

EFFICIENT COORDINATED POWER DISTRIBUTION ON PRIVATE INFRASTRUCTURE

by

Andrew Perrault

A thesis submitted in conformity with the requirements
for the degree of Master of Science
Graduate Department of Computer Science
University of Toronto

© Copyright 2013 by Andrew Perrault

Abstract

Efficient Coordinated Power Distribution on Private Infrastructure

Andrew Perrault

Master of Science

Graduate Department of Computer Science

University of Toronto

2013

Due to the design of current power distribution networks, it will become increasingly attractive for agents to generate their own power (distributed generation) and to construct private infrastructure (e.g., transmission lines) to exchange power with others nearby without using the main public grid. We show that such private transactions may increase overall load on the network because of the increased distance that power must flow from generation sources, thus increasing transmission loss. We present a coordination scheme that allows centralized control of private infrastructure while satisfying participation constraints and budget balance. Experiments show that our scheme reduces transmission losses by 4-5% when there are only a constant number of private lines and by 55%-60% when the number of private lines is proportional to the number of agents. These results hold even when there is only a small amount of distributed generation.

Contents

| | | |
|----------|---|-----------|
| 1 | Introduction | 1 |
| 2 | Background | 6 |
| 2.1 | The Power Grid | 6 |
| 2.2 | The Smart Grid | 7 |
| 2.3 | Game Theoretic Concepts | 8 |
| 2.4 | Network Games | 11 |
| 3 | Agent Behavior in the Power Distribution Game | 12 |
| 3.1 | Setting | 12 |
| 3.2 | Models of Agent Behavior | 15 |
| 3.3 | A Simple Incentive Scheme | 20 |
| 3.4 | No-Resistive-Losses Model | 23 |
| 4 | Experiments | 26 |
| 5 | Conclusion | 31 |
| | Bibliography | 34 |
| A | Market Games | 37 |
| B | Power Distribution Game Without Resistive Losses | 41 |

Chapter 1

Introduction

The *smart grid* can be roughly defined as an augmented electrical grid that gathers information about its own operation—specifically the behavior of consumers and producers—to automatically improve efficiency and reliability. Utility companies have been the primary driver of development of the smart grid thus far, particularly because they see the opportunity to replace monitoring tasks typically undertaken by people (e.g., reading electricity meters, looking for damage to transmission lines, monitoring transformers) with electronic sensing and control devices that are accessible from a central location. This stage of development is still ongoing—smart residential meters themselves are still only available in relatively few places—but access to comprehensive grid data will provide exciting research opportunities in many fields.

While much smart grid research has focused on the development of the infrastructure necessary to monitor operation, full realization of the smart grid requires intelligent control and incentive schemes to make the best use of information so gathered [17]. As a consequence, smart grid research should attract the attention of the computer science, artificial intelligence, machine learning and computational economics communities. We briefly list some of the interesting problems in the area.

Recommender systems [3] can use large amounts of user activity data to make behavior suggestions to users. Historically, they have been primarily used by web-based businesses that both have access to user activity data and communication capabilities that can be cheaply customized for each user. Data collected from the smart grid will allow utilities to make specific behavior recommendations to households and businesses; however, these systems will be challenging to design as there will be an immense amount of extremely high-dimensional data. Some actions that might be suggested by a recommender system are new appliance purchases, shifting loads to off-peak hours, and fault detection.

These systems will require machine learning to use data to conjecture actions that users have taken, identify appliance profiles for a household, and to infer user preferences from a limited number of observations. There are also opportunities for machine learning in sensor denoising and fault detection, at both the household and grid levels.

It might be desirable to augment a recommender system with a survey functionality, as it is possible to learn the shape of a user’s utility function, i.e., their preference for events that have not occurred, to make more helpful suggestions. This problem is in the domain of preference elicitation [16] or active learning [18].

Another major area of improvement is how *power exchanges*—the way agents can declare and bid for electricity—on the grid itself work, which is the focus of this thesis. We will describe in detail the problem of how electricity should be priced in order to make maximally efficient transactions. We view the problem as one of coordination, where we assume that there is some trust among the agents. From a more adversarial perspective, it can also be viewed as a *mechanism design* problem [13], where the agent types represent the electricity supply or demand of each agent, and the goal is to adjust the payoff vectors so that an equilibrium with optimal social welfare is implemented. There are several other interesting potential uses of mechanisms that fall under the category of contract design: energy generation and consumption prediction, consumer and producer pooling for electric utilities, and infrastructure and generation cost sharing.

In this thesis, we will abstract away some of the details of the current system and physics and focus on the key principles behind price design. Briefly, current electricity pricing regimes in many jurisdictions require that electric utilities report the expected demand of their consumers, and generators report their expected production, to an *independent system operator (ISO)* [7]. In the United States, ISOs are nonprofit corporations formed at the recommendation of the federal energy regulator in order to facilitate increased transaction volume between utilities resulting from electricity deregulation. ISO-controlled electricity markets are examples of *mechanisms*, where a central controller declares a protocol for mapping declarations by market participants to transactions and payments. The market participants in this context are utilities and large industrial consumers on the demand side and generators on the supply side.

The ISO typically solves a linear relaxation of the *electricity distribution problem*, telling generators how much power to produce (with prices set using dual variables of the optimization). The linear relaxation neglects the non-linear physics of the system—it will be more accurate in relatively uncongested networks or those with only short transmission lines than when congestion and lengthy transmission amplify the impact of the physical non-linearities—but it can be solved efficiently and transparently.

An underlying tension in market design is between easy to calculate, transparent prices and prices that incentivize optimal behavior. The ISO solution is relatively complex, and the specific implementation is analyzed and simulated by market participants to calculate ways they can misrepresent their supply or demand and production costs to increase the amount they are paid by the ISO. One area our perspective could be applied is to improve the stability guarantees of ISO algorithms. While it seems that electricity buyers and sellers have no option but to participate in ISO exchanges, markets that are not provably stable or approximately stable have historically been outperformed by their stable counterparts [21]. ISOs have existed for around 15 years.

At the household and small-business level, most consumers pay fixed rates to their electric utilities representing the average cost of generation and distribution on their utility’s network over a long period of time.¹ These prices do not give the consumer any information about the true cost of electricity at a particular moment. Electricity generation has roughly two “layers:” a base load and peaking plants. The base load is cheap to operate, but it cannot be adjusted to current demand. Coal, hydro and nuclear are three examples of base load power sources. Peaking plants are expensive and can be adjusted quickly in response to relatively rapid changes in demand. Oil and gas turbines are examples. The result of this system is that electricity has a low median price and a high average price.

In recent years, it has become increasingly easy for “household” consumers to generate electricity from renewable sources at cheaper cost and with lower emissions—most notably wind and solar power. Installation capacity is often higher in rural areas where space is available. Groups of such consumers can often limit their dependence on the public grid by exploiting local generation, e.g., in the form of *microgrids* [10]. Microgrids are groups of generators and consumers connected to the larger grid at a single point, which allows decoupling in the event of grid failures.

For instance, consider a rural area in a developing country where a significant amount of power is produced by wind farms, solar panels or biofuels. These tend to be located nearer to farms than to urban centers, so some of these local generators may strike deals with nearby consumers (e.g., farms) to construct private transmission lines and exchange power outside the public grid. Indeed, due to fixed rate-schedules that mediate transactions with the public grid, it is often more profitable for them to sell directly to their neighbors than to the grid. However, as we see below, such transactions can reduce overall efficiency. The network in Figure 1.1, explained further in Chapter 3, illustrates such a setting.

This dynamic creates tension between the publicly-managed grid and the private agents it serves. While the public grid may wish to construct additional transmission capacity, especially in the presence

¹In areas that have received smart meters, time-of-use rates are often used, which designate a separate price for “peak” and “off-peak” times.

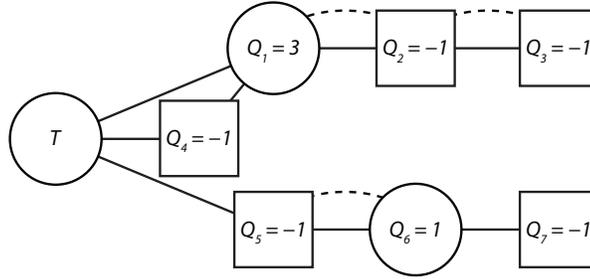


Figure 1.1: A network with private and public links. (a running example used in Secs. 3.1 and 3.2). Circles represent net generators and squares net consumers. Node T represents the connection to the outside power grid. Solid lines are publicly-owned links and dashed lines private links. We assume that public and private links have roughly the same voltage and resistance. Q_i denotes the supply of agent A_i .

of increasing distributed, unreliable generation, the benefit of doing so is limited by the degree of private control of transmission lines allowed. In this thesis, we develop a routing scheme that allows agents in a local “microgrid” to coordinate private distribution with the publicly-owned utility. Our scheme achieves optimal distribution at minimum cost, while satisfying participation constraints in the sense that each agent is at least as well off as in the uncoordinated regime. Our scheme can also be used as a basis for determining cost-sharing of new local generation and transmission infrastructure among agents by quantifying their individual benefits assuming optimal usage.

This rest of the thesis is organized as follows. In Chapter 2, we describe the current electricity grid, smart grid plans and related research and provide background on the relevant game theoretic concepts and briefly discuss previous research. In Section 3.1, we outline our model of the networks involved; including agent utility functions, the physics of electricity distribution, and admissible strategies. In Section 3.2, we describe several models of system behavior and describe how they are related to incentive schemes. In Section 3.3, we present an incentive/routing scheme that coordinates agents on the private network with the public network, maximizing aggregate agent utility across the distribution network. We perform experiments in Chapter 4 that compare the performance of our system under various models of agent behavior introduced in the previous chapter. We use random graphs to test performance under a variety of assumptions, and we find that there is a large gap in performance between current pricing schemes and the most sensitive ones. We present conclusions and other directions for future work in Chapter 5.

The electricity generation and distribution context we describe and analyze is necessarily speculative in certain aspects as it describes a system that does not yet exist and is still under debate by policymakers. From a high level, we assume that small producers and consumers will have increased control over trading decisions in future systems—that the most desirable way to increase the efficiency of electricity markets

is by deregulation made possible by smart grid engineering allowing for more fine-grained control of routing decisions across the grid. It is possible that the costs of deregulation due to increased cognitive burden on market participants will outweigh its efficiency benefits in the short-term, but we think it is important to investigate the value of greater decentralization given that technology will soon make this possible. Our model thus emphasizes that agents must be incentivized to participate in any centralized mechanism, i.e., it is possible for any agent to make outside arrangements which they will prefer if they are more beneficial.

Nearly all research in power distribution uses some kind of approximate model of electrical physics as the exact power flow equations are highly non-convex. We simplify the physics substantially to focus on how the dynamics of power grids as mediums for electricity exchange between pairs of agents. The most significant simplification we make is neglecting sympathetic flows which essentially impose constraints on the fraction of power out of a node on each outgoing edge. We believe that these constraints could be integrated into our optimization framework and priced, but this is an important area of future work. We describe our physics model in detail in Section 3.1 and Section 3.2.

Computational considerations are another important area for future consideration. We focus on the microgrid setting, where most networks have less than 150 nodes. In this context, our quadratic convex optimizations are not difficult to solve (took under a second), but scaling problems might be encountered in larger settings. We describe the computational issues in more detail in Section 3.3 and Chapter 4. The number of optimizations that must be performed in the case that the system operator does not have knowledge of previous transactions between agents is a potential computational bottleneck—it is expensive to compute the payments that must be given to agents to incentivize them to participate in the new exchange mechanism, as far as we know. It is likely that a heuristic or approximate algorithm would be necessary in practice.

Chapter 2

Background

2.1 The Power Grid

In this section, we briefly describe the current power grid and how we model it in our analysis. Since we do not have access to data describing the properties of current power networks, our approach to modeling them will be to abstract away some of the physical properties and test a wide variety of parameters for the others. Privacy and security concerns prevent most grid-related data from public release.

Our basic model for the power grid is a tree-like publicly-owned grid that has very little monitoring equipment on the lines and simple meters at the connection points to users. We will represent the public grid by a graph $G_{pub} = (V, E_{pub})$. Assume that all users are connected to the public grid. Our basic model for a user v 's *utility function* will be a single number Q_v representing the supply or demand for power of user v . A positive Q_v represents a net supply and a negative one represents a net demand. We use a single number for the sake of simplicity; the number can either represent a long-term average of the supply or demand of the user or demand over a small slice of time. Following empirical work in this area, we assume that v 's supply or demand is essentially unaffected by prices in the short term.

To explore the incentive properties of power exchange schemes, we will assume that there exist some *privately-owned transmission lines*—for example, as part of a microgrid. We will refer to these lines collectively as the *private grid* and represent it by a graph $G_{priv} = (V_{priv}, E_{priv})$. See Figure 1.1 for an example of a grid with public and private links represented by a graph. Not every agent needs to be connected to a private link. In our experiments, our model for constructing private edges is a random graph.

Our major sources for the physics of power distribution are Wood and Wollenberg (1996) [28] and

Saad et al. (2011) [22]. The former describes many of the key problems of power engineering, including predicting power flow through a network, which has been extensively studied. The most general version of the problem poses a difficult non-convex optimization where progress is still being made [11]. Nearly all commonly-used methods neglect *resistive losses*, which are an important part of our analysis. The general framework for these methods is to simplify the problem a lot or a little, perhaps making it convex, and then solve the simplified model either exactly if it's linear and convex or via a numerical method otherwise.

Following Saad et al. (2011), we make several simplifications to the physics of power distribution that go beyond simple power flows, while continuing to model resistive losses. We model alternating current as direct current to avoid having to deal with complex numbers. We also neglect Kirchhoff's voltage law, which causes current to flow sympathetically around loops. We assume that physics minimizes resistive losses; the resulting performance can be viewed as an upper-bound for real physics. It would not be difficult to use more complicated physics models that make use of more specific information about the properties of the grid, if we had access to that information.

Resistive losses are a key aspect of our model of the grid. When power is sent from agent u to agent v , some the power dissipates in transmission. The amount of power lost is quadratic in the amount transmitted and linear in the length of the transmission lines. We will denote the amount of power sent from u to v as $f_{u,v}$ and the amount of power received from u at v as $f'_{u,v}$ (after losses are deducted).

2.2 The Smart Grid

Smart grid research has largely focused on the development and implementation of sensing and actuating hardware, usually with the goal of being able to monitor and control the grid for an area from a single location. The objectives of this approach are to reduce resistive losses, efficiently manage generation sources, and increase reliability. In recent years, there has been increasing interest in smart grid applications in the artificial intelligence community [17], particularly focused on how market participants can cooperate to increase reliability and decrease cost. We will briefly describe several notable papers on smart grid applications from an artificial intelligence perspective.

The most similar approach to our work is that of Miller et al. (2012), which formulates optimal power distribution as a decentralized multiagent coordination problem, with the goal of reducing the amount of computation required [12]. Our objectives are different because we are primarily interested in the incentives of the agents to cooperate in the distribution scheme.

One common theme of algorithmic smart grid research is to shape demand to reduce maximum loads

and react to fluctuations in supply availability (e.g., by shifting some peak demand to off-peak hours). Peak demand is a critical driver behind supply costs because it is typically provided by finely adjustable oil or gas turbines that are much more expensive to operate than other generation sources. Large scale distributed battery capacity can help address this problem, as cheaply-generated power can be stored until peak hours and then resold. One potential source of battery capacity is idle electric cars, but any trading schemes will have to take into account when the car will be needed for transportation, and owners are unlikely to want to manage trading manually. Vytelingum et al. (2010) designs decentralized coordination algorithms that maximize the efficiency of distributed storage [27]. Their model of the market is simple, but shows that without coordination algorithms, storage can decrease market stability due to herding.

Similar coordination issues exist on the supply side of the market, e.g., among wind farms. Wind power is unpredictable, making it difficult for generators to negotiate contracts with ISOs. *Coalitions of generators* in geographically disparate areas can mitigate the impact of unpredictability, making it easier for wind generators to enter the market [19]. By using *scoring rules*, they can incentivize accurate prediction and reduce insurance costs.

Another approach is to form coalitions of consumers that have complementary demands. These coalitions can be cheaper for the generator to supply to the extent that they have flatter aggregate load profiles, accurate predictions of future use, and low consumption variance. Stable coalition formation allows the generator to pass on savings to agents that have “useful” consumption patterns (e.g., that improve overall efficiency) without having to calculate personalized prices for each agent. Vinyals et al. (2012) studies coalition formation for consumers under the assumption that larger coalitions are harder to manage [25].

Since consumers do not want to constantly make decisions about power usage based on supply availability, intelligent agents may be able to manage electricity usage and storage based on user preferences. Vytelingum et al. (2010) develops a continuous double-auction protocol along with trading strategies to maximize market efficiency [26]. They use the full DC power flow approximation and assume that no resistive losses occur. However, their analysis does not address participation constraints or market stability.

2.3 Game Theoretic Concepts

We will briefly define the relevant game theoretic concepts used to model electricity distribution markets. Our basic view is that the players are the producers and consumers on the network, and the actions

are the choices of where to send power over private infrastructure and what payments to make to other agents. We assume that individual agents cannot determine the routing that occurs on publicly-owned links, but they can reroute any portion of incoming power over whichever private edges they choose. This assumption is unrealistic if nodes on the graph are viewed as individual households and businesses on the current electricity grid because it would require sophisticated routing equipment. It is reasonable in the context of a community-owned microgrid.

The problem can be viewed as either a competitive or cooperative game depending on the modeling assumptions. We will describe both, but our primary analysis will focus on the cooperative perspective because it neatly deals with the issue of agreements between agents. A *competitive game* consists of a set N of *agents (players)*, a set A_i of actions for each player i , and a function that specifies a *payoff* for each player $u_i : A \rightarrow R$ depending on the the action selected by each player [15]. Since our players coincide with the nodes of a network, we will use the set V to represent the group of players.

We will model the game as having two stages, although only the first will be strategic. In the first stage, each agent connected to a private edge sets the flow over the outgoing private edges. In the second stage of the game, all remaining agent supplies or demands are sold to or bought by the public grid. In this model, we will assume that the each agent pays the grid a fixed sell price p_s or a fixed buy price p_b for each remaining unit of supply or demand. Note that flow does not need to be conserved in the first stage because the public grid will buy or sell any amount of electricity in the second stage.

As the mechanism designer, our goal will be to determine a set of payoffs such that the agents choose actions that minimize the aggregate cost to all agents. To predict the outcome of a game, we will refer to solution concepts from game theory. The most prevalent of these is the *Nash equilibrium*, which consists of an action for each player such that no individual agent has incentive to change her action given that the others do not change theirs. Formally, a Nash equilibrium is an action profile a such that $u_i(a) \geq u_i(a_{-i}, a_i), \forall a_i \in A_i, \forall i \in N$ where a_{-i} denotes the actions of all agents except i . A natural strengthening of this concept is the *Strong Nash equilibrium*, a *action profile* (an action for each agent) where no group of agents has incentive to change their actions given that the other agents do not, that is, an action profile a such that $\sum_{i \in S} u_i(a) \geq \sum_{i \in S} u_i(a_{-S}, a_S), \forall a_S \in \times_{j \in N} A_j, \forall S \subseteq N$.

One potential difficulty with this approach is that the set of valid moves for a player depends on the moves made by other players. Often, games will have multiple equilibria, which can make enforcing such constraints difficult, as agents do not know which equilibria is active, and they may have incentives to mislead other agents regarding the action they will play.

Viewing the problem as a cooperative game leads to a different perspective. In a *cooperative game*, each agent cooperates with a *coalition* of agents, forming a contract to pick the action that is best for

the group [15]. In a standard model of a cooperative game, the aggregate utility of any group of agents is independent of which coalitions are formed by the other agents and the actions they take. Thus, we can define a *characteristic function* $v : 2^N \rightarrow \mathbb{R}$ where $v(C)$ represents the maximum *aggregate utility* of a group/coalition C can guarantee themselves by operating effectively.

The coalition then decides how to distribute its aggregate utility among its participants—hopefully giving enough to each participant that they do not regret joining that coalition. To formalize this, the *core* solution has all agents cooperate and sets payments such that the net payment to each group of agents is greater than what they would receive if they were to join a different coalition or *defect* [15]. Specifically, let t be a vector of payments where t_i is the payment to agent i . A vector t is a core allocation if

$$\sum_{i \in N} t_i = v(N) \tag{2.1}$$

$$\sum_{i \in C} t_i \geq v(C) \tag{2.2}$$

In the electricity distribution problem, a coalition consists of a group of agents that exchange electricity among themselves. Finding a core payment vector corresponds to finding a “personalized” pricing scheme such that no group of agents defects and forms a separate “market” or “trading group”. In the case where transmission losses between two agents are independent of transmissions between other agents, the cooperative game theory framework can be applied directly to efficiently calculate prices. This can be accomplished by modeling it as a *market game* [24] or *market game with transaction costs* [20]. We review some of the basic theory of market games, including the proof that core payments always exist, in Appendix A.

Difficulties arise as a result of non-independent congestion losses on transmission framework. The total loss incurred will be quadratic in the amount transmitted on a given line; details will be given in Section 3.1. There are several modeling choices that can be made, depending on whether transmission lines are shared among all agents on the private network or owned solely by the agents at the endpoints of the lines. In the latter case, the lines connecting members of a coalition are owned by that coalition, making the value of the characteristic function independent of the behavior of the other agents—a typical cooperative game, but not a market game. If the lines are shared, coalitions can impose large externalities on each other by congesting the lines that are important for other coalitions, making a suitable characteristic function representation impossible. Instead, we model the competition between coalitions as a competitive game and argue that stability can be guaranteed if market coordinator is

aware of prior arrangements among market participants.

2.4 Network Games

In this section, we will give an overview of related work on network routing, sharing, and construction games. Previous work in *communication networks* has produced edge-based market-clearing prices that induce Nash equilibria in routing games in a variety of contexts. Work by Kelly et al. [9] proposes a decentralized resource sharing protocol that produces approximate market-clearing prices via the Nash equilibrium of a game, which was later extended to sharing of edges on a graph [8]. These approaches do not address congestion of resources, and, as a consequence, do not provide a suitable model for power distribution. Equilibrium prices in the case of competing network operators with congestion externalities was explored by Acemoglu and Ozdaglar [1, 2], but only in parallel-serial networks. Their primary result is that competition among network operators does not improve the quality of equilibrium behavior in the presence of congestion, a counter-intuitive result that has some similarities to our analysis. The communication networks framework is significantly different since each message has a fixed destination. In power distribution, it is not important where a particular unit of power is delivered.

Another avenue of analysis has investigated models of decentralized network construction, primarily focusing on the properties of the equilibrium networks [4, 6]. These papers are able to describe many specific properties of the equilibrium because the formation games analyzed are quite simple. Once congestion externalities are integrated, the analysis becomes much more difficult.

Chapter 3

Agent Behavior in the Power Distribution Game

In this chapter, we describe and analyze our model of electricity markets, including a discussion of how payment schemes affect the behavior of market participants. In Section 3.1, we introduce the formal specifications of our model of electricity markets and our approach to physical modeling and its assumptions. In Section 3.2, we present several models of behavior, which can be understood as solutions to optimization problems. In Section 3.3, we analyze the problem of setting prices to induce more efficient agent behavior. This approach determines payments that have useful properties, but the meaning of the payments is quite opaque. In the following section, 3.4, we develop a complete theory for payments for a simpler model (without resistive losses) that involves no optimization. These payments give an intuition for the payments produced by the more complicated algorithm.

3.1 Setting

We first introduce the notation needed to describe our model. The network $G = (V, E)$ is composed of two parts, a *public network* $G_{pub} = (V, E_{pub})$ that is connected and a *private network* $G_{priv} = (V_{priv}, E_{priv})$ where $V_{priv} \subseteq V$. A single node T represents the link to the outside grid, which will buy or sell arbitrary amounts of electricity; it is not connected to any agent via the private network. Let $f_{u,v}$ denote the power sent from agent u to agent v over edge (u, v) , and let $f'_{u,v}$ denote the power received by agent v over edge (u, v) after losses are deducted. These variables will be constrained to be non-negative. The $f_{u,v}$ variables are used to formulate the problem as an optimization, but the solutions will always have

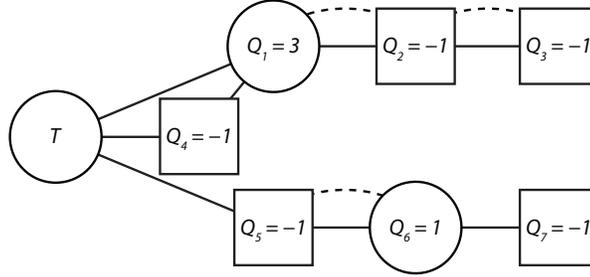


Figure 3.1: A network with private and public links. (a running example used in Secs. 3.1 and 3.2). Circles represent net generators and squares net consumers. Node T represents the connection to the outside power grid. Solid lines are publicly-owned links and dashed lines private links. We assume that public and private links have roughly the same voltage and resistance. Q_i denotes the supply of agent A_i .

at least one of $f_{u,v}$ and $f_{v,u}$ equal to zero, since power may only flow in one direction.

Each agent A_v (except T) has a *utility function* $U_v(x_v)$ for any allocation x_v of power representing their net consumption. Because demand for electricity is usually highly inelastic, we represent a utility function by a single constant Q_v , the agent's, steady-state (long-term) supply or demand for power.¹ Positive Q_v reflects a surplus of electricity, in which case we assume $U_v(x_v)$ is very steep for $x_v < 0$ and equal to zero thereafter. Negative Q_v reflects net demand, where $U_v(x_v)$ is very steep for $x_v < -Q_v$ and zero after. We will only consider settings in which all demands can be feasibly satisfied, so we treat these utility functions as if they were *constraints* on the final solution. These assumptions about utility functions simplify the analysis significantly. They also have the benefit of being simple to query agents about. We believe the core of the analysis can be extended the case of arbitrary utility functions.

Figure 3.1 illustrates the model: each circle is a net generator and each square a net consumer, with Q_v being net power generated by agent A_v . T represents the connection to the public grid, while public and private transmission lines are solid and dashed lines, respectively.

We consider several models of agent behavior below corresponding to different degrees of coordination and incentive structures. Each model, together with an agent's connections, determines an agent's strategy space. All models below share a common simplified physics model:

- Flow of electricity must be conserved: flow into any v must equal flow out of v less the agent's supply Q_v . This a natural requirement when the problem is viewed from the network-flow perspective. All physics models will have an analogous constraint.
- *Resistive losses* occur when power is transmitted between two nodes. Energy lost to heat and radiation is proportional to the *square* of the amount transmitted, and depends on the voltage and

¹There are obvious limits to the assumption of total demand inelasticity; studies of elasticity of electricity demand have been restricted to small price ranges. The problem making demand more sensitive to the cost of electricity supply, whether through price or other means (e.g. visual feedback) is an important one that is beyond the scope of our present analysis.

resistance of a particular link. Let $R_{u,v}$ be the resistance and $U_{u,v}$ the voltage on edge (u, v) . If u sends $f_{u,v}$ units of power over the line (u, v) , the amount v receives from u is:

$$f'_{u,v} = f_{u,v} - \frac{f_{u,v}^2 R_{u,v}}{U_{u,v}^2} \quad (3.1)$$

Note that $R_{u,v}$ is proportional to the length of (u, v) . Resistive losses are often assumed to be negligible in many simplified power flow models.

- Transformer losses occur when electricity is “stepped up/down” between high outside grid voltages and low local grid voltages at T . This loss is a constant fraction β of the power converted.
- Electricity can be disposed of at any point in the network. We make this assumption to make optimization simpler, but it should not have a large effect qualitatively. It could potentially affect the analysis if the generation is unpredictable and not easily adjusted, creating the potential for cascading failures due to overloading. We do not consider this problem.

From a physics standpoint, our simplified model treats alternating current (AC), the choice of most power systems worldwide, as direct current (DC), which is less complicated to deal with, and is a common simplification in power systems research (flow prediction in AC networks is often non-convex) [11]. We also neglect Kirchhoff’s voltage law, that causes power to flow sympathetically on loops. We will later introduce an energy-minimization problem that will approximate physics-induced flow on our simplified model.

To justify our model of electricity physics, we want to make two observations. The first is that the model we use gives an upper-bound for the efficiency of physics-based routing; in particular, AC power introduces interference when currents are out of phase, and Kirchhoff’s voltage law introduces “phantom” power flows that are not useful for achieving the routing objective. The second issue is that any physics model can be substituted in to the analysis without affecting calculations in other areas. Since we lack specific knowledge of the physical properties of grids, we use a high-level physics model that is optimistic to show that the “goals” of physics are not the same as those of a grid operator; in particular, physics-based routing does not take into account the cost of generation from a particular source. High-cost power can be overused if it is close to consumers, when it may have been better to send low-cost power, incurring higher losses in the process though still providing more cost effective consumption.

Our overall goal is to induce transmissions that minimize the amount of power purchased from the outside grid (and implicitly total generation), while satisfying the demands of all agents. This perspective

is valid in settings where private generation has lower marginal cost than that of the public grid—which is often true of renewable sources—and/or where public generation is a long distance away from our local agents. Agents can either transact exclusively with the public grid or join a coalition of other agents and coordinate their actions. The public grid offers a fixed *buy price* p_b and a lower fixed *sell price* p_s to the local network, which is how utilities currently engage small consumers and producers (it is impossible to do otherwise without smart meters). The prices can reflect average costs to the public grid (e.g., reflecting distance), but otherwise assumes that local consumption has little impact on the larger public network, which is true when local networks represent a small fraction of total demand. Flow on the public network is dictated by our simplified physics model in the following sense: it will be the flow that minimizes transmission losses. This can be viewed as the lowest energy state of the system.

For an example of how private routing can affect the network, consider Figure 3.1. Since the grid’s buy and sell prices are fixed, A_5 and A_6 will exchange power over the dashed edge—a privately-owned transmission line. The public grid will then have to send power over a long distance from T to A_7 . Given that A_5 and A_6 are cooperating, if the private edge $(5,6)$ did not exist, power would only need to be sent from T to A_5 , which is closer. Due to transmission losses, the existence of the private link between A_5 and A_6 combined with rational behavior of A_5 and A_6 *decreases* the social welfare of the resulting routing equilibrium.

Below we outline different ways in which *coalitions* of agents can coordinate their behavior to maximize their collective utility, assuming they can divide any surplus generated appropriately. The public network is centrally managed, but routing on any link in the private network can be arbitrarily controlled. We assume that agents in separate coalitions do not trade with each other on the private network (otherwise, they will be taken as part of the same coalition).

Two coordination issues arise in this setting: the first is finding optimal coordinated strategies, or a *routing scheme*, for the *grand coalition* of all agents: complete coordination will ensure that all agent demands are met at minimum cost. The second is finding an *incentive scheme*, or payments, that stabilize the grand coalition by aligning their interests to induce minimum cost electricity flow on the network. In this work, we first address the former, and assume that agents have the ability to find stabilizing payments. We discuss incentives in more depth in Section 3.3.

3.2 Models of Agent Behavior

We compare optimal agent routing behavior under four different models of cooperation. In addition to a mathematical specification for each model, we provide a description of the flows it induces on the

network in Figure 3.1. Induced flows in each model are computed sequentially in two stages. In the *private* stage, agents on the private network agree to exchange power. In the *public* stage, routing on the public network occurs given net demands after the private stage, and any payments are made between the grid and the private agents.

The Ad Hoc Model. The *ad hoc model* simulates system behavior when agents make decisions in a local, myopic manner without coordination or cooperation. Agents will only ever trade power with their neighbors, and they do so in order of closeness. Each trade is made to maximize the joint utility of the two traders. This model probably best reflects *status quo* behavior because private transmission lines are unregulated, and there are relatively few of them. It is unlikely that agents that are very distant from each other will be able to discover that it would be beneficial to build a private line between them. Although the random network generation model we use for experiments does not take distance into account, this behavior model favors the use of lines between agents that are closer to each other. The routing on the public network is induced by physics, which aligns with the primarily physics-based routing that occurs on the public grid.

We now formally describe the algorithm used to compute behavior under the ad hoc model. To compute exchanges on the private network, we order private edges by increasing length—agents that are close together are more likely to transact because they are generally more familiar with each other. For each $(u, v) \in E_{priv}$, set the non-negative flow on that edge such that the total utility of agents u and v is maximized:

$$\max_{f_{u,v}} \begin{cases} p_s(Q_u - f_{u,v}) & \text{if } Q_u - f_{u,v} \geq 0 \\ -p_b(Q_u - f_{u,v}) & \text{if } Q_u - f_{u,v} < 0 \end{cases} + \begin{cases} p_s(Q_v + f'_{u,v}) & \text{if } Q_v + f'_{u,v} \geq 0 \\ -p_b(Q_v + f'_{u,v}) & \text{if } Q_v + f'_{u,v} < 0 \end{cases} \quad (3.2)$$

The first term of this equation represents the utility of agent u , and the second term represents the utility of agent v . These utility terms describe the cost or payoff resulting from buying or selling the remaining power of each agent, after trading, to the public grid. We then update the demands of these two nodes by setting $Q_u^{new} = Q_u^{old} - f_{u,v}$ and $Q_v^{new} = Q_v^{old} + f'_{u,v}$. We continue through the ordered sequence of edges until no profitable trades remain. Note that two agents only exchange if one is a net generator and the other a net consumer, and agents only trade with their immediate neighbors in G_{priv} .

After these local transactions have occurred, remaining demand is met using physics-based flow. This

occurs without the agents choosing a strategy because they are unable to affect the flow on infrastructure that they do not own. Physics-based routing happens according to the physics model described in the previous section. We minimize transmission losses by solving the following optimization:

$$\begin{aligned} \min_{f_{u,v}, \forall (u,v) \in E_{pub}} \quad & \sum_{(u,v) \in E_{pub}} \frac{f_{u,v}^2 R_{u,v}}{U_{u,v}^2} \\ & + \frac{\beta}{1-\beta} \sum_{v:(T,v) \in E_{pub}} f_{T,v} + \beta \sum_{v:(v,T) \in E_{pub}} f'_{v,T} \end{aligned} \quad (3.3)$$

subject to, for each node $v \in V$:

$$Q'_v + \sum_{u:(u,v) \in E} f'_{u,v} \geq \sum_{u:(v,u) \in E} f_{v,u} \quad (3.4)$$

where Q'_v (in this case, Q_v^{new}) is the supply of node v induced by the flows from private trading. The first term of Objective 3.3 represents the resistive losses incurred on every edge in the public network, while the second and third terms represent the amount lost while changing the voltage of electricity from the grid at node T . Constraint 3.4 states that flow should be conserved at every node in the network.

We observe the following induced flows on the network in Figure 3.1 under this ad hoc model. On the upper branch, A_1 sends power to A_2 , but the private edge (2, 3) remains unused because A_1 is unable to trade with A_3 in the private stage. This causes public edge (2, 3) to carry more power than necessary, incurring higher losses. On the middle branch, the cost of supplying A_4 is higher than necessary because physics-induced routing equalizes the losses on $(T, 4)$ and $(1, 4)$ —wasting power from A_1 —as long as losses on $(1, 4)$ exceed β . On the lower branch, A_6 sends its power to A_5 , leaving A_7 to receive power from T . This is inefficient because A_5 is closer to T than is A_7 .

The Private Self-Interest Model. The *private self-interest model* reflects agents on the private network cooperating to maximize their aggregate utility under fixed public network pricing. The agents behave as if they solved a joint optimization with a term representing the utility of each. The variables in this optimization are the flow variables on the private network. The flow on the public network is determined by physics in a method analogous to that in the ad hoc section—via minimizing the amount of transmission losses under our basic physical constraints.

The agents connected on the private network solve the following optimization:

$$\min_{f_{u,v} : (u,v) \in E_{priv}} \sum_{v \in V_{priv}} \begin{cases} -p_s(Q_v + x'_v) & \text{if } Q_v + x'_v \geq 0 \\ p_b(Q_v + x'_v) & \text{if } Q_v + x'_v < 0 \end{cases} \quad (3.5)$$

where

$$x'_v = \sum_{u:(u,v) \in E_{priv}} f'_{u,v} - \sum_{u:(v,u) \in E_{priv}} f_{v,u} \quad (3.6)$$

Objective 3.5 represents the total cost to nodes on the private network to satisfy remaining demand from (or sell surplus to) the public grid after all (optimal) private transactions are made. Each term in the minimization represents the utility of a single agent after the routing on the private network described by the $f_{u,v}$ variables is carried out. The utility terms represent the cost or payoff of buying or selling to the public grid after all routing on the private network is carried out.

After the private agents have minimized their liability to the grid, physics governs flow on the public network as in the ad hoc model (Objective 3.3 subject to Constraint 3.4).

In this model, the flow on the middle and lower branches of Figure 3.1 is the same as in the ad hoc model. On the upper branch, A_1 sends power to A_3 during the private stage (exchanges between non-adjacent nodes is now permitted). However, this does not reduce the cost of power from the public network, since this merely shifts load from public to private edges, leaving the public edges virtually unused.

By solving the optimization, agents on the private network maximize their aggregate utility. However, they still must divide up the surplus in a way so that none of their agents defect. We will discuss this problem in Section 3.3.

The Cooperative Model. In the *cooperative model*, agents on the private network distribute electricity so that the load on the public network is minimized (again, flow on the public network is determined by physics). The agents solve a joint optimization that minimizes the overall cost to supply electricity to the entire network by minimizing the amount of externally generated power consumed by the *entire* set of agents (according to our assumption that all locally generated power is from renewable sources and thus has near zero marginal cost). Causing agents to cooperate will require suitable incentives, which we will discuss in Section 3.3. Minimizing the overall draw from the grid requires arranging flows on the

private network that solve the following optimization:

$$\min_{f_{u,v:(u,v) \in E_{priv}}} \frac{1}{1-\beta} \sum_{v:(T,v) \in E} f_{T,v} - (1-\beta) \sum_{v:(v,T) \in E} f'_{v,T} \quad (3.7)$$

subject to the usual flow and physical constraints, Equations 3.3 and Equation 3.4. Recall that T is the node representing the link to the larger grid. The first sum in Objective 3.7 represents the total amount of external power imported—the loss β incurred at the transformer acts as a percent additional cost on this power. The second sum represents the total amount of locally-generated power that is exported less transformer losses. Note that for every edge (v, T) , at least one of $f_{v,T}$ and $f_{T,v}$ will be 0.

This optimization is more difficult to solve than the others because it involves a constraint that the flows on the public network are set such that a separate objective (physics-based energy minimization) is minimized. This computation is explained at the end of this section.

The flows induced by the cooperative model in Figure 3.1 are quite different from those seen in the previous models. Flows on the upper and lower branches resolve the inefficiencies present in earlier models: the upper branch flow from A_1 to A_2 and A_3 uses the public and private links equally to minimize transmission losses, and on the lower branch, power from A_6 will flow to A_7 instead of A_5 . Despite the improvement, flow on the middle branch remains inefficient due to physics-based routing, causing A_4 to draw power from T as well as A_1 .

The Integrated Model. In the *integrated model*, we allow the flow on all edges, both private and public, to be controlled by the local agents, relaxing the constraint that flow on public edges be governed by the physics model. The objective minimizes the amount of externally generated power consumed, as in the previous model, but all flow variables are assumed to be controllable by the agents, not just those on the private network.

This can be viewed as “ideal” behavior for agents on a network that have arbitrarily precise routing equipment and some degree of control over (local) transmission on public lines. This reflects the most effective use of resources if transmission lines are expensive to construct and routing equipment is cheap. It requires solution of the following optimization:

$$\min_{f_{u,v:(u,v) \in E}} \frac{1}{1-\beta} \sum_{v:(T,v) \in E} f_{T,v} - (1-\beta) \sum_{v:(v,T) \in E} f'_{v,T} \quad (3.8)$$

subject to constraints Equation 3.4. The objective is identical to that in the cooperative case (Equation 3.7) except that *all flows* are controllable, not just those on private links. The first term represents the amount of externally generated power imported, and the second term represents the amount of local

power exported.

The flows induced by the integrated model in Figure 3.1 are the same as those of the cooperative model on the upper and lower branches. However, because no physics-based flows occur, the inefficiency on the middle branch disappears, and A_4 will draw power only from A_1 .

Computation. Except for the cooperative model, the optimization problems above are all convex, quadratically-constrained quadratic programs (QCQP) that can be directly solved by CPLEX or other off-the-shelf optimization tools.

In the cooperative case, we have a constraint that the flows on the public network must minimize the physics-based energy function given the flows set by the agents. Our overall objective of minimizing the consumption of externally generated power then depends on flows on both the public and private network. We approximate a solution to this problem by minimizing the objective over all the flow variables. We can then calculate the value of the objective by solving the physics-based minimization over the variables on the public network. The resulting solution is feasible, and proves to be within a few percent of the optimal value in our experiments. The cooperative model will always induce higher losses than the integrated model, but it is quite close in practice.

3.3 A Simple Incentive Scheme

While the methods above compute optimal coordinated behavior under the behavioral restrictions imposed, we must also ensure incentives can be put in place such that all agents effectively cooperate by actually adopting the strategy prescribed by the optimal solution. The aim of these prices is to induce the cooperative or integrated behavior described in the previous section. Which of the two models can be used in practice depends on what kind of routing equipment is available on the network.

Equilibrium Prices. We first present a pricing scheme that aligns the interests of each agent on the private network with socially optimal behavior. This induces self-interested agents to adopt the cooperative solution (or the integrated solution if public links are controlled locally). If the public links are controlled locally, the prices need to be charged to every agent in the network. If they are controlled centrally, the prices only need to be applied to agents that are connected to at least one private edge.

Our technique is relatively standard in network models [23, 14, 9, 1]: we compute agent-specific buy/sell prices using dual prices obtained from the optimization. These induce individual agents to implement the globally optimal solution by maximizing their own net utility at these prices.

We first solve the integrated optimization Equation 3.8. By the first-order optimality conditions,

there is a set of positive dual variables μ such that: (a) for each edge (u, v) , where $u, v \neq T$:

$$-\mu_v \left(1 - \frac{2f_{u,v}R_{u,v}}{U_{u,v}^2} \right) + \mu_u - \mu_{u,v} = 0; \quad (3.9)$$

(b) for each edge (T, v) :

$$\frac{1}{1-\beta} - \mu_v \left(1 - \frac{2f_{T,v}R_{T,v}}{U_{T,v}^2} \right) - \mu_{T,v} = 0; \quad (3.10)$$

and (c) for each edge (v, T) :

$$-(1-\beta) \left(1 - \frac{2f_{v,T}R_{v,T}}{U_{v,T}^2} \right) + \mu_v - \mu_{v,T} = 0. \quad (3.11)$$

The values of these dual variables “price” the direct transmission to and from each agent. We charge each agent v :

- $\mu_v + \mu_{v,u}$ for each unit sent to $u \neq T$ and $-\mu_u$ for each unit received (after losses) by $u \neq T$.
- $-\mu_u + \mu_{u,v}$ for each unit sent from $u \neq T$ and μ_v for each unit received (after losses) from $u \neq T$.
- $\frac{1}{1-\beta} - \mu_v \left(1 - \frac{f_{T,v}R_{T,v}}{U_{T,v}^2} \right)$, for each unit sent from T (if $(T, v) \in E$), where $f_{T,v}$ is the total amount sent by T to v .
- $\mu_v - (1-\beta) \left(1 - \frac{f_{v,T}R_{v,T}}{U_{v,T}^2} \right)$, for each unit sent to T (if $(T, v) \in E$), where $f_{v,T}$ is the total amount sent to T by v .

Note, in our models, there are no direct *private* links to T —we include payments for such links for completeness. It is possible to imagine a model where an agent may construct a transmission line to a grid substation (or other public infrastructure), but it is not clear how the line would be controlled. Thus, we omit it from our overall model.

Assume these prices are fixed. Each v in the local network has the following “personal” optimization problem to maximize its utility for the (net) power received minus the price paid:

$$\begin{aligned} & \max_{\substack{f_{u,v}:(u,v) \in E \\ f_{v,u}:(v,u) \in E}} U_v \left(Q_v + \sum_{u:(u,v) \in E} f'_{u,v} - \sum_{u:(v,u) \in E} f_{v,u} \right) \\ & + \sum_{\substack{u:(v,u) \in E \\ u \neq T}} ((\mu_v + \mu_{v,u})f_{v,u} - \mu_u f'_{v,u}) + \sum_{\substack{u:(u,v) \in E \\ u \neq T}} ((-\mu_u + \mu_{u,v})f_{u,v} + \mu_v f'_{u,v}) \end{aligned}$$

plus the following expression if the node is connected to T :

$$(1 - \beta)f'_{v,T} - (\mu_v - \mu_{v,T})f_{v,T} - \left(\frac{1}{1 - \beta} - \mu_{T,v}\right)f_{T,v} + \mu_v \left(1 - \frac{2f_{T,v}R_{T,v}}{U_{T,v}^2}\right)$$

The first-order optimality conditions for this problem are a subset of those for the global optimization. Thus, the optimal flow for any agent at the given prices coincides precisely with the globally optimal flows. Furthermore, no *coalition of agents* can improve their aggregate utility by changing their strategies at these prices. In other words, these prices induce a *strong Nash equilibrium* that coincides with the globally optimal outcome. However, these prices have some weaknesses. We do not have any guarantees about the total amount paid to the mechanism. Ideally, we would want it to be zero. We also do not know if the individual agents would choose to participate in this mechanism because we do not know what their previous contracts were. We will address these issues further below.

Core Stability. If the prices that arise as a by-product of the optimization are taken as given, they induce strongly stable optimal behavior in utility maximizing agents. However, if the market coordinator has no binding power over the agents to accept these prices, certain agents may refuse to “enter this market” if they had prior local arrangements (e.g., in the private self-interest model or in some other form) that gave them higher utility. For example, in a network that has a single private edge that connects a generator and a consumer, these two agents will likely benefit greatly from the operation of the edge, and thus be unwilling to agree to a change in pricing without inducement. Furthermore, while coalitional defections cannot be beneficial *given these fixed prices*, the allocation induced by these prices cannot generally be supported if one or more agents refuse to participate. Thus, we might need to consider payment schemes that prevent such defection.

Superficially our models seem to fit within the framework of market games, albeit with variable “transaction” costs induced by distance and load. Unfortunately, there is no natural characteristic function representation for this cooperative game. The value that one coalition of cooperating agents receives from a given strategy can be altered by the actions of another: if a second coalition routes power over the same transmission lines, the higher load increases power loss quadratically. Research on market games [24] shows that core transfers exist in a variety of exchanges by using dual variables of allocation optimization as prices, as we do above, making core transfers easy to compute. However, our model is not reducible to market games due to interaction effects. These require that, when reasoning about potential defections from the grand coalition, we should consider Nash equilibria that emerge when treating different coalitions as players.

Unfortunately, we have not yet been able to prove that transfers exist that support a strong Nash

equilibrium (or loosely, a *core imputation*). We strongly suspect that this is so: in *all instances* tested empirically, we are able to compute stability-inducing transfers. This remains an important topic of research. We note that if the market coordinator (system operator) knows of existing arrangements between agents prior to joining the coordinated local network, participation constraints that induce them to participate can be easily addressed. If the agents were to agree to participate in the pricing scheme, the social welfare of the system will increase by a non-negative amount, since the current transmission schedule is a feasible option in the global optimization. The operator can distribute this surplus to agents on the private network, guaranteeing that each receives utility at least that derived prior to coordination.

3.4 No-Resistive-Losses Model

Since it is not intuitively obvious what the payments calculated by our incentive look like, it is interesting to consider a simpler case where core payments can be explicitly calculated: the power distribution model without resistive losses. In this section, we will describe the model and the pricing that naturally emerges from it.

As in the general resistive losses model, we will assume the agents are represented by nodes on a graph G , and the transmission lines between them are represented by edges on that graph. We will assume that the graph is connected, which can be relaxed by decomposing a single disconnected graph into multiple instances of no-resistive-losses distribution. We will further assume that each agent can buy and sell power from the public grid at the same prices, p_b and p_s . We note that this is a natural assumption as long as all agents are connected to the public grid because there are no resistive losses in transmission.

Without loss of generality, we will normalize each utility function such that $U_v(-Q_v) = 0$ if $Q_v < 0$, i.e., each agent with demand for power receives zero utility at his stated demand, and $U_v(0) = 0$ if $Q_v \geq 0$, i.e., each agent with power surplus receives zero utility when allocated no power. We will also make an assumption about utility functions that is realistic, but more restrictive than we have previously: agents do not receive any utility for excess power other than p_s per unit and that they do not lose utility for a power deficit other than p_b per unit. Essentially, we assume that agents have a completely inelastic demand for power and the grid offers linear prices.

The global optimization version of no-resistive-losses distribution is to find the allocation of power to agents such that the aggregate utility of the agents is maximized. We note that the actual power flow used in distribution does not matter—there are no resistive losses so the capacity of each edge is

unlimited.² This model qualifies as a network game, but we will not model it as such because the network game formulation only provides payments as the solution to a linear program, not in closed form.

Let $v(C)$ be the net payment to the grid for optimal behavior for coalition C . Since the model is simple, optimal behavior consists only of sending surplus power to agents that need it. The agents in a coalition C will derive some benefit, i.e., a reduction in cost or increase in profit, from cooperating. Under this simple model, it is a core allocation to divide this benefit among the cooperating agents proportionally according to how much of the scarce resource that each provides, where the scarce resource is the resource—either supply or demand—which would increase the amount of benefit from cooperation if more of it were present in the coalition. We will call it the *proportional scarce allocation*.

Definition 1 (Proportional Scarce Allocation). The following \vec{t} is the proportional scarce allocation:

$$t_i = v(\{i\}) + s_i$$

$$= \begin{cases} Q_i p_s & \text{if } Q_i \geq 0 \\ Q_i p_b & \text{if } Q_i < 0 \end{cases} + \begin{cases} (p_s - p_b) Q_i & \text{if } Q(N) \geq 0, Q_i < 0 \\ (p_b - p_s) Q_i & \text{if } Q(N) < 0, Q_i \geq 0 \\ 0 & \text{otherwise} \end{cases}$$

where

$$s_i = \begin{cases} (p_s - p_b) Q_i & \text{if } Q(N) \geq 0, Q_i < 0 \\ (p_b - p_s) Q_i & \text{if } Q(N) < 0, Q_i \geq 0 \\ 0 & \text{otherwise} \end{cases}$$

Here s_i represents the share of the benefits of cooperation agent i receives based on the amount of scarce resource they have. Agent i 's reservation value $v(\{i\})$ represents a participation constraint: how well the agent could do if she did not cooperate.

Theorem 1. *The proportional scarce allocation \vec{t} is a core allocation.*

For example, in a coalition with 50 kWh of supply and 60 kWh of demand, supply would be the scarce resource, and it would be a core allocation to divide the benefits of cooperation among agents with net positive supply proportionally to the amount of supply they provide. A full proof of Theorem 1 and further discussion is included in Appendix B.

²The optimization typically solved by ISOs for pricing purposes is similar to this problem except transmission lines with finite capacity are used to represent resistive losses.

The simpler market game with no resistive losses can be thought of as an approximation to the full power distribution game when transmission losses are low. Since they are typically around 7% in practice, the approximation will be quite close. Therefore, we can think of a core allocation for the full power distribution game as roughly distributing profits according to the amount of scarce resource provided by each agent, with some separation of different parts of the network caused by the losses on each edge.

Chapter 4

Experiments

In this chapter, we simulate the effect of the four behavior models described in the previous chapter on various types of electricity grids. We do not have access to data on the topology of real electricity grids, and private infrastructure does not widely exist outside of self-contained microgrids that connect to the public grid at a single point. Thus, we will make some general assumptions about grid structure and test a variety of parameters for aspects of the grid we are unsure about. At a high level, we assume that the public grid is a minimum spanning tree that averages 7% in transmission losses, the private grid is a random graph, and that agent power supplies and demands are sampled from a normal distribution. We will describe our model formally, justifying the assumptions we make about specific values and then discuss some potential critique of our assumptions.

We test a variety of generation and distribution scenarios. Some reflect current practice, while others reflect future settings with potential increased usage of local, private generation and transmission infrastructure. Our proxy for social welfare is *transmission loss*, i.e., the amount of power originating in the public grid that is not consumed by agents in the local network (assuming all agent demands are satisfied). This correlates linearly with social welfare assuming local producers have linear cost curves for generation. Since cost curves are typically convex, our results will, in fact, underestimate true improvement in social welfare due to lowered losses. All local networks in our experiments consume many times more power than they produce (which reflects most realistic settings) and have an average transmission loss of around 7% (with standard deviation of 2.5%) when no private infrastructure or distributed generation is assumed: this is the current U.S. average for transmission losses.

In each trial, 100 agents are distributed uniformly over a square grid of size either 100×100 km. (low density) or 1×1 km. (high density). We generate a distance-minimizing spanning tree connecting these

agents (and T), reflecting how most public transmission networks are constructed. Links between agents have voltage 22kV, and links to T are 50 kV (links to T tend to have a higher load, so are typically higher voltage). All edges have 0.2 ohms of resistance per km. of length. Transformer losses at T are given by $\beta = 0.02$ (this determines the theoretical lower limit on transmission loss for a net-consumer network, as all publicly-supplied power must pass through T). We use a 10% gap between the grid buy price p_b and sell price p_s .

Agent demands are distributed normally with fixed mean μ , and a variable std. dev. σ which varies the amount of distributed generation: we set $\sigma = \mu/2$ (low degree of distributed generation), $3\mu/4$ (medium), μ (high) and $5\mu/4$ (very high), meaning that we expect 2.5%, 9%, 16%, and 21% of agents to be net producers, respectively. The private network was generated by sampling an Erdős-Renyi random graph using p to determine edge density in the network: each edge is included with probability p . We use edge densities of 0 (i.e., no private network), $\frac{1}{n^2}$, $\frac{1}{n}$, and $\frac{1}{\log(n)}$ (the last representing heavy investment in private infrastructure).

These levels of sampling probability have some significance in terms of emergent behavior in Erdős-Renyi random graphs [5]. From the perspective of expectation, at $p = \frac{1}{n^2}$ we expect a single edge on the private network since there are $\frac{1}{n^2}$ edges to sample from. At $p = \frac{1}{n}$, the expected degree for each node on the private network is 1. We can also make some observations about component size and number of components. For $p = \frac{1}{n}$, the private network will almost surely have a component of size of order $O(n^{2/3})$ —around 22 nodes in a 100 node network. At $p = \frac{1}{\log(n)}$, the graph will almost surely be connected.

Before we describe our results, we will discuss several potential weaknesses of our experimental setup. One issue with our modeling of resistive losses is that it is possible that most losses are incurred in transmission to or from remote places, rather than locally. This occurs naturally to an extent in our low density 100×100 km. scenario, but does not fully account for the case where a large amount of a network generation occurs at extreme distance from the agents it interacts with, e.g., rural hydroelectric dams or remote settlements. Our modeling of resistive losses will concentrate them around heavy producers and consumers and agents that are centrally located in the network.

Another area of weakness relates to our assumptions about the probabilistic distribution of electricity supply/demand. We assume a unimodal normal distribution with relatively high variance. This distribution could easily be bimodal in practice, as one might expect that those agents who have distributed generation capacity would likely be separable from those without it. We may also expect to see generation being built in areas of high demand in the case that price incentives are well-aligned.

Our induced-spanning tree model of the public grid does not exactly align with the reality of grid

| <i>Amt. of dist. gen.</i> | <i>Model</i> | <i>Edge density on priv. net.</i> | | | |
|-------------------------------|-----------------|-----------------------------------|-----------------|---------------|---------------------|
| | | 0 | $\frac{1}{n^2}$ | $\frac{1}{n}$ | $\frac{1}{\log(n)}$ |
| Low | Ad hoc | 7.14 | 7.25 | 7.05 | 7.13 |
| | Priv. self-int. | 7.14 | 7.25 | 7.05 | 7.13 |
| | Coop. | 7.14 | 6.95 | 2.91 | 2.16 |
| | Integrated | 7.14 | 6.95 | 2.92 | 2.16 |
| Med. | Ad hoc | 7.28 | 7.19 | 7.15 | 7.20 |
| | Priv. self-int. | 7.28 | 7.19 | 7.15 | 7.20 |
| | Coop. | 7.28 | 6.92 | 2.93 | 2.16 |
| | Integrated | 7.28 | 6.91 | 2.93 | 2.16 |
| High | Ad hoc | 7.19 | 7.37 | 7.34 | 7.24 |
| | Priv. self-int. | 7.19 | 7.37 | 7.34 | 7.24 |
| | Coop. | 7.19 | 7.07 | 2.97 | 2.15 |
| | Integrated | 7.15 | 7.05 | 2.96 | 2.15 |
| V. High | Ad hoc | 7.32 | 7.26 | 7.28 | 7.22 |
| | Priv. self-int. | 7.32 | 7.26 | 7.27 | 7.20 |
| | Coop. | 7.32 | 6.93 | 2.95 | 2.16 |
| | Integrated | 7.26 | 6.87 | 2.94 | 2.16 |

Figure 4.1: Avg. transmission loss as a percentage of net demand in a low (node) density network (1 agent per sq. kilometer).

| <i>Amt. of dist. gen.</i> | <i>Model</i> | <i>Edge density on priv. net.</i> | | | |
|-------------------------------|-----------------|-----------------------------------|-----------------|---------------|---------------------|
| | | 0 | $\frac{1}{n^2}$ | $\frac{1}{n}$ | $\frac{1}{\log(n)}$ |
| Low | Ad hoc | 7.25 | 7.28 | 7.31 | 7.33 |
| | Priv. self-int. | 7.25 | 7.28 | 7.31 | 7.33 |
| | Coop. | 7.25 | 6.97 | 2.93 | 2.16 |
| | Integrated | 7.24 | 6.97 | 2.93 | 2.16 |
| Med. | Ad hoc | 7.22 | 7.39 | 7.40 | 7.08 |
| | Priv. self-int. | 7.22 | 7.39 | 7.40 | 7.09 |
| | Coop. | 7.22 | 7.05 | 3.00 | 2.16 |
| | Integrated | 7.21 | 7.04 | 3.00 | 2.16 |
| High | Ad hoc | 7.42 | 7.30 | 7.31 | 7.37 |
| | Priv. self-int. | 7.42 | 7.30 | 7.31 | 7.36 |
| | Coop. | 7.42 | 6.99 | 2.99 | 2.16 |
| | Integrated | 7.40 | 6.97 | 2.98 | 2.16 |
| V. High | Ad hoc | 7.30 | 7.26 | 7.50 | 7.41 |
| | Priv. self-int. | 7.30 | 7.26 | 7.50 | 7.39 |
| | Coop. | 7.30 | 6.91 | 3.00 | 2.16 |
| | Integrated | 7.26 | 6.85 | 2.99 | 2.16 |

Figure 4.2: Avg. transmission loss as a percentage of net demand in a high (node) density network (100 agents sq. kilometer).

construction costs. From observation, agents are often connected to a single “terminal” node in an electricity grid, especially in urban settings. In terms of tree structure, it seems likely that the branching factor of parents of leaf nodes is higher on average than the branching factor of other nodes on the tree. We speculate that this structure emerges due to the costs of land-use rights.

Figures 4.1 and 4.2 summarize our results, with each data point showing average transmission loss over 1000 random instances.¹ At a high level, we see that the cooperative and integrated models provide significantly lower losses relative to the ad hoc and private-self interest models in nearly all cases. We see relatively small amounts of savings at private edge densities of 0 and $\frac{1}{n^2}$ and relatively large amounts at $\frac{1}{n}$ and $\frac{1}{\log(n)}$. Increased private network density correlates very strongly with savings in the more coordinated models: losses are roughly 60% less than in the ad hoc and private self-interest models when edge density is $\frac{1}{n}$ or greater (in both the low and high node density models). With edge density $\frac{1}{\log(n)}$, transmission losses approach the lower bound of 2% associated with purchase of all power from the grid with no transmission (only transformer) loss.

Theoretically, the integrated model should be more efficient than the cooperative model, which in turn should outperform the ad hoc and private self-interested models. This hierarchy is reflected in our experimental results. Integrated tends to not do much better than Cooperative in high private-edge density. In the low node density, low distributed generation and $\frac{1}{n}$ private-edge density, Integrated performed slightly worse than cooperative, which appears to be due to build up of rounding error. With low private-edge density and low or medium distributed generation cases, there was also little difference between the two models. The largest gaps appear in low private edge density and high distributed generation settings, with a largest gap of nearly one percent occurring in low node density, very high distributed generation and $\frac{1}{n^2}$ private edge density.

We do not have a theoretical result that dictates how the ad hoc and private self-interest models should compare. In practice, their performance is extremely similar, usually with a slight edge for Private Self-interest. The largest gap of 0.3% between the two models occurs in the low node density, very high distributed generation and $\frac{1}{\log(n)}$ private-edge density case. Private Self-interest outperforms Ad Hoc in all cases except high node density, medium distributed generation and $\frac{1}{\log(n)}$ edge density.

The main explanation for the similar performance is that it seems to make little difference what trading goes on between agents on the private network if the private network is not used to assist the public grid distributing power from the transmission link node T . These networks are all aggregate net consumers so the most important task of infrastructure is to efficiently distribute power coming from

¹Std. dev. is not shown, but is 2.5-3% for avg. losses in the 7% range, 0.5% for losses around 3%, and 0.01% for losses around 2%.

the outside grid. The internal power flows from distributed generation are relatively less important.

Our results assume a fixed buy-sell price gap of 10%. The smaller this difference is, the less trading will occur in the ad hoc and private self-interest models. A small gap mediates the effect of bad routing decisions on the private network, but also causes less usage of private edges, increasing load on public infrastructure. We observe in Figures 4.1 and 4.2 what could be a mild increase in losses as more distributed generation is deployed (although it is hard to be sure due to sampling noise). This would not be surprising since this also introduces higher demands in our model (we keep *net local demand* constant)—load becomes more focused around large producers and consumers, causing large losses at these points. The average loss in the zero private edge case is 7.22% in low distributed generation and 7.29% in high distributed generation. In $\frac{1}{\log(n)}$ edge density, the average loss is 4.67% in low distributed generation and 4.73% in high distributed generation.

Generating the random graphs was usually the most computationally expensive part of each trial—at least for 100 node graphs, solving the convex quadratically-constrained quadratic programs took under a second, but we have not investigated the scaling behavior of optimization. Note that we do not need to calculate strong Nash-stable prices to run these experiments. Calculating the prices is easy if you have knowledge of the previous trading arrangements between agents, but potentially computationally difficult if you do not, i.e., the scaling behavior is exponential in the number of agents, as far as we know, and would likely require an heuristic or approximate algorithm in practice. Section 3.3 gives a high-level description of the payment calculation process.

Chapter 5

Conclusion

We have presented a routing scheme that coordinates control of private infrastructure with the public grid, improving social welfare while satisfying participation constraints for the agents who control private infrastructure if we assume that coalitions of agents have the ability to calculate stabilizing payments. We presented several models of agent behavior under different incentive conditions and tested the efficiency of various power distribution networks under different models. Using private infrastructure in cooperation with public infrastructure was shown to be quite important.

Although we were always able to compute “core” stabilizing payments empirically, it remains an open question whether they always exist. The cooperative game theory model we analyze assumes that the prices on the public grid are fixed regardless of the behavior of agents on the private grid. Since these prices reflect long term costs of supplying costs of electricity, this is probably a reasonable assumption from a short term perspective or if there is a relatively small number of agents on the private grid. It makes the problem significantly easier to analyze and does not require consideration of the generation cost curves of the supplying utilities. If the generation cost curve were convex, it should make the problem easier because there is an “extra” incentive for agents to cooperate, but it would be interesting to further analyze this version of the problem.

We conjecture that there is a class of cooperative games that generalizes market games to allow non-independent congestion costs, while retaining the property that core payments always exist. We are not sure if these games have core payments that are easy to calculate, but if the proof of existence resembles that of market games, they should be. This class would include at least the power distribution game with fixed prices on the public grid as described above, but also would include games with other transmission losses (likely convex). We investigated this problem from many perspectives, but have not

yet been able to prove either core existence or construct a counter example.

Our experimental results show that coordination is critical—without it, private links can be constructed without decreasing the overall network load. Indeed, building a private link may be profitable for an agent in the short term; but if many agents construct links without coordinating, the return on investment tends to zero, even when a large number of private links are constructed. Essentially, these private links are only valuable under the ad hoc and private self-interest models to the extent that they allow agents to access the available distributed generation. Since the amount of distributed generation is limited, the construction of further links divides these benefits into smaller and smaller pieces.

Since the grids we simulate have relatively small amounts of distributed generation, the main objective for routing is to deliver power as cheaply as possible. Although the private grid can reduce the amount of demand from the public grid, the lack of coordination in the way this occurs prevents it from having a large impact on overall grid efficiency. When the private lines are used under the cooperative and integrated models, they can reduce delivery costs by taking additional load off of the public grid. The ad hoc and private self-interest models do not prevent them from performing this role, but there is no incentive for agents to do so. If the network were not a net consumer, the role of the grid becomes primarily to distribute power locally, and the role of the public grid is less critical. In this scenario, the private self-interest model is likely to perform decently once most agents are connected via the private network. One avenue for extension is to investigate higher levels of distributed generation to see where the thresholds for behavior change are.

Another outstanding question is the importance of the assumptions made about network formation. There are many different random graph models that match the structure of power grids more closely; for example, those with the *small world property*, i.e., low average node to node distance, reduce the average length of power exchanges.

Using our incentive schemes to determine pricing and capital cost-sharing of new (private and public) infrastructure is a natural future direction. Once a pricing scheme is set for distribution, it becomes possible to integrate prices into a cost-sharing scheme that in turn creates both a new network formation process and a paradigm for collective decision making for infrastructure construction—a goal of this line of inquiry. Collective decision making in this context raises questions about preference elicitation and truthfulness. This approach requires analysis of capital costs and time horizons, making the analysis quite different than that provided here for use of infrastructure.

Studying the mechanism design problem in power distribution is also important—here we assume that agent demands are known (e.g., from historical data) or that agents report their utility functions truthfully. Asking each agent for a complete utility function may be expensive, impractical, and subject

to manipulation by profit-seeking agents. From the perspective of the no-losses approximation, it is clear that there is a great deal of potential strategic behavior near the point where the amounts of supply and demand are similar. Since prices have to be provided in advance but agent supplies and demands are not certain until the exchange takes place, there is a risk that supply and demand predictions interact with proposed prices to create instability. For example, models of dynamic pricing of electricity often show herding behavior among agents: demand is low so prices are reduced, which causes demand to become high causing prices to increase, etc. Instability can cause capacity to be underused as well as system failures when stress is too high.

We focused on background research relating directly to power distribution or markets with transaction costs, whether in computer science, operations research or economics. There are some interesting connections in other areas to explore. For example, there is the network-flow perspective in operations research—particularly generalized network flows and minimum cost flows, which both relate directly to the problem. A study of these areas could clarify the connection between markets on networks and cooperative game theory, which we hope will shed light on the price calculation problem.

However, it is also possible that the proof of core existence or non-existence for market games with congestion does not lead to any computational insights. If this is the case, perhaps the pricing computation problem could be approached via heuristic search or approximation algorithms. It seems that prices are relatively hard to approximate, but this a question worthy of further investigation. It is also possible that real-world power distribution networks possess properties that simplify calculation of prices.

Bibliography

- [1] Daron Acemoglu and Asuman Ozdaglar. Competition and efficiency in congested markets. *Mathematics of Operations Research*, 32(1):1–31, 2007.
- [2] Daron Acemoglu and Asuman Ozdaglar. Competition in parallel-serial networks. *IEEE Journal on Selected Areas in Communications*, 25(6):1180–1192, 2007.
- [3] Gediminas Adomavicius and Alexander Tuzhilin. Towards the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE Transactions on Knowledge and Data Engineering*, 17(6), 2005.
- [4] Elliot Anshelevich, Anirban Dasgupta, Jon Kleinberg, Eva Tardos, Tom Wexler, and Tim Roughgarden. The price of stability for network design with fair cost allocation. In *Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science*, pages 295–304, Washington, DC, 2004.
- [5] Paul Erdős and Alfréd Rényi. On the evolution of random graphs. *Publications of the Mathematical Institute of the Hungarian Academy of Sciences*, 5:17–61, 1960.
- [6] Alex Fabrikant, Ankur Luthra, Elitza Maneva, Christos H. Papadimitriou, and Scott Shenker. On a network creation game. In *Proceedings of the Twenty-Second Annual Symposium on Principles of Distributed Computing*, pages 347–351, New York, 2003.
- [7] William W. Hogan. Electricity market restructuring: Reforms of reforms. *Journal of Regulatory Economics*, 21(1):103–132, 2002.
- [8] Ramesh Johari and John N. Tsitsiklis. Efficiency loss in a network resource allocation game. *Mathematics of Operations Research*, 29(3):407–435, 2004.

- [9] F. P. Kelly, A. K. Maulloo, and D. K. H. Tan. Rate control for communication networks: Shadow prices, proportional fairness and stability. *Journal of the Operational Research Society*, 49(3):237–252, 1998.
- [10] Robert H. Lasseter and Paolo Piagi. Microgrid: a conceptual solution. In *IEEE Power Electronics Specialists Conference (PESC 2004)*, pages 4284–4290, Aachen, Germany, 2004.
- [11] Javad Lavaei and Steven H. Low. Zero duality gap in optimal power flow problem. *IEEE Transactions on Power Systems*, 27(1):92–107, 2012.
- [12] Sam Miller, Sarvapali D Ramchurn, and Alex Rogers. Optimal decentralised dispatch of embedded generation in the smart grid. In *Proceedings of the Eleventh International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-12)*, pages 281–288, Valencia, Spain, 2012.
- [13] Noam Nisan. Introduction to mechanism design (for computer scientists). In Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani, editors, *Algorithmic Game Theory*, pages 209–241. Cambridge University Press, 2007.
- [14] Richard P. O’Neill, Paul M. Sotkiewicz, Benjamin F. Hobbs, Michael H. Rothkopf, and William R. Stewart, Jr. Efficient market-clearing prices in markets with nonconvexities. *European Journal of Operational Research*, 164(1):269–285, 2005.
- [15] Martin J. Osborne and Ariel Rubinstein. *A Course in Game Theory*. MIT Press, Cambridge, 1994.
- [16] Pearl Pu and Li Chen. User-involved preference elicitation for product search and recommender systems. *AI Magazine*, 29(4):93–103, 2008.
- [17] Sarvapali D. Ramchurn, Perukrishnen Vytelingum, Alex Rogers, and Nicholas R. Jennings. Putting the “smarts” into the smart grid: A grand challenge for artificial intelligence. *Communications of the ACM*, 55(4):86–97, 2012.
- [18] Al Mamunur Rashid, Istvan Albert, Dan Cosley, Shyong K. Lam, Sean M. McNee, Joseph A. Konstan, and John Riedl. Getting to know you: Learning new user preferences in recommender systems. In *International Conference on Intelligent User Interfaces*, pages 127–134, San Francisco, 2002.
- [19] Valentin Robu, Ramachandra Kota, Georgios Chalkiadakis, Alex Rogers, and Nicholas R. Jennings. Cooperative virtual power plant formation using scoring rules. In *Proceedings of the Twenty-sixth AAAI Conference on Artificial Intelligence (AAAI-12)*, pages 1165–1166, Toronto, 2012.

- [20] John D. Rogawski and Martin Shubik. A strategic market game with transactions costs. *Mathematical Social Sciences*, 11(2):139–160, 1986.
- [21] Alvin E. Roth. The economist as engineer: Game theory, experimentation, and computation as tools for design economics. *Econometrica*, 70(4):1341–1378, 2002.
- [22] Walid Saad, Zhu Han, and H. Vincent Poor. Coalitional game theory for cooperative micro-grid distribution networks. In *Proceedings of 2011 IEEE International Conference on Communications Workshops (ICC)*, pages 1–5, Kyoto, 2011.
- [23] Paul A. Samuelson. Spatial price equilibrium and linear programming. *American Economic Review*, 42(3):283–303, 1952.
- [24] Lloyd S. Shapley and Martin Shubik. On market games. *Journal of Economic Theory*, 1(1):9–25, 1969.
- [25] Meritxell Vinyals, Filippo Bistaffa, Alessandro Farinelli, and Alex Rogers. Coalitional energy purchasing in the smart grid. In *IEEE International Energy Conference & Exhibition (ENERGYCON 2012)*, pages 848–853, Florence, 2012.
- [26] Perukrishnen Vytelingum, Sarvapali D. Ramchurn, Thomas D. Voice, Alex Rogers, and Nicholas R. Jennings. Trading agents for the smart electricity grid. In *Proceedings of the Ninth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-10)*, pages 897–904, Toronto, 2010.
- [27] Perukrishnen Vytelingum, Thomas D. Voice, Sarvapali D. Ramchurn, Alex Rogers, and Nicholas R. Jennings. Agent-based micro-storage management for the smart grid. In *Proceedings of the Ninth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-10)*, pages 39–46, Toronto, 2010.
- [28] Allen J. Wood and Bruce F. Wollenberg. *Power Generation, Operation, and Control*. Wiley-Interscience, New York, 1996.

Appendix A

Market Games

In this section, we will provide a brief background on market games and their extensions [24, 20]. Power distribution games with certain types of losses that on each edge that depend only on the behavior of the agents at the endpoints can be represented as market games. Market games always have a core allocation, and it is easy to find (requires solution of a convex optimization problem). We reproduce the proof of core existence here because it provides a possible avenue for proving that the power distribution game with congestive losses has a core allocation.

A market game consists of n agents where each agent i has a concave utility function U_i , an initial allocation of goods a_i . We will also assume that the utility functions are increasing, which can be assured by inducing small slopes to flat parts of utility functions. A coalition of agents C chooses an a final allocation x_i to each agent that maximizes the aggregate utility of the coalition subject to the constraint the the amount of goods present in the final allocation is equal to the amount present in the initial allocation, i.e. agents can only trade with other agents in the same coalition. The value of the aggregate utility of the maximizing allocation for coalition C defines the characteristic function $v(C)$. We want to show that a core allocation always exists, i.e. a payment to each agent t_i such that the sum of the payments is equal to $v(N)$ and for every coalition C , $\sum_{i \in C} t_i \geq v(C)$.

Theorem 2. *Every market game has a core [24].*

Proof. For simplicity, we will present the case where only a single good is traded, so that a_i and x_i are scalars, not vectors. The value of a coalition C is determined by the value of the solution to the following

optimization problem:

$$\begin{aligned} & \max_{x_i, \forall i \in N} \sum U_i(x_i) \\ \text{subject to } & \sum_{i \in N} x_i = \sum_{i \in N} a_i \\ & x_i \geq 0, \forall i \in N \end{aligned}$$

We will show that there exist *market-clearing prices* in the following sense.

Definition 2. A price λ is market-clearing if there exists an allocation $x^* = \{x_1^*, \dots, x_n^*\}$ such that:

1. $U_i(x_i) - \lambda(x_i - a_i)$ subject to $x_i \geq 0$ is maximized at $x_i = x_i^*$ for all agents i .
2. $\sum_{i \in N} b_i = \sum_{i \in N} a_i$.

We will refer to x^* as an allocation induced by λ .

Lemma 1. 1. *Any allocation induced by a market-clearing price maximizes social welfare.*

2. *Any allocation that maximizes social welfare is induced by a market-clearing price.*

Proof. 1. Suppose that $B = \{b_1, \dots, b_n\}$ is an allocation that maximizes social welfare and x^* is an allocation induced by market-clearing price λ .

$$\begin{aligned} U_i(x_i^*) - \lambda(x_i^* - a_i) &\geq U_i(b_i) - \lambda(b_i - a_i) \\ \sum_{i \in N} U_i(x_i^*) - \sum_{i \in N} \lambda(x_i^* - a_i) &\geq \sum_{i \in N} U_i(b_i) - \sum_{i \in N} \lambda(b_i - a_i) \end{aligned}$$

using the definition of market-clearing prices. Because $\sum_{i \in N} x_i^* = \sum_{i \in N} a_i = \sum_{i \in N} b_i$ by definition,

$$\sum_{i \in N} U_i(x_i^*) \geq \sum_{i \in N} U_i(b_i)$$

Thus, the allocation induced by a market-clearing price maximizes social welfare.

2. Since we have assumed that the agent's utility functions are concave and increasing, for any price λ , $U_i(x_i) - \lambda(x_i - a_i)$ subject to $x_i \geq 0$ will have a unique maximum. Suppose that there is an allocation B' that maximizes social welfare that does not have a price associated with it. Since B' is a feasible allocation and we assume that $\sum_i a_i$, at least one agent has $b'_i > 0$. Since there is no

market-clearing price, $U'_i(b'_i)$ is not a market-clearing price. Therefore, there is some agent j such that either:

- (a) $U'_j(b'_j) > U'_i(b'_i)$. In this case, social welfare could be improved by transferring some of the commodity from agent i to agent j , so B' does not maximize social welfare.
- (b) $U'_j(b'_j) < U'_i(b'_i)$ and $b'_j > 0$. In this case, transferring some of the commodity from j to i would improve social welfare, so B' does not maximize social welfare.

□

To show that the core exists, we will give a payoff vector that is in the core, β , which is defined as:

$$\beta_i = U_i(b_i) - \lambda(b_i - a_i)$$

where B is an allocation that maximizes social welfare and λ is a market-clearing price that induces it. Consider any nonempty subset S of N , and let $Y_S = \{y_i : i \in N\}$ be a feasible allocation over S which achieves $v(S)$. Since b_i maximizes the user optimization problem,

$$\beta_i \geq U_i(y_i) - \lambda(y_i - a_i)$$

By summing over i , we will get that the total payoff to agents in S meets or exceeds $v(S)$.

$$\beta(S) \geq \sum_{i \in S} U_i(y_i) - \lambda \sum_{i \in S} (y_i - a_i) = \sum_{i \in S} U_i(y_i) = v(S)$$

because $\sum_{i \in S} (y_i - a_i) = 0$. This shows that β is a stable allocation. To show that it is efficient, we can sum over the payments for all agents.

$$\beta(N) = \sum_{i \in N} (U_i(b_i) - \lambda(b_i - a_i)) = \sum_{i \in N} U_i(b_i) = v(N)$$

□

We can extend market games to allow for a large family of transmission loss functions. The idea for doing this is to rewrite the new game with loss functions as a market game with more commodities—a different commodity for each different origin, e.g., a single commodity becomes n commodities, one for each agent. The loss functions can be different between every pair of agents, and they must be convex so that the new utility functions are still concave. Most critically for power distribution games, the

transmission loss incurred between a pair of agents must only depend on what is transmitted between those two agents. This condition prohibits the natural extension of market games to power distribution games, where the transmission loss incurred on a link depends on the total amount of power transmitted over that edge.

Appendix B

Power Distribution Game Without Resistive Losses

In this appendix, we present a simplified version of the power distribution game described in Chapter 3. Since this game has no resistive losses, it will be possible to explicitly construct core payments, which is useful as these payments will be approximately core stable for the full power distribution game if the losses incurred in practice are low. Therefore, these payments can be viewed as intuitively similar to core payments in the full game, if they exist.

As in Appendix A, we will use slightly different notation than we use in the rest of the analysis to avoid conflict between standard notation for graphs and for cooperative games. Let N be the set of agents. Since there are no resistive losses, it is not necessary to describe the graph structure the agents reside on as long as the agents are connected and cooperating. Each agent i has a net supply of Q_i . Let $Q(C)$ be the net power surplus of coalition (subset of agents) C , i.e., $\sum_{i \in C} Q_i$. Let $Q^+(C) = \sum_{i \in C: Q_i \geq 0} Q_i$ and let $Q^-(C) = \sum_{i \in C: Q_i < 0} Q_i$. $Q^+(C)$ is the total demand of the net demanding agents in C , and $Q^-(C)$ is the total demand of the net supplying agents in C . Then,

$$v(C) = \begin{cases} Q(C)p_s & \text{if } Q(C) \geq 0 \\ Q(C)p_b & \text{if } Q(C) < 0 \end{cases}$$

The value of a coalition is defined as the payment it makes to the grid in order for its total demand to be satisfied, which could be positive or negative depending on whether it has net negative demand or net positive demand. We will divide this value into two parts: the *underlying cost* of satisfying the

agents and the amount by which that cost is reduced when agents cooperate—the *excess utility*. Define the excess utility generated by C as,

$$\begin{aligned}
e(C) &= v(C) - \sum_{i \in C} v(\{i\}) \\
&= -Q^+(C)p_s - Q^-(C)p_b + \begin{cases} Q(C)p_s & \text{if } Q(C) \geq 0 \\ Q(C)p_b & \text{if } Q(C) < 0 \end{cases} \\
&= -Q^+(C)p_s - Q^-(C)p_b + \begin{cases} Q^+(C)p_s + Q^-(C)p_s & \text{if } Q(C) \geq 0 \\ Q^+(C)p_b + Q^-(C)p_b & \text{if } Q(C) < 0 \end{cases} \\
&= \begin{cases} Q^-(C)p_s - Q^-(C)p_b & \text{if } Q(C) \geq 0 \\ Q^+(C)p_b - Q^+(C)p_s & \text{if } Q(C) < 0 \end{cases}
\end{aligned}$$

We will define the *scarce resource* to be whichever of demand and supply is limiting the amount of utility received by the coalition, which is determined by whether $Q^+(C)$ or $Q^-(C)$ has a smaller absolute value. We want to show that dividing the excess utility among the agents who have the scarce resource proportionally to how much of the scarce resource they provide is a core allocation. Agent i 's share of the excess utility will be denoted s_i :

$$s_i = \begin{cases} (p_s - p_b) Q_i & \text{if } Q(N) \geq 0, Q_i < 0 \\ (p_b - p_s) Q_i & \text{if } Q(N) < 0, Q_i \geq 0 \\ 0 & \text{otherwise} \end{cases}$$

Paying s_i plus their underlying costs to each agent i will be the *proportional scarce allocation*, which we will prove is a core allocation.

Definition 3 (Proportional Scarce Allocation (PSA)). The following \vec{t} is the *proportional scarce allocation*.

$$\begin{aligned}
t_i &= v(\{i\}) + s_i \\
&= \begin{cases} Q_i p_s & \text{if } Q_i \geq 0 \\ Q_i p_b & \text{if } Q_i < 0 \end{cases} + \begin{cases} (p_s - p_b) Q_i & \text{if } Q(N) \geq 0, Q_i < 0 \\ (p_b - p_s) Q_i & \text{if } Q(N) < 0, Q_i \geq 0 \\ 0 & \text{otherwise} \end{cases}
\end{aligned}$$

Theorem 3. *The proportional scarce allocation \vec{t} is a core allocation for the power distribution game with no resistive losses.*

Proof. We need to show that the allocation is efficient, i.e. $\sum_{i \in N} t_i = v(N)$, and that each subcoalition's constraint on the sum of the payments to its members is satisfied, i.e. $\sum_{i \in C} t_i \geq v(C), \forall C \in 2^N$. To see that the allocation is efficient, we can show that the sum of the payments equals the value of the grand coalition.

$$\begin{aligned} \sum_{i \in N} t_i &= \sum_{i \in N} v(\{i\}) + \sum_{i \in N} s_i = \sum_{i \in N} v(\{i\}) + \sum_{i \in N} \begin{cases} (p_s - p_b) Q_i & \text{if } Q(N) \geq 0, Q_i < 0 \\ (p_b - p_s) Q_i & \text{if } Q(N) < 0, Q_i \geq 0 \\ 0 & \text{otherwise} \end{cases} \\ &= \sum_{i \in N} v(\{i\}) + \begin{cases} (p_s - p_b) Q^-(N) & \text{if } Q(N) \geq 0 \\ (p_b - p_s) Q^+(N) & \text{if } Q(N) < 0 \end{cases} = \sum_{i \in N} v(\{i\}) + e(N) = v(N) \end{aligned}$$

To show that each subcoalition's constraint is satisfied, consider an arbitrary subcoalition C . The sum of the payments to members of C is

$$\begin{aligned} \sum_{i \in C} t_i &= \sum_{i \in C} v(\{i\}) + \sum_{i \in C} s_i = \sum_{i \in C} v(\{i\}) + \sum_{i \in C} \begin{cases} (p_s - p_b) Q_i & \text{if } Q(N) \geq 0, Q_i < 0 \\ (p_b - p_s) Q_i & \text{if } Q(N) < 0, Q_i \geq 0 \\ 0 & \text{otherwise} \end{cases} \\ &= \sum_{i \in C} v(\{i\}) + \begin{cases} (p_s - p_b) Q^-(C) & \text{if } Q(N) \geq 0 \\ (p_b - p_s) Q^+(C) & \text{if } Q(N) < 0 \end{cases} \\ &= \sum_{i \in C} v(\{i\}) + \begin{cases} e(C) & \text{if } (Q(N) \geq 0, Q(C) \geq 0) \text{ or } (Q(N) < 0, Q(C) < 0) \\ (p_s - p_b) Q^-(C) & \text{if } Q(N) \geq 0, Q(C) < 0 \\ (p_b - p_s) Q^+(C) & \text{if } Q(N) < 0, Q(C) \geq 0 \end{cases} \end{aligned}$$

Consider the three cases separately.

1. $\sum_{i \in C} t_i = \sum_{i \in C} v(\{i\}) + e(C) = v(C)$. In this case, the sum of the payments paid to the members of C exactly equals the value of C .

2. $\sum_{i \in C} t_i = \sum_{i \in C} v(\{i\}) + (p_s - p_b) Q^-(C)$. Since $Q(C) < 0$,

$$(p_s - p_b) Q^-(C) > (p_b - p_s) Q^+(C) = e(C)$$

Therefore, $\sum_{i \in C} t_i > \sum_{i \in C} v(\{i\}) + e(C) = v(C)$.

3. $\sum_{i \in C} t_i = \sum_{i \in C} v(\{i\}) + (p_b - p_s) Q^+(C)$. Since $Q(C) \geq 0$,

$$(p_b - p_s) Q^+(C) \geq (p_s - p_b) Q^-(C) = e(C)$$

Therefore, $\sum_{i \in C} t_i \geq \sum_{i \in C} v(\{i\}) + e(C) = v(C)$.

□

In this allocation, the agents on the scarce side of the market (i.e., supply or demand) divide the benefits of cooperation proportionally amongst themselves according to how much of the scarce resource each member has. It is the unique core allocation in many instances of this coalitional game. The following observations provide some conditions that must hold for other (non-PSA) core allocations to exist.

Observation 1. *The following are necessary conditions for a non-singleton core in the no losses game:*

1. *Existence of veto player. The surplus side of the market has at least one player who would reduce the excess utility of the grand coalition if he unilaterally refused to participate, i.e. his endowment is greater than the difference between the amounts of scarce and surplus resource.*
2. *Existence of insecure player. The scarce side of the market has at least one player whose resource endowment is less than that of the non-veto players on the surplus side.*

Proof. 1. Consider an arbitrary market that does not satisfy the veto player property. Removing any player from the surplus side does not decrease the excess utility of the grand coalition. Thus, $\sum_{i \in N \setminus \{j\}} s_i = e(N)$ for all players j on the surplus side. Therefore, $s_j = 0$ for all players j on the surplus side. Since the surplus side has more resources than the scarce side, every player on the scarce side can threaten to defect with the entire surplus side, thereby guaranteeing that they receive all surplus that they generated. Therefore, the unique core allocation is the PSA if there are no veto players.

2. Consider an arbitrary market that does not satisfy the insecure player property. Non-veto players cannot receive a positive payoff by the logic in the preceding proof. Thus, each player on the scarce

side can threaten to defect with all non-veto players, thereby assuring that each receives all of the excess utility that he generates.

□

Observation 2. 1. *If there is a single veto player and at least one insecure player, the core is non-singleton.*

2. *If there is a single insecure player and at least two veto players, the core may be singleton or non-singleton.*

Proof. 1. Once the insecure players have received their minimum payments given the magnitude of the non-veto players, the remaining excess utility can be given to the veto player.

2. Consider a game with one player of supply 3, $2n - 3$ players of supply 1, one player of demand 1, and n players of demand 2. Supply 3 is insecure with a minimum payment of one unit of excess utility. The supply 1 players are all secure, receiving one unit each. This leaves 2 units left to distribute, which must be received by every pairing of supply 3 and demand 2. Therefore, supply 3 receives 3 units, and the core is singleton for this class of examples.

Consider a game with one player of supply $2n - 1$ and n players of demand 2. Allocate $2n - 2$ units of excess utility to the player of supply $2n - 1$, then divide the remaining unit equally among the n players of demand 2. If there are $n - 1$ or fewer players of demand 2 in a coalition with the supply player, they will have received enough units of excess utility, and the grand coalition is satisfied because the allocation is efficient. Therefore, the core in this class of examples is non-singleton.

□