Q}^{D,i}$
return $(\mathbf{P}^*, \mathbf{Q}^{D*})$

In this procedure, we alternate between solving each problem given the parameters for demand or price set by the other. At line 9, the procedure re-optimizes the supply side program for the new parameter values for $\mathbf{Q}^{D,i}$, aggregated from the demand side problem solved at line 7, where i denotes the current iteration of the algorithm. Step 9 and 10, the supply update, are required in order to internalize the cost of other players' actions into prices. Utilizing the endogenous expression (4.10) for hourly prices, each player n's true payoff becomes a function of other player's consumption in each hour, which enters the expression through the constraints on flow variables $\mathbf{f}$. This, in effect, produces a nonlinear payoff function. In this algorithm, we seek to iteratively update this component of player payoffs that depends on other players' consumption profiles by communicating updated values for price from the supply side, which incorporate the required information after every computational iteration. Ordinarily, we cannot expect household n to have knowledge about any other household's consumption habits in setting its own $\mathbf{x}_n$, but given our proposed market structure of two-way communication between producers and households, we may utilize the information provided to the supply side, in terms of aggregate hourly demand, in order to provide a proxy for this hidden information.

If Algorithm 1 converges and we are able to determine a vector of equilibrium prices, electric utilities may set hourly retail prices accordingly such that each household's best response to one another will be to play the joint action profile that defines the equilibrium. By utilizing the mechanism of best response in our algorithm, we are testing for convergence to Nash equilibrium. If such an equilibrium is shown to exist, we may utilize the equilibrium prices to set hourly prices for households in the retail market.

4.4.3 Convergence for `MYOPIC_BEST_RESPONSE`

However, given our formulations for the supply side and demand side optimization programs, we can prove that `MYOPIC_BEST_RESPONSE` fails to converge. In the following section, we provide a proof using arguments from congestion game theory and argue that by tabulating price after letting all households re-optimize consumption, we actually fail to internalize the congestion effects of shifting one's consumption from period to period. Analyzing the structure of our algorithm as a congestion game motivates the development of other algorithms for computing equilibrium.

4.5 Congestion Game Formulation

Congestion games are useful in practice as a framework for modeling the effects that a single player's action in a game can have on other players' payoffs. In a multi-resource congestion game, each player chooses some subset from a set of resources, where the cost of each resource depends on the number of other players who also select it [14]. We extend this framework to model the Retail Market Energy Consumption Game, with household payoffs defined by (4.11), in order to build intuition as to why our `MYOPIC_BEST_RESPONSE` procedure fails to converge.

For our problem, we treat consumable electricity during each hour as a distinct resource. Players may substitute between consumption in one hour for another, so long as the constraints formulated in our network for each appliance and task to be scheduled are still met. As we recall, the player payoff to playing an action that does not satisfy every constraint is $-\infty$. Our game lends itself to the congestion game framework because the cost of a resource, consumption in hour h, increases with the number of players who choose to consume in that hour. Given our endogenous expression for prices, the inclusion of the hourly production function as a factor in each household's payoff expression is indicative of an underlying congestion game.

Another appealing property of congestion games that applies to our model is *anonymity:* players care about how many others use a given resource, but do not care about which others do so [14]. In our communication protocol, we communicate prices that incorporate the aggregate quantity of electricity consumed in each hour, but this does not indicate *which* households are the ones actually consuming it. The vector $\mathbf{Q}^D$ is a sufficient statistic for the congestion effect of other players' actions on a given household's cost of consuming electricity.

The supply update in `MYOPIC_BEST_RESPONSE` does not occur until after all households have already shifted consumption to a new $\mathbf{x}^*$, and therefore consumption decisions made jointly by households in each iteration are never based on the true cost of consumption. This wedge between the price of consumption as households see it when they move and the true price that incorporates the full congestion effect, which is not realized until after all households best respond, prevents the system from converging to equilibrium.

To make this point clearer, we provide an illustrative example. Take the payoff matrices in Figures 4.4a and 4.4b. Without a loss of generality, we simplify our game to a consumption decision with just two players, A and B, in which each player must choose an hour from the set $\{h_1, h_2\}$ in which to schedule all of his appliances, which total to a single unit of consumption. Prices for each hour are either high (H) or low (L), $L < H$, corresponding to the total demand Q_h^D for that hour:

$$P_h = \begin{cases} H, & \text{if } Q_h^D = 2 \\ L, & \text{if } Q_h^D < 2 \end{cases}, \quad h \in \{h_1, h_2\}. \tag{4.12}$$

If both players want to consume in the same hour, the price for that hour becomes H. From this setup, it is intuitive that the optimal solution is for each player to select a different hour, a demand smoothing outcome, in order to maximize player utility. However, in Figure 4.4a, given prices from a previous iteration, both players select h_1 in iteration i. The best response for both players in iteration i, given these prices $\mathbf{P}^i$, is bolded in Figure 4.4a: $\mathbf{x}^{i*} = (h_1, h_1)$.

Between the left and right panels in the figure, we perform a supply update, or the producer's best response, by tabulating aggregate demand for each hour and then update prices according to (4.12) from above. Though players propose consumption based on the static prices from the previous iteration, the prices they actually pay are the new prices determined from the supply update, which is intended to incorporate the congestion effect of other households consuming electricity during the same hour into the price. The new prices are designed to impact household payoff functions as if changes in consumption by other households were to enter the expression directly as a change in the hourly production cost function, which would affect the endogenous price expression. By utilizing this intermediate supply update step, we are able to maintain linearity in each expression.

Returning to our example, given that both players choose to consume in h_1, the price for

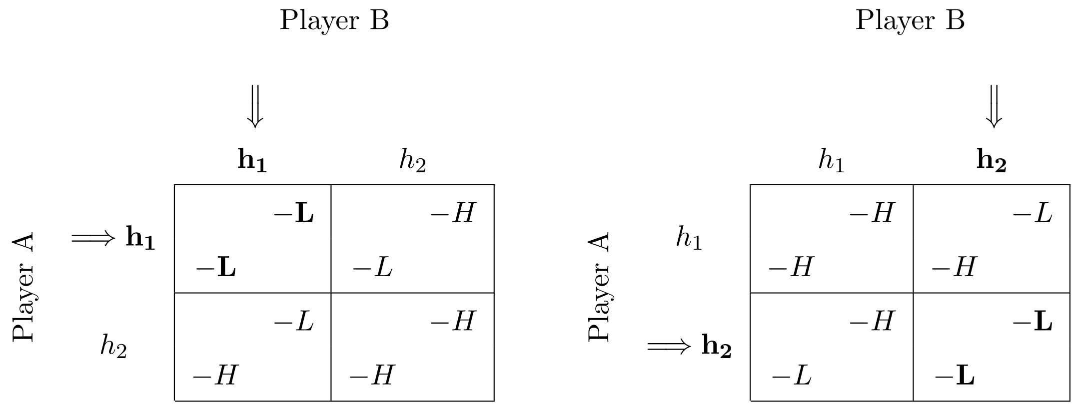

(a) Hourly prices and $\mathbf{x}^{i*}$ at iteration i before supply update

(b) Hourly prices and $\mathbf{x}^{i+1*}$ at iteration $i+1$ after supply update

Figure 4.4: Payoff Matrices Before and After Supply Update in `MYOPIC_BEST_RESPONSE` Algorithm

this hour increases to H and the new price for h_2 is now L. The algorithm terminates only if no player finds it optimal to deviate from his proposed strategy given the supply update. When players are again permitted to update in the next iteration, the payoff matrix is now the opposite of what it was in iteration i, such that H and L prices have flipped, and so both players again find it advantageous to deviate. The best response for the households in Figure 4.4b is now $\mathbf{x}^{i+1*} = (h_2, h_2)$. When players make this move, prices again flip and h_1 now becomes the non-peak hour once both players have made their move in iteration $i+1$. At the next iteration, payoffs for which look similar to Figure 4.4a, optimal player strategy again oscillates back to playing h_1.

By following a purely `MYOPIC_BEST_RESPONSE` procedure, in which all households react simultaneously during the demand update steps and prices do not change until after, we are unable to capture the congestion effect of players switching consumption profiles concurrently with player movement. Consequently, when players move, the static payoffs in the matrix do not reflect the full cost imposed to the system. As a result and as shown here, the `MYOPIC_BEST_RESPONSE` algorithm fails to converge. By noting the shortcomings of this procedure, we can begin to develop more complex algorithms that do converge to equilibrium or at least provide a robust approximation.

4.5.1 ASYNCH_CONVERGENCE Algorithm

The main problem with our proposed algorithm rests with its treatment of demand updates. In reducing the complex nonlinear expression for utility, which takes aggregate consumption as a parameter into a vector of hourly prices, we are unable to capture the cost implication on each household when other players shift their consumption as a best response. Instead, prices will incorporate this congestion effect only after every household has already shifted its consumption on the basis of the previous iteration's prices. Prices should dynamically update such that we are able to converge to the fixed point at which players no longer find it beneficial to continue to shift consumption, indicative of a Nash equilibrium. Nash equilibrium, by definition, is this stable strategy profile.

To improve upon our MYOPIC_BEST_RESPONSE procedure, we must find a way to simulate dynamic updating to prices such that congestion effects are incorporated immediately while subsets of players continue to make moves. We define the ASYNCH_CONVERGENCE algorithm as a stochastic method, that as we will demonstrate through simulation, in the limit approaches an equilibrium. The procedure is based on randomly selecting subsets of players to provide a best response and then incorporating the change to prices before other households are allowed to optimize again. Our procedure is outlined here:

Algorithm 2 ASYNCH_CONVERGENCE($\mathbf{x}^0$)

$i \leftarrow 0$
$\mathbf{Q}^{D,0} \leftarrow$ evaluate $(4.8)\big|_{\mathbf{x}^0}$
$(\mathbf{f}^{0*}, \sigma^{0*}) \leftarrow \arg\min_{\mathbf{f},\sigma}$ (4.3)
$\mathbf{P}^0 \leftarrow$ evaluate $(4.10)\big|_{\mathbf{f}^*,\sigma^*}$
while $\exists\, n \in \mathcal{N} : P_n(\mathbf{x}_n^*) \ngeq P_n(\mathbf{x}_n) - \epsilon, \forall\, \mathbf{x}_n \in X_n, \mathbf{x}_n \neq \mathbf{x}_n^*$ **do**
  $i \leftarrow i + 1$
  $S \leftarrow$ select $S \in \wp_k(\mathcal{N})$
  $\mathbf{x}^{i*}(S) \leftarrow \arg\min_{\mathbf{x},\sigma}$ (4.7) on network $G^D(S)$
  $\mathbf{Q}^{D,i} \leftarrow$ evaluate $(4.8)\big|_{\mathbf{x}^*}$
  $(\mathbf{f}^{i*}, \sigma^{i*}) \leftarrow \arg\min_{\mathbf{f}}$ (4.3)
  $\mathbf{P}^i \leftarrow$ evaluate $(4.10)\big|_{\mathbf{f}^*,\sigma^*}$
end while
$\mathbf{P}^* \leftarrow \mathbf{P}^i$
$\mathbf{Q}^{D,*} \leftarrow \mathbf{Q}^{D,i}$
return $(\mathbf{P}^*, \mathbf{Q}^{D*})$

The procedure for ASYNCH_CONVERGENCE follows from our decomposition results for the aggregate household appliance scheduling problem earlier in this chapter. In this algorithm, we generate random elements from the power set $\wp_k(\mathcal{N})$ of $\mathcal{N}$, of limited cardinality k. That is, $\wp_k(\mathcal{N})$ denotes the set of all subsets of $\mathcal{N}$ of cardinality less than or equal to k. In line 8 above, we optimize for the decomposed appliance scheduling problem over the network $G^D(S)$, holding all other flows as fixed from the previous iteration. This method allows us to incorporate incremental changes in player strategy into prices after each iteration such that hourly prices always reflect the true cost of consumption for when the next subset of players is permitted to shift their strategy. The problem we saw in our previous algorithm was that when players were permitted to make changes to their strategy, the static prices that are intended to served as a proxy for their true payoff function did not actually reflect the true cost of consumption. Prices under MYOPIC_BEST_RESPONSE omit the congestion effect induced by other players' actions on a household's payoff.

The ASYNCH_CONVERGENCE procedure at best approximates equilibrium values where agents are indifferent to small gains in payoffs up to some bound $\epsilon > 0$, which we introduce to the while loop condition in Algorithm 2. An important result we must show is that while we have changed the formulation of Game 1 to introduce household payments as payoffs, any solution found by our procedure, if it is stable, is still a Nash equilibrium for the original Retail Market Energy Consumption Game where player utility depends only on whether an action satisfies the constraints to the household problem. For households to have bounded payoffs in our modified game, they must play a consumption profile that is a feasible solution, which would always maximize household utility in the original Retail Market Energy Consumption Game. We introduce cost minimization into our mechanism as a heuristic to align household decision-making with the interests of social welfare, where we seek to minimize the aggregate total cost of production, which is equivalent to the objective value for economic dispatch and the aggregate household appliance scheduling program. By minimizing cost, households play strategies in line with the social objectives we outline as our motivation.

4.6 Summary

In this chapter, we have extended the model and problem formulation introduced in Chapter 2 to construct a practical and tractable system to solve for optimal production and consumption schedules. Utilizing network flows, we described the computational properties of our formulation and the advantages it brings to solving large-scale instances of the economic dispatch and aggregate household appliance scheduling problems. Our formulation is scalable and addresses the practical need to be able to generate hourly schedules for households utilizing a demand side management system and for coordinating production. Beyond allowing us to verify the theoretical claims made in Chapter 3 with computer simulations, the results of which are discussed in the next chapter, our model is implementable as currently formulated. Though the programs formulated in this chapter are important to confirming whether our theoretical results for the market under complete information still hold, they are also contributions in and of themselves to the literature as efficient formulations for the economic dispatch and household scheduling problems.

We also motivated and introduced two algorithms that build on our programs in order to compute equilibrium for the market under our proposed two-way communication of prices and hourly demand. Extending our game theoretic formulation from the previous chapter, we provided an explanation of why a myopic best response procedure to updating prices and player strategies will fail to converge. Consequently, we proposed a stochastic extension to better approximate the congestion effects on price from households selecting and shifting between consumption profiles. We address the fact that Nash equilibria in practice, for games with large player and action sets, are not easily computable and utilize our linear programming formulations to approximate equilibrium prices and quantities for use in a mechanism to design socially efficient production outcomes.

In the next chapter, we apply the programs and algorithms developed from this chapter to simulations verifying convergence. We also provide evidence for the theoretical claims made in previous chapters on the nature of equilibrium quantities and prices as welfare-maximizing outcomes.

Chapter 5

Simulation Results

In this chapter, we present results of computer simulations for our optimization models. We code and test the economic dispatch and aggregate household appliance scheduling programs, as well as our complete model for two-way communication between the supply side and households. We intend to test and demonstrate convergence for the mechanisms proposed in the previous chapter for determining prices. Our simulations also provide insight into the computational feasibility of our programs and the rate at which we converge to equilibrium, testing the practicality of our models for real-world application. We begin by considering simulation of our standalone programs to determine optimal production and household scheduling.

5.1 Simulation Overview

We first pose the following questions as ones we seek to answer through simulation of our model:

- What is the total social welfare gain from correcting the problem of information asymmetry in the electricity generation market? Do these results justify the infrastructure investment necessary in order to implement this system?
  - What can we expect in terms of a reduction in the *peak-to-average ratio* (PAR)[1]

[1] PAR is measured as the ratio of the maximum power level to the time-averaged power level.

for energy consumption throughout the day by households? What is the nature of our optimal solution for households?

  - What is the expected cost savings for households and for the total cost of production? How does this system reduce the dependence on expensive backup peaker plants or over-investment in generation capacity?

- How can this system, as a mechanism for determining retail prices, improve expectations of hourly demand in the problem of economic dispatch?

  - Does this system push prices toward their equilibrium values?

  - How do prices in this system reflect the true cost of consumption for households?

- To what extent can we solve for equilibrium prices and quantities in order to set retail prices? Comparing our `MYOPIC_BEST_RESPONSE` algorithm with the `ASYNCH_CONVERGENCE` procedure we propose, do either of these algorithms converge?

  - Can they be utilized as an approximation for congestion pricing in order to compute equilibrium?

  - What is the rate of convergence for these algorithms, and given computational time, are they practical for implementation in large-scale optimizations?

In the remainder of this section, we present an overview of the model assumptions used in our simulations for this chapter.

5.1.1 Household and Appliance Assumptions

For our market simulations, we model the household problem for 100 households. Each household is randomly selected to have between 1 to 20 appliances from a master set $\mathcal{A}$ and between 1 to 5 jobs with *soft* energy consumption scheduling constraints submitted for each appliance. For each job submitted, we randomly select a power mode from the predefined set $\mathcal{M}(a)$ for appliance a, the scheduling time range $[\alpha, \beta] : \beta \geq \alpha$ selected within a 24-hour period, and a duration d for the task that meets the condition: $\beta - \alpha + 1 \geq d$. In scheduling multiple jobs on a single appliance, we select time ranges such that at least one feasible solution will always exist. We have discussed the necessary and sufficient conditions for this in Chapter 2 based on Hall's Marriage Theorem [6] and incorporate this check to guarantee feasibility. For each household, we also randomly select between 1 to 10 appliances with *hard* energy consumption scheduling constraints, appliances with fixed energy consumption in each hour that cannot be rescheduled: $\beta - \alpha + 1 = d$. Appliances with hard constraints help model the energy use by refrigerators, lighting, heating, etc.; in general, these are appliances incompatible for rescheduling that still must be taken into account. Conveniently, for every household, we may collapse consumption due to these into a single appliance with hard constraints. Additionally, hard constraints on scheduling guarantee at least some baseline load of aggregate consumption in every hour since it is unrealistic to expect rescheduling for all consumption. The simulation time frame is 24 hours, starting from midnight 00:00:00 until 23:59:59 on the same day. Elements in set $\mathcal{H}$ are bounded by the interval $[0, 24)$. Our assumptions for appliance energy usage (in watts) in our model are based on averages taken from the GE Ecomagination: Home Appliance Energy Use application.

For every simulation dataset we generate, we also compile summary statistics for reference. Optimally, we adjust parameters to scale our dataset to match population averages from the U.S. Energy Information Administration's 2009 Residential Energy Consumption Survey (RECS). For the figures compiled in the chapter, the following statistics characterize our dataset:

Number of Households:	100 households
Average Daily Household Consumption:	24.77 kWh
Average Job Duration:	2.71 hours
Average Number of Jobs per Household:	10.85 jobs/household
Average Number of Appliances per Household:	3.54 appliances/household
Average Number of Jobs per Appliance:	3.06 jobs/appliance

5.1.2 Cost Assumptions

On the supply side, we assume that the aggregate cost functions are increasing and strictly convex, consistent with earlier assumptions for the model. This requires that we generate marginal capacities and costs accordingly for each power plant.

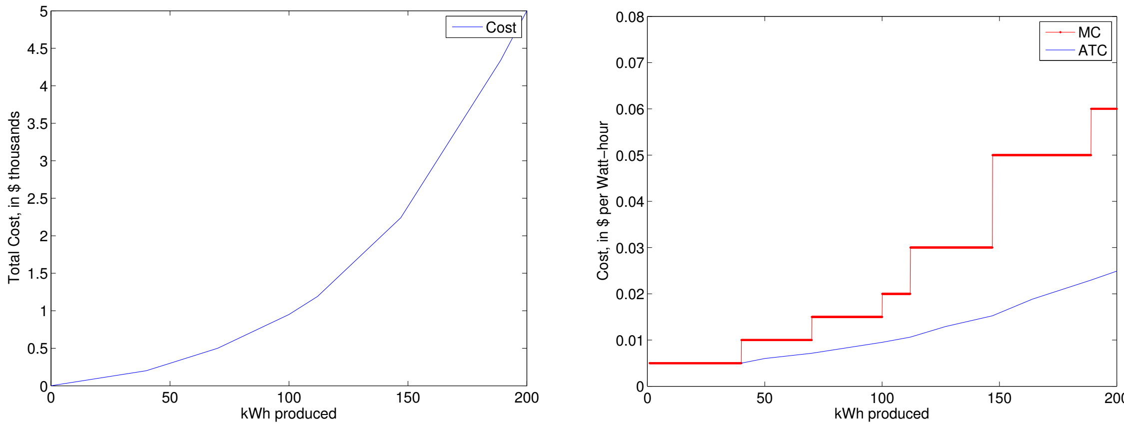

(a) Hourly Production Cost Curve (b) Average and Marginal Production Cost Curves

Figure 5.1: Hourly Production Cost Curve $C(Q_h^S)$ Assumptions

The figures generated above are the respective cost curves used in the simulations for this chapter. The x-axis in both Figures 5.1a and 5.1b denotes Q_h^S, aggregate hourly production. In our model, we had assumed that costs are the same for every hour and fixed costs are incurred uniformly during the day. The thicker marginal cost curve, in the figure to the right, is a piecewise function representative of all marginal production levels in merit order in our model. Production occurs at constant marginal cost for some incremental capacity before the system is faced with an increase in marginal cost to produce at the next incremental mode. The dispatch center always allocate production at lower marginal cost until the capacity constraints bind.

In our model, the true fixed costs of the energy technologies we simulate, like nuclear or coal power plants, are scaled by the total generation capacity, and added to the cost if the producer selects the plant to be online for the day. Scaling fixed costs permits us to restrict the problem to generation for 100 households and maintain meaningful values for cost. We utilize data on the leveled cost of generation resources projected for 2016 from the Energy Information Administration, Annual Energy Outlook 2011 (DOE/EIA-0383) to simulate cost curves.

Finally, assuming a budget balanced system, we set gross profit margin $\kappa \equiv 1$, for simplicity.

5.2 Simulation Results

Our network flow programs from the previous chapter and simulation tests are coded using A Mathematical Programming Language (AMPL) and utilize the IBM ILOG CPLEX Optimization Studio (CPLEX) commercial solver.

In this section, we compare results of our simulations using `MYOPIC_BEST_RESPONSE`, `ASYNCH_CONVERGENCE`, and for comparison, the optimal solution for the market under complete information, found by solving the joint program set up in Chapter 3. In our plots, the label with each of these algorithms represents the iteration at which the solution is taken (e.g. MyopBR-5000 is the state of the market at 5000 iterations of `MYOPIC_BEST_RESPONSE`).

5.2.1 Demand Smoothing

Figure 5.2 compares aggregate consumption over the day the three different simulations for the market with hidden information.

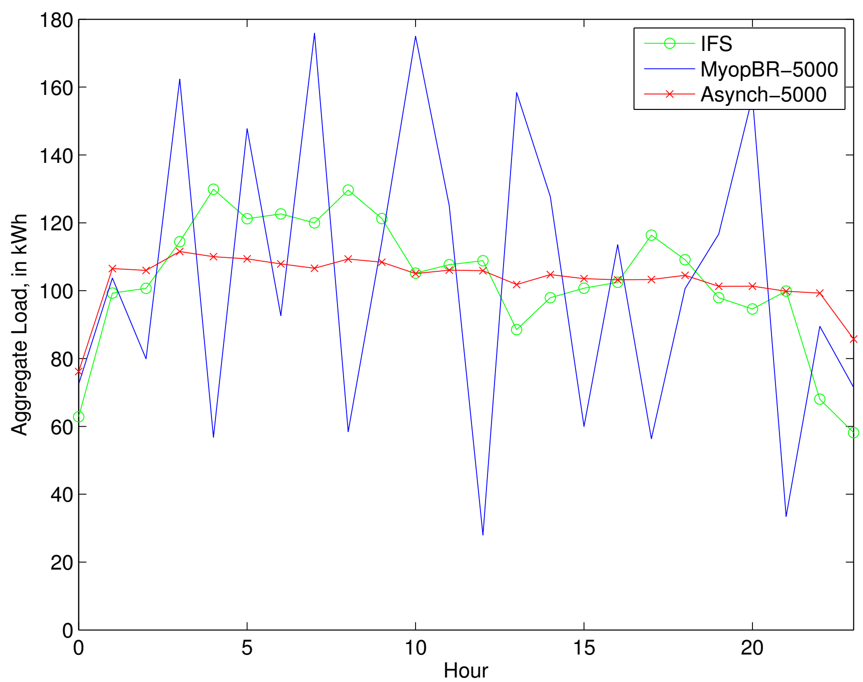

Figure 5.2: Comparison of Aggregate Household Consumption

The line labeled IFS is the initial feasible solution $\mathbf{x}^0$ used to seed both our algorithms.

To generate this solution, the aggregate household program is fed exogenously-determined hourly prices. Having generated hourly demand from a uniform distribution (time ranges for jobs were selected uniformly), we apply the same block price for each hour to mirror real-world block prices, which are set proportionally to load distributions. Because consumption changes are made in response only to relative discrepancies between hourly prices, the IFS solution is representative of a market outcome where no power scheduling is implemented (the initial price we set also does not matter). The solution is an outcome under block pricing.

The line labeled MyopBR-5000 is the consumption vector $\mathbf{Q}^{5000*}$ after 5000 iterations of `MYOPIC_BEST_RESPONSE`. The solution reflects consumption under real-time pricing, where households adjust scheduling in response to published hourly prices, but congestion effects are not felt. The wider fluctuations in demand between each hour in the solution are consistent with the result on real-time pricing and grid stability proven in Roozbehani *et al.* [11]; the algorithm does not converge to a stable solution. By providing households with real-time prices we introduce greater volatility in demand for the market, producing an outcome that is more unstable than even block pricing. With real-time pricing without producer response, we see an increase in peak-to-average ratio (PAR) to 1.70 from a PAR of 1.26 for the block pricing outcome. The total daily production cost of $32,444 for real-time pricing is 25.2% greater than the IFS cost of $25,915.

Finally, we utilize `ASYNCH_CONVERGENCE` to simulate the result of two-way coordination between households and the supply side to convergence. The line labeled Asynch-5000 plots the solution after 5000 iterations. The result in Figure 5.2 provides evidence of a demand smoothing effect in aggregate consumption, even though perfect smoothing is infeasible given the scheduling constraints in the dataset. With the `ASYNCH_CONVERGENCE` procedure, the solution tends toward the aggregate consumption profile expected at equilibrium, and is consistent with our theoretical results in Chapter 3. There is a reduction in PAR to 1.08 (14.3% reduction compared to PAR without demand side management implemented), and a total daily production cost of $24,546 (5.3% reduction compared to cost under block pricing).

By implementing power scheduling, we observe a 14.2% reduction in the maximum generation capacity necessary and a 50.7% reduction in peak-to-minimum spread in hourly demand. There is also an 83.7% reduction in the variance in hourly production by implementing power scheduling. This would allow the dispatch center to drastically reduce utilization of peaker plants and depend more heavily on cheaper base load power plants to provide more constant levels of power throughout the day. In general, the consumption result for our mechanism is smoother, suggesting our pricing mechanism does succeed in shaping demand, and the profile begins to resemble the theoretical optimum we derived in Chapter 3 for perfect smoothing.

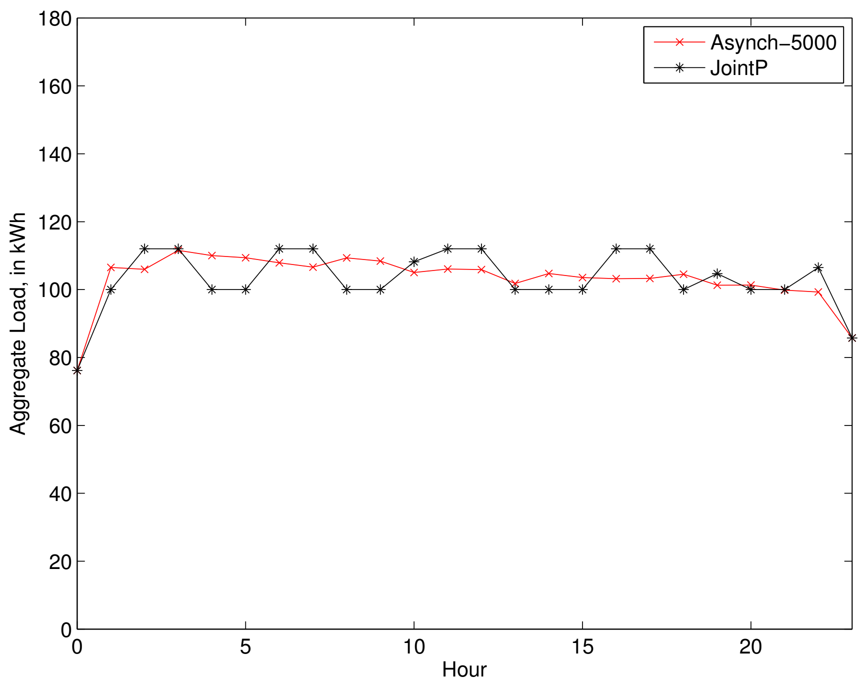

Figure 5.3: Hourly Load Comparison with Perfect Information

Figure 5.3 provides a comparison of the optimal solution found using our mechanism with the optimal solution to the joint program described in Chapter 3, which models the market under complete information. The solution labeled JointP is a solution set by the producer by optimizing over household constraints in an ideal market with complete information and is a pure-strategy Nash equilibrium for the producer and households. Ordinarily, however, the producer does not have this information. As a result, we depend on our mechanism to approximate this complete information solution using demand response to prices. It is interesting that the JointP solution appears less smooth than the Asynch-5000 solution. It is possible that the mechanism solution needs to be smoother hour-to-hour to reduce deviations by households as they re-optimize at each iteration. Nonetheless, the optimality gap between the market solution under complete information and the solution found by our mechanism is only 0.069%. The total daily production cost for the equilibrium solution under complete information is $24,529 (compared to $24,546 for Asynch-5000) and the peak-to-average ratio is also 1.08. Our mechanism is very effective in approximating the welfare-maximizing market equilibrium we found analytically in the market under complete information.

In a later section in this chapter, we analyze the stability of our result, since the solution found with `ASYNCH_CONVERGENCE` is an ϵ-Nash equilibrium.

5.2.2 Social Welfare

As our measure for social welfare, we look at the total aggregate cost of production, given that all constraints for scheduling are met. Figure 5.4a compares the hourly costs of production for the solution found after 5000 iterations of `ASYNCH_CONVERGENCE` to the solution after 5000 iterations of `MYOPIC_BEST_RESPONSE`. The smaller variance in production for the Asynch solution indicates a reduced dependence on backup generators, given that demand is more predictable. Overall capacity utilization (hourly production as a fraction of online generation capacity) for the solution in red is also consistently higher throughout the day. As we discussed above, overall cost for the solution found using our `ASYNCH_CONVERGENCE` mechanism is lower. Plotted against the hourly production curve, we see that the solution begins to resemble the perfect production smoothing strategy from our result in Chapter 3. Our simulation results provide evidence that reducing variance in daily production is an optimal strategy for improving efficiency.

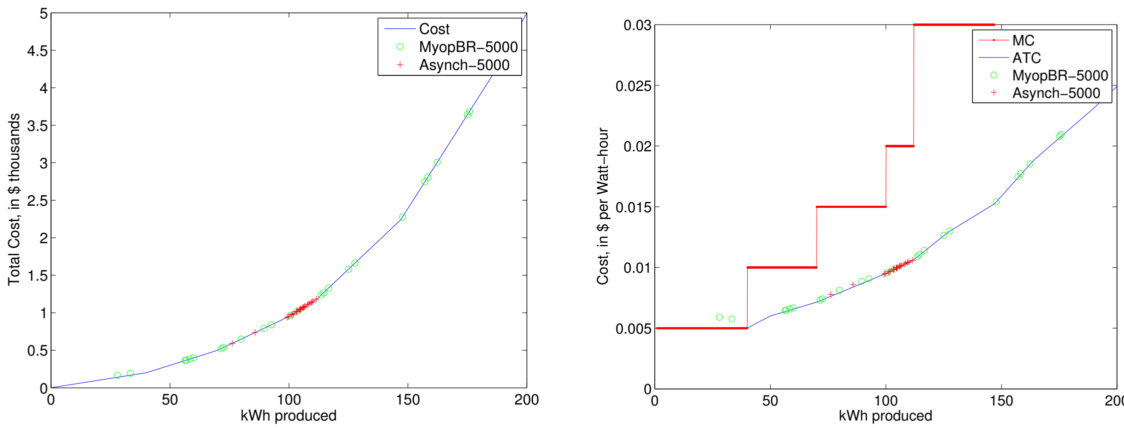

(a) Hourly Production Costs

(b) Average and Marginal Production Costs

Figure 5.4: Cost of Production Comparison

The plot to the right in Figure 5.4b shows the corresponding per unit cost and marginal production function for our simulation. Given the convexity of these functions, the solution points in red are optimal, and tend toward the theoretical perfect smoothing solution. In terms of maximizing social welfare, our mechanism does achieve its result in reducing the hourly costs of production and need for excess generation capacity throughout the day. The main result achieved is the reduction in variance in hourly production through demand load shaping.

We also plot the hourly production costs as time series for the solutions from our simulations in Figure 5.5 and notice the substantially higher costs during peak swings for the MyopBR-5000 solution due to the greater marginal costs at higher production modes. Cost convexity demonstrates the greater need to reduce peak periods of consumption as cost scales nonlinearly with production in each hour. Prices are more uniform in the Asynch-5000 solution and there is less incentive for households to shift consumption between hours; the consumption profile for households is more stable and cost minimizing for the producer.

From the perspective of households, the solution found through `ASYNCH_CONVERGENCE` leads to a 17.8% reduction in average household payment and 93% of households saw a reduction in daily payment as compared to the aggregate consumption profile after 5000 iterations of `MYOPIC_BEST_RESPONSE`.

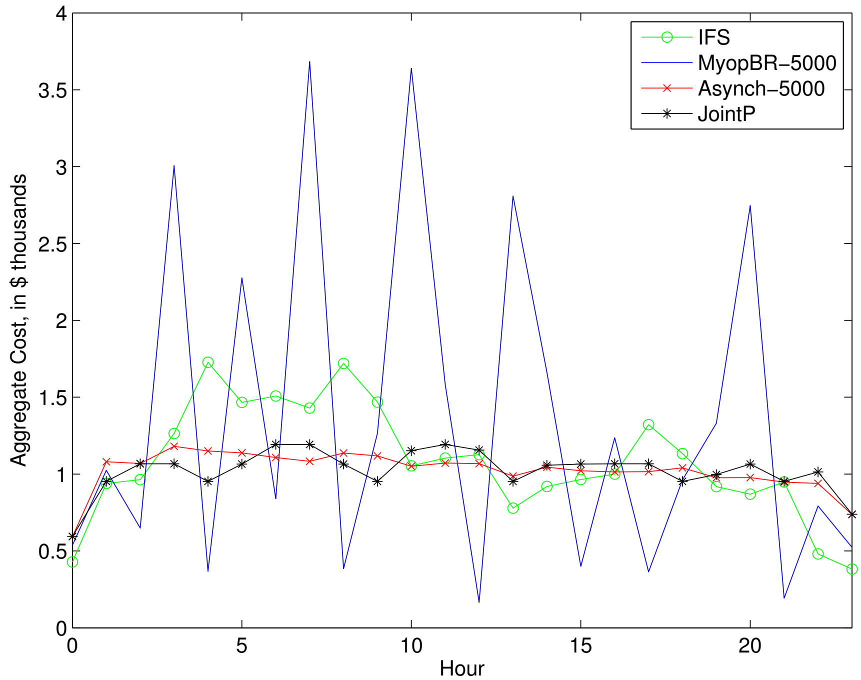

Figure 5.5: Aggregate Cost of Production Comparison

5.2.3 Convergence Results

In this section, we look for evidence of convergence in our mechanisms. Figure 5.6a plots the objective values for the economic dispatch (labeled Supply) and aggregate household appliance scheduling (labeled Demand) programs over iterations of the MYOPIC_BEST_RESPONSE procedure. It is clear that after 3-4 iterations, the mechanism begins cycling between solutions. The plot here provides evidence of our congestion game proof from the previous chapter in which rebound peaks appear. The MYOPIC_BEST_RESPONSE mechanism fails to converge as we anticipated because prices fail to incorporate the congestion effects of household actions. Thus, the production costs are actually greater than the perceived value of the objective for the household problem. The oscillations correspond to best responses as households continue to switch back and forth between time periods that are myopically cheapest, forming what the literature labels "rebound peaks" [3]. The simulation result in this plot helps illustrate our theoretical result explained in Figures 4.4a and 4.4b. Figure 5.6b corresponds to exactly these two figures from Chapter 4, plotting aggregate demand at the 4999th and 5000th iterations of `MYOPIC_BEST_RESPONSE`. As it is evident, aggregate demand continues to cycle between peaks and troughs as households myopically respond to prices that do not incorporate congestion.

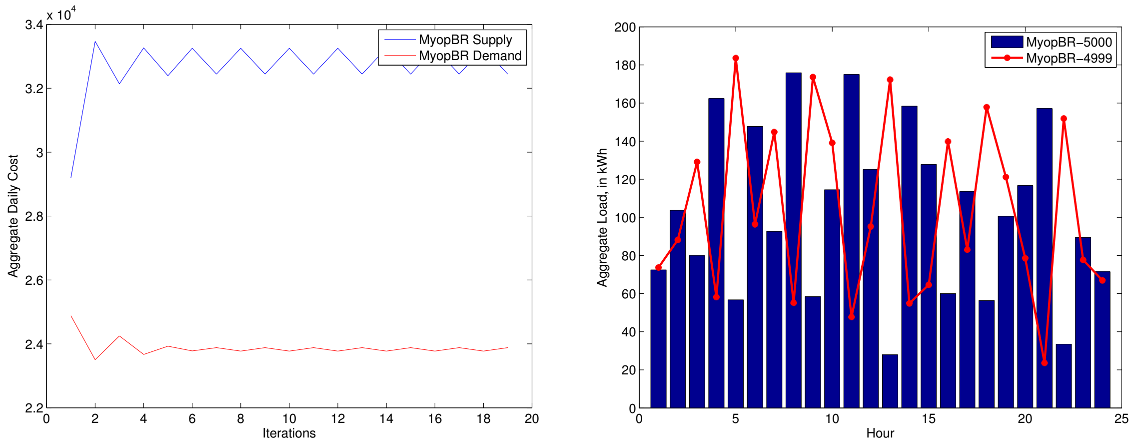

(a) Iterations of `MYOPIC_BEST_RESPONSE` vs. Objective Values

(b) `MYOPIC_BEST_RESPONSE` Rebound Peaks

Figure 5.6: `MYOPIC_BEST_RESPONSE` Convergence Results

It is worth noting that `MYOPIC_BEST_RESPONSE` enters this fixed oscillation of best responses relatively quickly. In Figure 5.6a, the objective values also never converge, implying that the social cost of production and individual cost of consumption are never the same, a socially inefficient outcome representative of the market when out of equilibrium. From this plot, we can see it is cheaper on average for households to consume electricity than it is for the producer to produce it, on a per unit basis. Ordinarily, this discrepancy is corrected by the producer's profit margin, but from a social welfare perspective, represents a wedge between the true cost of consumption and the cost faced by households in their consumption calculus.

Figure 5.7 plots the objective values of the supply and demand side programs over computational iterations of the `ASYNCH_CONVERGENCE` procedure. The two objective values closely track one another, as our mechanism succeeds in incorporating congestion effects into the prices published to households. The cost faced to households closely matches the true cost of consumption once congestion effects are felt. After two hundred or so iterations, the

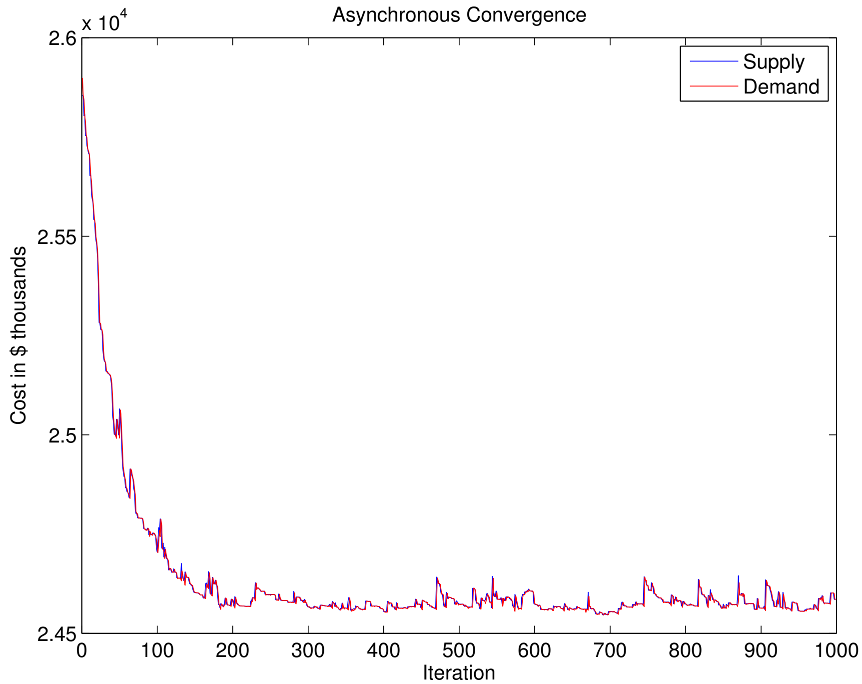

Figure 5.7: Iterations of `ASYNCH_CONVERGENCE` vs. Objective Values

procedure continues to asymptotically oscillate about the Nash equilibrium objective value found using the joint program. As we defined in the procedure, we expect convergence to an ϵ-equilibrium as subsets of households randomly selected to move will continue to best respond. In our simulation, we restrict the parameter k to allow no more than 10% of households to best respond at each iteration of the mechanism and note that fluctuations after 200 or so iterations are minimal as a percent of total production cost. The graph here provides evidence of the rate of convergence in our mechanism in approximating the market equilibrium found with complete information.

Computationally, each iteration of the mechanism is easy to solve given our formulation from the previous chapter. Requiring just over a couple hundred iterations to reach our ϵ-equilibrium for 100 households, the procedure converges quickly and it is not very computationally intensive to determine prices for the market using our algorithm.

5.3 Discussion

The goal of this chapter is mainly to provide verification of our theoretical results from previous chapters through simulation. The results here also test the programs we set up in Chapter 4. Constructing a dataset of 100 households, we simulate market interactions under various pricing schemes to analyze outcomes and compare performance.

We first compare welfare results for consumption and production with and without power scheduling implemented. Our simulations provide empirical evidence of a demand smoothing effect in aggregate consumption rendered by our mechanism, suggesting prices and consumption tend toward the socially efficient outcomes we derived analytically in Chapter 3. We analyze the implications to cost, consumption, and online generation capacity, confirming the motivations behind our model for power scheduling. By utilizing our mechanism for prices with the restricted information transfer set up by our complete market model, we are able to achieve a 5.3% reduction in production costs using our test dataset and a 14.3% reduction in peak-to-average ratio in demand, compared to block pricing. In the long term, this also reduces the need to expand generation capacity ahead of demand growth, improving capacity utilization and production efficiency.

We find that the market outcome with our mechanism differs only by 0.069% compared to the cost of production for the optimal solution found using the joint program from Chapter 3, which simulates the market under complete information. This result provides evidence that our mechanism, while imperfect, is still able to produce an outcome within a very small optimality gap of the market outcome under complete information. Though our solution is still suboptimal for production, given only the restricted information transfer of prices and aggregate demand, we are able to approximate efficient market outcomes comparable to those under complete information.

We also turn to a discussion of our mechanisms for determining equilibrium prices. We test both of the procedures proposed in the previous chapter and confirm the congestion game result for the `MYOPIC_BEST_RESPONSE` algorithm. Our results in this chapter confirm expectations set in Chapter 4 that the procedure will not converge, and our intuition proves true given the oscillatory cycling we see in our simulation results for iterations of the MYOPIC_BEST_RESPONSE procedure. Our simulation results provides evidence that our ASYNCH_CONVERGENCE mechanism does however converge asymptotically, as anticipated, to an approximation of market equilibrium and we build further intuition through simulations on the rate of convergence to equilibrium. For both of these procedures, the rate of convergence or to oscillations is relatively quick, making our ASYNCH_CONVERGENCE procedure practical for setting equilibrium prices in real-time.

Our simulation results are very promising and ideally we would extend our analysis to measure robustness and sensitivity as future work. Though we depend on only a restricted set of simulations to confirm our theoretical work, it sufficiently provides evidence for the effects we anticipated, in terms of demand smoothing and congestion effects, and demonstrates the promising outcomes from our mechanism design, even when compared again the market outcome with complete information.

Chapter 6

Conclusion

Our simulations in the previous chapter are a promising assessment of our work and validate the theoretical results proven earlier in this thesis. Here, we conclude with a brief review of our motivations, methods, and results, as well as with a critique of assumptions made in our model and directions for future work.

We end the chapter with some concluding remarks about our work and its place in the growing literature on demand response.

6.1 Brief Review

In this thesis, we developed a pricing mechanism in order to coordinate supply and demand in the market for electricity generation. We proposed a retail market structure that enabled producers to communicate day-ahead hourly prices to households, in exchange for each household's proposed load schedules as implicit payment. We formulated the optimization problems solved by each side of the market and provided a game theoretic model for understanding the interaction between participants in the ideal market under complete information.

The mechanism for prices we developed instead incorporates a restricted transfer of information across the true market with hidden information, in order to better optimize electricity production and appliance use in the home. Our objective through implementing our mechanism has been to achieve more socially efficient production, through which we may reduce the dependence on expensive peaker power plants and improve capacity utilization. In Chapter 3, we prove the existence of a Nash equilibrium in the market under complete information and through our simulation results in Chapter 5, confirm that our mechanism succeeds in approximating this social welfare maximizing outcome. More work is nonetheless left to analyze parameter values in our mechanism to measure the rate of convergence as well as to consider market outcomes under a relaxation of our modeling assumptions. Nevertheless, our work is promising for future research in utilizing demand response to improving production outcomes in the market for electricity generation.

6.2 Future Work

The ideas we developed can be extended in several directions. We review a few of our key assumptions here and propose extensions for future work:

- In our model, we have assumed the total daily consumption for each household is set at midnight and no new tasks can be submitted for scheduling once the day begins. We may however measure sensitivities to these parameters for our solutions to analyze the impact of unexpected shocks in consumption throughout the day. While we rationalize that most large appliance use of consequence can likely be anticipated, it would prove insightful to still understand the robustness of our production schedules to changes in total household demand. Our framework could be restructured to recourse production decisions every hour to incorporate the introduction or omission of household tasks throughout the day.
- All aspects of the household's parameters for each task (i.e. release and deadline times) are actually taken to be deterministic when submitted for scheduling. Instead, we may model the program with uncertainty on release and deadline times, updating posterior expectations by tracking each household's consumption patterns before the mechanism is in place. Given this modeling challenge, we may extend our model to update appliance schedules in other households for later in the day to balance out unexpected changes in real-time consumption.

- The perfect production smoothing result in Chapter 3 is largely enabled by our assumption that the production cost curve is the same for each hour. We may lift this assumption and model the market as variable costs for generating plants change throughout the day. We may also introduce uncertainty in costs or model shocks to production in terms of unexpected plant failure.

- More complex behavior for generating plants can be modeled, as we have made some simplifying assumptions in formulating the economic dispatch problem. We may incorporate ramp constraints restricting rapid fluctuations in production levels at a plant or impose restrictions on the minimum number of time periods a plant is allowed on and/or off. A spatial model for the power grid may also be useful in scheduling production at generating plants to minimize transmission distance to demand. Network flows provide an ideal framework for modeling this latter extension.

- We may also test different parameter values for fixed-to-variable cost ratios to develop insight into when high-fixed/low-variable cost technologies may be preferred to low-fixed/high-variable cost production at different types of generating plants across the system.

- In our model, we assumed uniform preference for scheduling across time ranges $[\alpha, \beta]$ submitted by households for each task. Our motivation was to design a system as minimally intrusive to households as possible: the only information we required for setting up a task for scheduling was the power setting, duration, release, and deadline times. Our work can be extended to analyze scheduling outcomes given other distributions for preference over time ranges $[\alpha, \beta]$, such as a truncated normal. In line with the literature on optimal delays for start times, we may also query users for time ranges to indicate preference for job start (or completion) times and assume exponentially-distributed preference over ranges specified. Our result for demand smoothing is made more feasible by our assumption of uniform preference over intervals, though as the number of households and tasks submitted increase, smoothing is still viable under other distributions for time preference.

- In the ASYNCH_CONVERGENCE procedure, we set an upper limit on the number of households allowed to move in each iteration as the parameter k. This value restricted the cardinality of subsets in $\wp_k(\mathcal{N})$ from which we sampled households to provide best response in each iteration. By varying the parameter k we may measure the effect on the rate of convergence. In our simulation, we held k constant at 10% of households. Reducing k limits the number of households that move before prices update to incorporate congestion effects. At higher values of k, we may see our ASYNCH_CONVERGENCE procedure begin to break down as prices will begin to no longer reflect the true cost of consumption.

6.3 Conclusion

In this thesis, we formulated models for the wholesale and retail electricity markets to explore the problem of information asymmetry. Our solution proposes the transfer of prices and aggregate demand across the market, compiling these values from centralized optimization programs engineered for the producer and households. We designed a mechanism for prices to improve production outcomes and demonstrated that even with a minimal transfer of information in the market, improvements can be seen in social welfare.

Simulation results provide support of our theoretical work in earlier chapters, illustrating the success of our proposed mechanism in reducing the costs of production. We have shown that incorporating the demand side of the market into a solution for minimizing production waste is almost essential. Demand response provides the key toward improving efficiency in the market. By aligning production into a load-shaping strategy as opposed to load-following, we are able to mitigate the current risks borne by the system in terms of demand uncertainty. Though information asymmetry is at the core of the problem in electricity markets, our work proves promising at finding a solution for a century-old problem.

References

[1] D. Bertsimas and J.N. Tsitsiklis. *Introduction to Linear Optimization.* Belmont, MA: Athena Science, 1997.

[2] S. Borenstein, M. Jaske, and A. Rosenfeld. Dynamic Pricing, Advanced Metering, and Demand Response in Electricity Markets. UC Berkley: Center for the Study of Energy Markets, 2002.

[3] C. Chen, S. Kishore, and L. Synder. An Innovative RTP-Based Residential Power Scheduling Scheme for Smart Grids. *ICASSP,* May 2011, Invited Paper, to appear.

[4] G. Constable and B. Somerville. *A Century of Innovation: Twenty Engineering Achievements That Transformed Our Lives.* Washington, D.C.: Joseph Henry Press, 2003.

[5] U.S. Department of Energy. *The Smart Grid: An Introduction.* 2009.

[6] P. Hall. On Representatives of Subsets. *J. London Math. Soc.* 10 (1), pp. 26-30, 1935.

[7] A. Ipakchi and F. Albuyeh. Grid of the Future. *IEEE Power Energy Magazine,* vol. 8, no. 4, pp. 52-62, Mar. 2009.

[8] S. Kishore and L. Snyder. Control Mechanisms for Residential Electricity Demand in SmartGrids. In *Proceedings of IEEE Conference on Smart Grid Communications,* 2010.

[9] A. Mohsenian-Rad and A. Leon-Garcia. Optimal Residential Load Control With Price Prediction in Real-Time Electricity Pricing Environments. *IEEE Transactions on Smart Grid,* vol. 1, no. 2, pp. 120-133, 2010.

[10] A. Mohsenian-Rad, V.W. Wong, J. Jatskevich, and R. Schober. Optimal and Autonomous Incentive-based Energy Consumption Scheduling Algorithm for Smart Grid. In *Proceedings of IEEE PES Conference on Innovative Smart Grid Technologies,* Gaithersburg, MD, Jan. 2010.

[11] M. Roozbehani, M.A. Dahleh, and S.K. Mitter. Volatility of Power Grids under Real-Time Pricing. LIDS Report, 2011.

[12] J. B. Rosen. Existence and Uniqueness of Equilibrium Points for Concave n-Person Games. *Econometrica,* vol. 33, pp. 347-351, 1965.

[13] P. Samadi, H. Mohsenian-Rad, R. Schober, V. Wong, and J. Jatskevich. Optimal Real-time Pricing Algorithm Based on Utility Maximization for Smart Grid. In *Proceedings of IEEE Conference on Smart Grid Communications,* 2010.

[14] Y. Shoham and K. Leyton-Brown. *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations.* New York: Cambridge University Press, 2009.

[15] A. Tucker. *Applied Combinatorics.* New York: John Wiley & Sons, 2007.

[16] G. Xiong, C. Chen, S. Ishore, and A. Yener. Smart (In-home) Power Scheduling for Demand Response on the Smart Grid. In *Proceedings of IEEE PES Conference on Innovative Smart Grid Technologies,* 2011.