

New Tools for Econometric Analysis of High-Frequency Time Series Data - Application to Demand-Side Management in Electricity Markets

Maximilian Balandat



Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/EECS-2016-203

<http://www2.eecs.berkeley.edu/Pubs/TechRpts/2016/EECS-2016-203.html>

December 15, 2016

Copyright © 2016, by the author(s).
All rights reserved.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission.

**New Tools for Econometric Analysis of High-Frequency Time Series Data –
Application to Demand-Side Management in Electricity Markets**

by

Maximilian Balandat

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Engineering - Electrical Engineering and Computer Sciences

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Claire J. Tomlin, Chair

Professor Kameshwar Poolla

Professor Jean C. Walrand

Professor Duncan S. Callaway

Fall 2016

**New Tools for Econometric Analysis of High-Frequency Time Series Data –
Application to Demand-Side Management in Electricity Markets**

Copyright 2016
by
Maximilian Balandat

Abstract

New Tools for Econometric Analysis of High-Frequency Time Series Data –
Application to Demand-Side Management in Electricity Markets

by

Maximilian Balandat

Doctor of Philosophy in Engineering - Electrical Engineering and Computer Sciences

University of California, Berkeley

Professor Claire J. Tomlin, Chair

This thesis develops novel methods for econometric analysis of time series data, and applies these methods in the context of demand-side management in California’s electricity markets. The first method is an optimization framework for simulating the behavior of electricity consumers with realistic dynamic consumption models under various dynamic electricity tariffs. This allows to generate new insights into the effect of intertemporal substitution on individual and social surplus under various existing and hypothetical pricing schemes, including Real-Time pricing, Time-of-Use pricing, Critical Peak Pricing, and Critical Peak Rebates. By introducing the concept of a baseline-taking equilibrium, it is possible to quantify the welfare implications of the manipulation of Demand Response (DR) baselines. A second contribution of this thesis is the formulation of the optimal contract design problem a DR aggregator faces when employing inter-temporal flexibility of commercial buildings’ HVAC-related power consumption in order to participate in the regulation capacity spot market. The associated bilevel optimization problem can be cast as a mixed-integer optimization problem that can be solved efficiently. The third contribution of this thesis is a novel methodology for causal inference on time series data based on ideas from Machine Learning. This two-stage approach allows to evaluate conditional individual treatment effects in experimental and non-experimental settings with repeated treatment exposure. The setup is very general and agnostic to the specific regression models used. Besides establishing core theoretical results, this thesis presents associated multi-stage bootstrapping techniques, which allow to perform statistical inference even under computationally challenging prediction models. Based on the developed inference methodology, the empirical contribution of this thesis is to study the effect of receiving DR notifications on energy consumption of residential electricity customers in California, including causal estimates of individual and average treatment effects. These results inform the design of a large-scale randomized controlled trial of much broader scope, which is currently being implemented in collaboration with a DR provider in California and involves more than 10,000 customers over a 14 month duration.

Contents

Contents	i
List of Figures	iii
List of Tables	v
1 Introduction	1
2 Welfare Effects of Dynamic Electricity Pricing	5
2.1 Introduction	5
2.2 Electricity Tariffs	13
2.3 Consumption Model	15
2.4 Simulation Setting and Evaluation Metrics	20
2.5 The Principal Determinants of Tariff Efficiency	22
2.6 The Effects of Demand Response and DR Distortions	38
2.7 Discussion	47
3 Contract Design for Provision of Frequency Regulation Capacity via Aggregation of Commercial Buildings	49
3.1 Preliminaries	50
3.2 A Building’s Scheduling Problem	54
3.3 The Optimal Contract Design Problem	58
4 Machine Learning Methods for Causal Inference on Time Series Data	70
4.1 Introduction	70
4.2 Background	72
4.3 Individual Treatment Effects Under Binary Treatments	77
4.4 Conditional Individual Treatment Effects Under General Treatments	94
4.5 Average Treatment Effects	98
4.6 Bootstrapped Confidence Intervals	103

5	Estimating Individual Treatment Effects in a Residential Demand Response Program	109
5.1	Related Work	110
5.2	A Large-Scale Randomized Controlled Trial	113
5.3	The Pre-Experiment Data Set	117
5.4	Synthetic Experiments on Real Data	124
5.5	Empirical Results on Observational Data	148
A	Supplementary Material for Chapter 2	161
A.1	Formulation of the Optimization Problem	161
A.2	Dynamical System Models	165
A.3	Economics Appendix	169
A.4	Data	172
A.5	Tables of Results	173
B	Supplementary Material for Chapter 5	187
C	Proofs	190
C.1	Proofs from Chapter 3	190
C.2	Proofs from Chapter 4	192
	Bibliography	196

List of Figures

2.1	Simulation DWL vs. time-separable DWL	27
2.2	Normalized social surplus (QU) for different battery sizes and tariffs	28
2.3	VSEAR and VPEAR for different tariffs and battery sizes	31
2.4	QU model, changes in surplus from A-1 tariff benchmark; actual tariffs	33
2.5	QU model, changes in surplus from A-1 tariff; hypothetical tariffs	35
2.6	HVAC model, changes in surplus from A-1 tariff	36
2.7	Baseline manipulation behavior in the QU model	40
2.8	Baseline manipulation behavior in the HVAC model	41
2.9	QU social surplus changes, DR and distortions	44
2.10	HVAC social surplus changes, DR and distortions	47
3.1	The aggregator as an interface between AS market and buildings	52
3.2	Different time scales of the aggregation problem	53
3.3	Disturbance forecasts	64
3.4	CAISO DAM regulation capacity prices for August 2012	65
3.5	Shares of the overall regulation capacities between buildings	66
3.6	Temperature and control for Building 1	67
3.7	Sample temperature trajectories for Building 1	67
3.8	Effect of the parameter ζ	68
4.1	Splitting observations into training and estimation data sets (illustration)	81
5.1	Geographic location of customers and weather stations	118
5.2	Histogram of length of observation histories	119
5.3	Frequency of DR notifications by hour of day	120
5.4	Heterogeneity of LMP across pricing nodes during DR periods	122
5.5	Heterogeneity of LMP across DR periods	123
5.6	Predicting the counterfactual consumption during DR periods	124
5.7	ITE estimates vs. true effects, uniform response	131
5.8	Empirical distribution of ITE estimation errors, uniform response	132
5.9	QQ-plot ITE estimates $\hat{\mu}_i$ against Normal, uniform response	135
5.10	Sample size vs. MDE, uniform response	136

5.11	Estimates and CIs for different group sizes, RFR model, uniform response . . .	137
5.12	QQ-plot parameter estimates $\hat{\beta}_i$ against Normal, linear response	138
5.13	Sample size vs. MDE, linear response	139
5.14	Estimates and CIs for different group sizes, RFR model, linear response	139
5.15	Identifying quantiles using ITE estimates, uniform response	141
5.16	Effect of ranking on conditional EATT, uniform response	142
5.17	Identifying quantiles using ITE estimates, linear response	143
5.18	Effect of ranking on conditional EATT, linear response	144
5.19	Illustration of the ITE for low and high marginal reducers	145
5.20	True and estimated parameters of the treatment effect	146
5.21	Empirical Results: ITE Estimates (semi-log)	151
5.22	Consumption of the “outlier” around DR periods	152
5.23	Empirical Results: ITE and ATT estimates (semi-log)	154
5.24	Empirical Results: ITE and ATT estimates (level)	155
5.25	ITE Estimates (semi-log) - Placebo Test	156
5.26	ITE Estimates (level) - Placebo Test	157
5.27	ITE estimates (level) vs. mean local air temperature during DR events	158
5.28	ITE estimates (level) vs. hourly peak consumption	159
5.29	Heterogeneity in ITE estimates (semi-log) by automation	160
B.1	Parameter estimates vs. true parameters, linear response	187
B.2	Empirical distribution of parameter estimation errors, linear response	188

List of Tables

2.1	Bias-variance decomposition of time-separable DWL ($E_d = -0.1$, A-1 load) . . .	24
2.2	PG&E load-weighted average NP-15 SMC and retail prices	25
2.3	Load-weighted average A-1 ToU prices, non-gen component, and NP-15 SMC . .	42
3.1	Building model parameters	65
5.1	Customers by electric utility	117
5.2	Penetration of home automation devices	121
5.3	Estimation errors $\hat{\mu}_i - \mu_i$ under different methods	129
5.4	Properties of the ITE estimation error, uniform response	133
5.5	Empirical coverage of 95% ITE confidence intervals, uniform response	133
5.6	Empirical coverage of 95% ATT confidence intervals, uniform response	134
5.7	Average width of 95% ATT confidence intervals, uniform response	136
5.8	Empirical coverage of 95% ATT confidence intervals, linear response	137
5.9	Average width of 95% ATT confidence intervals, linear response	137
5.10	Effect of targeting on the EATT in the synthetic experiment	147
5.11	Summary statistics on observations and DR notifications	149
5.12	Summary statistics on consumption and temperature	149
5.13	Penetration of home automation devices in estimation sample	159
A.1	HVAC model parameters	168
A.2	PG&E “Smart Days”	173
A.3	QU results; $E_d = -0.05$, Battery size = None	174
A.4	QU results; $E_d = -0.05$, Battery size = Medium	175
A.5	QU results; $E_d = -0.05$, Battery size = Large	176
A.6	QU results; $E_d = -0.1$, Battery size = None	177
A.7	QU results; $E_d = -0.1$, Battery size = Medium	178
A.8	QU results; $E_d = -0.1$, Battery size = Large	179
A.9	QU results; $E_d = -0.2$, Battery size = None	180
A.10	QU results; $E_d = -0.2$, Battery size = Medium	181
A.11	QU results; $E_d = -0.2$, Battery size = Large	182
A.12	QU results; $E_d = -0.3$, Battery size = None	183

A.13 QU results; $E_d = -0.3$, Battery size = Medium	184
A.14 QU results; $E_d = -0.3$, Battery size = Large	185
A.15 HVAC results	186
B.1 Properties of the ITE estimation error, linear response	188
B.2 Empirical coverage of 95% ITE confidence intervals, linear response	189

Acknowledgments

I would like to thank Clay Campaigne and Lillian Ratliff, with whom much of the research contained in Chapter 2 of this thesis has been performed collaboratively. I have also benefited greatly from collaborating with James Gillan and Datong Zhou on designing and implementing the field experiment described in Chapter 5. Finally, I would like to thank my PhD advisor, Claire Tomlin, for her excellent and consistent support and guidance.

Chapter 1

Introduction

Econometrics on High-Frequency Time Series Data

It is hardly news that we live in a world of data. But of course data itself has no inherent value, it is the insights that can be derived from it that are valuable. For this reason some have called data the “new oil”. Given the pace of the progress in Machine Learning research in recent years, it may be tempting to think that generating such insights will soon amount to little more than applying large-scale learning algorithms to ever increasing amounts of data. This is of course a fallacy – while Machine Learning is extremely good at finding correlations in large-scale and high dimensional data sets, doing so may often not be the goal worth aspiring to. In many situations we want to perform causal inference rather than just uncovering correlations, so as to generate insights from which we can extrapolate to situations in which the underlying data generating process may be different. An important example of high social relevance is the ongoing debate about so-called predictive policing, and to what extent the algorithms employed by law enforcement agencies in fact identify high crime areas and individuals, rather than merely skewing crime discovery rates through the reallocation of policing resources.

Causal inference has been the basic goal of econometricians for a long time, and their field has developed many sophisticated inference tools. While data sets used in econometrics may well be large, they are often either static or at least of relatively low frequency, with the amount of data on any single unit of observation (e.g. an individual person) being typically quite limited. However, this situation is changing quickly: The ubiquity of smart phones, the emergence of completely new areas such as personalized health or autonomous driving, and data-centric designs in many other domains (such as the smart grid) all generate high-frequency time series data on individuals. The amount of information contained in such data sets offers the opportunity to understand behavior on the level of individuals to a potentially much larger extent than is possible with conventional econometric tools. For instance, one can employ a wealth of sophisticated non-parametric Machine Learning methods to capture detailed aspects of the the behavior of individuals.

However, if we want to perform causal inference, we cannot use the conventional approach

to Machine Learning of “everything is just a feature” – instead, we must ensure that the interventions or “treatments” whose causal effect we are interested in understanding receive special attention. If individuals are exposed to these interventions repeatedly, then in some situations we may be able to perform inference by comparing the associated observations with the predictions of the models of their nominal behavior. In other words, we can use non-treatment periods of individuals as a form of “synthetic control,” which provides the potential for estimating the causal effect of the intervention on an individual by using data on that individual alone. While intriguingly simple at the surface, such an approach is subject to a number of significant challenges, including confoundedness of treatment assignments, non-stationarity of the outcome distribution, and the question of how to train Machine Learning models on such data with significant serial correlation. We address these challenges in Chapter 4, where we develop a causal inference framework based on this approach that allows for using generic Machine Learning models.

But it is not only the task of causal inference for which we can utilize high-frequency data to generate new economic insights. Another possibility is to use this data as input to detailed simulations of the consumption behavior of individual economic agents, through which we can evaluate the effects of existing or proposed policies. The benefit of such an analysis is that it is straightforward to consider a many different scenarios, as well as scenarios that for various reasons might be hard or impossible to test in practice. We develop approaches of this kind in Chapters 2 and 3 of this thesis.

While the ideas and methods developed throughout this thesis are relatively universal and easily adopted to other domains, we will focus on their application in the context of electricity markets.

Demand-Side Management in Electricity Markets

The essence of reliable electricity grid operation is matching supply and demand at all times. Today, this is commonly achieved by having generators, load-serving entities (LSEs), and other actors participate in wholesale electricity markets. Markets for electric power, as a place where balancing of demand and supply occurs, are at their core not fundamentally different from markets for any other good, be it milk, cars, or insurance policies. However, electricity is a very special good with a number of peculiarities, the most important ones being the highly variable cost of generation (both temporally and spatially), the high investment cost for increasing generation capacity, the dedicated infrastructure required for transport, and the very high cost of storage. Combined, these factors lead to very high temporal and spatial variability in the marginal cost of consuming electricity, which in conjunction with traditional retail pricing schemes pose major challenges for the market to operate efficiently.

In particular, retail electricity customers have traditionally paid flat electricity tariffs, which by design cannot reflect the different marginal costs of consumption at different times. Since under such tariffs electricity demand on an aggregate level can be forecasted with relatively high accuracy, the traditional way of matching supply and demand has been to have “supply follow load,” i.e., to schedule generation so as to match demand. From a

technical perspective, this approach has been working quite well, but it comes at a potentially high cost, as consumers do not face the correct marginal incentives for consuming electricity. Moreover, with the increasing penetration of renewable generation, additional uncertainty is being introduced on the supply side, as the output of renewable generation (such as wind and solar power) is both stochastic and intermittent. This impact is particularly large in states like California, which have formulated aggressive goals for adopting renewable generation. As a result, in order to be able to ensure reliable supply under the traditional "supply follows load" approach, the capacity of conventional generation (such as coal and gas fired power plants) must be designed to be able to meet peak demand, which often is highly economically inefficient, as this capacity is idle during most hours of the year.

A complementary approach is, instead, to try to modify consumer demand, either through financial incentives or through information and education aimed at inducing a changes in consumption behavior. Broadly, this idea is referred to as "demand-side management," and encompasses a wide range of policies. From the point of view of economic efficiency, the first-best policy is Real-Time Pricing (RTP), in which consumers are directly exposed to the price variation on the electricity wholesale market. However, such a policy has a number of practical issues, including primarily the high volatility in the consumer's electricity bill (Alexander, 2010). A variety of less ambitious policies have been designed and implemented to approximate RTP, including Time-of-Use (ToU) pricing, tariffs that include demand charges, Critical Peak Pricing (CPP), and Critical Peak Rebates (CPR), which are sometimes also referred to simply as Demand Response (DR).

In this thesis, we study the welfare effects of these different dynamic pricing schemes under the behavior of (boundedly) rational electricity consumers, using simulations based on historical high-frequency data. We also propose a new pricing scheme for regulation capacity based on the formulation of an aggregator's contract design problem. To complement these theoretical analyses, we also empirically evaluate the effect of incentivizing consumers to reduce consumption in an existing Residential DR program.

Outline and Main Contributions

To understand the effects of dynamic electricity pricing policies on consumers' consumption behaviors and subsequently on social welfare, one approach is to assume that consumers are (boundedly) rational, utility-maximizing agents. In Chapter 2 we take this view and study the welfare effects of different dynamic electricity pricing schemes (including hypothetical benchmark tariffs as well as various existing tariffs offered by PG&E) by simulating consumer behavior under different realistic scenarios based on historical data from California. Allowing for realistic dynamic consumption models, we gain novel insights into the effect of intertemporal substitution on individual and social surplus. Defining the concept of a baseline-taking equilibrium we are able to quantify the welfare implications of the adverse incentives associated with manipulating the Demand Response baseline.

Maintaining the rationality assumption, and given enough information about the utility function, we can also go one step further and ask what an *optimal* dynamic pricing

scheme would look like. Working within the context of the principal-agent framework, we in Chapter 3 formulate the optimal contract design problem an aggregator faces when using the inter-temporal flexibility of buildings’ HVAC consumption in order to participate in the regulation capacity spot market. We show how to cast the resulting bilevel optimization problem as a mixed-integer program that can be solved efficiently.

While the assumption of rational consumers with known utility functions is convenient for analysis, in reality we typically lack information about consumers’ utility functions, and more often than not face consumers that do not act in a fully rational way. In order to still obtain economic insights into the effects of different dynamic pricing schemes, we thus need to determine from data how consumers actually respond to different policies in the real world. That is, we must perform causal inference, i.e., estimate the causal effect of different interventions (or “treatments”), such as sending a Demand Response notification asking a consumer to reduce their consumption for a limited duration. In Chapter 4 we develop a general causal inference framework for time series data based on generic Machine Learning models. In the context of demand-side management, the high-frequency data collected by smart meters can, in combination with other data sources, be used to estimate a consumption model for an individual in non-treatment periods. Based on this framework, we in Chapter 5 evaluate the effect of a Critical Peak Rebate program operated by our industry partner, a residential Demand Response provider in California.

The fundamental challenge with our generic estimators is that the assumptions under which they are consistent are typically hard (sometimes even impossible) to verify in practice. To validate our estimators in the context of the evaluation, we therefore first perform a number of synthetic experiments on real data. We then investigate the potential for *adaptive targeting*, that is, improve efficiency of DR dispatch by exploiting heterogeneity in the conditional individual treatment effect. Finally, we estimate individual and average treatment effects based on a non-experimental data set and explore heterogeneity in the effects.

In collaboration with our industry partner, we are also currently conducting a large-scale Randomized Controlled Trial (RCT) of much broader scope, which will involve at least 10,000 customers over a 14 month duration. The primary goals of this trial are to estimate the effect of home automation technology on participants’ measured reduction during DR events and to estimate a demand curve, i.e. the average response to a DR event as a function of the reward level provided. Within this field experiment, we also benchmark the average treatment effect estimates we obtain using our methodology against the “gold standard” treatment effect estimates based on random assignment of participants to treatment and control groups. Finally, we study targeting in a randomized experiment, which allows us to obtain causal estimates of the effect of targeting on the average reductions during DR periods. As this RCT has launched only recently, we can at this point only report non-experimental results.

Chapter 2

Welfare Effects of Dynamic Electricity Pricing

This chapter is an adaptation of our current working paper (Campaigne et al., 2016).

2.1 Introduction

Many economists have stressed the importance of dynamic retail pricing for the efficient and reliable functioning of electricity markets. Dynamic tariffs include general Time-of-Use (ToU) tariffs, Critical Peak Pricing (CPP), historical-baseline-dependent Demand Response (DR),¹ and Real-Time Pricing (RTP). Economists are often particularly critical of baseline-dependent DR programs, for giving consumers adverse incentives (double-payment and baseline inflation) that may make their individually optimal behavior detrimental to social welfare. Policy-oriented discussions of dynamic pricing programs often stress the importance of load-shifting behavior, but economic evaluations of dynamic pricing typically assume a time-separable economy, precluding the possibility of intertemporal consumption substitution. Economic studies also typically assume simplified representations of retail tariffs. For example, the standard time-separability assumption precludes the representation of baseline-dependent DR rebates or demand charges.

In this chapter, we model prototypical rational electricity consumers, and formulate their consumption decisions under various dynamic pricing schemes as mathematical optimization problems. Based on historical data from California, we simulate a large number of different realistic scenarios and quantify the impact of various real and idealized electricity tariffs on social welfare. Our framework explicitly models intertemporal substitution of consumption. We introduce the concept of a “baseline-taking equilibrium,” and compute these equilibria, so

¹ The term Demand Response is often used loosely to describe any type of effort to invoke short-term flexibility in electricity consumption. In this thesis we use it much more narrowly to refer to so-called Critical Peak Rebate (CPR) programs (Borenstein, 2005), in which consumers are rewarded according to their “reduction” with respect to a historical baseline.

that we can calculate the welfare impacts of baseline manipulation. The methods and results of this chapter should be seen in contrast to those developed in Chapter 4 and applied in Chapter 5, which focus on empirically measuring aspects of the behavior or actual electricity consumers in a DR program.

The organization of this chapter is as follows. In the remainder of this section we review the relevant literature, discuss our contributions to this literature, and preview our main results. In Sections 2.2 and 2.3, we describe the tariffs we simulate, and the consumer models that face these tariffs, respectively. In Section 2.4, we describe the data setting and parameters that determine certain aspects of the tariffs we simulate, and also the welfare metrics according to which we evaluate tariffs. Sections 2.5 and 2.6 contain our main findings, which divide respectively into a broad discussion of the determinants of efficiency and comparison across real and hypothetical tariffs in general, and the effects of DR and DR distortions in particular. Each of these sections begins with a theoretical overview, followed by a summary of the results of our relevant simulations. Finally, in Section 2.7, we discuss policy implications, modeling limitations, and possible future research directions. Further details are provided in Appendix A.

2.1.1 Related Literature

Our research relates to three main strands of literature. Two strands are in economics: treating the efficiency of various retail electricity pricing schemes in general, and distorted incentives from baseline-dependent demand response programs in particular. The third strand studies engineering models of energy consumers.

Efficiency of Retail Pricing in General

The problem of economically efficient retail pricing of electricity is one of the core instances of the “peak-load pricing” problem: how to optimally price a non-storable good subject to fluctuating demand, produced by a regulated monopolist that faces a production capacity constraint and a break-even revenue requirement. Crew et al. (1995) provide a classic survey of this literature. They characterize the optimal markups of retail prices over marginal operating costs that may be required to pay for capacity costs and other fixed costs under linear pricing, and discuss extensions and related settings.

The most fundamental conclusion of the economics of electricity pricing is that for consumers who behave according to the standard economic model, the most efficient (or “first-best”) outcome occurs when they face a Real-Time Price (RTP) equal to the time-varying social marginal cost (SMC) of generating electricity, including the costs of externalities like Greenhouse Gases (GHGs) and other pollutants.

Borenstein and Holland (2005) discuss the effects of real-time metering and pricing on the efficiency of retail competition in restructured electricity markets, particularly when some fraction of customers remain in flat tariffs. They give a theoretical argument that, while retail competition results in the efficient outcome when all customers face real-time prices,

when some or all remain in a flat tariff, competition fails to achieve the second-best outcome; and nor does it provide optimal incentives for the marginal adoption of real-time metering.² More relevant to our concerns, they also provide simulation-based estimates of the welfare gains and cost savings from three different penetration levels of Real-Time Pricing in a long run competitive equilibrium simulation model, using data from 1999-2003 from California. Borenstein (2005) elaborates further on these simulation results and the underlying data and methodology, and also discusses the much smaller gains that can be obtained from ToU pricing in this model, under various rules determining how fixed costs are collected through volumetric adders.³ In energy and capacity cost terms (disregarding operating reserves, producer market power, and other complicating factors), he estimates the gains from RTP to be on the order of hundreds of millions of dollars annually, or 5-10% total customer bills, and those from ToU to be about 20% as large.⁴

Joskow and Tirole (2006) analyze several economic environments, including those of Borenstein and Holland (2005), and challenge some of the latter's modeling assumptions together with their corresponding conclusions. For our purposes, most relevant is their demonstration that Borenstein and Holland (2005)'s theoretical inefficiency results stem from the restriction to linear (i.e. volumetric) tariffs. Joskow and Tirole conclude that retail competition with flat two-part pricing (a lump-sum access charge plus a linear, per MWh charge that is constant across hours) *can* achieve the second-best optimum.

Borenstein (2005) provides a less formal, more policy-oriented discussion of various types of retail tariffs, including RTP, ToU, demand charges, CPP, interruptible service contracts, and baseline-dependent Demand Response. He estimates, based on wholesale price statistics, that ToU prices can reflect at most 6-13% of wholesale price variation in California (see his footnote 8). Hogan (2014) observes that this fraction of wholesale price variation that is "explained" by hourly or ToU indicator variables (the R-squared from a linear regression model) is an approximate index for the fraction of welfare gains that can be obtained by switching a group of consumers from a flat tariff to a ToU tariff, as compared to switching from flat to RTP. For the PJM market, Hogan reaches a more pessimistic estimate of the gains achievable by ToU than Borenstein (2005). Jacobsen et al. (2016), studying second-best taxation of environmental externalities, establish conditions under which formulas for deadweight loss itself, rather than the ratios of deadweight losses given by Hogan's index, can be expressed as functions of summary statistics from such regression analyses. As a preliminary step in their analysis, they present a standard expression for the deadweight loss due to suboptimal linear prices, based on Harberger (1964)'s seminal "welfare triangle" analysis:

²Second-best settings are settings where some constraints on policy make the otherwise unconstrained socially optimal solution infeasible. In this case the constraints are that consumers are subject to linear pricing, and that some fraction of customers are on flat-rate pricing instead of RTP.

³Transmission and distribution (T&D) costs are assumed to be passed through to customers as a time-invariant \$40/MWh charge.

⁴Borenstein and Holland (2005)'s constant-elasticity demand model implies that total surplus is infinite, so they consider absolute gains in surplus and cost savings, and as a fraction of customer bills (system cost plus \$40/MWh T&D costs, in their model).

The deadweight loss is the demand-derivative-weighted sum of squared differences of retail prices from social marginal costs (eq. (2.6) in our Section 2.5.1).⁵ Hogan’s index corresponds to the special case in which demand derivatives are constant but unknown, and the average markup in each ToU period is zero. The formula of Jacobsen et al. (2016) has the advantage of being applicable to any linear pricing scheme, whereas Hogan’s index is applicable only to comparing the just mentioned second-best scheme, ToU with zero average markup, with the first-best: RTP with zero markup.⁶ The assumption of zero average markup is restrictive, because political and other normative constraints seem to prevent utilities from collecting transmission and distribution (T&D) costs entirely through fixed “meter” charges.⁷

Demand Response Incentives in Particular

Many economists have argued that baseline-dependent Demand Response is a poor substitute for Real-Time Pricing in terms of economic efficiency (Wolak et al., 2009; Chao and DePillis, 2013; Borenstein et al., 2002; Bushnell et al., 2009; Borenstein, 2014). Borenstein (2014) criticizes it for giving incentives that vary drastically about the baseline quantity: if the participant’s demand is too great during a DR event, then it faces no incentive to conserve at all. This is a consequence of the fact that DR is typically designed as a risk-free, “carrot-only” incentive program, rather than a “carrot-and-stick” incentive (Alexander, 2010): customers are encouraged to change their behavior, but they face only an “upside” incentive from the status quo and no downside risk.

DR programs also give consumers two distorted incentives that are a principal focus of this chapter. The “double-payment” distortion is the excessive incentive for demand reduction during DR events that results from the fact that DR participants not only receive the wholesale price per unit reduction, but also avoid paying the retail price, which already includes an estimate of the wholesale price. The “baseline-inflation” distortion is the perverse incentive that consumers are given to increase their consumption in hours that they anticipate may determine the baseline for an upcoming DR hour, in order to increase their DR payment.

Chao and DePillis (2013) analyze these two incentive effects by characterizing the stationary Markov equilibrium of a dynamic model in which the consumer’s utility is a sum

⁵This standard Harberger triangle analysis assumes linear demand functions and constant marginal costs. Therefore, it does not take into account long-term equilibrium effects on the capital stock, as Borenstein and Holland (2005) do.

⁶Jacobsen et al. (2016)’s primary focus is the same as the assumed setting of Hogan’s index. However, Ito (2014)’s observation that consumers seem to respond to average prices rather than marginal puts throws this conventional “marginal” approach into question. For this and similar reasons, we interpret our results as pertaining to idealized rational consumers, rather than consumers as they currently are.

⁷Joskow and Tirole (2006) argue that some fixed costs are already collected through lump-sum charges, so that the need to recoup fixed costs should not prevent efficient, marginal-cost-based pricing. However, Borenstein (2016) points out that determining the appropriate share of system-wide fixed costs for each consumer seems to require arbitrary cost-allocations that are hard to square with normative principles. This is particularly the case for business customers, since companies can have radically different sizes. This difficulty provides an argument in favor of volumetric collection of some portion of fixed costs, despite the resulting economic inefficiency.

over concave, temporally independent stage utility functions, and DR participation is compulsory once enrolled.⁸ They show that both double-payment and baseline inflation result in inefficient consumption levels in their model. In a case of static linear demand and supply, they demonstrate that without baseline inflation (i.e. under a “contractual baseline”), the effect of the double payment distortion is that a Demand Response policy only improves efficiency when DR events are called when the wholesale market price is at least twice the retail rate.

Engineering / HVAC Literature

In the engineering literature, authors have proposed and studied relatively sophisticated consumers and developed algorithms for computing optimal behavior in the face of different pricing schemes. Zavala (2013) focuses on buildings as consumers and gives a comprehensive overview of real-time optimization strategies under dynamic prices. The problem of optimally scheduling different loads of a single consumer, such as electric appliances, is particularly well studied; see e.g. Chen et al. (2012) and Tsui and Chan (2012) and references therein. However, while authors consider a variety of dynamic pricing schemes (Vardakas et al., 2015), results on the welfare effects of existing dynamic policies under historical data are hard to find. Also, there does not seem to be any study of the impact of adverse incentives on social welfare. A number of authors have focused on developing new pricing schemes based on maximizing social welfare, see for example Shi and Wong (2011); Singh et al. (2011); Dong et al. (2012); Samadi et al. (2012); Yang et al. (2013). Others considers the relationship between electricity retailer and consumers in a principal-agent framework (Zugno et al., 2013; Balandat et al., 2014). However, these approaches typically results in very complicated pricing mechanisms that are very far from currently existing policies.

2.1.2 Main Contributions of this Chapter

Relation to the Extant Literature

Our study examines the welfare effects of a number of different real and hypothetical tariffs, for two principal electricity consumer models. Our Quadratic Utility (QU) model represents a generic consumer with a separate demand curve for each time stage, like those from Borenstein and Holland (2005) and Chao and DePillis (2013); but by incorporating a physical model of a battery, we further endow this consumer with the ability to engage in intertemporal substitution. The second model represents an agent operating a commercial building’s Heating, Ventilation and Air Conditioning (HVAC) system, who seeks to minimize total expenditures, subject to maintaining the building’s internal temperature within time-dependent comfort constraints.

⁸When we say that DR is compulsory, we mean that the consumer receives a DR payment which is the reduction from baseline times the wholesale price, even if this quantity is negative. That is, Chao and DePillis’ formulation assumes away the problem of discontinuous incentives noted by (Borenstein, 2014).

We complement the simulation analysis of Borenstein and Holland (2005) by studying the welfare effects of a range of real and realistic tariffs, represented in fine-grained detail. The type of analysis in Borenstein and Holland (2005) does not incorporate Critical Peak Pricing, demand charges, and baseline-dependent Demand Response, and since the latter two features involve intertemporal coupling (as opposed to simple linear prices), the model from Borenstein and Holland (2005) could not be extended to incorporate them. Using the preliminary results in Jacobsen et al. (2016), we show that making the assumption of zero average markup (Hogan, 2014) can be quite misleading, particularly given the large markups over social marginal costs in the real tariffs we study. We assess the gains from Real-Time Pricing, various ToU tariffs, CPP, and DR; and show that under the standard model (particularly without intertemporal substitution), high volumetric markups are a much greater contributor to deadweight loss than is the absence of Real-Time Pricing, at least for realistic tariffs and data drawn from California. But our simulation results also indicate that Real-Time Pricing becomes more important as the capacity for consumption substitution increases. We complement the analysis of Chao and DePillis (2013) with a more detailed and accurate representation of Demand Response revenues, which, due to the voluntary nature of participation, are non-convex in the consumption quantity.

Perhaps the most significant advance in our approach, vis-à-vis the literature described above, is that our models incorporate realistic intertemporal consumption substitution: shifting energy consumption through time either by means of a battery, or via the inherent thermal inertia of a building and its air volume.⁹ This is especially significant because one of the key policy rationales for DR and other time-varying tariffs programs is to incentivize “load shifting” (Faruqui et al., 2012): incentives reflecting scarcity might not merely prevent an act of consumption, but might result in it being rescheduled.¹⁰ We think it is important, especially given advances in automation technology,¹¹ to consider how the dynamic nature of *consumption* may interact with dynamic tariffs.

A shortcoming of our approach, compared to Borenstein and Holland (2005), is that we have no representation of the supply side. We take historical market prices as fully representing the supply side, whereas Borenstein and Holland model the supply mix and market equilibrium. This means that our simulation results are best interpreted as showing the marginal effects of shifting a single consumer, or a small group of consumers, between tariffs, without thereby affecting the supply side. Another limitation of our approach, particularly in comparison to Chao and DePillis (2013), is our assumption of perfect foreknowledge of wholesale prices and thus the timing of DR events.

Our work bridges the gap between the economics and engineering literatures, while making important contributions to both fields: On the one hand it enriches the economics and

⁹While Chao and DePillis (2013) analyze a dynamic equilibrium, dynamics only enter into their model through the baseline-formation process itself, rather than in the internal state of the consumer.

¹⁰Herter and Wayland (2010) provide empirical evidence for load shifting among residential customers in the California Statewide Pricing Pilot, a CPP field experiment.

¹¹Bollinger and Hartmann (2016) and Harding and Lamarche (2016) both find that automation technology, in particular “smart” thermostats, provides significant reductions of peak load.

policy literature by extending existing analyses to more realistic consumption models that allow true inter-temporal substitution, and by defining the novel baseline-taking equilibrium concept that allows to evaluate the cost of manipulation of the DR baseline. On the other hand, our work contributes to the engineering literature by developing a novel optimization formulation that allows us to endogenize realistic DR baselining methodologies currently in use (in particular by CAISO), and by making available a software package that allows to easily apply our economic analyses to a variety of engineering-style consumption models.

Executive Summary of Results

Using historical data from California including real-time wholesale market prices, weather, and representative consumption to calibrate our models (see Appendices A.2 and A.4), we provide estimates of the welfare effects various dynamic pricing policies, and in our Quadratic Utility model, assess their dependence on the elasticity and substitution capacity of demand.

In Section 2.5.1, we show that in our data setting, according to the standard analysis of tariff efficiency from Borenstein and Holland (2005), which ignores intertemporal substitution (essentially, our QU model without battery), the deadweight loss is mostly due to the high average level of markups, rather than tariffs’ failure to co-vary in real time with social marginal costs. However, as we introduce and increase the capability of intertemporal substitution, the average markup has less of an impact on total welfare, and Real-Time Pricing becomes relatively more important (Section 2.5.2). We also show how ToU and RTP tariffs whose price ratios do not reflect the ratios of social marginal costs create inefficient load-shifting incentives for customers who can intertemporally substitute, with the result that having a battery can be detrimental to social welfare.

In Section 2.5.3 we perform a comprehensive comparison of the simulation results for general tariff types. The most salient patterns are that welfare effects generally scale approximately linearly with elasticity, because the effects of tariff differences are mediated by their effects on consumption quantities;¹² and that the larger efficiency effects are typically across tariffs, rather than from adding DR or Peak Day Pricing (PDP)¹³ to a tariff (except for PDP in the A-6 ToU tariff). For example, the A-6 ToU has quite low social welfare in our QU model without a large battery, mostly because the extremely high prices over-penalize consumption during peak ToUs. But with a large battery, the A-6 ToU tariff becomes more efficient; not because the consumer is using the battery efficiently, but because the battery is encouraging it to take advantage of low off-peak prices to consume *more*.

In the QU model, real time pricing tariffs are generally much more efficient than all actually existing tariffs. For a typical business consumer with an annual bill of \$4,010 and elasticity $E_d = -0.1$, our “A-1 RTP” tariff achieves welfare gains of \$66 annually with no battery, and \$357 annually with a medium battery. The greatest gains achieved by (arguably) non-hypothetical tariffs in those settings are \$8 and \$69 respectively, from the A-1 ToU tariff with baseline-taking DR. (All benefits are all calculated relative to the benchmark of the

¹²This result is exact for the Quadratic Utility model in absence of a battery: see eq. (2.6).

¹³PDP is PG&E’s branding for Critical Peak Pricing.

“vanilla” A-1 tariff, on which this consumer model is calibrated.) In the QU model, the effects of PDP and DR are quite small without a battery. For example, for the same typical business consumer with an annual bill of \$4,010 and elasticity $E_d = -0.1$, the benefits are positive but negligible, from \$1-\$10 annually. The no-battery welfare effects scale approximately linearly in elasticity. With a medium battery (with similar specifications as a Tesla PowerWall battery (Tesla Motors, 2016)), the benefits of DR are between \$40 and \$80 annually, with baseline manipulation cases typically falling toward the lower end of the range; and PDP saves nothing in the A-1 ToU, and \$28 annually in the A-6 ToU.

But for the HVAC model, aggressive ToU tariffs (the A-6 and E-19 ToU) are competitive with some RTP tariffs in terms of social welfare, due to large capacity cost savings. DR and Critical Peak Pricing programs typically have beneficial welfare effects, and the latter significantly reduce consumers’ contribution to long-run capacity costs. Tariffs from the A-1 ToU with Peak Day Pricing, E-19 ToU tariffs, and our hypothetical A-1 RTP tariff all achieve cost savings between \$12,000 and \$22,000 annually compared to the “vanilla” A-1 tariff, which has baseline generation cost of \$68,100, and consumer expenditure of \$146,795. These large effects reflect the high degree of flexibility of our HVAC system. DR typically achieves cost savings of \$2,000-\$7,000 annually depending on the base tariff, and PDP delivers smaller benefits, from \$0-\$300, except in the A-1 ToU, where it achieves a surprisingly large benefit of \$12,439, which is entirely due to reduced capacity costs.

In Section 2.6 we focus on the welfare effects of DR and DR distortions. We compute the welfare effects of the double-payment incentive in our simulation scenarios; and, by formulating the concept of a “baseline-taking equilibrium,” we similarly compute the welfare effects of DR baseline manipulation. In the Quadratic Utility model under a realistic tariff, Demand Response has negligible effects without a battery. With a medium battery, DR generates welfare improvements on the scale of 1-2% of annual customer expenditures (or 3-6% of capacity plus generation costs), but the adverse incentives reduce the benefits toward the lower end of that range. With a large battery, the welfare benefits are slightly larger without manipulation, but with manipulation, DR becomes destructive, making demand response destructive to social welfare overall. In the HVAC model, Demand Response creates much larger benefits, on the order of 10% of the social cost of generation plus capacity. Surprisingly, the “adverse incentives” of DR can actually be *beneficial* in a realistic tariff; but in a theoretical case with zero average markup, the adverse incentives are destructive, so much so that the net effect of DR on social welfare becomes negative.

Finally, we observe that a battery may be a worthy investment for the QU consumer under an RTP tariff, but that under currently realistic ToU and DR tariffs, the benefits of a battery are not likely to justify the cost, and may even be negative. Therefore, we would argue that, at this point, programs to subsidize battery investment are unadvisable, until tariffs are reformed to give economically efficient incentives.

2.2 Electricity Tariffs

Economists typically advocate for Real-Time Pricing on the basis of economic efficiency (Borenstein, 2005). However, very few consumers seem to have the sophistication and motivation to make consumption decisions based on real-time prices, such that exposure to unpredictable prices and bills would be worthwhile — as of 2012, only two utilities offered retail RTP plans (Faruqui et al., 2012). Regulators and consumer advocates are wary of requiring or defaulting their constituents into programs with high volatility and unpredictable bills, or exposing sub-populations to retail rates that would be higher than under the status quo (Alexander, 2010; Faruqui et al., 2012). As a result, a number of alternatives have been introduced that can be seen as striking a risk-reward tradeoff that is intermediate between traditional flat-rate pricing and RTP, capturing some of the variability in the cost of energy, while avoiding the unpredictability (Faruqui et al., 2012). Borenstein (2005) and Faruqui et al. (2012) provide overviews of the different types of time-varying pricing programs.

Time-of-Use (ToU) tariffs are one such alternative, under which customers pay different rates in different periods, classified by season, business day vs. holiday or weekend, and time of day. In theory, ToU prices can be interpreted as composed of an estimate of the conditional expectation or conditional weighted average of wholesale prices during the respective ToU period (Hogan, 2014), plus a markup to cover fixed costs that do not vary with the customer’s quantity of energy consumption.

Demand charges are charges proportional to the customer’s maximum demand (in kW), typically averaged over a fifteen minute period. There are several possible rationales for applying demand charges, although Borenstein (2005) is very skeptical that any is economically satisfactory. One justification is that demand charges help manage peak demand, since ToU pricing does not capture any of the considerable residual cost variation during the peak ToU period (Borenstein, 2005). However, this problem would be better addressed with Critical Peak Pricing. Another rationale is that peak demand is a proxy for the cost a customer imposes on the system for distribution capacity. Perhaps the best explanation for the existence of demand charges is simple historical entrenchment: Arthur Wright invented the “electric maximum demand indicator” in 1902 (Wright, 1902), and advocated vigorously on behalf of demand charges. His technology offered an approximate solution to managing peak electric load almost a century before the widespread adoption of real-time meters (Faruqui, 2015).

Critical Peak Pricing (CPP) is another alternative: under CPP, higher prices are charged in a small subset of hours, but the particular times are determined and communicated to the consumer on relatively short notice (e.g., the 20 hours of the year with the highest anticipated prices, with day-ahead notice to the customer). Under standard CPP, the peak price is known at the beginning of the season; under variable peak pricing (VPP), the critical peak prices determined close to real-time, and are related to locational marginal prices (LMPs). CPP is often combined with ToU pricing.

Finally, in Demand Response (DR) or Critical Peak Rebate (CPR) programs,¹⁴ consumers

¹⁴We use the terms DR and CPP interchangeably in this thesis, see footnote 1 in the previous section.

are rewarded for their “reduction” in consumption with respect to some baseline computed based on historical consumption values. A DR policy can be combined with any of the above tariff types.

We use the terms “markup,” “volumetric adder,” and, in Section 2.5.1, “pricing error,” almost interchangeably. Generally the markup is the retail price minus the wholesale price, and volumetric adder is a common term in the electricity industry for the markup per unit energy. The pricing error is the difference between the retail price and the social marginal cost, which, technically speaking, is the markup minus externality costs.

Specific Tariffs Used in our Simulations

In our analysis, we focus on a number of commercial tariffs offered by PG&E in Northern and Central California, as well as several hypothetical benchmark tariffs. The actually existing tariffs include the A-1, A-1 ToU, A-6 ToU, A-10, A-10 ToU, and the E-19 ToU tariffs, which we briefly outline in the remainder of this section. We refer the reader to PG&E’s documentation¹⁵ for more detail.

The A-1 tariff is a “small general service” flat rate tariff. It charges one same energy charge (i.e. per kWh) for all winter periods, from November 1 to April 30, and another energy charge that is approximately 50% higher for all summer periods. However, the A-1 is not open to customers with a maximum demand of 75kW or more for three months in a row, or to newly connecting customers with smart meters. The A-1 ToU is a new small general service Time-of-Use tariff, meant to replace the A-1. Like all PG&E ToU tariffs, in the summer A-1 ToU has on-peak, part-peak, and off-peak energy charges, and for the winter it has part-peak and off-peak charges. Within each season, the difference between the highest and lowest energy rates is about 20%, while the ratio of average summer to winter rates is similar to the original A-1.

The A-6 ToU is a ToU tariff that more aggressively incentivizes load shifting. The summer peak price is \$0.60/kWh, which is four times the summer off-peak price. Off-peak prices are lower than those under A-1 ToU.

The A-10 tariff is similar to the A-1, except that it also has a demand charge. In turn the energy charges are lowered, relative to the A-1. The demand charge is meant to be a simple proxy for the contribution to system peak, which, in theory, would ensure that consumers face the appropriate price signals for contributing to the need for marginal system capacity expansion, although it has been criticized as being ill-suited to that goal (Borenstein, 2005).

Finally, the E-19 tariff combines the aggressive ToU pricing of the A-6 – with a summer peak energy rate more than four times the summer off-peak rate – with a demand charge. It has the lowest off-peak rates of any of the tariffs. Further, it has a more elaborate demand charge formula: in the summer, there are separate charges proportional to the highest 15 minute power draw in a on-peak period and part-peak period respectively, and there is also an additional charge proportional to the maximum of the two power draws just mentioned.

¹⁵Pacific Gas and Electric Company (b), <http://www.pge.com/tariffs/electric.shtml>

The ToU tariffs (A-1 ToU, A-6 ToU, A-10 ToU, and E-19 ToU) all allow for CPP. In particular, PG&E’s Peak Day Pricing (PDP) program is an optional rate offered to consumers already on one of the ToU tariffs that provides a discount on regular summer electricity rates in exchange for higher prices during a small number of peak pricing event days per year. Under PDP, PG&E has the right to call between 9 and 15 PDP events, on a day-ahead basis. On a PDP day, a substantial adder, between \$0.60 and \$1.20, is added to energy charges between 2-6 pm. In exchange, customers are offered reductions in both their off-peak energy charges and in their demand charges. Our simulations include PDP events on the days that they actually occurred.¹⁶ For our analysis we treat Peak Day Pricing and Demand Response as exclusive alternatives.¹⁷

We also consider three hypothetical tariffs: the SMC RTP tariff, the A-1 RTP tariff, and the Opt Flat tariff. The SMC RTP tariff consists only of an energy charge, equal to the time-varying social marginal cost (SMC), that is, the LMP plus the Social Cost of Carbon (see Appendix A.3.2). Capacity costs and non-GHG externalities are not included in these SMCs. The A-1 RTP tariff is a more realistic RTP tariff, which adds LMPs to the A-1 “non-generation rate.” The non-generation rate is the surcharge charged to customers of Community Choice Aggregator as an estimate of the transmission and distribution cost allocation that those customers must pay to PG&E to maintain the transmission and distribution infrastructure. However, we should note that the A-1 RTP tariff has a much lower average price than actually existing tariffs, because the imputed generation portion that we remove from the A-1 tariff to get the non-generation rate is actually much larger than average LMPs.¹⁸ The Opt Flat tariff is a flat tariff equal to the average SMC. This is the optimal flat tariff for a time-separable demand system with identical demand derivatives in every period.

We consider DR in the Opt Flat tariff. We do not consider DR or peak-day pricing in either of the RTP tariffs.

2.3 Consumption Model

In this section we provide a brief overview of the consumption models in our study. We think of such models as consisting of two parts: a basic expenditure model and an electricity consumption model. The expenditure model describes the generic costs associated with consuming electricity under various tariffs,¹⁹ while the electricity consumption model captures the specifics of the consumer’s utility function, constraints, and system dynamics.

¹⁶See Appendix A.4 for a list of these days.

¹⁷PG&E allows simultaneous (“dual”) participation in both programs, but only if the DR is a “day-of” capacity program, rather than a day-ahead energy program. (Pacific Gas and Electric Company, 2011).

¹⁸See Tables 2.2 and 2.3.

¹⁹Accounting for the base rate as well as other components such as demand charges, PDP charges and credits (rate discounts during non-event days in CPP), and DR reward payments.

This modular framework allows us to easily incorporate different types consumers and to analyze how this drives the welfare effects.

We assume that utility is quasi-linear, so that the overall utility of a risk-neutral consumer is $V = U - E$, where U is the total consumption utility (given by the electricity consumption model) and E is the total expenditure.

2.3.1 Expenditure Model

A customer’s total expenditure over T periods under a given retail tariff is given by

$$E = FC + \sum_{t=1}^T \left(p_t^R q_t - \mathbf{1}_{\{t \in \mathcal{E}\}} p_t^{DR} DR_t \right) + DC \tag{2.1}$$

where FC are the tariff’s total fixed charges over all periods,²⁰ and q_t , p_t^R , and p_t^{DR} are the electricity consumption (in kWh), retail energy charge, and demand response reward²¹ (in \$/kWh) in period t , respectively. Further, \mathcal{E} is the set of DR periods, and DC are the total demand charges accrued over all periods. For each period $t \in \mathcal{E}$ the quantity $DR_t = q_t^{BL} - q_t$ is the “reduction” in electricity consumption with respect to the baseline value q_t^{BL} , based on which the consumer is compensated if it participates.²² For the sake of simplicity, our expenditure model assumes that the revenues that a DR provider receives are passed on directly to the DR participant.²³

Various baselining methodologies for DR have been proposed and are used by different ISOs. In our analysis in this chapter, we will focus on the so-called “10 in 10” baseline used by CAISO and detailed in Appendix A.1, under which q_t^{BL} essentially is the average consumption during the same hour of the day over a number of recent non-event days.²⁴ Demand charges, if part of the tariff, are typically high linear prices on the customer’s peak power consumption during each month (or they may be specific to each ToU period of each month: see Section 2.2), averaged over an hourly or quarter-hourly period.

In general, the times at which DR events take place are unknown to the consumer a priori, at least until a certain period (e.g. 24 hours for a day ahead warning) before the event. Moreover, if the DR rewards depend on the real-time or hour-ahead LMP, there is

²⁰e.g. daily meter charges, processing and billing charges, etc.

²¹Here $p_t^{DR} = p_t^W$ for standard DR rewards, and $p_t^{DR} = p_t^W - (p_t^R - p_t^{T\&D})$ for LMP-G rewards, which have been proposed to reduce the double payment distortion (see Section 2.6).

²²In this section we are not yet concerned with how to model the consumer’s decision process of participating in a DR event. See Appendix A.1 for details on how this problem can be formulated as a mixed-integer optimization problem.

²³According to FERC Order 745 (FERC, 2011b), DR is to be compensated at the LMP.

²⁴For simplicity, we do not perform a so-called Load Point Adjustment (LPA) (Coughlin et al., 2008), which aims to improve precision of the baseline by taking into account day-of consumption levels in the early hours of an event day. A multiplicative LPA would be difficult to implement within our optimization framework, because it would introduce a ratio of decision variables into the constraints. An additive LPA, however, would be straightforward to include.

uncertainty about the marginal benefit of reducing consumption during a DR event, even if the period of the event is known. Thus in reality a utility-maximizing consumer faces a stochastic optimization problem that includes her beliefs about both the probability of a DR event occurring and the marginal reward in every period. As such a problem appears intractable without making additional modeling assumptions, we for simplicity consider the benchmark case where the periods and rewards of the DR events are known a priori for the entire simulation horizon.

Assumption 1 (A priori knowledge of DR events). *The set $\mathcal{E} \subset \{1, \dots, T\}$ of DR events as well as the associated Demand Response rewards p_t^{DR} for $t \in \mathcal{E}$ are known to the consumer in period $t = 0$.*

Under Assumption 1, the consumer has perfect knowledge of the effect of its consumption choices on the amount of DR rewards received. Intuitively speaking, this will over-emphasize a consumer’s potential to benefit from artificially inflating their baseline in order to maximize DR payoffs, as doing so in the presence of uncertainty is typically a much less compelling strategy.²⁵

2.3.2 Electricity Consumption Model

We capture the dynamics of the consumption model (and thus the potential for inter-temporal substitution) using the language of dynamical systems. In order to obtain a tractable optimization problem, we restrict ourselves to linear dynamical systems. To simplify notation we focus on time-invariant systems, noting that the extension to time-varying systems is straightforward. Furthermore, since our optimization formulation already includes integer variables, it would be relatively easy to extend our framework to piece-wise affine (PWA) dynamical systems. This increased generality would allow to include approximations to non-linear models that better capture the dynamics of certain systems. For example, Aswani et al. (2012b) argue that PWA systems can provide a more accurate representation of the dynamics of HVAC systems in different operating regimes.

To simplify the exposition in the following, we will restrict our attention to linear dynamical systems for the purpose of this chapter. Specifically, we consider a generic electricity consumer with an internal state $x_t \in \mathbb{R}^{n_x}$ that evolves over time according to a discrete-time linear time-invariant (LTI) system of the form:

$$x_{t+1} = Ax_t + Bu_t + Ev_t \tag{2.2a}$$

$$y_t = Cx_t + Du_t \tag{2.2b}$$

$$q_t = c_q u_t \tag{2.2c}$$

²⁵However, we cannot easily claim the solution under Assumption 1 as an upper bound on the effects of artificial baseline inflation (at least in an almost sure sense), as suboptimal decision-making due to false beliefs in the presence of uncertainty potentially may yield better outcomes for the individual than the expectation-maximizing strategy. We plan to investigate this question further in future work.

Here $u_t \in \mathbb{R}^{n_u}$ denotes the vector of inputs, $y_t \in \mathbb{R}^{n_y}$ the vector of outputs, and $v_t \in \mathbb{R}^{n_v}$ the a vector of disturbances. We assume that $v_t = \hat{v}_t + \nu_t$, where \hat{v}_t is the disturbance forecast and ν_t is a random vector representing the forecast error. The initial state x_0 at time $t = 0$ is known. The system matrices $A \in \mathbb{R}^{n_x \times n_x}$, $B \in \mathbb{R}^{n_x \times n_u}$ and $E \in \mathbb{R}^{n_x \times n_v}$, which describe the state evolution, and the output matrices $C \in \mathbb{R}^{n_y \times n_x}$ and $D \in \mathbb{R}^{n_y \times n_u}$ are also known. Finally, $c_q \in \mathbb{R}^{1 \times n_u}$ is a known vector mapping control input to power consumption. Note that the latter assumption implies that the energy consumption q_t in period t is linear in the control u_t . While this assumption is somewhat restrictive, it still allows for a wide range of interesting and relevant consumption models.²⁶ Stacking states, controls, and outputs, respectively, we can write $\mathbf{x} := [x_0^\top, \dots, x_T^\top]^\top$, $\mathbf{u} := [u_0^\top, \dots, u_{T-1}^\top]^\top$, and $\mathbf{y} := [y_0, \dots, y_T]^\top$.

Typically there will be some hard constraints on the system's control input u , due for example to actuator limits. Moreover, physical limits as well as safety considerations will generally impose hard constraints on the state x and the output y . We assume that these constraints are linear in state and control and thus can be expressed as

$$F\mathbf{x} + G\mathbf{u} \leq \mathbf{0} \tag{2.3}$$

where F and G are appropriate matrices.²⁷ Note that this formulation allows for a wide range of constraints, from simple box constraints over complicated polytopic constraint sets to intertemporal constraints, for example in the form of budget constraints on the control input, or an overall target production quantity in a production model. As with the uncertainty about DR events, we in this initial work for simplicity choose to ignore the forecast errors:

Assumption 2. *There are no errors in the disturbance forecast, i.e., $\nu_t \equiv 0$.*

Under Assumption 2, it is straightforward to see that the dynamics (2.2a) can also be included into the constraints (2.3).

Compared with the fidelity and generality of the consumption models that have been used in the relevant economics literature, our formulation allows for a broad range of more realistic models. In the following we describe the two specific models that we use in our simulations.

Quadratic Utility (QU) with Battery

The first model we consider is a consumer that derives a quadratic utility (QU) from electricity consumption in each time period, giving rise to a standard linear demand curve for each period. We augment this system with a battery that allows for energy storage and thus enables intertemporal substitution of consumption. The consumer who consumes quantity \tilde{q}_t in period t at price p_t^R derives stage utility

$$U_t(q_t) = a_t \tilde{q}_t - \frac{1}{2} b_t \tilde{q}_t^2 - p_t^R \tilde{q}_t \tag{2.4}$$

²⁶Requiring the consumption to be linear in the control is governed by how we formulate the participation decision during Demand Response hours, which relies on this linearity.

²⁷Due to linearity it is clear that constraints on the output \mathbf{y} can be written in this way as well.

so that the total utility is $U = \sum_{t=1}^T U_t(\tilde{q}_t)$. The parameters a_t and b_t are calibrated based on observed consumption levels of a sample of consumers under the A-1 tariff, positing a particular elasticity of demand (see Appendix A.2.1 for details). We model the battery dynamics as a simple continuous-time first-order linear system given by the ODE

$$\dot{x}_\tau = -\frac{1}{T_{\text{leak}}}x_\tau + \eta_c u_{1,\tau} - \frac{1}{\eta_d}u_{2,\tau} \tag{2.5}$$

were x_τ is the battery charge (in kWh) and $u_{1,\tau}$ and $u_{2,\tau}$ are the charge and discharge power (in kW) at time τ , respectively. Further, T_{leak} is the leakage time constant and η_c and η_d are the charging and discharging efficiencies of the battery, respectively. The discrete-time battery model of the form (2.2) is obtained by discretizing (2.5) under zero-order hold sampling. The total amount of energy drawn from the grid in period t is $q_t = u_{1,t} + u_{3,t}$, where $u_{3,t}$ is the energy that is consumed directly, and $u_{1,t}$ is the energy used for charging the battery. The total amount of energy consumed in period t is $\tilde{q}_t = u_{2,t} + u_{3,t}$.

In addition to the base case of no battery, we consider a “medium” and a “large” battery with 10kWh and 25kWh capacity, respectively. We assume simple lower and upper bounds (conditional on battery size) of the form $0 \leq u_{i,t} \leq u_i^{\text{max}}$ on charging and discharging rates. We assume a unidirectional interface to the grid, in which the consumer is unable to discharge stored energy to the grid.²⁸ All parameters and the discrete-time model are given in Section A.2.1 of the Appendix.

Commercial Building HVAC Model

We also consider a simple model of the Heating, Ventilation and Air-Conditioning (HVAC) system of a commercial building. Commercial building HVAC makes up about 14% of total electricity consumption in the U.S.²⁹ Such HVAC systems are a natural candidate for the provision of Demand Response, due to their high level of consumption and the intrinsic thermal inertia of buildings, which allows to shift heating and cooling inter-temporally (Oldewurtel et al., 2013). While commercial HVAC – even when participating in Demand Response programs – tends to be governed by relatively simple heuristic control strategies (Oldewurtel et al., 2013), we contend that an optimization-based approach is well-motivated for the comparison of the economic outcomes under many different policy settings.

The form and parameters of our simplified model are taken from Gondhalekar et al. (2013). The model has three states, which describe aggregates of indoor air temperature, interior wall surface temperature, and exterior wall core temperature (all in °C). The two control inputs u_1 and u_2 are the electric power (in kW) used for heating and for cooling,

²⁸Allowing for a bidirectional interface, while easy to model, would add a second dimension of energy arbitrage to the problem, which would distract from the problem we aim to study in this chapter.

²⁹HVAC accounts for about 40% of commercial building electricity consumption (Fagilde), and commercial buildings comprise about 35% of total U.S. electricity consumption (U.S. Department of Energy, 2012; U.S. Energy Information Administration).

respectively.³⁰ The electric energy drawn from the grid in period t is $q_t = u_{1,t} + u_{2,t}$. Exogenous disturbances are outdoor air temperature, solar radiation, and internal heat sources, and are taken from publicly available data sources (see Appendix A.4 for details).

We impose “comfort constraints” on the interior air temperature $x_{1,t}$ as well as actuation constraints on heating and cooling power consumptions $u_{1,t}$ and $u_{2,t}$ (see Appendix A.2.2 for details). We assume that the utility generated from consuming electricity is independent of the particular temperature profile, so long as it satisfies the comfort constraints. Hence effectively we have that $U = C$ for some constant C if the comfort constraints are satisfied, and $U = -\infty$ otherwise.³¹ By representing the preferences of the occupants by hard comfort constraints, we avoid the issue of estimating the occupants’ dollar value of discomfort incurred by slight deviations from a most-preferred set-point.

Note that while, unlike the QU model, the HVAC model does not include an electric battery, the thermal capacity of the building also enables inter-temporal substitution of consumption, for example by pre-cooling the building during the morning hours.

2.4 Simulation Setting and Evaluation Metrics

2.4.1 Simulation Parameters

For both the Quadratic Utility and HVAC models, we simulate the behavior of the consumer under the different pricing schemes for a range of different parameters. We consider data for the following five geographic regions:³² San Francisco East Bay, San Francisco Peninsula, California Central Coast, Fresno, and Sacramento. For each of these areas, we consider as simulation periods the years 2012, 2013, and 2014, each taken separately. To simplify our exposition and to obtain a metric that is, in some sense, representative for the consumption in recent years in all of California, most of our results are reported in form of the average over both geographical areas and simulation periods.

The periods during each simulation run that are potential DR periods are those whose real-time LMP exceed the threshold determined by CAISO’s so-called net benefits test (NBT) (Xu, 2011). We artificially limit the number of DR events, since simply applying the NBT results in thousands of DR events per year, which we judge to be unrealistic. To simulate n_{DR} DR events during the simulation period, we determine the n_{DR} hours with the highest LMPs, subject to the constraint that there are no more than two events in a single day. While we theoretically can run simulations for an arbitrary number of DR events, for large n_{DR} the

³⁰We acknowledge that most buildings in California are not electrically heated. The point here is not to have a model as accurate as possible, but to understand the effect of intertemporal substitution capability based on the thermal inertia of the building. Moreover, most periods with high LMPs fall in the hot summer months, which means that the effect of heating plays a limited role for Demand Response anyway.

³¹See Section 2.4.2 for additional discussions of consumer utility effects.

³²These map to so-called Sub-Load Aggregation Point (SLAP) nodes defined by CAISO (California Independent System Operator Corporation, 2013).

problem size of the baseline manipulation case quickly becomes intractable.³³ We report results for $n_{DR} = 75$, which appears relatively high given the number of events that are typically called in existing CPP and DR programs.³⁴

2.4.2 Welfare Measures

We evaluate welfare effects of retail tariffs under both the Quadratic Utility (QU) and HVAC consumption models described in Section 2.3. For both prototypical consumers, we evaluate tariffs according to variants of standard welfare measures: consumer surplus, retailer surplus, and the sum of these: social surplus, or total welfare. The consumer surplus is the consumption utility minus the consumer expenditure. The retailer surplus is the consumer expenditure, treated as revenue, minus LMP-weighted consumption, capacity costs (which we break out separately), and greenhouse gas (GHG) externality costs.³⁵ By netting externality costs from the retailer surplus, we are in a sense partitioning society into the consumer on the one hand, and everything else on the other. This is a reasonable scheme, because the consumer is the only optimizing agent in our setup; and in any case, California electric utilities are subject to revenue regulation, such that their allowed revenues are “uncoupled” from sales volume (Migden-Ostrander et al., 2014).³⁶

Because we take historical wholesale prices as given rather than depending on the consumption, these measures give us the *marginal* welfare impact, to the consumer and to the rest of society, of moving a small group of consumers onto one or another tariff.³⁷ The (marginal) social surplus is the sum of the consumer and retailer surpluses: consumption utility, minus procurement and environmental costs.

In any consumption model, ignoring capacity costs, if the consumer faces a tariff equal to the LMP plus externality costs, then the consumer’s objective is identical with the social welfare objective.³⁸ This is the best case for society, and we simulate this situation with our

³³The complexity of this problem does not grow linearly in the number of events, and depends heavily on the number of potential DR events during the 10 day period before the event that is used to determine the 10 in 10 CAISO baseline. See Appendix A.1.1 for details.

³⁴For example, no more than 15 events per year are called in PG&E’s *SmartRate* critical peak pricing plan (Pacific Gas and Electric Company, 2016c).

³⁵The CPUC (California Public Utilities Commission) requires that load-serving entities in California procure sufficient long term capacity to cover their peak loads. We discuss the calculation of environmental costs and capacity costs in Appendices A.3.2 and A.3.3.

³⁶Other presentations might break out externalized environmental costs or DR revenues separately, since they clearly do not accrue to the retailer.

³⁷If we estimated historical cost curves instead of taking historical LMPs as given, we could study the aggregate impact of moving a larger number of consumers between tariffs. We restrict ourselves to the “marginal” setting for the sake of simplicity. This partly accounts for our use of the term “retailer surplus” instead of the more standard “producer surplus,” since it is more realistic to assume that the retailer would procure the bulk of its energy at the LMP, in expectation. In addition, we circumvent certain data challenges, in particular that of attributing spatial heterogeneity to the historical bid curves.

³⁸This is to say that the consumer’s contribution to social cost can be well approximated as a linear expression with a coefficient for energy consumed in each hour. In principle, the consumer’s marginal

SMC RTP tariff. The deadweight loss (DWL) under a given tariff is the total welfare in this hypothetical best case, minus the total welfare under the tariff under consideration.³⁹ Our calculations of deadweight loss are relative to the particular demand models – in particular, battery size and elasticity. This deadweight loss can be interpreted as the amount society loses by suboptimal pricing, assuming that consumer preferences and technology are fixed.

In the HVAC model, the consumption utility is taken to be an arbitrary constant (see Section 2.3.2). In our analysis of welfare impacts we always consider changes in surplus from some benchmark tariff, so that this constant cancels out.

2.5 The Principal Determinants of Tariff Efficiency

2.5.1 “Classical” Time-Separable Analyses

A common refrain among electricity market economists is that real time pricing is the most efficient retail pricing scheme, and that ToU pricing and DR are very inadequate approximations of it (Borenstein, 2005; Hogan, 2014). The latter policies may even be counterproductive distractions, some authors argue, by competing for limited attention and political capital (Bushnell et al., 2009; Hogan, 2014). Hogan (2014) makes this argument with respect to ToU pricing in the context of second-best pricing.⁴⁰ He observes that taking the optimal flat tariff as a benchmark, the optimal ToU tariff can only capture about 11% of the welfare gains achievable by the optimal RTP tariff.⁴¹

The optimal flat tariff has an energy price equal to the demand-derivative-weighted average of social marginal costs, and similarly, the optimal ToU tariff sets the price in each ToU period equal to the demand-derivative-weighted expectation of the SMC conditional on that ToU period.⁴² However, this argument must be qualified by the fact that this form of second-best pricing is itself difficult or impossible to achieve. This is because the utility has substantial fixed administrative, transmission, and distribution costs, and, for the time being, it seems that fixed tariff charges sufficiently high to recoup these costs are politically

contribution to production cost also includes its contribution to ancillary service costs (Tsitsiklis and Xu, 2015).

³⁹In fact, we treat capacity costs in a somewhat inconsistent manner. On the one hand, we do not incorporate them into Social Marginal Costs, because the available data are of questionable quality; there is no definitive methodology for their calculation (see Appendix A.3.3); and the SMC data are of central importance, as an input to both our simulations and the conceptual and statistical analyses in Sections 2.5.1 and 2.5.2. On the other hand, we do depict capacity costs in the summary descriptions of the simulation results in Sections 2.5.3 and 2.6.2.

⁴⁰A second-best policy is a policy that is suboptimal, but optimal subject to some policy constraint under consideration. Here, the constraint is that prices not vary, or not vary within each ToU period.

⁴¹We report Hogan (2014)’s figure for the seemingly favorable assumption that the price can differ for each hour of the day, and that prices are updated annually (see his footnote 4). For hourly ToU prices updated every month, the achievable welfare gains increase to about 20%. According to our data, conditioning on year and ToU period can achieve about a 2-3% reduction in deadweight loss.

⁴²See Borenstein and Holland (2005) p. 475, or Joskow and Tirole (2006).

infeasible, so that a substantial portion must be recovered through volumetric adders to the tariff (Borenstein, 2016). In fact, we observe that the average markups embedded in PG&E tariffs are large enough that, according to standard time-separable consumption models, they account for the great majority of the deadweight loss, so that the failure to co-vary dynamically with social costs pales in comparison.⁴³

However, we describe in Section 2.5.2 that when we allow for intertemporal consumption substitution, the importance of the average markup is diminished, and our results are more consistent with the arguments of economists mentioned above, including their lack of emphasis on average markups.⁴⁴

Jacobsen et al. (2016) present a general formula for Harberger (1964)’s standard characterization of DWL as a function of the mis-pricing “errors,” in a system with linear demand and constant marginal costs:

$$-2 \times \text{DWL} = \sum_{j=1}^J \sum_{k=1}^J e_j e_k \frac{\partial x_j}{\partial e_k} = \sum_{j=1}^J e_j^2 \frac{\partial x_j}{\partial e_j} + \sum_{j=1}^J \sum_{k \neq j}^J e_j e_k \frac{\partial x_j}{\partial e_k} \quad (2.6)$$

In general, the “error” e_i is the difference between the retail price and the social marginal cost for commodity i , and x_j is the demand for commodity j , for $i, j \in \{1, \dots, J\}$. In our setting, the commodities are electricity delivery in particular hours, and the “errors” are the hourly markups, which we also refer to as “volumetric adders.”⁴⁵

In the time-separable consumption model, the second term on the right hand side of (2.6) is zero, and the DWL is a weighted least squares objective, with the weights being the demand derivatives. Treating the retail price as a statistical predictor of the social cost, we can decompose this mean-squared-error loss function into bias and variance components.⁴⁶ The bias component of a tariff’s DWL is the mean tariff error: the average markup, less externality costs. The variance component is the average squared difference between the error and the bias. The variance component is zero if and only if the tariff differs from the SMC by a constant, namely, the bias. Such a tariff is an RTP tariff (reflecting both internal and externalized costs) with a constant volumetric markup.

⁴³Borenstein (2005) addresses this issue, arguing that disregarding the volumetric adder is unlikely to have a substantial effect (see his page 5). One factor explaining the discrepancy between Borenstein’s conclusions and our own is that he considers markups on the order of 10% or 20% of wholesale prices, whereas the markups we observe are on the order of several hundred percent, and none as small as 100% (see Table 2.2).

⁴⁴The reader should bear in mind the caveat that in the present section, we consider only short-run costs, and ignore the cost of peak capacity. We incorporate capacity costs into the social welfare measures in Sections 2.5.3 and 2.6.

⁴⁵We should note that equation (2.6) only holds when nonnegativity constraints are not active for the consumer’s optimal consumption vector. This condition does not hold for our QU consumer under the A-6 ToU PDP tariff with elasticity $E_d \geq 0.2$, since the resulting energy prices are about four times the prices on which that consumer model is calibrated.

⁴⁶Technically, a bias-variance decomposition requires scaling the demand derivatives so that they sum to one, thereby scaling the DWL as well, and then treating them as a notional probability measure (see Appendix A.3.1 for details). Whenever we refer to expectations or variances, the corresponding probability measure incorporates demand-derivative weighting.

In a ToU tariff, the variance component is a weighted average of the SMC variance within each ToU period: $\sum_i w_i \cdot \text{Var}(\text{SMC}|\text{ToU} = \text{ToU}_i)$, where ToU_i ranges over the ToU period types (i.e. summer peak, summer part-peak, summer off peak, winter part-peak, winter off peak), and w_i incorporates both the frequency of ToU periods and their average demand derivatives. The law of total variance implies that $\mathbb{E}[\text{Var}(X|Y)] \leq \text{Var}(X)$, which guarantees that the optimal ToU tariff reduces DWL as compared to the optimal flat tariff, as both have only variance components. In this zero-average-markup second-best setting, the fraction of welfare gained from the optimal ToU tariff, compared to the optimal RTP tariff, is equal to the R-squared from a linear regression of the SMC on ToU indicator variables. This R-squared is the fraction of SMC variance “explained” by the ToU, and when demand-derivatives are the same for all time periods, it is Hogan (2014)’s index, mentioned above.⁴⁷

In Table 2.1 we see that using this decomposition, for a time-separable consumption model on a PG&E tariff, the average markup makes a much larger contribution to deadweight loss than does the failure of retail prices to covary with the SMC.⁴⁸ In the A-1 and A-1 ToU tariffs, the bias component contributes approximately 90% of the DWL. In the A-6 ToU tariff, whose price difference between summer peak and winter off peak ToU periods is approximately 10 times the average summer peak LMP, both the bias and the variance contributions are much greater than those of the A-1 tariffs.⁴⁹

Tariff	DWL	Bias Portion	Variance Portion
A-1	\$112	\$98	\$14
A-1 TOU	\$115	\$101	\$14
A-1 TOU PDP	\$137	\$104	\$33
A-6 TOU	\$278	\$166	\$112
A-6 TOU PDP	\$320	\$154	\$165
A-1 RTP	\$45	\$44	\$1
Opt Flat	\$7	\$0	\$7
SMC RTP	\$0	\$0	\$0

Table 2.1: Bias-variance decomposition of time-separable DWL ($E_d = -0.1$, A-1 load)

In Table 2.2, we display the mean NP-15 pricing node SMC, as well as prices under two tariffs, to give an idea of the magnitude of the average markup.

⁴⁷Some of these statements are made somewhat more nuanced by demand-weighting, but the same intuitive principles apply.

⁴⁸Table 2.1 assumes a constant elasticity of -0.1, and calibrates demand derivatives based on historical load data from the A-1 tariff. These quantities are linear in elasticity.

⁴⁹The A-1 RTP tariff’s nonzero variance component reflects the fact that the volumetric adders equal to PG&E’s non-generation-rate, which are used as volumetric adders on top of an LMP pass-through, are different in the summer and the winter.

	2012	2013	2014	A-1	A-1 ToU	A-6 ToU
Summer peak	60	56	67	242	262	612
Summer part-peak	50	52	61	242	253	286
Summer off peak	40	44	55	242	225	158
Winter part-peak	47	54	66	164	175	181
Winter off peak	40	48	59	164	155	148

Table 2.2: PG&E load-weighted average NP-15 SMC for three years, and retail prices, all in \$/MWh, by ToU period

2.5.2 Substitution Effects Under Linear Energy Pricing

In this section, we focus on incentives that result from linear energy prices – that is, per-unit-energy prices, rather than demand charges or Demand Response – particularly in models like our QU and HVAC models, in which consumers are able to substitute intertemporally. First we explore how the relative contributions of the average markup vs. time-invariance change as substitution capacity changes, either “directly,” via cross-price elasticity, or “indirectly,” by load-shifting using either existing means of storage (HVAC model) or a battery (QU model).⁵⁰ Then we draw a distinction between “level effects” and “load shifting effects” of tariffs on consumption patterns, which helps us explain why some tariffs have the efficiency effects that they do.

For a consumer with the ability to intertemporally substitute, the bias-variance decomposition introduced above no longer exhausts the deadweight loss. Nevertheless, we can still consider markups and a lack of Real-Time Pricing as two principal factors impacting tariff efficiency, and compare their effects. We present two arguments to demonstrate that, as we increase cross-price elasticity directly or indirectly, the high level of markups diminishes in importance, and the lack of Real-Time Pricing – which is in a sense the same thing as high markup variance – becomes more important.

First we consider changing cross-price elasticity directly in a linear demand model. Examining the cross terms in (2.6), we see that, roughly speaking, the more highly correlated tariff errors are for pairs of periods which serve as substitutes (i.e., have large positive cross-price elasticities), the more the substitution effect reduces deadweight loss. On the other hand, if two pricing errors have opposite signs in substitute hours, then they induce inefficient substitution between their respective hours. Using (2.6), we can derive a condition for a two-good linear demand system under which, even if both pricing errors are positive,

⁵⁰In our QU model, we use the battery model as an indirect means of introducing cross-price elasticity. The QU-with-battery demand system is piecewise linear, rather than linear, and so equation (2.6) is only an approximation to the DWL. For the HVAC model, if there were no heat dissipation, then as long as constraints are not binding, the consumer will shift consumption to the cheapest period. This means that, effectively, the cross price-elasticity would be infinite between two periods with different price as long as consumption can be shifted without violating the constraints. In reality, heat dissipation renders it finite, although potentially very high.

increasing the smaller of them can reduce deadweight loss:⁵¹

$$\frac{\partial}{\partial e_1} \left(e_1^2 \frac{\partial x_1}{\partial e_1} + e_2^2 \frac{\partial x_2}{\partial e_2} + e_1 e_2 \left(\frac{\partial x_1}{\partial e_2} + \frac{\partial x_2}{\partial e_1} \right) \right) > 0 \Leftrightarrow \frac{e_2}{e_1} > -2 \frac{\frac{\partial x_1}{\partial e_1}}{\frac{\partial x_1}{\partial e_2} + \frac{\partial x_2}{\partial e_1}} \quad (2.7)$$

This inequality shows that, if the markup of good 2 is high compared to that of good 1, and the cross-price elasticities are large compared to good 1’s own-price elasticity, then increasing the magnitude of e_1 can actually *reduce* deadweight loss, by diminishing exaggerated incentives to substitute good 1 for good 2 (recall that $\partial x_1 / \partial e_1 < 0$, and generally, the cross-price elasticities are positive). The lesson is that equalizing markups across time becomes more important as cross-elasticity increases.⁵²

Now we consider the effect of changing cross-price elasticity “indirectly,” by varying the size of the Quadratic Utility consumer’s battery between None, Medium, and Large.⁵³ We see how this variation in elasticity affects the relative contributions of average markup and correlation with RTP change by comparing the DWL under two hypothetical tariffs. The A-1 RTP tariff has a constant markup (no markup variance), so that its DWL is entirely attributable to markups.⁵⁴ On the other hand, the “Opt Flat” tariff does not track SMC variation at all, but has an average markup of zero, so that its DWL is entirely attributable to the lack of Real-Time Pricing.⁵⁵

In Figure 2.1, we present the result of this analysis, for demand elasticity $E_d = -0.1$. On the x -axis, we plot the deadweight loss in each tariff that results from a time-separable model, such as those assumed by Borenstein and Holland (2005) and Hogan (2014); this is calculated directly from equation (2.6), without intertemporal cross-terms, using tariff data and utility function parameters. On the y -axis, we plot the deadweight loss from our simulation results relative to the social surplus under the SMC RTP tariff, with the same elasticity and battery technology.⁵⁶ In Figure 2.1, circle markers represent no substitution (no battery), diamond markers represent moderate substitution (Medium battery), and inverted triangle markers represent high substitution (Large battery). Each tariff is represented as a vertical stack of three markers, one of each shape, because the x -axis quantity does not

⁵¹The derivation relies on the linearity of the demand system, i.e., the fact that higher-order derivatives of demand quantities with respect to price are zero.

⁵²The argument that substitution between goods drives their optimal markups together has a long history in the taxation literature. Hatta and Haltiwanger (1986), for example, give sufficient conditions on the “strength” of substitutes, which guarantee that “squeezing” their tax rates toward each other would be welfare-improving.

⁵³See Appendix A.2.1 for details on the battery parameters.

⁵⁴The tariff rate is the SMC, plus a constant volumetric adder equal to the A-1 non-generation rate.

⁵⁵The optimal flat tariff, assuming time-separable consumption utility, weights SMCs by their demand derivatives: see eq. (5) in Borenstein and Holland (2005). However, as we use the same tariffs for several different consumer types, we reflect our agnosticism about demand by using an arithmetic average. This discrepancy accounts for the fact that the bias component is not exactly zero. Another reasonable choice might be to use system load weighting, to assure that energy costs are recovered by the LSE.

⁵⁶This comparison assumes that technology is fixed. If we interpret the battery as a proxy for other kinds of substitution preferences, then the comparison would hold those constant as well.

account for substitution capability. The fact that the circle markers lie on the $y = x$ line shows that our consumption model is correctly calibrated.⁵⁷

In the left portion of the figure, we see that without intertemporal substitution, the Opt Flat tariff has a much lower DWL than the A-1 RTP (see also Table 2.1). But as we allow for substitution by introducing and increasing the size of the physical battery model, the Opt Flat tariff induces much larger DWLs. This is because the Opt Flat tariff fails to encourage efficient intertemporal substitution, while the A-1 RTP tariff promotes it. The corresponding results for several actual PG&E commercial tariffs appear in the right portion of Figure 2.1. These tariffs are less efficient than the hypothetical tariffs described so far.

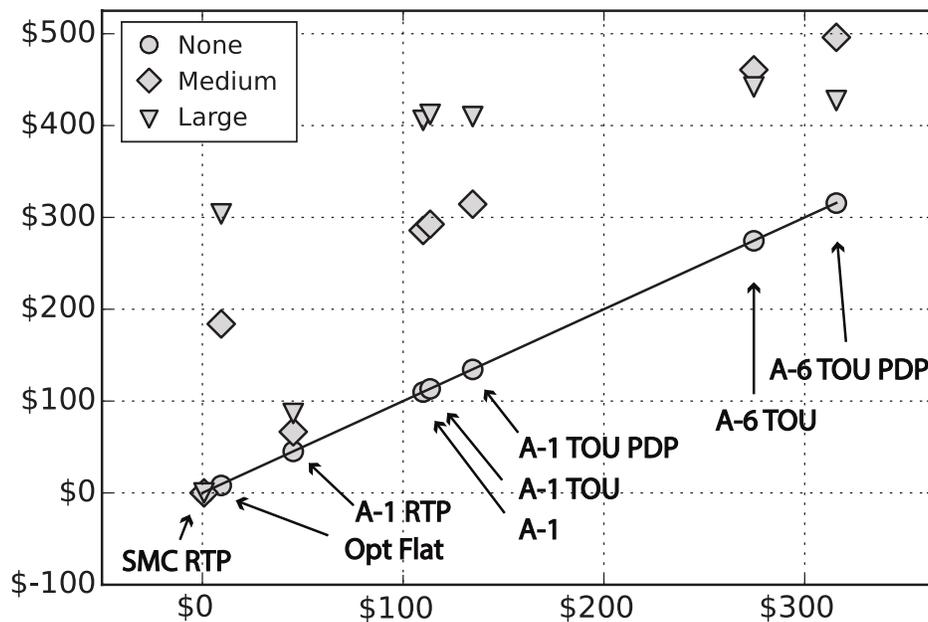


Figure 2.1: Simulation DWL (y -axis) vs. time-separable DWL (x -axis) for actual and idealized tariffs, for three battery sizes.

The A-1 tariff is also flat throughout the day, with a price depending only on season, and therefore provides essentially no incentive to use battery storage.⁵⁸ Its pattern of results is the same as the Opt Flat, except for a translation representing lower efficiency overall, due to consumption-suppression effects because of the higher price level of A-1. The effect that the deadweight loss increases with battery size for all tariffs is primarily due to the reference value: the larger the battery, the more efficient consumption under the SMC RTP becomes, which means that more social value is “left on the table.”

⁵⁷However, if we displayed the same plot for $E_d \leq -0.2$, the DWL predicted by equation (2.6) would overstate the actual DWL for PDP tariffs, because that equation is only valid for “interior solutions,” whereas PDP prices get so high that they drive the optimal unconstrained consumption quantities negative for elastic consumers. See also footnote 45 in this chapter.

⁵⁸Except for a few hours each year when the price levels changes between seasons.

The A-6 ToU tariffs, on the right, show a much different pattern: the DWL for the large battery is less than that for the medium battery. This reflects the fact that, in the A-6 ToU tariff, the medium battery provides little or no welfare benefit, while the marginal benefit of switching to the large battery is even greater in the A-6 ToU tariff than it is in the SMC RTP. This pattern is displayed in Figure 2.2, in which we plot the effect of increasing battery size on social welfare, for several elasticities. Since the y -axis values for the A-6 ToU tariff in Figure 2.1 are scaled differences between the SMC RTP values and the A-6 ToU values in Figure 2.2, the fact that the marginal benefit of the large battery is greater in the A-6 than in the RTP explains the decrease in DWL mentioned above.

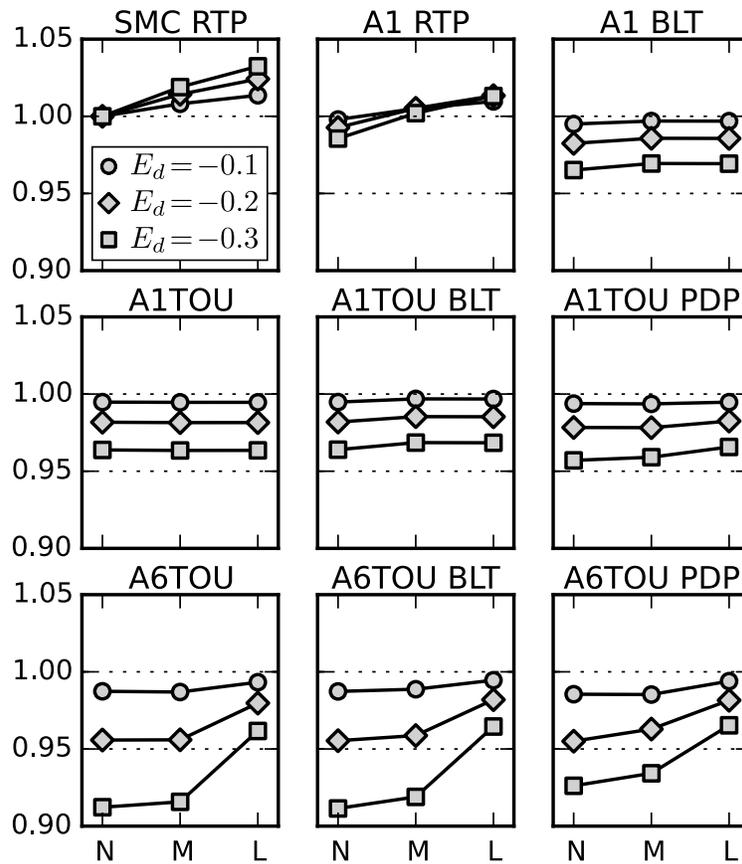


Figure 2.2: Quadratic Utility: Normalized social surplus (disregarding capacity costs) for battery size = (N)one, (M)edium, and (L)arge, for 9 tariff \times DR type combinations (social surplus normalized by value for SMC RTP, no battery)

In Figure 2.2 we omit plots for tariffs, such as the A-1, where the figure would be indistinguishable from constant; but we retain some that one may expect to show variation, but do not. Under RTP tariffs, the welfare increases nearly linearly. A large battery increases social surplus by \$435-\$441 annually in the SMC RTP tariff, and \$340-\$345 in the A-1 RTP tariff, as compared to an annual bill of \$4,010 (or surpluses between \$9,000 and \$40,000,

depending on elasticity). We see that the battery makes very little difference in the A-1 ToU tariff. Counterintuitively, the social surplus slightly *decreases* when a medium battery is added under the A-6 ToU tariff, for elasticities -0.1 (circles) and -0.05 (not pictured), by \$10 and \$13 respectively.⁵⁹

One observation we can make from Figure 2.2 is that the battery does not raise social welfare to very high levels, except in the RTP tariffs. In fact, we will now give an argument that even the substantial social welfare gains from increasing battery size in the A-6 ToU tariff are not due to efficient use of the battery itself.⁶⁰

To make this argument, we decompose the deadweight loss into “level effects” and “load-shifting effects.” That is, we distinguish between (i) whether the consumption levels in each period are efficient, and (ii) whether, given those levels, the use of the battery is efficient.⁶¹ This decomposition is expressed in the equality between (2.8) and (2.9) below.

Defining the “virtual social energy arbitrage revenue” (VSEAR) as the social benefit from shifting energy across time without changing consumption levels,⁶² the inefficiency of suboptimal use of storage can be expressed as the efficient VSEAR, minus the VSEAR under individually optimal behavior, resulting in the equality between (2.9) and (2.10) below:

$$DWL = \text{Efficient Surplus} - \text{Actual Surplus} \tag{2.8}$$

$$= \underbrace{(\text{Efficient Surplus} - \text{Actual Levels Efficiently Sourced})}_{\text{DWL from consuming at inefficient levels: } \geq 0} + \underbrace{(\text{Actual Levels Efficiently Sourced} - \text{Actual Surplus})}_{\text{DWL from inefficient battery use given actual consumption levels: } \geq 0} \tag{2.9}$$

$$= \underbrace{(\text{Efficient Surplus} - \text{Actual Levels Efficiently Sourced})}_{\text{DWL from consuming at inefficient levels: } \geq 0} + \underbrace{(\text{Actual Levels Efficiently Sourced} - \text{Actual Levels No Battery})}_{\text{VSEAR given actual consumption but socially optimal shifting: } \geq 0} + \underbrace{(\text{Actual Levels No Battery} - \text{Actual Surplus})}_{-\text{VSEAR under individually optimal behavior}} \tag{2.10}$$

The first two summands in (2.10) are nonnegative by construction, but their calculation requires an auxiliary optimization which we do not perform. The last term is, in a sense, the impact of actual (individually optimal) battery use on economic efficiency.

⁵⁹This trend is too small to see in the plots, but the reader can refer to the tables in Appendix A.5.

⁶⁰Neubauer and Simpson (2015) make a similar argument, that demand charges give consumers inefficient incentives to exploit on-site storage.

⁶¹These are features of consumption decisions given the ability for intertemporal substitution, rather than tariffs, and are distinct from the bias-variance decomposition of DWL that is applicable for time-separable consumption models.

⁶²i.e., $VSEAR = \sum_t (u_{2,t} - u_{1,t}) SMC_t$ according to the notation from Section 2.3.2. Similarly, VPEAR, defined below, is $VPEAR = \sum_t (u_{2,t} - u_{1,t}) p_t^R$.

For intuition about the cause of inefficient substitution, consider a consumer that will consume a unit of energy in hour i , and has the option of drawing that unit from the grid either in hour i , or drawing a slightly larger quantity to provide for that consumption in hour $j < i$ (that is, we hold the consumption quantity constant, as in the latter components of the decompositions above). The consumer’s battery has charge, discharge, and leakage inefficiencies, η_c , η_d and T_{leak} respectively. Using the parameter values from Appendix A.2.1, the consumer saves the following dollar quantity per unit if it chooses to draw the power in period $j < i$:

$$p_i^R - \frac{p_j^R}{\eta_d \eta_c e^{(i-j)/T_{\text{leak}}}} \approx p_i^R - 1.11 \cdot 1.01^{(i-j)} p_j^R \approx p_i^R - (1.11 + 0.01(i - j)) p_j^R \quad (2.11)$$

We refer to (2.11), summed over time indices, as the virtual private energy arbitrage revenue, or VPEAR.⁶³ The middle expression of (2.11) plugs in our battery model parameters. Because $(1.01)^k \approx 1 + .01 k$ for small k , battery storage effectively increases the price and cost by a fixed 11% for a charge-discharge cycle, plus 1% per hour in storage, compared to the price and cost in the actual production hour.

The social cost savings from that substitution is the same expression with the corresponding social marginal costs in place of retail prices:

$$\text{SMC}_i - \frac{\text{SMC}_j}{\eta_d \eta_c e^{(i-j)/T_{\text{leak}}}} \quad (2.12)$$

The summation of (2.12) over time is the VSEAR, introduced above. When (2.11) $> 0 >$ (2.12), the consumer is given a socially inefficient incentive to substitute intertemporally with storage, and for each unit of energy drawn in j and consumed in i , the quantity (2.12) is incurred as deadweight loss.

The price statistics in Table 2.2 suggest that the consumer is often given inefficient substitution incentives. In particular, the ratio of summer peak to part- and off peak prices is exaggerated in the A-6 ToU tariff (although the consumption-suppression effects are much greater). The fact that wholesale prices have very high variance implies that the static tariffs generate inefficient substitution incentives more often than the means would suggest.

In Figure 2.3, we plot the virtual social energy arbitrage revenues (VSEAR) under three tariffs, as well as the corresponding private arbitrage revenue (VPEAR). Elasticity has very little effect on the results, so we only display the result for elasticity of demand $E_d = -0.1$. We can see that the use of the battery itself is on average destructive of value in the A-6 ToU tariff. This is surprising, when contrasted with the increases in social surplus between the medium and large battery (Figure 2.2). The implication is that the battery increases efficiency in the A-6 ToU tariff by encouraging the consumer to consume *more*, but not by getting the consumer to draw power at more socially efficient times. The “load-shifting”

⁶³This measure only includes energy charges, and is thus not an accurate measure of expenditure savings for tariffs with DR or demand charges. However, it is valid for Peak Day Pricing, since we model PDP as part of the energy charge of the tariff.

effect increases generation costs, but its effect on social surplus is outweighed by the beneficial level effect. Society would be even better off if the consumer’s consumption quantities were held at the levels chosen when it has a battery, without it actually using the battery.

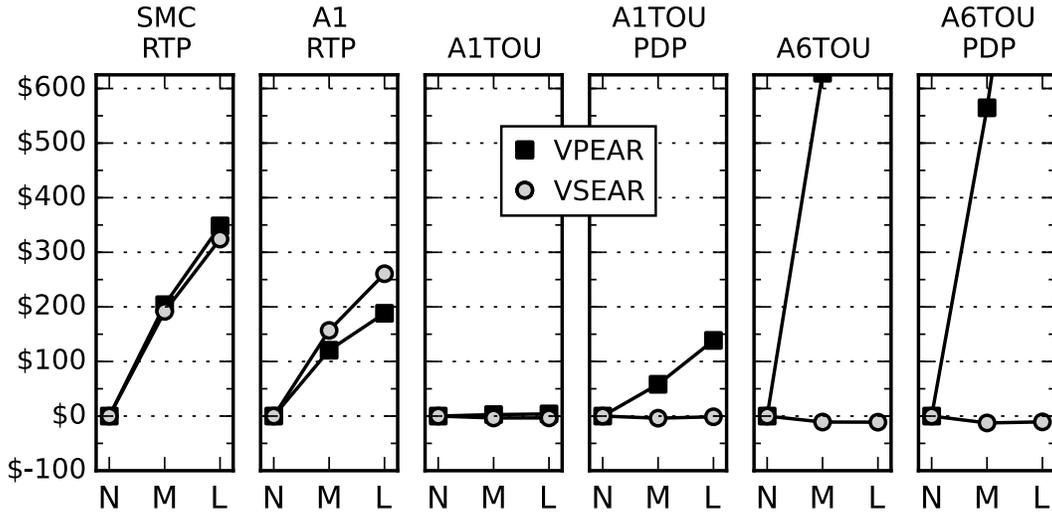


Figure 2.3: Virtual social energy arbitrage revenue (VSEAR) and virtual private energy arbitrage revenue (VPEAR) for 6 tariffs, for Battery Size = (N)one, (M)edium, and (L)arge.

Under existing RTP tariffs, the VSEAR is quite large, whereas under existing ToU tariffs, it is negative, and quite small. In the A-1 ToU case, this is also reflected in the fact that the private benefits from load shifting, as captured in the VPEAR, are also quite small. But in the case of the A-6 ToU, the customer realizes extremely large private benefits – about \$1,300 annually – from load shifting which is in itself socially destructive. This pattern of results helps explain the trends in social surplus depicted in Figure 2.2 above.

These observations, that existing retail tariffs do not align private incentives with social welfare, make us skeptical of the case for public subsidization of on-site battery storage. Noting the often small, mixed or unpredictable effects of battery storage in existing tariffs, we believe that any subsidies should be conditioned on the development of retail tariffs that give consumers reliably efficient price signals.

2.5.3 Simulation Results: Comparison Across Tariffs

Building on the preliminary analysis above, in this section we summarize the cross-tariff comparison of welfare measures.

Quadratic Utility Consumption Model

First we continue discussing the results for the Quadratic Utility model. In Figures 2.4 and 2.5, we plot dollar changes in economic surpluses, under real and hypothetical tariffs

respectively, using the A-1 tariff as a benchmark. The social surplus under the A-1 tariff does not depend on battery size, and is \$8,855 for elasticity $E_d = -0.3$, \$12,000 for $E_d = -0.2$, \$21,435 for $E_d = -0.10$, and \$40,305 for $E_d = -0.05$. For a more concrete benchmark, the consumer expenditure is \$4,010, and the total of SMC and capacity cost is \$1,209, regardless of elasticity or battery size.⁶⁴ All simulation data is presented in Tables A.3 - A.14. We plot the changes in social surplus (thick black arrows) as the sum of three components: change in consumer surplus on top (blue arrows), change in retailer surplus ignoring capacity costs (“retail energy surplus” – red arrows), and *negative* change in capacity costs on the bottom (purple arrows). We represent the summation of these components into the total change in social surplus in the style of “tip-to-tail” vector sum diagrams.

We do not simulate the Quadratic Utility model under tariffs with demand charges (the A-10 and E-19 tariffs), as calibrating their utility parameters by assuming the optimality of historical load data is much more complicated under such tariffs. Moreover, the consumer classes associated with these tariffs seem quite different from those associated to the A-1 and A-6 tariff. For the sake of completeness, we also plot the effect of DR under “baseline-taking equilibrium” (we explore the effects of DR in greater detail in Section 2.6).

The most salient trends in the results reported in Figures 2.4 and 2.5 are that with low elasticities, efficiency effects are quite small, because prices have smaller effects on consumption levels;⁶⁵ and that the larger efficiency effects are typically across tariffs, rather than between the various dynamic variation of each tariffs (except for PDP in the A-6 ToU tariff). As discussed in Section 2.5.1, with consumption substitution the Real-Time Pricing tariffs are generally much more efficient than all other tariffs, particularly with intertemporal substitution. The efficiency effects of DR under “baseline-taking equilibrium” are generally positive but often small.

A-1 and A-1 ToU

The A-1 tariff without DR is not represented in Figure 2.4, because it is the baseline against which other tariffs are compared. The “vanilla” settings of the A-1 and A-1 ToU tariffs have nearly the same results, and are within \$10 of each other for every metric, for every elasticity and battery size.

With no battery, the various settings of the A-1 and A-1 ToU tariffs make almost no difference, particularly in terms of total welfare. The largest difference in total surplus between any two such settings is nearly proportional to elasticity, with \$5 annually at elasticity $E_d = -0.05$, and \$30 annually at $E_d = -0.3$ (recall that for the QU model without intertemporal substitution and with a linear tariff, DWL is linear in elasticity).

With a medium battery, DR and PDP start to have beneficial effects. DR increases social surplus by about \$57 annually in the A-1 tariff, and \$74 annually in the A-1 ToU. Elasticity

⁶⁴Expenditure and cost are constant because the consumer parameters are calibrated to reproduce a given reference consumption trace for each elasticity, and the battery plays almost no role under the A-1 tariff.

⁶⁵The Harberger equation for DWL, (2.6), shows that without intertemporal substitution, deadweight loss is linear in elasticity.

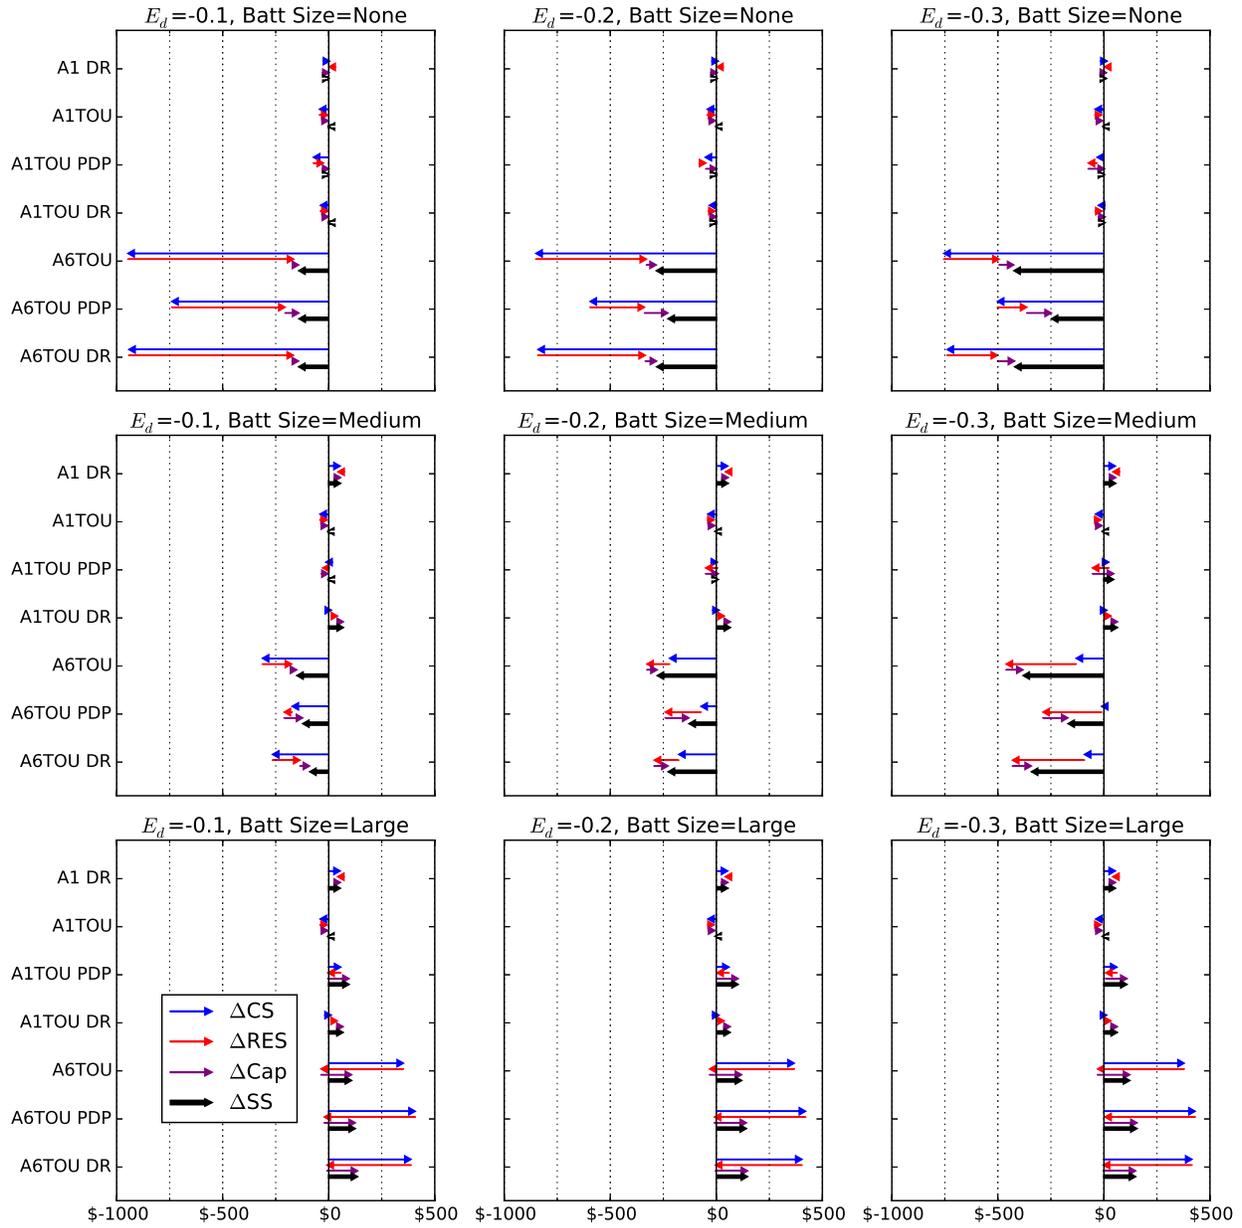


Figure 2.4: QU model, changes in surplus from A-1 tariff benchmark; actual tariffs

does not change these tariff effects by more than a dollar within the range we simulate. The effect of PDP is much smaller and typically less than \$20 annually, except with elasticity $E_d = -0.3$, in which it is \$56 annually. Also, more of the benefits from DR accrue to the consumer (with all parties benefiting when capacity costs are accounted for), which may make DR more viable than PDP as a voluntary program.

With a large battery, the effect of DR is the same as it is in the Medium battery, plus or

minus two dollars. However, PDP becomes more efficient, resulting in social surplus benefits \$30-\$40 greater than DR. In this case, the benefits mostly accrue to the consumer, and the retailer sees a reduced energy surplus, but after accounting for capacity costs, all parties are better off.

A-6 ToU

With no battery, the A-6 ToU tariff is strikingly less efficient than the A-1 tariffs. We discussed this above in Section 2.5.1: the very high markups during peak ToU periods suppress consumption during those periods; and without load shifting, the A-6 ToU's consumption-suppression effect outweighs the effect of the lower prices during the off-peak ToU periods. At elasticity $E_d = -0.1$, the difference is about \$150 annually, and this effect is linear in elasticity. With low elasticities, the price increase in the A-6 ToU causes a large monetary transfer from the consumer to the retailer (\$140 with $E_d = -0.1$), but as elasticity increases, the consumer cuts back, and the size of the transfer diminishes.

With a medium battery, the efficiency effects of the A-6 ToU are similar as without a battery, but the allocation of surplus is much more favorable to the consumer. The effects of elasticity on the allocation of these losses are similar as above, except that the retailer shares in the losses for larger elasticities.

With a large battery, the A-6 ToU becomes more efficient than the A-1, by approximately \$110 annually. We have discussed part of the explanation above, particularly in connection with Figures 2.2 and 2.3: disregarding capacity costs, the beneficial effect is due to the fact that the consumer consumes greater amounts, enjoying higher consumption utility, not because the battery usage is itself efficient. However, we also see that the A-6 ToU generates substantial capacity cost savings with the large battery, since most system peak hours occur in the peak ToU period. We note that elasticity has very little effect with the large battery.

Peak Day Pricing has a much more pronounced effect in the A-6 ToU tariff than in the A-1 tariffs with the small and medium batteries, although the resulting efficiency is still much less. PDP results in substantial capacity cost savings, as well as smaller consumer losses.

With no battery and with the medium battery, DR has a smaller effect in the A-6 ToU tariff than it does in the A-1 tariffs. With a large battery, neither PDP nor DR has an appreciable effect in the A-6 ToU tariff.

Hypothetical Tariffs

The social surplus results for the hypothetical tariffs is largely explained in Sections 2.5.1 and 2.5.2: with no battery, the Opt Flat tariff is more efficient than the A-1 RTP. However, as the consumer becomes able to intertemporally substitute, the Opt Flat tariff is greatly surpassed by the RTP tariffs.

The SMC RTP and Opt Flat tariffs induce huge transfers from the retailer to the consumer, because they do not include volumetric adders. This prevents us from plotting the

hypothetical tariff results on the same scale as the real tariffs above. The A-1 RTP tariff also transfers money to the consumer, but less than the others do, reflecting the fact that the implied generation rate that is subtracted off from the A-1 to arrive at the A-1 non-generation rate is actually much greater than the average LMP.

In all cases, the Real-Time Pricing tariffs are much more efficient than actually existing tariffs. It is easier to compare the efficiency of real and hypothetical tariffs above in Figure 2.1, where they are all on the same scale.

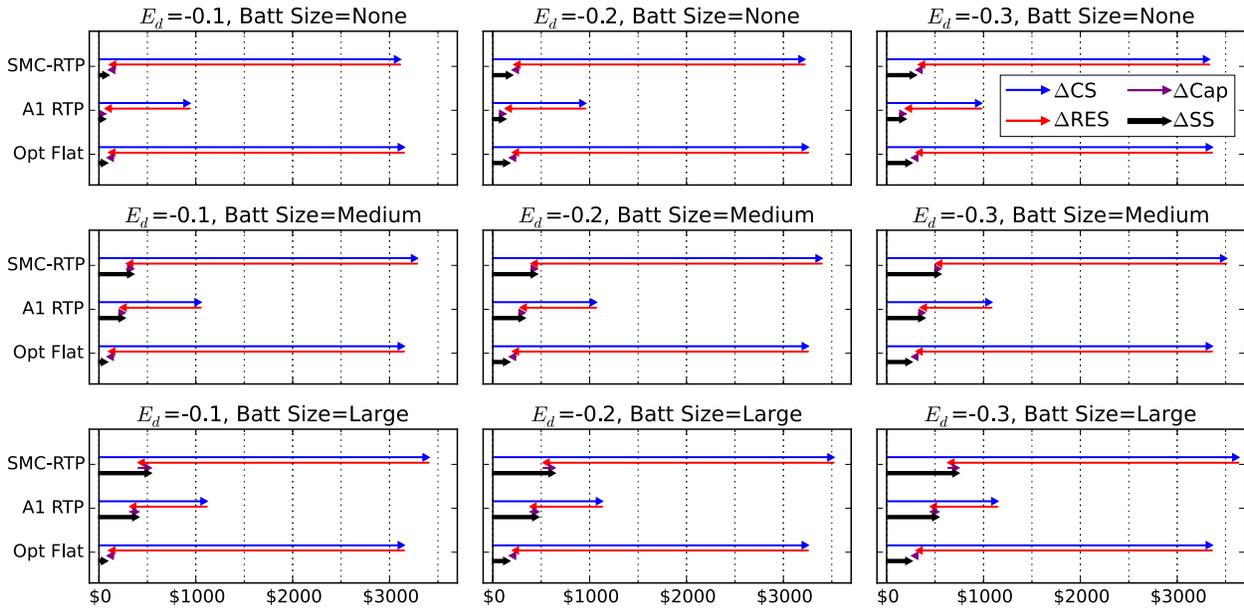


Figure 2.5: QU model, changes in surplus from A-1 tariff; hypothetical tariffs

HVAC consumption model

The comparison across tariffs is less daunting for the HVAC case, since we only consider a single consumption model. In Figure 2.6, we display changes in welfare relative to the benchmark of the A-1 tariff. The A-1 tariff induces a social surplus of $-\$68,100$, which is negative because the value of meeting the comfort constraints is normalized to zero. Consumer expenditure under the A-1 tariff is $\$146,795$.⁶⁶ This data is also presented in Table A.15.

In the HVAC consumption model, the decompositions we introduced above, distinguishing between level effects and load shifting effects, do not apply. Nevertheless, it is clear from

⁶⁶For ease of reading, we truncate the plot at $\$35,000$. The increase in consumer surplus from changing to the SMC RTP tariff is $\$113,491$, and for changing to the Opt Flat tariff, that increase is $\$107,534$. These tariffs entail huge transfers to the consumer, because they do not include volumetric adders. Any realistic implementation would need to include some kind of lump-sum transfer from the consumer, which would arbitrarily change the right endpoints which are not visible here. Also note that by construction, the endpoint of the red arrow is the change in energy generation cost from the A-1 benchmark.

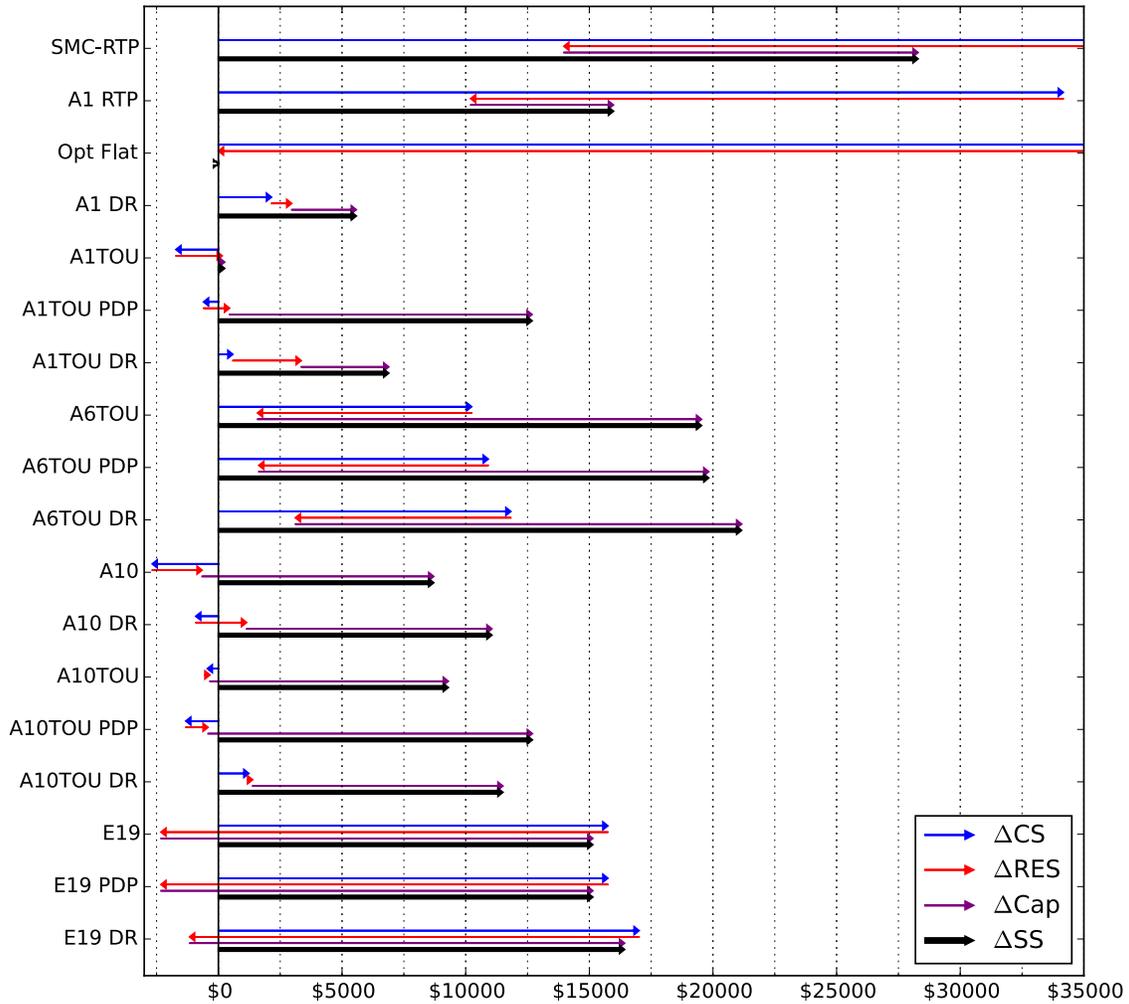


Figure 2.6: HVAC model, changes in surplus from A-1 tariff

the problem formulation that this demand system is subject to load shifting effects but not level effects, in the sense that multiplying all tariff prices by a positive scalar has no effect on the consumer’s optimal solution.⁶⁷

The HVAC model results are very different from those for the Quadratic Utility consumer, presumably largely because of the absence of level effects. The A-1 ToU tariff is very similar to the flat A-1, but with Peak Day Pricing, the A-1 ToU achieves large cost savings, of about

⁶⁷This is because the consumer’s optimization problem is to minimize expenditure, subject to comfort constraints. This problem can be reformulated so that the objective is linear in the vector of retail prices, with no prices showing up in the constraints. Then scaling the vector of prices scales the objective function (by linearity) without affecting the constraints, so that the optimal solution is unaffected. If a tariff includes complicated elements like demand response or demand charges, prices show up in auxiliary constraints. But these constraints can be eliminated by substitution into the objective, at the cost of no longer having a standard-form LP or MIP, which only matters for computational reasons.

\$12,500. Every other tariff induces a substantial improvement over the A-1 tariff, mostly because of large capacity savings. This is especially the case for the A-6 and E-19 tariffs, which have very aggressive ToU pricing. Demand Response (without baseline manipulation) is always beneficial, with the effects are largest in the A-1 and A-1 ToU tariffs. In all tariffs, PDP has positive effects – by far the most beneficial in the A-1 ToU tariff, where it saves approximately \$12,500 in capacity costs.

SMC RTP is, reassuringly, the most efficient tariff. But it is striking, and surprising, that the A-6 ToU tariffs are more efficient than the hypothetical A-1 RTP, and that the E-19 tariffs are only slightly less efficient than the A-1 RTP. The benefits from these aggressive ToU tariffs are primarily due to reductions in capacity costs, which more than compensate for less beneficial energy cost effects (as compared to A-1 RTP). This can be seen by noting that the endpoint of the red arrow represents the change in social energy costs, and the length of the purple arrow represents the reduction in capacity costs. The A-10 tariff has similar energy costs as the A-1, but, presumably due to demand charges, has lower capacity costs. The E-19 tariff has consistently *higher* generation costs than the A-1 benchmark, but the capacity cost savings more than compensate, so that it is also more efficient than A-1.

The fact that the hypothetical A-1 RTP tariff does not compare as favorably against realistic tariffs as it does in the Quadratic Utility model, due to its smaller reductions in capacity costs than those achieved by the realistic tariffs, probably indicates that RTP tariffs should include contribution to capacity costs, as our simulated RTP tariffs do not. It appears that LMPs alone do not provide a sufficient incentive to reduce system capacity costs, particularly when their effect is diluted by flat volumetric adders. However, we should note that our methodology for computing capacity costs is subject to noise, being derived from such small “samples” of hours, and is based on a public available dataset which we do not regard as a particularly reliable measure of actual capacity costs.⁶⁸

The most dramatic pattern in private expenditures in the comparison of real tariffs is that the A-6 ToU and E-19 ToU tariffs are the cheapest for the consumer. This is because our HVAC system is very capable of intertemporal substitution, perhaps particularly in the weather regime we consider, so that the tariffs with the lowest off-peak price are the cheapest. This substitution entails a loss of retailer energy surplus, because the retail price differences are greater than the average wholesale price differences. The hypothetical tariffs all induce large transfers from the retailer to the consumer, because actual tariffs include such high markups over LMPs.

⁶⁸One way to reduce the variance of capacity cost estimates would be to average effects across heterogenous consumers. Another would be to adopt something like the “probabilistic” capacity charge allocation from Boomhower and Davis (2016).

2.6 The Effects of Demand Response and DR Distortions

In this section, we examine the welfare effects of Demand Response, with a focus on two economic distortions commented on in the literature: “baseline manipulation” and “double payment” (Chao and DePillis, 2013; Hogan, 2010; Borlick et al., 2012).

2.6.1 Theoretical Overview

Baseline Manipulation

A fully rational consumer who understands the method used for determining the DR baseline may have an incentive to artificially inflate her consumption during certain periods in order to increase the rewards from DR “reductions” during periods of high reward p_t^{DR} . In our model the baseline values are determined endogenously as part of the optimization problem, so these incentives are captured correctly and we can indeed observe this behavior in our simulations (see below for examples). To evaluate the effects of baseline manipulation compared to the behavior of a non-strategic customer, we consider a no-manipulation benchmark that we refer to as “baseline-taking equilibrium”:

Definition 1 (Baseline-Taking Equilibrium). *Let $\beta : \mathbf{q} \mapsto \mathbf{q}^{BL}$ denote the function mapping a consumption sequence $\mathbf{q} = (q_1, \dots, q_T)$ to a sequence of baseline values $\mathbf{q}^{BL} = (q_1^{BL}, \dots, q_T^{BL})$, and let \mathcal{C} denote the set of constraints on state and control variables of the consumer. Then $(\mathbf{x}^*, \mathbf{u}^*, \mathbf{q}^*)$ is a baseline-taking equilibrium if $\mathbf{q}^* = \beta(\mathbf{q}^*)$ and $(\mathbf{x}^*, \mathbf{u}^*) \in \arg \max_{\mathbf{x}, \mathbf{u} \in \mathcal{C}} V(\mathbf{x}, \mathbf{u}; \mathbf{q}^*)$.*

In words, a baseline-taking consumer regards the DR baseline values as exogenously given data, rather than decision variables as in the “fully rational” model. In equilibrium, the consumer’s optimal response to these baseline values results, as computed by the given baselining methodology (in our case, CAISO “10 in 10”), in the given DR baseline. That is, the vector of baseline quantities is a fixed point of the operator $\beta(\cdot)$ in Definition 1. Algorithm 5 in Appendix A.1.2 describes the fixed-point iteration we use to compute a baseline-taking equilibrium.

The contrast between such a baseline-taking equilibrium and the strategic, or “baseline-manipulation” optimum is analogous to the contrast between a price-taking equilibrium on the one hand, and monopoly or Cournot oligopoly pricing outcomes on the other.⁶⁹

⁶⁹To solve for a price-taking equilibrium, the economist characterizes producers’ optimal quantity response as a function of an exogenously determined price over which the producer has no strategic control. Then the economist uses a market-clearing condition relating prices and total production quantities to determine the equilibrium price that supports these quantity decision. By comparing prices and quantities in both economic environments, one can, arguably, capture the effects of strategic “manipulation” of baselines, and prices, respectively.

To illustrate some of the baseline manipulation resulting from the distorted incentives we show some simulation results that highlight the effect of DR with and without baseline manipulation. These plots show how a fully rational (i.e. strategic) DR participant would behave with advanced knowledge of DR event days. They also demonstrate that optimal behavior in baseline-taking equilibrium corresponds to an intuitive understanding of how a boundedly rational DR participant, who ignores the incentive to inflate the baseline, would behave instead.

Baseline manipulation in the QU model A sample solution from our simulation of the Quadratic Utility model under the Opt Flat tariff with elasticity $E_d = -0.3$ is shown in Figure 2.7 for the strategic agent (“CAISO”), the baseline-taking agent (“BLT”), and the “Nominal” agent, who is not exposed to any DR incentives. We highlight the DR event as well as the associated “BL-relevant” periods, i.e. the periods that are considered for determining the baseline value for the DR period.

The top panel of Figure 2.7 shows the charging (u_1) and discharging (u_2) rates of the consumer’s battery.⁷⁰ We see that immediately before and during the DR event, both the strategic agent and the BLT agent act the same: before the event they charge the battery, and then they consume from the battery during the event. The middle panel shows the total energy drawn from the grid ($u_1 + u_3$). From this we see that during the DR event both the strategic and BLT agent in fact consume exclusively from their battery and draw no power from the grid (the Nominal agent ignores the event). But at 24 and 48 hours before the event, the BLT agent acts the same as the Nominal agent, while the strategic agent charges its battery as rapidly as possible during the baseline-setting hour, and discharges during the subsequent two hours. The evolution of the battery charge,⁷¹ shown in the bottom panel of Figure 2.7, reflects the behaviors described above.

Baseline manipulation in the HVAC model We present a similar sample solution for the HVAC model under the A-1 ToU tariff in Figure 2.8. The top panel shows the cooling power, while the bottom panel shows the evolution of the temperature in the building. Since this is a hot summer day, the building has to use a significant amount of energy to cool the building in order to satisfy the comfort constraints.

In the 12 hours leading up the DR event, the strategic and the BLT agent both behave in the same way: they pre-cool the building considerably, so that during the DR event, they can forgo the use of cooling. In the temperature evolution we can see that after the DR event the temperature hits the upper constraint. This can be contrasted with the behavior of the Nominal agent, who does not pre-cool the building and hence needs to use considerable cooling power during the DR event in order to satisfy the comfort constraint. 24 hours

⁷⁰We do not include the respective inputs of the nominal agent, who under the OptFlat tariff does not utilize the battery and exclusively consumes energy directly from the grid.

⁷¹Note that in our discrete-time model the charging and discharging during period t is reflected in the battery charge only in period $t + 1$.

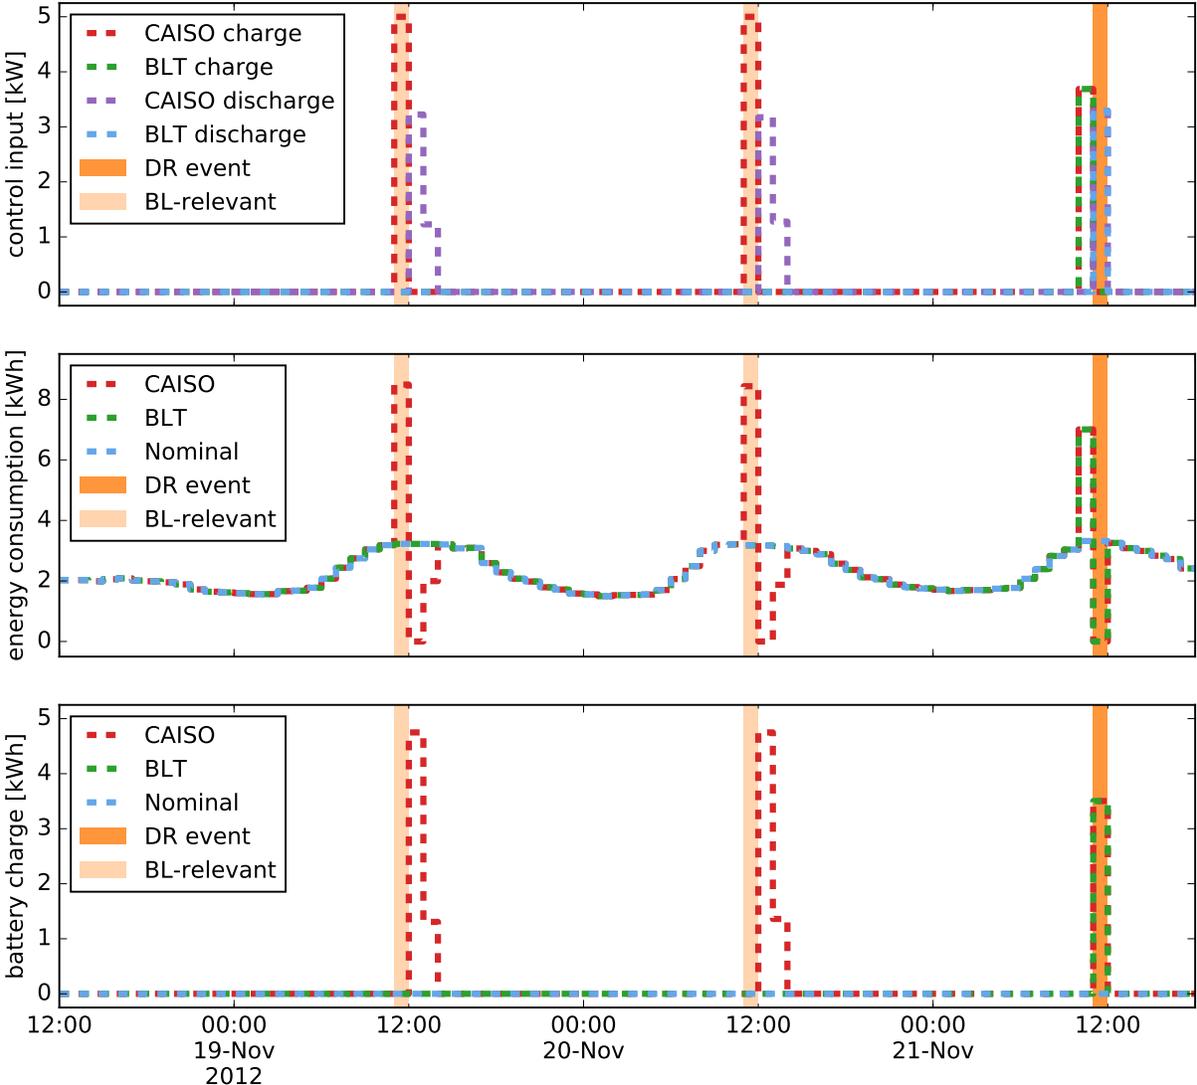


Figure 2.7: Baseline manipulation behavior in the QU model, OptFlat, medium battery

before the DR event, the BLT agent’s behavior is indistinguishable from that of the Nominal agent, while the strategic agent “overcools” for an hour by running the HVAC at a higher power level than necessary for satisfying the comfort constraints (we can see this because the associated temperature trace drops slightly below the other temperature traces in the hour after the baseline-relevant hour, whereas at other times it is mostly hidden behind them). Then the strategic DR participant allows the temperature to drift back up to the upper constraint. Given our modeling assumption that the building’s occupants are indifferent to temperature as long as it satisfies the comfort constraints, this overcooling behavior is wasteful, because the energy needed to maintain a certain temperature level is greater the

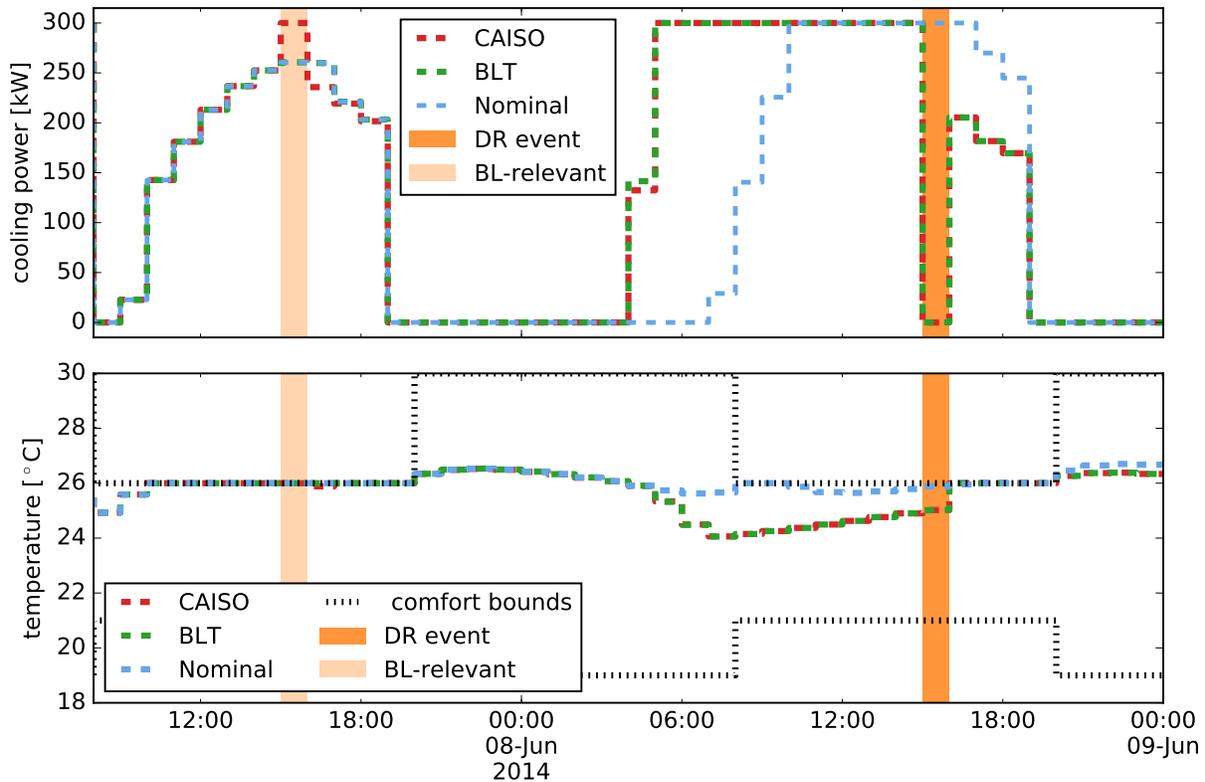


Figure 2.8: Baseline manipulation behavior in the HVAC model

further it is from the temperature the building would be without actuation.⁷²

Double Payment

The other economic distortion we consider is “double payment.” Under the compensation scheme mandated by FERC Order 745 (FERC, 2011b), providers of Demand Response are to be paid for reductions from their historical baseline at the LMP. But by reducing consumption, the DR participant also avoids paying the retail price p_t^R . The retail price can be decomposed into a component that reflects the average cost of energy procurement, plus a markup intended to recoup the retailer’s additional costs, particularly their fixed costs. We can write this decomposition as $p_t^R = \mathbb{E}[p_t^W] + T\&D_t$, where the first term denotes the average wholesale price, and the second denotes fixed costs such as Transmission and Distribution.

⁷²If the only exogenous determinants of building temperature were outside temperature, this would be the simple result of Newton’s law of cooling, which states that the rate at which a body dissipates heat into its surroundings is proportional to the difference between the body’s temperature and that of its surroundings. In a model with higher order terms relating power draw and HVAC cooling output, overcooling might be wasteful also because it would encourage the building to run the HVAC at an inefficiently high level; but in our linear model, this is not an issue.

The result is that, in a DR hour, the avoided expenditure, or effective “net price” per MWh reduced, is $p_t^R + p_t^W$, which exceeds the efficient price by p_t^R . But in an average hour, the retail price exceeds the efficient price by only T&D_t.⁷³

Many economists have concluded on the basis of this and other arguments that this effective price is too high (Borlick et al., 2012; Chao and DePillis, 2013), and that if baseline-dependent DR is to exist, the DR payment should be the LMP *minus* some quantity, usually referred to as “G” (for “generation”), to correct the effective net price. However, there is a lack of clarity in the literature about what exactly this “G” should be. Some authors take it to be the retail rate (Chao and DePillis (2013); Borlick et al. (2012); Borlick (2010); Shanker (2010)), so that the effective price is p_t^W in DR hours. This is the efficient price for a single hour considered alone. But Hogan (2010) (also approvingly cited by Chao and DePillis (2013)) argues that “G” refers to “the imputed generation portion of retail rates.” If the imputed generation component $\mathbb{E}[p_t^W]$, then the LMP-minus-G payment is $p_t^W - \mathbb{E}[p_t^W]$, and the effective net price is T&D_t + p_t^W .

When we simulate the elimination of the double payment distortion by making LMP-minus-G payments in our study, we adopt a variant of this latter approach subtracting an imputed generation component of the retail price from the LMP payment. This seems sensible because it equalizes average markups across DR and non-DR hours (see Section 2.5.2), and DR does not necessarily abate fixed costs (although that point is arguable).

To derive the imputed generation component of a retail tariff, we subtract the “non-generation rate” for that tariff from the tariff itself. The non-generation rate is a surcharge paid to PG&E by customers of Community Choice Aggregation customers, to cover PG&E’s infrastructure costs (see Table 2.3, which includes data from Pacific Gas and Electric Company (a)). In our hypothetical Opt Flat tariff, we take the imputed generation component to be the average LMP itself, since the Opt Flat tariff represents the average generation rate.

	Tariff	Non-Gen	Imputed Gen	Average SMC
Summer peak	262	128	134	61
Summer part-peak	253	128	125	55
Summer off peak	225	128	97	46
Winter part-peak	175	97	77	55
Winter off peak	155	97	58	48

Table 2.3: Load-weighted average A-1 ToU prices, non-gen component, and NP-15 SMC, all in \$/MWh

While we consider this to be a reasonable and realistic way of decomposing PG&E’s tariffs into generation and non-generation components, this decomposition results in generation

⁷³This analysis disregards externality costs. To account for externality costs, the reader can subtract them from T&D_t in what follows. This is straightforward in our setting, because externality costs are essentially time-invariant in California, as natural gas is the marginal fuel in the majority of hours (see Appendix A.3.2 and Callaway et al. (2015)).

components that are much higher than the average LMP.⁷⁴ The resulting value of “G” is a sort of compromise, intermediate between the retail rate itself endorsed by Chao and DePillis (2013) on the high side, and the average wholesale price, which is perhaps the lowest value suggested by Hogan (2010), on the low side. It seems the fact that PG&E’s imputed generation component is so much higher than the average LMP is accounted for by CPUC mandates to procure various expensive renewable resources, such as wind, solar, and biogas energy, using out-of-market feed-in tariffs.⁷⁵

2.6.2 Simulation Results: DR Distortions

We now present our findings on the effects of DR with and without the two economic distortions just introduced, for both the Quadratic Utility model and the HVAC model.

In Figures 2.9 and 2.10, we plot, for both consumption models, the welfare effects of the four DR variants: standard DR with both the manipulation and double payment distortions (“DR”); DR without double payment (“LMP-G”); DR in a baseline-taking equilibrium (“BLT”); and DR without double payment *and* in a baseline-taking equilibrium (“BLT & LMP-G”). The welfare effects are represented in terms of their changes (in \$) from the no-DR benchmark, as in the previous figures.

We explore the welfare impacts of these different demand response environments under two tariffs: A-1 ToU and Opt Flat. We choose the A-1 ToU tariff as representative of extant tariffs, as flat tariffs such as the A-1 are being phased out for commercial and industrial tariffs, and are not available to customers with smart meters (Pacific Gas and Electric Company, b). We choose the Opt Flat tariff because it is common in the DR and dynamic pricing literature to treat DR as superimposed on a second-best tariff, i.e. in which the retail price is the average social marginal cost (Chao and DePillis, 2013, Appendix A).

Note that, in our model, DR can never have a detrimental effect on the consumer surplus, because participation is voluntary and we assume that the consumer has perfect foresight.

Quadratic Utility Model

In general, the changes in surplus due to DR for the Quadratic Utility consumption model are on the order of several hundred dollars at most. This is very small as a fraction of total social surplus since calibration of inelastic utility functions results in very large social surpluses (ranging from \$8,855 for elasticity $E_d = -0.3$ to \$21,435 at elasticity $E_d = -0.1$). However, in absolute terms, or as a fraction of the annual electricity bill (calibrated to be equal to \$4,013 annually in the A-1 case), the changes are more substantial.

⁷⁴A reasonable alternative would be some kind of load-weighted average LMP. Since PG&E publishes representative consumption data for each customer class (Pacific Gas and Electric Company, 2016a), a separate load-weighting measure could be used for each tariff.

⁷⁵See Senate Bill 1122 (California State Senate, 2012), as well as Pacific Gas and Electric Company (2016d).

In Figure 2.9, we plot the results for a demand elasticity of $E_d = -0.1$.⁷⁶ Since the effects of DR are almost negligible without battery storage, this lack of dependence on demand elasticity suggests that the economically significant effects of DR are a result of battery storage arbitrage, and are mostly financial from the consumer’s perspective, without much effect on end-use consumption quantities.

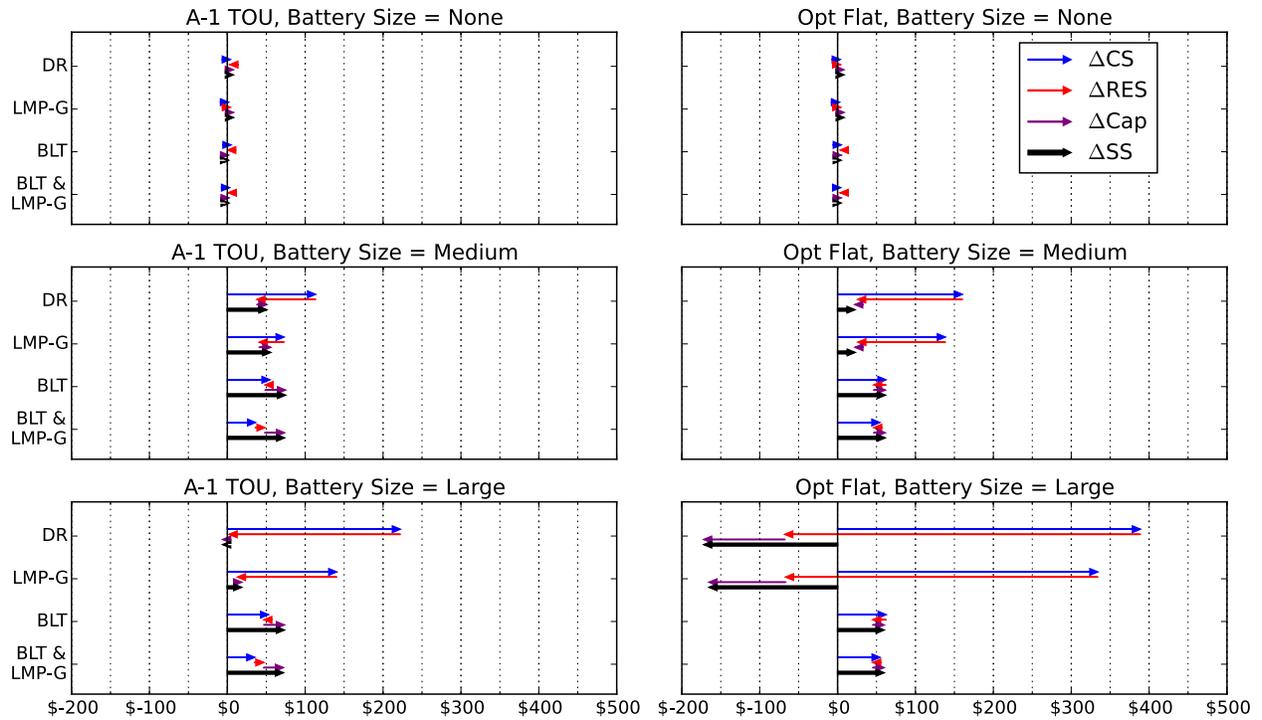


Figure 2.9: Quadratic Utility social surplus changes, DR and distortions, $E_d = -0.1$

DR has an appreciable effect on the surpluses whenever there is battery storage. In these cases, manipulation of the DR baseline has a comparatively large and detrimental effect on social surplus. That is, “BLT” always has a higher surplus than “DR,” and “BLT & LMP-G” always has a higher social surplus than “LMP-G.” Whenever DR with baseline manipulation has a non-negligible effect, the consumer benefits are much greater than for the corresponding BLT case, but the aggregate social benefits are smaller, which implies a financial transfer from the retailer (really, the rest of society) to the consumer, which is sometimes quite large.⁷⁷ In baseline-taking equilibria, the retailer surplus increases under

⁷⁶The plots for elasticities in the range from -0.05 to -0.2 are very similar for both tariffs, so we only include one. In the tables in Appendix A.5 we also present results for $E_d = -0.3$, but this case is more extreme, and perhaps not very realistic.

⁷⁷Ignoring capacity costs, this transfer is the length of overlapping blue and red arrows. Accounting for capacity costs, one would replace the red arrows with the sum of the red and purple arrows. Note that for any tariff, the transfer between the consumer and the retailer could be changed to any desired values by changing the fixed charges, i.e., the meter charges.

the A-1 ToU tariff, and is unaffected under the Opt Flat tariff. The retailer benefits are entirely due to reduction in capacity costs (purple arrows).

Without a battery, DR is beneficial to society, but the scale is negligible: at most \$15 annually (a little over \$1 a month) under the very high elasticity $E_d = -0.2$, and at most \$3 annually under the more plausible elasticity of $E_d = -0.05$. The largest effect of DR without a battery (and the only one greater than \$15 annually) is a reduction in generation cost of \$36 annually, when $E_d = -0.3$.

With a medium battery, all DR variants increase social surplus, with an increase between \$50 and \$75 in the A-1 ToU tariff, and between \$19 and \$60 in the Opt Flat tariff. The lower ends of these ranges come from the baseline manipulation cases, and the upper ones are from baseline-taking equilibria. Eliminating the double payment distortion (“LMP-G” vs. “DR,” and “BLT & LMP-G” vs. “BLT”) has a small positive effect (sometimes none) on total surplus, and has a more substantial effect of redistributing part of the change in social surplus from the consumer to the retailer (reducing the size of the transfer from the retailer to the consumer in the baseline manipulation case, and magnifying the retailer’s revenue increase under baseline-taking). In the basic DR setting (with both manipulation and double payment), DR increases the social surplus by \$20 annually, but involves a transfer from the retailer to the consumer of \$60 to \$65 in the A-1 ToU and \$140 in the Opt Flat tariffs, annually. In the real world, much of such a transfer would likely involve cross-subsidization from consumers who do not participate in DR. But *without* baseline manipulation, both the retailer and the consumer see improvements from DR, which is what one would hope for from a DR program.

With a large battery, the effects of DR are quite large and mixed. Strikingly, increasing the battery size from Medium to Large has a negligible effect on total surplus in baseline-taking equilibria, and has a large *negative* effect when manipulation is possible (comparing results for the Large vs Medium battery). With both manipulation and double payment, DR has a small negative effect on social surplus in the A-1 ToU tariff (\$10 annually), and a larger negative effect in the Opt Flat tariff (\$190 annually). These totals include large transfers from the retailer to the consumer: about \$220 and \$380 respectively. Without manipulation, society gains \$60 to \$70 annually from DR, and the retailers revenues either increase or are unaffected, as compared to no DR.

These results compare DR tariffs against the A-1 ToU benchmark, which changes slightly with battery size. So we also summarize the effects of increasing the battery size on social surplus in DR tariffs: With Demand Response, increasing battery size can increase social welfare by between \$23 and \$36 annually in the A-1 tariff, and \$44-\$48 in the A-1 ToU (the high end of the ranges are from lower elasticities). The private gains are smaller. With baseline manipulation, a medium battery realizes a value to society of \$14-\$26, but with a large battery, the surplus *decreases* by \$21-\$26 annually. In conjunction with our discussion above, this suggests that under our modeling assumptions RTP tariffs may make battery investments worthwhile, but DR programs do not justify them, and may even make the operation of a battery economically destructive.

To sum up, in the Quadratic Utility model, the effects of DR on total surplus are highly

dependent on the base tariff (A-1 ToU vs. Opt Flat), the size of the battery, and the type of economic distortions allowed for, but are not dependent on elasticity. Baseline manipulation is always worse for society than baseline-taking behavior, and can in fact make the net social effects of DR deleterious. DR with baseline manipulation results in large transfers toward the consumer whenever there is any effect on efficiency. Double payment has smaller efficiency effects, mostly redistributing changes in surplus from the retailer to the consumer.⁷⁸ A large battery as compared to a medium battery is good for the consumer, but neutral for society in baseline-taking equilibrium, and bad for society when baseline manipulation is possible.

As a side-note, the non-monotonicity of the effect of the battery size on social surplus reflects the complexity of tariff incentive effects when consumers can intertemporally substitute, an issue we explored earlier in Section 2.5.2, particularly Figures 2.2 and 2.3, with respect to the A-6 ToU tariff.

HVAC Model

Figure 2.10 shows the changes in surpluses for the HVAC model, from the no-DR benchmarks for the A-1 and Opt Flat tariffs respectively. For the A-1 tariff, the benchmark social surplus (negative generation cost minus capacity cost with no DR) is -\$68,100, and the individual expenditure is \$146,795. For the Opt Flat tariff, the social surplus benchmark is nearly same as that of the A-1 tariff (-\$68,097), and the individual expenditure is \$39,261.⁷⁹

Under the A-1 ToU tariff, the effect of Demand Response is always beneficial for society in aggregate. It is also beneficial in every component of social surplus, except in the case with both DR distortions, in which case the retailer is negatively impacted in terms of energy costs, but in turn sees a larger benefit in terms of reduced capacity costs. The pattern of results stemming from distortions is hard to explain: eliminating baseline manipulation is good for society in the presence of the double payment distortion, but it is deleterious for society under LMP-G compensation. Similarly, eliminating double payment is beneficial under baseline manipulation, but it is harmful in baseline-taking equilibria. And the “undistorted” outcome with baseline-taking and LMP-G, is the worst of all four combinations of distortions.

In the Opt Flat tariff, the effects are much more concordant with economic intuition. This might be expected, since the Opt Flat tariff is the standard setting under which DR is studied, and we do not have to worry about the interaction of the DR incentives with high markups and potentially distorted time-of-use price ratios. Demand Response is deleterious for society with baseline manipulation, but beneficial in baseline-taking equilibrium. What is striking is the extremely high transfer from the retailer to the consumer under baseline

⁷⁸The notable exception to this is in the A-1 ToU tariff, with the large battery and baseline manipulation, where eliminating double payment creates a \$25 annual gain for society, transforming effects that were slightly or negligibly deleterious into effects that are noticeably positive.

⁷⁹Since the A-1 tariff is also flat, the only difference in incentives between the two tariffs is the incentive to substitute between summer and winter seasons in the A-1 tariff, which is only possible for a few hours annually. The Opt Flat individual expenditure is lower than the social cost of generation because the Opt Flat tariff does not reflect capacity costs.

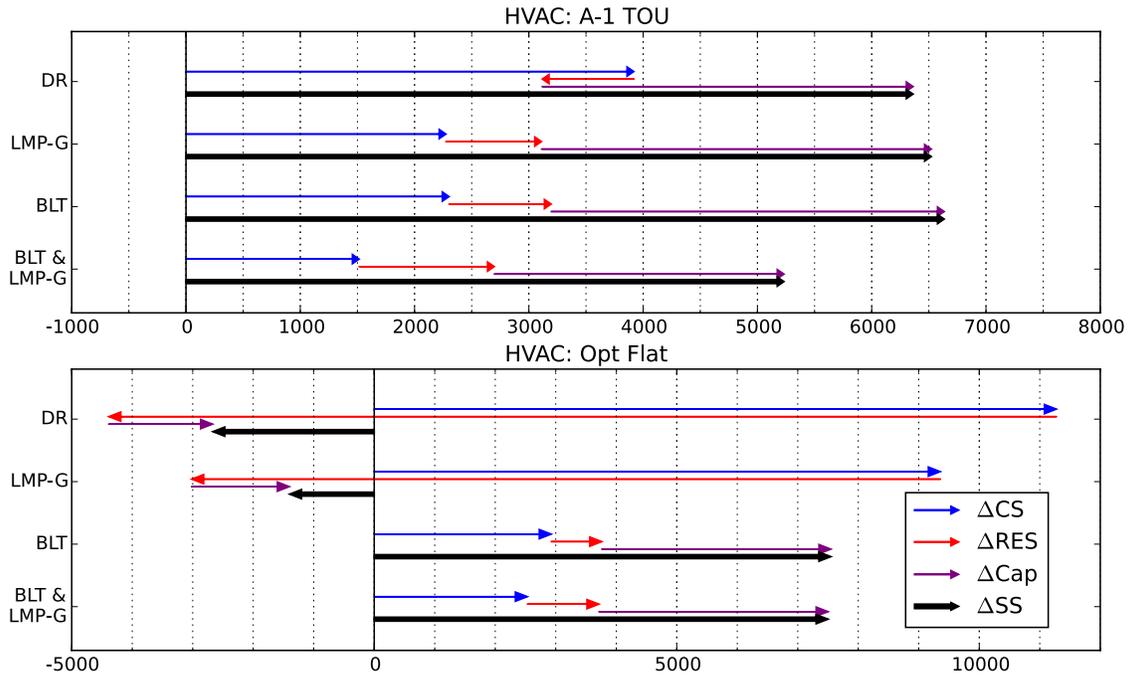


Figure 2.10: HVAC social surplus changes, DR and distortions (benchmark social cost of generation and capacity: \$68,100)

manipulation (note though that reduced capacity costs still have a significant positive effect on retailer surplus). Eliminating the double payment distortion is beneficial under baseline manipulation. In baseline-taking equilibrium, it has a small effect on total welfare, which is almost entirely a redistribution of revenue from the consumer to the retailer. This suggests that double payment does not have a large effect on consumption behavior in baseline-taking equilibrium.

2.7 Discussion

Throughout this chapter, we have observed that, not all too surprisingly, the efficiency effects of different tariffs depend strongly on the underlying consumption model. Moreover, they are subject to some rather unintuitive interactions of various aspects of the tariffs (in particular PDP and DR) that are hard to explain, yet alone to foresee. This provides an argument for policymakers to design simple tariffs whose efficiency properties can be easily understood, rather than trying to address specific shortcomings of tariffs by complicated designs.

The main limitation of our approach is that it considers data – LMPs, periods eligible for DR, and weather – to be known to the consumer in advance. This perfect knowledge assumption is clearly restrictive, and we would expect consumer behavior in DR tariffs to be somewhat different in a more realistic setting with limited information about future wholesale

market prices and DR events.⁸⁰ A central concern of ours are the efficiency effects of baseline manipulation. Intuitively, we would expect that perfect foresight gives the consumer the greatest ability to artificially inflate her baseline, and hence expect that our estimates of the social cost of baseline manipulation over-estimate the true cost. We view the current investigation as a first step in this direction.

Ideally, we would be able to extend our formulation to such more realistic cases by adopting methods such as approximate dynamic programming (Powell, 2011), stochastic multi-stage programming (Defourny et al., 2011), or stochastic MPC (Mesbah, 2016). However, if the goal is to retain the ability to treat complex consumption models and to formulate the DR baseline endogenously as part of the optimization problem, the extension to such methods poses very significant challenges, at least if the resulting algorithm should be of any reasonable computational complexity.⁸¹ The “perfect information relaxation” that we examine in the present work would be an important first step in benchmarking any such approximately optimal policies under more realistic information structures (Brown et al., 2010).

We hope that our open source software package `pyDR` (Balandat et al., 2016a) provides other researchers with a useful tool to study welfare effects of dynamic electricity pricing under a broad range of consumption models and tariffs.

⁸⁰For standard tariffs without DR (including ToU and PDP variants), information about LMPs is irrelevant to the consumer except for their correlation with CPP events.

⁸¹Our baseline-taking equilibrium concept and associated fixed-point algorithm translate to the stochastic multistage programming setting in a relatively straightforward way. However, the combination of hourly consumption decisions and the long-range dependence of the DR baseline on consumption decisions would result in prohibitively large scenario trees, with or without baseline manipulation. Perhaps heuristic methods could be applied to generate tractable scenario trees with a satisfactory level of modeling fidelity.

Chapter 3

Contract Design for Provision of Frequency Regulation Capacity via Aggregation of Commercial Buildings

The dynamic retail tariffs we investigated in the previous chapter aim to help load serving entities to balance supply and demand on a relatively slow time-scale, where imbalances are typically resolved in *energy* markets. These markets run on different time-scales (day-ahead, hour-ahead, real-time) in order to benefit from the increasingly accurate demand and supply forecasts as lead time decreases. At some point, however, running such an energy spot market becomes impractical. Instead, the grid operator procures generation *capacity*, so-called Ancillary Services (AS), that can be dispatched on much faster time-scales. The highest quality AS is frequency regulation (up and down), over which the operator has near real-time control.

Traditionally, these frequency control reserves have been provided by conventional generators, primarily hydro and natural gas fired power plants. Due to the increasing penetration of renewables, more and more generation is becoming uncertain and intermittent, which results in an increasing need for frequency regulation reserves (Halamay et al., 2011; Makarov et al., 2009). Instead of providing these capacity reserves solely using conventional generators, the mitigation of frequency deviations can also be supported by demand-side resources. Compared to rather crude approaches such as load shedding during DR periods, the performance requirements for a demand side resource to serve as a AS resource are much stricter.

As we have argued in the previous chapter, buildings have temporal flexibility in parts of their power consumption due to their inherent thermal storage capacity. In this chapter we consider the problem of grouping several mid-to-large size commercial buildings together by an *aggregator*, which uses the buildings' inter-temporal consumption flexibility in order to provide frequency regulation capacity in the AS market. As Hao et al. (2013) point out, commercial buildings seem particularly well suited for this task due to their comparatively large power consumption, advanced building management systems, and their HVAC systems' potential to follow a high-frequency regulation signal.

In order to avoid load shedding or violating user comfort one generally needs to plan ahead, i.e., back off from comfort and actuator constraints to increase flexibility in power consumption. This can for example be achieved by using Model Predictive Control (MPC), see (Henze et al., 2005; Oldewurtel, 2011; Ma et al., 2012) and references therein. The idea of using MPC to have commercial buildings provide frequency regulation has been proposed earlier by Vrettos et al. (2014) and Maasoumy et al. (2014). Like most of the initial work in this domain, Vrettos et al. (2014) assumes that the objectives of aggregator and buildings are perfectly aligned. However, unless aggregator and buildings are under the same ownership, this will usually not be the case. Indeed, having buildings provide regulation capacity typically requires less aggressive control strategies and hence results in a more costly operation. An aggregator therefore will need to compensate the building for the service it provides. Hence, while there may be significant potential for using buildings to provide frequency regulation, a large part of this potential can only be tapped by providing the correct financial incentives. While Maasoumy et al. (2014) take some steps in investigating this problem, they do so in a rather ad-hoc way, and do not provide a structured approach to determining incentives that are optimal in some sense.

In this chapter, we formulate the aggregator’s optimal contract design problem. For simplicity, we limit ourselves to the case without information asymmetries between buildings and aggregator. We propose a contract structure in which a building agrees to adjust its power consumption according to a regulation signal within time-varying capacity bands around a nominal power schedule. In exchange, the building receives a monetary reward from the aggregator. The financial risk for the building due to following the regulation signal can be absorbed by the aggregator by charging the building for its energy consumption according to its nominal schedule. The resulting problem is a bilevel optimization problem, in which the aggregator jointly optimizes over nominal schedule, capacities and rewards in order to maximize its expected profit, subject to the buildings’ individual rationality constraints. Using techniques from integer programming we cast this bilevel problem as an equivalent mixed-integer program. Moreover, we show that if buildings do not impose externalities on the aggregator, determining the resulting “first-best contract” (Laffont and Martimort, 2002) reduces to solving a Linear Program.

3.1 Preliminaries

Wholesale Markets for Ancillary Services

We consider Ancillary Services (AS) markets in deregulated wholesale electricity spot markets. Specifically, we focus on the frequency regulation capacity market run by CAISO.¹ The market rules specify bidding, clearing and settlement processes. Certified market participants, called Scheduling Coordinators (SC), bid into the AS market and commonly rep-

¹While details of the implementation vary, the basic market setup is quite similar among different ISOs in North America, so our formulation can easily be extended to other markets.

resent multiple resources (generators or participating loads). For simplicity we will in the following speak of resources bidding in the market, with the understanding that is in fact the associated SC who places the bids.

Like energy markets, frequency regulation capacity markets run on different timescales, namely the Day-Ahead Market (DAM) and the Real-Time Market (RTM). Each market operates as follows: After submitting its bids (consisting of maximum capacities and prices for both regulation up and regulation down), a resource gets awarded a regulation capacity schedule based on a uniform price auction conducted by the ISO. The market clearing process is a large-scale optimization problem, in which capacity prices are determined from the dual variables (commonly referred to as “shadow prices”). At run-time, the resource then receives a Load Frequency Control (LFC) signal ω from the ISO, specifying how much it should deviate from its nominal power schedule (which is determined in the energy market). The LFC signal is constrained to the awarded regulation capacities. The ISO has direct access to measurements of the power output of the resource, and hence can verify whether it has fulfilled its obligations. If that is the case, then the resource receives a capacity payment according to the market clearing price.²

The Role of the Aggregator

We are interested in using the temporal flexibility that buildings have in terms of their HVAC power consumption in order to offer regulation capacity to the grid. The role of the aggregator is to provide an interface between the electricity spot market on one side and the individual buildings on the other, as illustrated in Figure 3.1. There are many reasons why individual buildings cannot easily participate as resources in the wholesale market themselves. Firstly, some markets have minimum bid size requirements (Cutter et al., 2009) that often exceed the power consumption of single buildings. Moreover, participating in the market creates additional overhead — not only does it involve acquiring the necessary certifications from the ISO (California Independent System Operator Corporation, 2015a), but it also requires understanding market processes and developing bidding strategies, tasks that building managers are typically not familiar with.

An aggregator with expertise in these areas, on the other hand, can aggregate different buildings and bundle their individual flexibilities in such a way that the resulting product is suitable for sale in the spot market. In order for the individual buildings to be willing to participate in such a scheme, the aggregator must provide financial incentives. While the market interface is well defined and clearly specified, the aggregator has large freedom in choosing this incentive structure.³ Here we take a contract-theoretic point of view and

²Recently, ISOs have been ordered by the federal regulator to implement “pay for performance” compensation (FERC, 2011a), which rewards resources for performance (e.g. mileage payments or signal tracking accuracy) in addition to capacity. For simplicity we will restrict our attention to capacity payments.

³This situation is the same as in residential DR programs, in which DR providers can decide relatively freely how they want to dispatch and reward their customers.

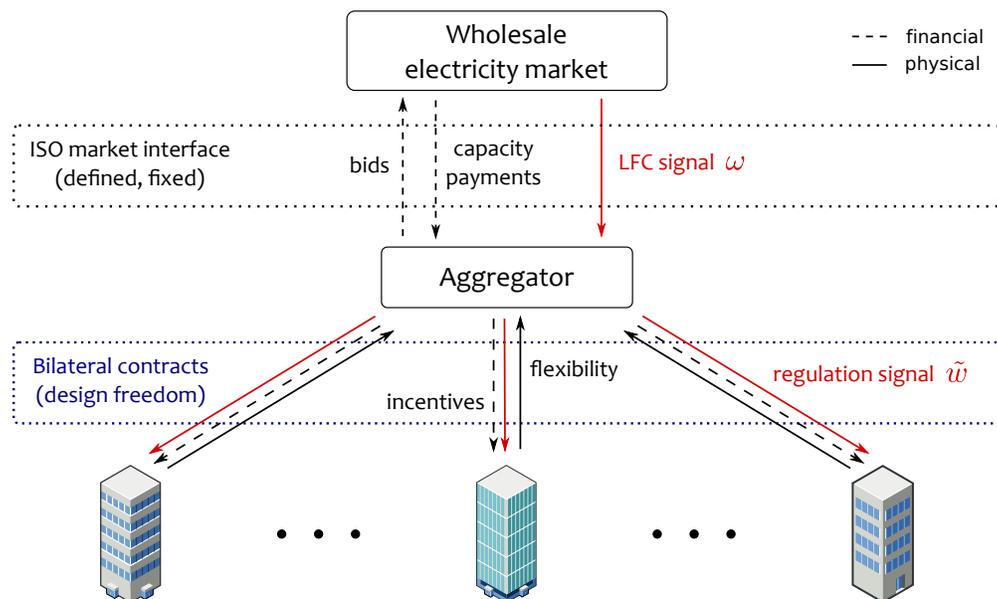


Figure 3.1: The aggregator as an interface between AS market and buildings

consider the problem of maximizing the aggregator’s profit over all contracts within a certain class. It should be noted that the resulting scheme is generally not an energy-saving solution.

The Basic Contract Structure

The basic structure of the contract we propose between aggregator and an individual building is the following: For a given scheduling horizon,⁴ the building agrees to adjust its HVAC power consumption with respect to a nominal power schedule, based on a regulation signal $\tilde{\omega}$ it receives from the aggregator that is guaranteed to lie within certain bounds. The building is charged for its energy consumption according to this nominal schedule, and the aggregator pays the building a one-time monetary reward R . Nominal schedule, capacity bounds and reward are all part of the contract specification.

Importantly, such a contract need not directly involve the building’s power provider, e.g., the local electric utility.⁵ In fact, the provider simply bills the building for its actual energy consumption under the regulation signal and need not even know of the contract.⁶ The difference between the realized cost and the hypothetical cost under the nominal schedule can then be settled ex-post between building and aggregator. Charging the building only for its nominal consumption is an important feature of our contract, as it does not expose the

⁴Since the aggregator is bidding in the spot market the scheduling horizon will not exceed 24 hours.

⁵The legality of retail customers interacting directly with an aggregator depends on the local legislation. For example, nonresidential customers in California can enter such contracts (California State Senate, 2011).

⁶Here we assume that the power provider is either oblivious or indifferent to the relationship between aggregator and customer. Without this assumption the question of incentives would become much more challenging in this context.

building to any financial risk associated with the randomness of the regulation signal. That is, such a contract is ex-post individually rational by design.

Addressing the Different Timescales of the Problem

An interesting characteristic of the multi-level problem we consider here is that it takes place on three different timescales, as illustrated in Figure 3.2. On the market level, bids,

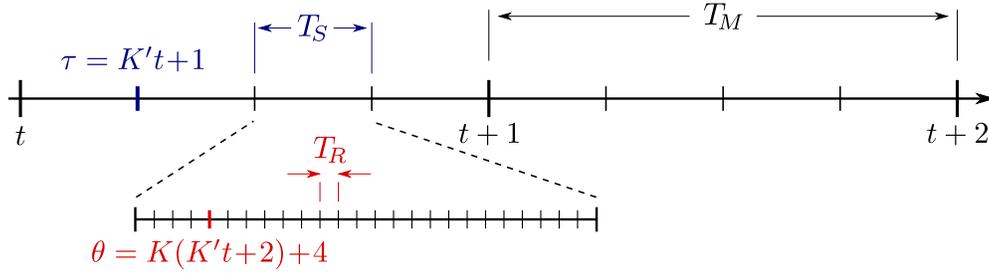


Figure 3.2: Different time scales of the aggregation problem

prices and awarded AS capacities are determined for market intervals of length T_M (typically $T_M = 1$ h for DAM and $T_M = 15$ min for RTM). On the other extreme, at the level of frequency control, the ISO sends a LFC signal ω that is piece-wise constant over intervals of length $T_R \ll T_M$ ($T_R = 4$ s in CAISO AS regions). Given that the time constants of building temperature dynamics are typically much larger than T_R , it is not necessary to schedule the nominal control on the timescale of the regulation signal. Moreover, the associated optimization problem would be computationally challenging, if not intractable. Therefore the control schedule for a building is determined for an intermediate interval length T_S which depends on the time-constants of the building, the frequency with which set points can be changed, and the possible computational challenges with the solving the resulting optimization problem. Commonly, T_S will lie in the range of 5-15 minutes. Here we make the following Assumption:

Assumption 3. *There exist constants $K, K' \in \mathbb{N}$ such that $T_S = KT_R$ and $T_M = K'T_S$, respectively.*

We consider a contract design problem over a scheduling horizon of N_M market periods (e.g., $N_M = 24$ for the Day-Ahead market) and let $N_S = K'N_M$, $N_R = KN_S$ and

$$\begin{aligned} \mathbb{T}_S(t) &= \{\tau \in N_S : tT_M \leq \tau T_S < (t+1)T_M\} \\ \mathbb{T}_R(\tau) &= \{\theta \in N_R : \tau T_S \leq \theta T_R < (\tau+1)T_S\} \end{aligned}$$

We also write $t_\tau = \mathbb{T}_S^{-1}(\tau)$ and $\tau_\theta = \mathbb{T}_R^{-1}(\theta)$.

Our contract will be based on a robust scheduling problems solved for the sampling time T_S . This implicitly assumes that the regulation signal \tilde{w} for a building, derived by the

aggregator from the ISO's LFC signal ω , is constant over intervals of length T_S . However, in reality the signal ω varies with frequency $1/T_R$. We thus need to make sure that the obtained schedule is robust also for the actual building model on the faster time-scale. Fortunately, this is the case for most building models, as we will show in the following section.

3.2 A Building's Scheduling Problem

The Building Model

Consider a building b whose HVAC dynamics for sampling time T_S are described by the discrete-time LTI system⁷

$$x_{\tau+1}^b = A^b x_{\tau}^b + B^b(u_{\tau}^b + \tilde{w}_{\tau}^b) + E^b v_{\tau}^b \quad (3.1a)$$

$$y_{\tau}^b = C^b x_{\tau}^b + D^b(u_{\tau}^b + \tilde{w}_{\tau}^b) + F^b v_{\tau}^b \quad (3.1b)$$

with known initial condition $x_0^b \in \mathbb{R}^{n_x}$. Here $u_{\tau}^b \in \mathbb{R}^{n_u}$ represents the control inputs to the HVAC system, $y_{\tau}^b \in \mathbb{R}^{n_y}$ is the system output and $v_{\tau}^b \in \mathbb{R}^{n_v}$ is the effect of weather and occupancy in time interval t . The signal $\tilde{w}_{\tau}^b \in \mathbb{R}^{n_u}$ is a proxy for the regulation signal w_{θ}^b received by the aggregator. We write $\mathbf{u}^b = (u_0^b, \dots, u_{N_S}^b)$ and will use analogous notation for other variables throughout this chapter.

Under some mild conditions, which are satisfied by most building models, one can show that robustness of (3.1) with respect to $\tilde{w}_{\tau}^b \in [-r_{\tau}^{b\downarrow}, r_{\tau}^{b\uparrow}]$ implies robustness of the "true" fast-sampled system with respect to $w_{\theta}^b \in [-r_{\tau_{\theta}}^{b\downarrow}, r_{\tau_{\theta}}^{b\uparrow}]$. Specifically, suppose the model of a building is given by the continuous-time LTI system

$$\dot{x}(t) = A_c x(t) + B_c(u(t) + w(t)) \quad (3.2a)$$

$$y(t) = Cx(t) + D(u(t) + w(t)) \quad (3.2b)$$

Sampling this system with sampling times T_R and T_S using zero-order hold yields the discrete-time LTI systems

$$\tilde{x}_{\theta+1} = \tilde{A}\tilde{x}_{\theta} + \tilde{B}(\tilde{u}_{\theta} + \tilde{w}_{\theta}) \quad (3.3a)$$

$$\tilde{y}_{\theta} = C\tilde{x}_{\theta} + D(\tilde{u}_{\theta} + \tilde{w}_{\theta}) \quad (3.3b)$$

$$x_{t+1} = Ax_t + B(u_t + w_t) \quad (3.4a)$$

$$y_t = Cx_t + D(u_t + w_t) \quad (3.4b)$$

For the remainder of this chapter, we will adhere to the following notational conventions: For a matrix M let $(M)_{i\cdot}$ and $(M)_{\cdot j}$ denote the i -th row and j -th column of M , respectively. Further, define $M^+ := \max(M, 0)$ and $M^- := \max(-M, 0)$ (element-wise). By \times we denote element-wise vector multiplication, and by $\pi_A(x)$ the Euclidean projection of x onto A .

⁷As in the previous chapter, the assumption of time-invariance is only made for notational convenience; the extension to linear time-varying systems is immediate.

Proposition 1. *Let $\tilde{u}_{Kt+j} = u_t$ for all $t = 0, \dots, N_R$ and $j = 0, \dots, K-1$. Then, if for all $i = 1, \dots, n_y$ and $k = 1, \dots, n_u$ the sign of $\eta_{ijk} := (C)_i \tilde{A}^{K-1-j} (\tilde{B})_{:k}$ is the same for all $j = 0, \dots, K-1$, the robust schedule u determined for the system (3.4) is robust also for system (3.3) at the corresponding time instances.*

Proof. See Appendix C.1. □

Intuitively, the conditions of Proposition 1 mean that the qualitative effect of a particular (scalar) input signal on a particular (scalar) output signal is consistent over time. This is commonly the case for temperatures dynamics of a building: For instance, decreasing the supply air temperature should never lead to an increase of a room temperature. The linear and linearized building models that we have worked with all satisfy the conditions of Proposition 1.

We make the following additional important assumption:

Assumption 4. *For each building, the power consumption of each actuator is linear in the associated control input.*

Under Assumption 4, building b 's overall HVAC energy consumption Q_t^b in interval τ can be expressed as $Q_\tau^b = (q^b)^\top u_\tau^b = \sum_{i=1}^{n_y} q_i^b u_{\tau,i}^b$, where $q_i^b \geq 0$ is the per-unit energy consumption of actuator i over an interval of length T_S .

As weather and occupancy effects are usually uncertain, we need to account for this uncertainty in the scheduling problem. There are different ways of accomplishing this; here we choose to formulate chance constraints that require the building's comfort constraints to be satisfied with given levels of probability. The weather and occupancy effect v_τ^b can be split into a known forecast \hat{v}_τ^b and the forecast error $\tilde{v}_\tau^b := v_\tau^b - \hat{v}_\tau^b$. We assume that – based on the analysis of historical data – a linear uncertainty model of the form

$$\tilde{v}_{\tau+1}^b = G^b \tilde{v}_\tau^b + H^b e_\tau^b \quad (3.5)$$

has been identified, where $e_\tau^b \stackrel{d}{=} \mathcal{N}(0, I)$ are i.i.d. Gaussian and $\tilde{v}_0^b \stackrel{d}{=} \mathcal{N}(0, \Sigma_0^b)$ with e_τ^b and \tilde{v}_0^b independent for all τ .

Some elementary algebra allows us to write the system output y_τ^b as

$$y_\tau^b = \hat{y}_\tau^b(\mathbf{u}^b) + \sum_{\iota < \tau} W_{\tau-1-\iota}^b \tilde{w}_\iota^b + D_\tau^b \tau_\tau^b + \mathcal{G}_\tau^b \tilde{v}_0^b + \sum_{\iota < \tau} \mathcal{H}_{\tau,\iota}^b e_\iota^b \quad (3.6)$$

where

$$\hat{y}_\tau^b(\mathbf{u}^b) = C^b D^b u_\tau^b + C^b F^b \hat{v}_\tau^b + C^b (A^b)^\tau x_0^b + \sum_{\iota < \tau} C^b (A^b)^{\tau-1-\iota} (B^b u_\iota^b + E^b \tilde{v}_\iota^b) \quad (3.7)$$

is the nominal system output, $W_\tau^b = C^b (A^b)^\tau B^b$, and

$$\mathcal{G}_\tau^b = \sum_{\iota < \tau} C^b (A^b)^{\tau-1-\iota} E^b (G^b)^\iota + F^b (G^b)^\tau \quad (3.8)$$

$$\mathcal{H}_{\tau,\iota}^b = \sum_{\theta < \tau} C^b (A^b)^{\tau-1-\theta} E^b (G^b)^{\theta-1-\iota} H^b + F^b (G^b)^{\tau-1-\iota} H^b \quad (3.9)$$

The Robust Optimal Scheduling Problem

Denote by $\mathbf{r}^{\text{b}\uparrow} = (r_0^{\text{b}\uparrow}, \dots, r_{N_S}^{\text{b}\uparrow})$ and $\mathbf{r}^{\text{b}\downarrow} = (r_0^{\text{b}\downarrow}, \dots, r_{N_S}^{\text{b}\downarrow})$ the sequences of non-negative vectors of “regulation up” and “regulation down” capacities for building b , respectively.⁸ Suppose that at each time τ , the building receives a regulation signal $\tilde{w}_\tau^{\text{b}} \in [-r_\tau^{\text{b}\downarrow}, r_\tau^{\text{b}\uparrow}]$. At this point we do not assume anything about the signal other than these worst-case bounds. For the purpose of this section suppose that the capacity vectors $\mathbf{r}^{\text{b}\uparrow}$ and $\mathbf{r}^{\text{b}\downarrow}$ are fixed and known. It may be infeasible for a building to tolerate all possible sequences of the regulation signal within $[-\mathbf{r}^{\text{b}\downarrow}, \mathbf{r}^{\text{b}\uparrow}]$ without jeopardizing occupant comfort. Hence we allow the building to shrink the interval within which it is required to robustly follow the regulation signal to $[-r_\tau^{\text{b}\uparrow} + s_\tau^{\text{b}\uparrow}, r_\tau^{\text{b}\downarrow} - s_\tau^{\text{b}\downarrow}]$, where $s_\tau^{\text{b}\uparrow}$ and $s_\tau^{\text{b}\downarrow}$ with

$$0 \leq s_\tau^{\text{b}\downarrow} \leq r_\tau^{\text{b}\downarrow}, \quad 0 \leq s_\tau^{\text{b}\uparrow} \leq r_\tau^{\text{b}\uparrow} \quad \forall \tau \leq N_S \quad (3.10)$$

are the *slacks*. We define the *effective regulation signal*

$$\check{w}_\tau = \pi_{[-r_\tau^{\text{b}\uparrow} + s_\tau^{\text{b}\uparrow}, r_\tau^{\text{b}\downarrow} - s_\tau^{\text{b}\downarrow]}(\tilde{w}_\tau^{\text{b}}) \quad (3.11)$$

In exchange for the reduced uncertainty resulting from this, the building pays a penalty $\mathcal{P}^{\text{b}}(\mathbf{s}^{\text{b}})$ to the aggregator, where $\mathbf{s}^{\text{b}} = (\mathbf{s}^{\text{b}\uparrow}, \mathbf{s}^{\text{b}\downarrow})$. Given \mathbf{r}^{b} and \mathbf{s}^{b} , the building seeks the nominal schedule of minimum expected cost that ensures constraint satisfaction for all possible sequences of $\check{\mathbf{w}}^{\text{b}}$.

Input Constraints The actual control inputs $u_\tau^{\text{b}} + \check{w}_\tau^{\text{b}}$ at time τ are restricted by physical limits on the actuators, which translate into the following constraints⁹:

$$\underline{u}_\tau^{\text{b}} \leq u_\tau^{\text{b}} + \check{w}_\tau^{\text{b}} \leq \bar{u}_\tau^{\text{b}} \quad \forall \tau \leq N_S \quad (3.12)$$

Since the regulation signal w_τ^{b} is unknown at scheduling time, the above constraints have to hold for any regulation signal $\check{w}_\tau^{\text{b}} \in [-r_\tau^{\text{b}\uparrow} + s_\tau^{\text{b}\uparrow}, r_\tau^{\text{b}\downarrow} - s_\tau^{\text{b}\downarrow}]$. This can be achieved by imposing the following robustified version of (3.12):

$$\underline{u}_\tau^{\text{b}} + r_\tau^{\text{b}\uparrow} - s_\tau^{\text{b}\uparrow} \leq u_\tau^{\text{b}} \leq \bar{u}_\tau^{\text{b}} - r_\tau^{\text{b}\downarrow} + s_\tau^{\text{b}\downarrow} \quad \forall \tau \leq N_S \quad (3.13)$$

Importantly, (3.13) deals with the uncertainty introduced by $\check{\mathbf{w}}^{\text{b}}$ in a robust fashion, i.e., it accounts for the worst case, so the building need not make any assumptions on the statistical properties of $\check{\mathbf{w}}^{\text{b}}$. This allows the building to make firm commitments on what signals (regulation signal minus slack) it can tolerate.

⁸Our convention is to use the classic point of view of generation markets in which “regulation up” means that a generator increases its power output and “regulation down” means that it decreases it. Thus, from the point of view of a load, “regulation up” means a *decrease* while “regulation down” means an *increase* in power consumption with respect to the nominal schedule.

⁹For simplicity we assume box constraints on the inputs, the extension to polyhedral constraints is straightforward.

Output Constraints The building aims to ensure that, for all possible sequences of the signal $\tilde{\mathbf{w}}^b$, the output \mathbf{y}^b satisfies the building's comfort constraints. Since small violations of comfort constraints are usually non-critical, we assume the building formulates chance constraints, which ensure that the constraints at any point in time are satisfied with a given level of probability. Specifically, we suppose that, for all possible sequences of $\tilde{\mathbf{w}}^b$, the comfort constraints $y_{\tau,i}^b \geq \underline{y}_{\tau,i}^b$ and $y_{\tau,i}^b \leq \bar{y}_{\tau,i}^b$ are to be satisfied with a probability of at least $1 - \underline{\alpha}_{\tau,i}^b$ and $1 - \bar{\alpha}_{\tau,i}^b$, respectively.¹⁰

Let $\mathcal{W}_\tau^b = \{(\tilde{w}_0, \dots, \tilde{w}_\tau) : -r_\iota^{\text{b}\uparrow} + s_\iota^{\text{b}\uparrow} \leq \tilde{w}_\iota \leq r_\iota^{\text{b}\downarrow} - s_\iota^{\text{b}\downarrow}, \forall \iota \leq \tau\}$. The constraints on the output then read

$$\forall \tau, \forall i, \forall \tilde{w} \in \mathcal{W}_\tau^b, \quad \begin{cases} \mathbb{P}(y_{\tau,i}^b \geq \underline{y}_{\tau,i}^b) \geq 1 - \underline{\alpha}_{\tau,i}^b \\ \mathbb{P}(y_{\tau,i}^b \leq \bar{y}_{\tau,i}^b) \geq 1 - \bar{\alpha}_{\tau,i}^b \end{cases} \quad (3.14)$$

Using techniques from robust optimization and chance-constrained programming, the above chance constraints for $\tau = 0, \dots, N_S$ and $1 \leq i \leq n_y^b$ can be expressed as

$$\hat{y}_{\tau,i}^b(\mathbf{u}^b) \geq \underline{y}_{\tau,i}^b + \sum_{\iota < \tau} \left[(W_{\tau-1-\iota}^b)_{i,:}^+ (r_\iota^{\text{b}\uparrow} - s_\iota^{\text{b}\uparrow}) + (W_{\tau-1-\iota}^b)_{i,:}^- (r_\iota^{\text{b}\downarrow} - s_\iota^{\text{b}\downarrow}) \right] \\ + (D^b)_{i,:}^+ (r_\tau^{\text{b}\uparrow} - s_\tau^{\text{b}\uparrow}) + (D^b)_{i,:}^- (r_\tau^{\text{b}\downarrow} - s_\tau^{\text{b}\downarrow}) + \sigma_{\tau,i}^b \Phi^{-1}(1 - \underline{\alpha}_{\tau,i}^b) \quad (3.15a)$$

$$\hat{y}_{\tau,i}^b(\mathbf{u}^b) \leq \bar{y}_{\tau,i}^b - \sum_{\iota < \tau} \left[(W_{\tau-1-\iota}^b)_{i,:}^+ (r_\iota^{\text{b}\downarrow} - s_\iota^{\text{b}\downarrow}) + (W_{\tau-1-\iota}^b)_{i,:}^- (r_\iota^{\text{b}\uparrow} - s_\iota^{\text{b}\uparrow}) \right] \\ - (D^b)_{i,:}^+ (r_\tau^{\text{b}\downarrow} - s_\tau^{\text{b}\downarrow}) - (D^b)_{i,:}^- (r_\tau^{\text{b}\uparrow} - s_\tau^{\text{b}\uparrow}) - \sigma_{\tau,i}^b \Phi^{-1}(1 - \bar{\alpha}_{\tau,i}^b) \quad (3.15b)$$

where $(\sigma_{\tau,i}^b)^2 = (\mathcal{G}_\tau^b)_{i,:} \Sigma_0^b (\mathcal{G}_\tau^b)_{i,:}^\top + \sum_{\iota < \tau} (\mathcal{H}_\tau^b)_{i,:} (\mathcal{H}_\tau^b)_{i,:}^\top$ and Φ is the cdf of a standard normal random variable. Note that, importantly, both (3.15a) and (3.15b) are linear in the control \mathbf{u}^b , the capacities $\mathbf{r}^{\text{b}\uparrow}$ and $\mathbf{r}^{\text{b}\downarrow}$, and the slack variables $\mathbf{r}^{\text{b}\uparrow}$ and $\mathbf{r}^{\text{b}\downarrow}$, respectively, while all other parameters are known constants.

The Robust Scheduling Problem for Fixed Capacities

Let $\mathbf{c}^b = (c_0^b, \dots, c_{N_S}^b)$ be the vector of electricity prices for the building over the scheduling horizon.¹¹ As discussed earlier in this section, in the contract we propose a building pays only for its power consumption according to nominal schedule. That is, the energy cost (excluding \mathcal{P}^b) of building b is

$$\varphi_{\text{nom}}^b(\mathbf{u}^b) = \sum_{\tau=1}^{N_S} c_\tau^b (q^b)^\top u_\tau^b$$

In practice the building will still pay the power provider

$$\varphi_{\text{act}}^b(\mathbf{u}^b, \tilde{\mathbf{w}}^b) = \sum_{\tau=1}^{N_S} c_\tau^b (q^b)^\top u_\tau^b + \sum_{\tau=1}^{N_S} c_\tau^b \frac{(q^b)^\top}{K} \sum_{\theta \in \mathbb{T}_R(\tau)} \tilde{w}_\theta^b$$

¹⁰Here $\underline{\alpha}_{\tau,i}^b, \bar{\alpha}_{\tau,i}^b < 0.5$ are given parameters, and typically $\underline{\alpha}_{\tau,i}^b, \bar{\alpha}_{\tau,i}^b \ll 1$.

¹¹We assume c^b is known over the whole scheduling horizon, but may otherwise be arbitrary (i.e. c^b can describe both flat and Time-of-Use tariffs).

for its actual consumption, while the difference

$$\Delta\varphi^b(\tilde{\mathbf{w}}^b) = \sum_{\tau=1}^{N_S} c_\tau^b \frac{(q^b)^\top}{K} \sum_{\theta \in \mathbb{T}_R(\tau)} \tilde{w}_\theta^b$$

is settled ex post with the aggregator. Here the signal $\tilde{\mathbf{w}}^b$ is the effective regulation signal received with frequency $1/T_R$. The building's overall robust scheduling problem for given regulation capacity vectors $\mathbf{r}^{b\downarrow}$ and $\mathbf{r}^{b\uparrow}$ is therefore

$$\begin{aligned} (\mathbf{u}^{b*}, \mathbf{s}^{b*})(\mathbf{r}^{b\downarrow}, \mathbf{r}^{b\uparrow}) &= \arg \min_{\mathbf{u}^b, \mathbf{s}^b} \varphi_{\text{nom}}^b(\mathbf{u}^b) + \mathcal{P}^b(\mathbf{s}^b) \\ &\text{s.t.} \quad (3.10), (3.13), (3.15) \end{aligned} \quad (3.16)$$

We make the following assumption:

Assumption 5. *The building penalties \mathcal{P}^b are of the form*

$$\mathcal{P}^b(\mathbf{s}^b) = \sum_b \sum_{\tau=1}^{N_S} (p_l^\uparrow s_\tau^\uparrow + p_q^\uparrow (s_\tau^\uparrow)^2 + p_l^\downarrow s_\tau^\downarrow + p_q^\downarrow (s_\tau^\downarrow)^2)$$

where $p_l^\uparrow, p_l^\downarrow \geq 0$ and $p_q^\uparrow, p_q^\downarrow \geq 0$ are fixed coefficients.

Under Assumption 5, (3.16) is a Quadratic Program (QP).

3.3 The Optimal Contract Design Problem

We now formulate the aggregator's optimal contract design problem for a given collection \mathcal{B} of buildings. To this end we require some additional assumptions.

Assumption 6. *The aggregator is risk-neutral.*

Assumption 7. *The aggregator has no market power in the wholesale market for frequency regulation capacity.*

Assumption 8. *Buildings do not possess any private information. That is, for each $b \in \mathcal{B}$ the aggregator has full knowledge of the system model (3.1), constraints, energy consumption rates q^b and tariff \mathbf{c}^b . Moreover, it has access to measurements of the initial state x_0^b and the HVAC power consumption on the frequency $1/T_R$ of the LFC signal ω .*

Assumption 6 is standard and made primarily for mathematical convenience. An extension to a risk-averse aggregator, while desirable, appears quite challenging. Assumption 7 holds if the total capacity that can be provided by all buildings is small enough that the aggregator's bid does not significantly affect the overall supply curve. Then the aggregator acts as a price-taker and places its bid at true cost. Due to Assumptions 4 and 6 it suffices that the aggregator has access to an unbiased estimate of the expected capacity prices over the scheduling horizon. Assumption 8 is more restrictive. In practice, an aggregator will not know the exact building models, which causes information asymmetries. It would be interesting to analyze how much of a so-called information rent a building could earn in this case, but this question is outside the scope of this chapter.

Objective Function

Denote by $r_t^{M\uparrow}$ and $r_t^{M\downarrow}$ the regulation up and down capacities offered in the market in period t , respectively. Furthermore, let $\mathbf{r}^M = (r^{M\uparrow}, r^{M\downarrow})$ and, similarly, $R = \{R^b : b \in \mathcal{B}\}$, $\mathbf{r} = \{(\mathbf{r}^{b\downarrow}, \mathbf{r}^{b\uparrow}) : b \in \mathcal{B}\}$, and $\mathbf{s} = \{(\mathbf{s}^{b\downarrow}, \mathbf{s}^{b\uparrow}) : b \in \mathcal{B}\}$. The deviations $\epsilon_\theta^b := \check{w}_\theta^b - w_\theta^b$ of building b 's control from the dispatch signal $u_{\tau_\theta} + w_\theta^b$ are given by

$$\epsilon_\theta^b = (w_\theta^b + r_{\tau_\theta}^{b\uparrow} - s_{\tau_\theta}^{b\uparrow})^- - (w_\theta^b - r_{\tau_\theta}^{b\downarrow} + s_{\tau_\theta}^{b\downarrow})^+$$

Observe that $\epsilon_\theta^b = 0$ if the slacks $s_{\tau_\theta}^{b\uparrow}, s_{\tau_\theta}^{b\downarrow}$ are zero. With

$$w_\theta^M = \begin{cases} -\omega_\theta r_{t\tau_\theta}^{M\uparrow} & \text{if } \omega_\theta \geq 0 \\ -\omega_\theta r_{t\tau_\theta}^{M\downarrow} & \text{if } \omega_\theta < 0 \end{cases} \quad (3.17)$$

denoting the aggregator's regulation signal derived from the LFC signal $\omega_\theta \in [-1, 1]$, the deviation ϵ_θ^M of the buildings' aggregate consumption from the ISO dispatch signal is

$$\epsilon_\theta^M = \sum_b (q^b)^\top w_\theta^b - w_\theta^M + \sum_b (q^b)^\top \epsilon_\theta^b$$

Using the above we can write

$$\mathbb{E}[\sum_b \Delta\varphi^b(\check{\mathbf{w}}^b)] = \sum_b \sum_{\tau=1}^{N_S} c_\tau^b \sum_{\theta \in \mathbb{T}_R(\tau)} \frac{(q^b)^\top}{K} \mathbb{E}[w_\theta^b + \epsilon_\theta^b]$$

Under Assumption 6 the aggregator aims to maximize its expected profit

$$\psi(\mathbf{r}^M, \mathbf{r}, \mathbf{s}, R) = \mathbb{E} \left[\sum_{t=1}^{N_M} (\rho_t^\uparrow r_t^{M\uparrow} + \rho_t^\downarrow r_t^{M\downarrow}) - \mathcal{P}^M(\boldsymbol{\epsilon}^M) - \sum_b (R^b + \Delta\varphi^b(\check{\mathbf{w}}^b) - \mathcal{P}^b(\mathbf{s}^b)) \right] \quad (3.18)$$

which consists of the expected revenue generated in the spot market, a penalty $\mathcal{P}^M(\boldsymbol{\epsilon}^M)$ incurred for deviations from the market dispatch signal and, for each building, the reward R^b , the cost difference $\Delta\varphi^b$ and the building's penalty \mathcal{P}^b . A typical market penalty would for example be

$$\mathcal{P}^M(\boldsymbol{\epsilon}^M) = \sum_{\theta=1}^{N_R} p_h(\epsilon_\theta^M)^+ + p_l(\epsilon_\theta^M)^- \quad (3.19)$$

where $p_h, p_l \geq 0$ are the per-unit penalty factors.¹²

So far we have not discussed how the aggregator would determine the signals w_θ^b from the ISO's LFC signal ω_θ . While the question of how to do this in an optimal fashion is intriguing,¹³ we for the purpose of this chapter make the following simplifying assumption:

¹²Alternatively, $-\mathcal{P}^M$ can also be interpreted as a "pay for performance" payment (FERC, 2011a).

¹³The overall demand modulation must be split not only between the different buildings but also between different actuators within the buildings.

Assumption 9. *The contributions of the different buildings to the provided overall modulation of the demand are determined in a “pro rata” fashion, i.e., proportional to their respective contracted capacities. Specifically, we assume that*

$$w_\theta^b = \begin{cases} -\omega_\theta r_{\tau_\theta}^{\text{b}\uparrow} & \text{if } \omega_\theta \geq 0 \\ -\omega_\theta r_{\tau_\theta}^{\text{b}\downarrow} & \text{if } \omega_\theta < 0 \end{cases} \quad (3.20)$$

for all $b \in \mathcal{B}$ and $\theta = 0, \dots, N_R$.

Not surprisingly, the aggregator’s payoff will depend on the statistical properties of the LFC signal ω . Since this signal depends directly on the mismatch of supply and demand in the grid, which the ISO consistently tries to minimize, it is quite natural to make the following assumption:

Assumption 10. *For all $\theta = 0, \dots, N_R$ the LFC signal ω_θ is a random variable supported on $[-1, 1]$ with a continuous distribution that is symmetric around zero.*

Using Assumptions 9 and 10 we can compute

$$\mathbb{E}[w_\theta^b] = \zeta_\theta (r_{\tau_\theta}^{\text{b}\downarrow} - r_{\tau_\theta}^{\text{b}\uparrow}) \quad (3.21a)$$

$$\mathbb{E}[\epsilon_\theta^b] = \mathbb{E}[\left((1 - \omega_\theta)r_{\tau_\theta}^{\text{b}\uparrow} - s_{\tau_\theta}^{\text{b}\uparrow}\right)^- \mathbf{1}_{\{\omega_\theta \geq 0\}}] - \mathbb{E}[\left(-(1 + \omega_\theta)r_{\tau_\theta}^{\text{b}\downarrow} + s_{\tau_\theta}^{\text{b}\downarrow}\right)^+ \mathbf{1}_{\{\omega_\theta < 0\}}] \quad (3.21b)$$

where $\zeta_\theta := 0.5 \mathbb{E}[\omega_\theta | \omega_\theta \geq 0]$. Hence $\mathbb{E}[w_\theta^b]$ is linear in the capacities, and this is independent of the distribution of ω_θ . This is not the case for $\mathbb{E}[\epsilon_\theta^b]$, i.e., the effect of $r_{\tau_\theta}^{\text{b}}$ and $s_{\tau_\theta}^{\text{b}}$ on $\mathbb{E}[\epsilon_\theta^b]$ on is distribution-dependent. E.g., if $\omega_\theta \sim U[-1, 1]$ one can find that

$$\mathbb{E}[\epsilon_\theta^b] = \frac{1}{4} \left[(s_{\tau_\theta}^{\text{b}\uparrow})^2 / r_{\tau_\theta}^{\text{b}\uparrow} - (s_{\tau_\theta}^{\text{b}\downarrow})^2 / r_{\tau_\theta}^{\text{b}\downarrow} \right]$$

From now on we will assume for simplicity that $\zeta_\theta = \zeta$ for all θ , which is for example the case if the ω_θ are identically distributed. From (3.21) and Assumption 10 it follows that $\mathbb{E}[w_\theta^b + \epsilon_\theta^b] = 0$ if capacities and slacks are symmetric.

Individual Rationality Constraints

Observe that for all buildings $b \in \mathcal{B}$ we have that

$$\varphi_{\text{oo}}^{\text{b}*} := \varphi_{\text{nom}}^{\text{b}}(\mathbf{u}^{\text{b}*}(0, 0)) \leq \varphi_{\text{nom}}^{\text{b}}(\mathbf{u}^{\text{b}*}(\mathbf{r}^{\text{b}\downarrow}, \mathbf{r}^{\text{b}\uparrow}))$$

for any capacity vectors $\mathbf{r}^{\text{b}\downarrow}, \mathbf{r}^{\text{b}\uparrow}$. We refer to $\varphi_{\text{oo}}^{\text{b}*}$ as the value of the *outside option* of building b , since it represents the minimum (nominal) energy cost without contract. The aggregator must then ensure that the overall cost incurred by building b (including the monetary reward R^b) when participating in the contract does not exceed $\varphi_{\text{oo}}^{\text{b}*}$.

The Aggregator's Optimization Problem

The aggregator's optimal contract design problem can now be stated as the following bilevel optimization problem (Colson et al., 2007):

$$\max_{\mathbf{r}^M, \mathbf{r}, \mathbf{s}, R} \psi(\mathbf{r}^M, \mathbf{r}, \mathbf{s}, R) \quad (3.22a)$$

subject to

$$r_t^{M\uparrow} \geq 0, r_t^{M\downarrow} \geq 0 \quad \forall t \leq N_M \quad (3.22b)$$

$$r_\tau^{b\uparrow} \geq 0, r_\tau^{b\downarrow} \geq 0 \quad \forall b \in \mathcal{B}, \forall \tau \leq N_S \quad (3.22c)$$

$$(\mathbf{u}^b, \mathbf{s}^b) \text{ solves (3.16)} \quad \forall b \in \mathcal{B} \quad (3.22d)$$

$$\varphi^b(\mathbf{u}^b, \mathbf{s}^b) - R^b \leq \varphi_{oo}^{b*} \quad \forall b \in \mathcal{B} \quad (3.22e)$$

In the above problem, (3.22d) requires each building to behave optimally with respect to its capacity vectors $\mathbf{r}^{b\downarrow}$ and $\mathbf{r}^{b\uparrow}$. Participation of each building is ensured by the individual rationality constraint (3.22e).

Observing that R^b does not affect (3.22d) and that (3.18) is strictly decreasing in R^b , it is clear that the individual rationality constraint (3.22e) will be tight at the optimum. Substituting the resulting equality into the objective we get

$$\begin{aligned} \tilde{\psi}(\mathbf{r}^M, \mathbf{r}, \mathbf{s}) = \mathbb{E} \left[\sum_{t=1}^{N_M} (\rho_t^\uparrow r_t^{M\uparrow} + \rho_t^\downarrow r_t^{M\downarrow}) - \mathcal{P}^M(\boldsymbol{\epsilon}^M) \right. \\ \left. - \sum_b \varphi_{\text{nom}}^b(\mathbf{u}^b) - \varphi_{oo}^{b*} + \Delta\varphi^b(\check{\mathbf{w}}^b) \right] \end{aligned} \quad (3.23)$$

and with that (3.22) is equivalent to

$$\max_{\mathbf{r}^M, \mathbf{r}, \mathbf{s}} \tilde{\psi}(\mathbf{r}^M, \mathbf{r}, \mathbf{s}) \quad \text{s.t.} \quad (3.22b), (3.22c), (3.22d) \quad (3.24)$$

If the nominal scheduling problems for all buildings are feasible, then (3.22) will always have a solution: if $\mathbf{r}^M = \mathbf{r} = \mathbf{s} = 0$ and $\mathbf{u}^b = \mathbf{u}^{b*}(0, 0)$ for all b then (3.22b) and (3.22c) hold with equality, and (3.22d) holds by definition of \mathbf{u}^{b*} .

3.3.1 Solution Methodology

The aggregator's optimization problem (3.22a) is a bilevel optimization problem, for which different solution approaches exist. These include vertex enumeration techniques, branch-and-bound (BNB) algorithms, penalty function methods, descent algorithms and trust-region methods (Colson et al., 2007). We focus on BNB techniques since they can provide certificates for global optimality of the solution based on duality theory.

The General Case

The basic idea of BNB methods for bilevel optimization problems is to replace the lower level problem by a set of suitable optimality conditions. In our case the lower level problem is a QP,

so the corresponding Karush-Kuhn-Tucker (KKT) optimality conditions are necessary and sufficient. Hence the bilevel problem can be written equivalently as a standard optimization problem in which the KKT conditions of the lower level problem appear as constraints, where the dual variables of the lower level problem now appear as primal variables.

Since for QPs the primal feasibility, dual feasibility and Lagrangian stationarity conditions are all affine constraints, the only challenge are the complementary slackness (CS) conditions. For example, the corresponding CS conditions for the (upper) robustified input constraints (3.13) are given by

$$\underline{\lambda}_\tau^b \times (-u_\tau^b + \underline{u}_\tau^b + r_\tau^{\text{b}\uparrow} - s_\tau^{\text{b}\uparrow}) = 0, \quad \forall \tau = 0, \dots, N_S$$

where $0 \leq \underline{\lambda}_\tau^b \in \mathbb{R}^{n_u^b}$ are the dual variables.

There are two main ways of dealing with the CS conditions. One is to directly apply a specialized BNB algorithm, the other one is to reformulate the problem by introducing auxiliary binary variables and bounding primal constraints and dual variables. We take the latter approach, which is the same “big-M” method we used in formulating the optimization problem in Chapter 2. Introducing variables $\underline{z}_\tau^b \in \{0, 1\}^{n_u^b}$, we may express primal feasibility, dual feasibility and CS conditions jointly by

$$-M_{\underline{u}_\tau^b} \times \underline{z}_\tau^b \leq -u_\tau^b + \underline{u}_\tau^b + r_\tau^{\text{b}\uparrow} - s_\tau^{\text{b}\uparrow} \leq 0 \quad (3.25a)$$

$$0 \leq \underline{\lambda}_\tau^b \leq (1 - \underline{z}_\tau^b) \times M_{\underline{\lambda}_{u,\tau}^b} \quad (3.25b)$$

where $M_{\underline{u}_\tau^b}, M_{\underline{\lambda}_\tau^b} \in \mathbb{R}_+^{n_u^b}$ are sufficiently large constants. All other CS conditions can be handled analogously.

Thus, the bilevel problem (3.24) is equivalent to

$$\max_{\mathbf{r}^M, \mathbf{r}, \mathbf{u}, \mathbf{s}, \underline{\lambda}, \mathbf{z}} \tilde{\psi}(\mathbf{r}^M, \mathbf{r}, \mathbf{s}) \quad \text{s.t.} \quad (3.22b), (3.22c), \text{ (CS), (LS)} \quad (3.26)$$

where **(CS)** denotes a set of constraints of the form (3.25), and **(LS)** represents the Lagrangian stationarity constraint, which is an affine equality constraint.

In general, (3.26) is a Mixed-Integer optimization problem with linear constraints and nonlinear objective, the specific form of which depends on the kind of market penalty and the shape of the distribution of the LFC signal ω_θ . Note that (3.26) is highly structured. In particular, the constraints are completely decoupled. This suggests that the problem should be amenable to decomposition techniques that aim at separating the coupling term \mathcal{P}^M in the objective. An important feature of this formulation is that, since the KKT conditions of the buildings’ problems appear explicitly as constraints, any feasible solution of (3.26) describes a situation in which all buildings behave optimally with respect to their regulation capacities. Hence it is not necessary to find the global optimum in order to ensure individual rationality.

The First-Best Contract

In general, the buildings' decisions impose externalities on the aggregator via their effect on the distributions of ϵ^M and $\tilde{\mathbf{w}}^b$. If the aggregator wants to avoid being penalized for deviations from schedule in the market and therefore insists that $\mathbf{s} = 0$, i.e., that all buildings are robustly providing their full capacities \mathbf{r}^b , then the problem can be transformed into a single Linear Program.

Indeed, if $\mathbf{s}^b = 0$ for all b , then $\mathcal{P}^b(\mathbf{s}^b) = 0$, $\tilde{\mathbf{w}}^b = \mathbf{w}^b$ for all b , and $\epsilon^M = 0$. Hence we see from (3.23) that the only term in the objective that depends on building b 's decision variables is φ_{nom}^b . Moreover, it is clearly optimal for the aggregator to commit all available capacity in the market:

$$\left. \begin{aligned} r_t^{M\uparrow} &= \sum_b (q^b)^\top r_\tau^{b\uparrow} \\ r_t^{M\downarrow} &= \sum_b (q^b)^\top r_\tau^{b\downarrow} \end{aligned} \right\}, \quad \forall t \leq N_M, \tau \in \mathbb{T}_S(t) \quad (3.27)$$

Writing

$$\begin{aligned} \mathcal{U}^b(\mathbf{r}^b) &= \{\mathbf{u}^b : (3.13), (3.15) \text{ hold with } \mathbf{s}^b = 0\} \\ \mathcal{R} &= \{(\mathbf{r}^M, \mathbf{r}) : \mathbf{r}^M, \mathbf{r} \geq 0, (3.27) \text{ holds}\} \end{aligned}$$

we have

$$\begin{aligned} & \max_{(\mathbf{r}^M, \mathbf{r}) \in \mathcal{R}} \mathbb{E} \left[\sum_{t=1}^{N_M} (\rho_t^\uparrow r_t^{M\uparrow} + \rho_t^\downarrow r_t^{M\downarrow}) + \sum_b \varphi_{\text{oo}}^{b*} \right. \\ & \quad \left. - \sum_b (\min_{\mathbf{u}^b \in \mathcal{U}^b(\mathbf{r}^b)} \varphi_{\text{nom}}^b(\mathbf{u}^b) + \Delta\varphi^b(\mathbf{w}^b)) \right] \\ &= \max_{\substack{(\mathbf{r}^M, \mathbf{r}) \in \mathcal{R} \\ \mathbf{u}^b \in \mathcal{U}^b(\mathbf{r}^b), \forall b}} \left[\sum_{t=1}^{N_M} (\mathbb{E}[\rho_t^\uparrow] r_t^{M\uparrow} + \mathbb{E}[\rho_t^\downarrow] r_t^{M\downarrow}) + \sum_b \varphi_{\text{oo}}^{b*} \right. \\ & \quad \left. - \sum_b (\varphi_{\text{nom}}^b(\mathbf{u}^b) + \mathbb{E}[\Delta\varphi^b(\mathbf{w}^b)]) \right] \end{aligned} \quad (3.28)$$

where

$$\mathbb{E}[\sum_b \Delta\varphi^b(\mathbf{w}^b)] = \sum_b \sum_{\tau=1}^{N_S} c_\tau^b \sum_{\theta \in \mathbb{T}_R(\tau)} \frac{(q^b)^\top}{K} \zeta_\theta (r_{\tau\theta}^{b\downarrow} - r_{\tau\theta}^{b\uparrow})$$

is linear in the decision variables. Finally \mathcal{R} and \mathcal{U}^b are described by a set of linear constraints, so (3.28) is a Linear Program. Economically speaking, there is no Moral Hazard, so the aggregator is able to implement what is known as the first-best contract (Laffont and Martimort, 2002) in contract theory.¹⁴ However, the first-best contract in this setting need not necessarily be optimal for the aggregator. This is the case if the benefit from the buildings' reduced costs for adjusting their schedule under less strict robustness requirements outweighs the expected deviation penalties incurred in the market.

¹⁴Note that there is also no Adverse Selection as there are no information asymmetries.

3.3.2 A Numerical Example

We now apply our methodology for the first-best contract to a simple numerical example with four buildings. For each building, we use a scaled version of a simple one-dimensional model identified by Aswani et al. (2012b) using semiparametric regression techniques on a testbed on the UC Berkeley campus.

We emphasize that our simulations are primarily for illustration and for showcasing some important qualitative properties of the optimal contract. In order to get meaningful quantitative results on the economic feasibility of such a contract scheme, a more thorough study is necessary. For our example we consider a scheduling horizon of $T_M = 24$ and scheduling intervals of length $T_S = 15$ min, so that $N_S = 96$ and $K' = 4$.

Model Parameters

As in the previous chapter, we use PG&E’s commercial electricity tariffs in our simulation. For simplicity, we assume that the monthly peak demand is not affected by the modified schedules of the contract, which means the demand charge does not change.

We generate four slightly perturbed versions of the model identified in Aswani et al. (2012b), where the disturbance v is given by a weighted sum of outside temperature and heating load from occupants, equipment, and solar heating. Figure 3.3 shows the disturbance forecasts we used, which reflect different outside temperatures and usage patterns.

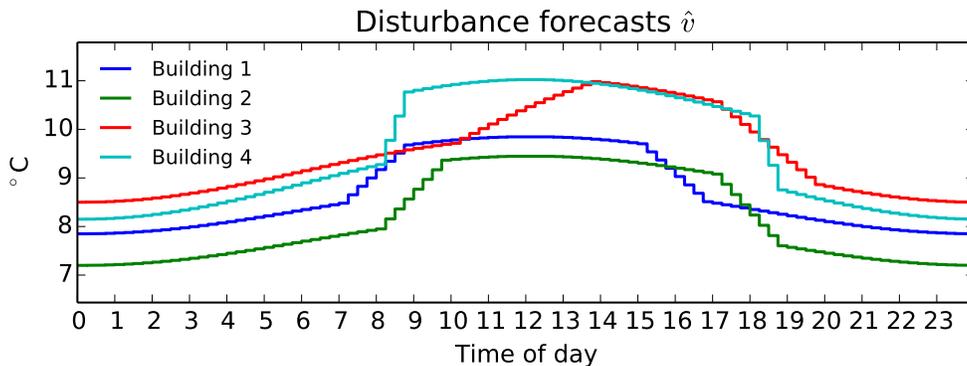


Figure 3.3: Disturbance forecasts

For each model we assume a maximum power consumption typical for mid-to-large size commercial buildings, from which we determine the consumption factors q . For all models we have $C = E = 1$ and $D = F = 0$, with the remaining parameters given in Table 3.1. Comfort constraints are chosen so as to have a tight temperature band during common working hours, and looser band outside common working hours. Finally, we assume that $\underline{\alpha}_{t,i} = \bar{\alpha}_{t,i} = 5\%$, $\Sigma_0 = 0.5$, $\kappa = 0$ and $\underline{u} = 0$ for all buildings.

Figure 3.4 shows CAISO DAM regulation up and down prices for August 2012, with the respective mean for each hour across the whole month. We find that with these capacity

b	A	B	x_0 [°C]	q [kWh]	\bar{u}	G	H	tariff
1	0.640	-2.64	23.0	43.75	0.5	0.20	0.15	A10 TOU
2	0.680	-2.55	22.0	62.50	0.75	0.25	0.20	E19 TOU
3	0.635	-2.70	24.0	31.25	0.65	0.15	0.10	A10 TOU
4	0.620	-2.65	22.5	50.00	0.5	0.10	0.15	A10

Table 3.1: Building model parameters

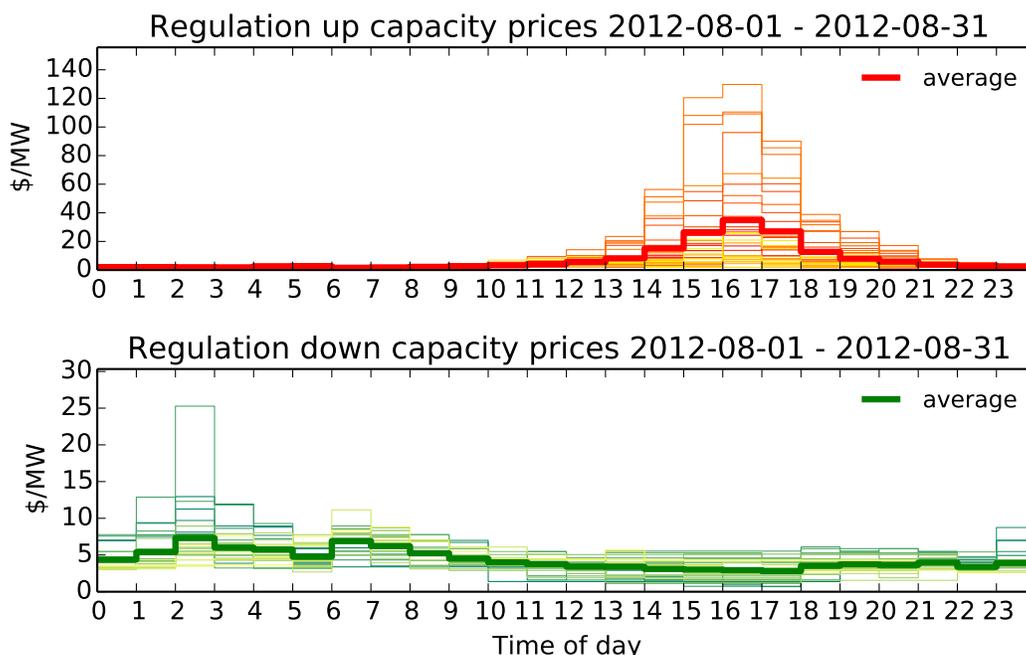


Figure 3.4: CAISO DAM regulation capacity prices for August 2012

prices the optimal strategy for the aggregator is to not bid any capacity in the market. This is primarily because of the small thermal capacity and inefficient HVAC system of our simple model. As a result the increased cost for deviating from the cost-optimal power schedule outweighs the potential revenue in the capacity market. Under the HVAC model and one of the benchmark tariffs with lower markups from Chapter 2 we would likely see the aggregator bid capacity in the market. However, the point of this illustrative example is to show some qualitative properties of the optimal contract, so we choose to simply scale the expected capacity prices ρ^\uparrow and ρ^\downarrow by a factor 20 and 15, respectively, for our simulations. Finally, we assume that the ISO's LFC signal ω_θ for each θ is uniformly distributed on $[-1, 1]$, which implies that $\zeta = 0.25$.

Simulation Results

The optimal contract design problem for the DAM for the above parameters is solved also with GUROBI. The aggregator’s expected profit is \$20.7, the rewards R^b are \$28.4, \$24.1, \$9.0 and \$17.5 for buildings 1-4. Hence for the simple building model used in this simulation, the aggregator must spend a large share of its revenue on incentivizing the buildings.

Figure 3.5 shows, using a stacked plot, how the individual buildings contribute to the overall regulation capacities offered in the market. Observe how the contributions vary on the faster scheduling frequency on the building level to provide capacity on the slower market time scale in the most cost-effective way.

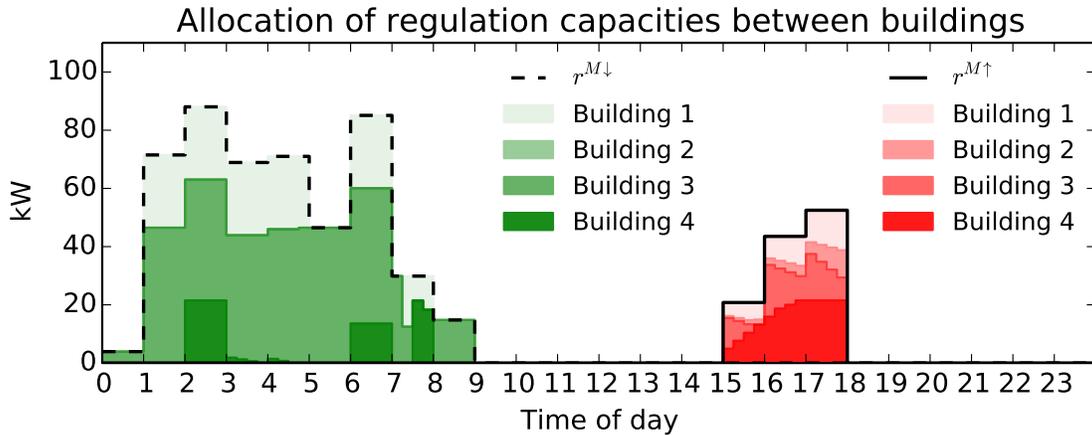


Figure 3.5: Shares of the overall regulation capacities between buildings

Temperature evolutions and control inputs for Building 1 are shown in Figure 3.6, where u_{oo} and u^* are the optimal schedules of outside option and optimal contract, respectively. The regulation signal w has been sampled uniformly from within the contracted bounds. In order to be able to clearly observe the behavior of the optimal control, we have set $\tilde{v} = 0$ in Figure 3.6. The shaded areas in the lower plot represent the amounts by which the actuator constraints have been tightened in order to be able to ensure that the resulting nominal control schedule is robustly feasible (observe that the control $u^* + w$ indeed never violates the actuator constraints).

Figure 3.7 shows the temperatures for Building 1 for 30 randomly sampled errors sequences \tilde{v} , again assuming a uniformly distributed LFC signal ω . The individual comfort constraints are violated in 0.57% of all time intervals.

Finally, Figure 3.8 shows the strong dependency of the total amount of contracted capacity over the scheduling horizon and the aggregator’s expected profit on the parameter ζ , and hence on the shape of the distribution of ω . Noting that higher values of ζ correspond to a high variance of ω , this is not surprising. Indeed, when contracting regulation down capacity this would result in the cost difference $\Delta\varphi(w)$ likely being large, so in expectation the aggregator would need to pay a large amount to ensure that buildings are only charged

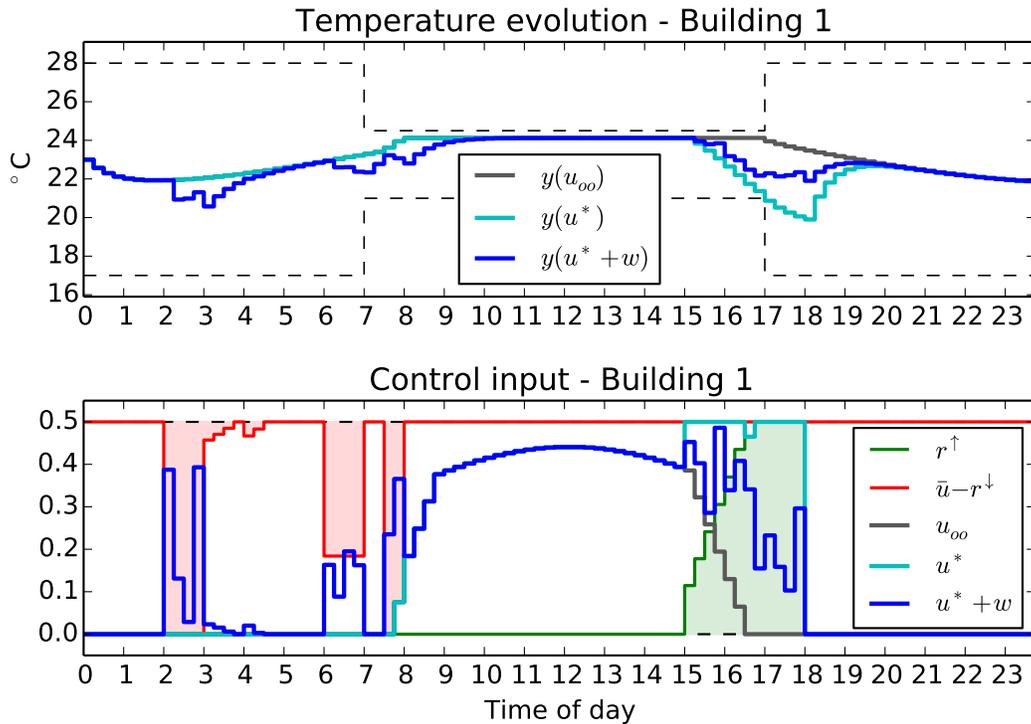


Figure 3.6: Temperature and control for Building 1

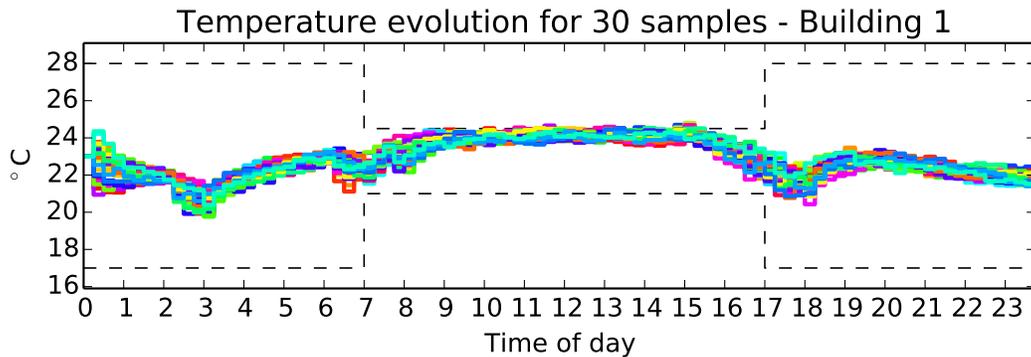


Figure 3.7: Sample temperature trajectories for Building 1

according to nominal schedules. For small ζ the distribution of ω is concentrated around zero and this issue is less pronounced. The situation is mirrored for regulation up capacity; in this case large values of ζ are beneficial for the aggregator. Figure 3.8 also shows that the overall amount of capacity offered is much higher for regulation down. This is due to the fact that in many cases regulation down capacity can be offered without having to deviate from the optimal schedule of the outside option.

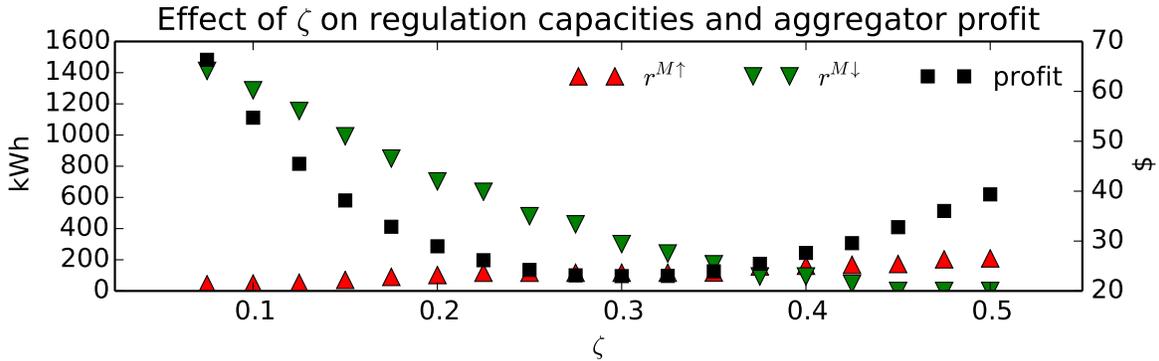


Figure 3.8: Effect of the parameter ζ

3.3.3 Discussion

We made the simplifying assumption that a building’s power consumption is linear in the HVAC control input u_t^b , which is quite restrictive. For example, the power consumption of a fan is often modeled as quadratic or cubic in its air mass flow. At the expense of having to solve a harder optimization problem, our approach can be extended to the case when a building’s total HVAC energy consumption per interval includes a convex quadratic of the control signal. Notably, in this case the structure of the KKT conditions of the building’s problems does not change, so the buildings’ individual rationality constraint can be incorporated in the same way. However, (3.22d) and (3.22e) now involve quadratic constraints including bilinear terms between scheduled control \mathbf{u}^b and the capacity vectors $\mathbf{r}^{b\uparrow}$ and $\mathbf{r}^{b\downarrow}$. The reason for this is that, unlike in the linear case, the amount of regulation capacity (in terms of power, not control) now also depends on u_t^b . By employing rank-relaxation techniques (see e.g. Recht et al. (2010)), the problem can then be “lifted” to a mixed-integer semidefinite program (SDP).

Under common simplifying assumptions many building models are bilinear,¹⁵ which can be approximated by linearizing the model at various operating points to obtain a piece-wise affine model. As we have remarked in the previous chapter, such models can be included into mixed-integer programs in relatively straightforward fashion. However, while this poses no major challenges to the building’s optimization problem from Section A.1, using the same approach to reformulating the bilevel optimization problem would require explicit and tractable optimality conditions for the resulting MILP, which unfortunately are not known.

A higher-level question is whether the aggregator actually requires a fully detailed building model for the purpose of determining a nominal energy schedule. In order to be able to aggregate a larger number of buildings, it would be preferable to use rather generic models that are easily parametrized for different buildings. This approach would have the added benefit of protecting the privacy of buildings by masking their true detailed models. Ideally, models for this purpose are sufficiently detailed to capture the buildings’ temporal flexi-

¹⁵see e.g. Aswani et al. (2012a), Oldewurtel et al. (2012) and Anthony and Borrelli (2011).

bilities, but simple enough so they are easily communicated and render the aggregator's optimization problem tractable. This raises questions pertaining to model reduction that go far beyond this particular application.

Finally, our assumption that there is no private information at all will generally not hold in practice. On the one hand, measuring HVAC power consumption at high frequency can be achieved by using a dedicated meter, so our problem is one of perfect monitoring. On the other hand, an aggregator in practice will not know the exact model of the building's dynamics, which causes information asymmetry. Therefore it is important to understand how much a building could benefit from misreporting some of its private parameters. Determining this information rent, however, seems quite challenging.

Chapter 4

Machine Learning Methods for Causal Inference on Time Series Data

4.1 Introduction

The goal of causal inference is to draw conclusions about a causal connection between variables, based on the conditions of the occurrence of some effect under a change of a subset of the variables (the *cause*). For instance, in Chapter 5, we are interested in estimating the causal effect of the rebate level during Demand Response events on the electricity consumption of the program participants during those events.

The well-known phrase “correlation does not imply causation” rightly indicates that causal inference can be a very challenging task, in particular if the variation in the cause is not controlled by the researcher. Holland (1986) coined the maxim of “no causation without manipulation,” emphasizing that it is impossible to perform causal inference without manipulation of the cause. However, even if the cause is being manipulated, a basic issue remains: the outcome of a unit of observation (an *individual*) is observed exactly under a single value of the cause. The outcome under different values of the cause is necessarily unobserved. This is commonly referred to as the “fundamental problem of causal inference” (Holland, 1986).

In order to estimate the causal effect, one therefore needs a good proxy for the missing observations, usually referred to as the *controls*. Consider the simple example in which the cause is a binary treatment (e.g. “prescribe drug” vs. “do not prescribe drug”). In an ideal world, the researcher conducts a well-designed experiment, in which multiple individuals from some experiment population are randomly assigned to two different groups (“treatment” and “control”). Comparing the outcomes (in this case patient health) between treatment and control groups then provides a causal estimate of the *Average Treatment Effect* (ATE), as the randomization renders the two groups statistically comparable in every way. In other words, the average outcome in the control group provides an estimate of the average *counterfactual* outcome of the treated (i.e. the outcome of the subjects in the treatment group had they not received treatment).

Such an experiment is referred to as a *Randomized Controlled Trial* (RCT), and is attractive in that it provides causal estimates under relatively weak assumptions (Angrist and Pischke, 2008). If one is willing to make stronger assumptions, it is also possible to obtain causal estimates through what is usually referred to as *observational* studies, in which the cause is manipulated not by the researcher, but through some other mechanism that is potentially unknown to the researcher. The problem in this case is of course that the value of the cause may depend on inherent properties of the subjects, so that different groups are not statistically comparable anymore. Under certain conditions, a proper re-weighting of the observations of the subjects can still yield a causal estimate of the effect (see e.g. Rosenbaum and Rubin (1983)).

Both in experimental and observational settings, causal inference is traditionally performed by comparing observations of outcomes across different subjects.¹ We will refer to this approach as “across-subject matching” (ASM). The reason for across-subject matching being the standard evaluation technique is simply that most studies have traditionally involved only few (often only a single) observation of the outcome variable per subject.² However, this situation is changing, and recently we have seen a growing availability of high-frequency time series data on the level of individuals, for example in personalized health and fitness tracking (with devices such as the fitbit), social networks, power systems (where consumption data from smart meters on a 15 minute or hourly resolution has become the norm), and many other domains whose future importance we are likely not yet aware of (e.g. autonomous driving). If such high-frequency data includes periods of repeated exposure to some treatment, it potentially can be used to perform causal inference on the level of individuals. Specifically, it may be feasible to estimate an individual treatment effect (ITE) by essentially comparing the outcomes in treatment periods with those in non-treatment periods, an approach we will refer to as “within-subject matching” (WSM). While such an approach does address some of the problems of conventional techniques (primarily that of implementing a randomized experiment), it also raises a number of new challenges.

This chapter develops a two-stage estimation framework for estimating individual treatment effects, by combining techniques from econometrics, Machine Learning, and computational statistics to perform within-subject matching. The basic idea of these methods is the following: Assuming a scalar outcome variable, we first split the observations on an individual in treatment (“estimation”) and non-treatment (training or “control”) periods. On the training data, the problem reduces to a standard supervised learning or prediction problem. Using cross-validation, we train a Machine Learning model on the training data, which can be thought of as capturing the behavior of this individual in the absence of treatment.

¹A subject here generally refers to a unit for which one observes outcomes, and could be an individual person, a group, a state, a country, etc. For the purpose of this thesis we will consider as subjects individual economic actors.

²Data with repeated observations per individual are usually referred to as “panel data.” Typically, panel data in econometrics is of relatively low frequency. The most important exception is financial data (Andersen et al., 2003), but there has been work on inference on high-frequency data also in other fields, such as neuroscience (Valdés-Sosa et al., 2006) and energy (Jesoe et al., 2015).

Our framework is model-agnostic, and allows for choosing basically any prediction model (ranging from simple Generalized Linear Models over Decision Trees and Neural Networks to ensemble methods), whichever is best suited for the task at hand. For each treatment observation, we then generate an estimate of the counterfactual outcome in the absence of treatment by generating a prediction from our trained model using the respective covariate vector. This completes the first stage of our procedure. In the second stage, we then compare the predicted counterfactuals with the observed outcome under treatment and characterize the individual treatment effect. In the simplest case we just compute the sample mean to obtain the marginal ITE. With sufficient treatment observations, we may also estimate the conditional ITE as a function of the covariates or a treatment parameter in order to characterize heterogeneity of the treatment effect on the level of individuals. Since in essence the second stage estimation is just another regression problem, we can again employ a variety of regression methods for this task, including both parametric and non-parametric models. Finally, marginalizing over the individuals of the group we can also easily obtain an estimate of what is known as the average treatment effect on the treated (ATT).

The methodology described above is modular, separating the causal inference problem into a prediction and an estimation phase. This allows for using the prediction method best suited to the problem, while still providing theoretical guarantees of the estimation procedure under reasonable assumptions. Utilizing the massive amounts of data that are being collected in today's world, our methodology can be used to obtain estimates with causal interpretation. In particular, it can provide a good proxy for treatment effect estimates whenever RCTs are infeasible, and in some situations actually provide more precise causal estimates.³ It is trivially parallelizable across individuals and hence avoids some of the computational challenges of performing conventional analyses based on linear regression models with fixed-effect specification. Most importantly, by characterizing the heterogeneity of the treatment effect by means of the ITE estimates, it offers new opportunities for both economic analysis and operational decision-making.

4.2 Background

The presence of heterogeneity in treatment effects has long been recognized by economists. However, the majority of methods developed in the program evaluation literature that follow the work by Rubin (1974) and Rosenbaum and Rubin (1983), while allowing for heterogeneity in the treatment effect, do not attempt to test for or characterize this heterogeneity.

One type of heterogeneity investigated in the literature pertains to the the distribution of the treatment effect across the population. For example, one may be interested in the effect of a treatment on the dispersion of an outcome, or in its effect on the lower tail of the outcome distribution. The quantity generally considered in characterizing this kind of heterogeneity is the quantile treatment effect (QTE), first defined by Doksum (1974). It corresponds, for any fixed percentile, to the horizontal distance between the cumulative

³See Jesso et al. (2015) for a similar observation.

distribution functions of outcomes of under treatment and non-treatment. Firpo (2007) proposes an estimator for quantile treatment effects under the identifying restriction that selection to treatment is based on observable characteristics. The estimator is based on a two-step method that combines nonparametric estimation of the propensity score with the computation of the difference between the solutions of two separate minimization problems. The objective functions of these minimization problems involve the sum of so-called check functions, an approach that was first suggested by Koenker and Bassett (1978).

Maybe the most basic notion of heterogeneity of the treatment effect is that of a difference in the effect across different subpopulations in the sample characterized by covariates. Such heterogeneity has long been studied in both empirical and methodological work. For instance, Crump et al. (2008) develop two nonparametric tests for testing the hypotheses that (i) the treatment has zero average effect for all subpopulations defined by covariates and that (ii) the average effect conditional on the covariates is identical for all subpopulations. One of the main challenges with having large amounts of data available is the very basic task of how to organize and make sense of data in a very large covariate space. In a study with only a few covariates per individual, it is typically relatively easy to manually examine heterogeneity along different dimensions. However, as the dimensionality of the problem increases this approach quickly becomes infeasible. This is one situation in which, by detecting patterns in the data, Machine Learning can, in principle, add value to econometric analyses. However, the lack of a common understanding of the respective goals between the two fields can pose a challenge.

4.2.1 Machine Learning and Econometrics

Econometrics and Machine Learning are disciplines between which interaction, at least historically, has been rather limited. Traditionally, econometrics has been concerned primarily with causal inference, while (supervised) Machine Learning has focused primarily on prediction. In causal inference, one performs statistical estimation in order to answer a question about the counterfactual impact of a change in policy, usually referred to as a "treatment" (Athey, 2015). Typically, the goal is to extrapolate from the measured treatment effect, i.e., to estimate the impact of adopting this change in policy on a wider scale and in a potentially different environment. In contrast, the goal in prediction problems is to forecast outcomes based on observed covariates (or "features") in a setting where the underlying data generating process remains unchanged. In particular, in prediction one is typically not concerned with potential correlations between "treatment features" and other features. Therefore, one has to be careful when trying to use methods designed for prediction as a tool for causal inference in econometrics.

Nevertheless, there is a significant opportunity for both fields to benefit from the approaches and tools in the respective other. Econometricians may isolate prediction problems in their work and utilize modern Machine Learning algorithms to greatly improve performance (over simpler methods such as Linear Regression), or go a step further and modify tools from Machine Learning in order to perform causal inference. Conversely, Machine

Learning researchers can broaden the reach of their methods by adopting certain points of view from econometrics, for example by making a distinction between internal and exogenously varied features, and thereby expanding the applicability of their methods beyond standard prediction. The connections between econometrics and Machine Learning as well as the potential for cross-pollination between the fields have been discussed in recent survey papers by Einav and Levin (2013), Varian (2014), and Athey (2015).

Machine Learning Tools in Econometrics

Machine Learning has been and continues to be remarkably successful in a wide range of practical applications, from recommender systems (Ricci et al., 2011), natural language processing (Manning and Schütze, 1999) and computer vision (Szeliski, 2010) to personalized medicine (Cruz and Wishart, 2006) and epidemiology (Lee et al., 2010).

With this success, some economists are starting to lose their reservations and are beginning to embrace these tools and apply them in various contexts in economic research. A straightforward way to do so is to simply use existing predictions as an input to a more involved economic analysis. For instance, in the context of credit and insurance markets Bundorf et al. (2012) and Einav et al. (2013) use off-the-shelf credit and health risk scores to account for the probability of default or likely health expenditures of individual consumers. In their work predictive modeling is used as an “off-the-shelf” tool to stratify the population into different risk levels, based on which heterogeneity in the consumer behavior can be investigated.

Rather than simply using the result of Machine Learning algorithms as an input, some researchers have started to use modern optimization-based regression methods in order to boost the performance of existing analysis tools. For example, Belloni et al. (2012) use Lasso (L1-norm regularization) methods to form first-stage predictions and estimate optimal instruments in linear instrumental variables (IV) models with many instruments. They show that their estimator is \sqrt{n} -consistent and asymptotically normal when the first stage of the IV model is approximately sparse, and, moreover, that it is semiparametrically efficient (Newey, 1990) when the structural error is homoscedastic.

A route more challenging than employing “out-of-the-box” tools is to instead modify Machine Learning algorithms in such a way that they become directly applicable to causal inference. Maybe the best example in this area is the work by Athey and Imbens (2015), who propose methods for estimating heterogeneity in causal effects in experimental and observational studies. They provide a data-driven approach to partitioning the data into subpopulations which differ in the magnitude of their treatment effects. Their “honest” approach to estimation uses one sample to construct the partition and another independent sample to estimate treatment effects for each subpopulation. The method they employ is a modification of regression tree methods popular in Machine Learning, optimized for goodness of fit in treatment effects and to account for honest estimation. Their model selection criteria focus on improving the prediction of treatment effects conditional on covariates, anticipating that bias will be eliminated by honest estimation, but also accounting for the change in the

variance of treatment effect estimates within each subpopulation as a result of the split. In a closely related article, Wager and Athey (2015) propose non-parametric causal forests, a modification of random forests (Breiman, 2001), for estimating heterogeneous treatment effects. They show that these forests are point-wise consistent with an asymptotically Gaussian centered sampling distribution. Moreover, based on their algorithm they provide a method for constructing asymptotic confidence intervals.

Researchers have also started to develop their own methods that, while not directly using existing Machine Learning algorithms, are inspired by the approaches in the field. One such example is the work by Abadie et al. (2010), which performs causal inference using a synthetic control group, whose observations are generated by weighting the observations of multiple groups in order to form a group that is “more similar” to the treatment group. The weighting is determined by minimizing the error between the synthetic observations and those of the treatment group in a pre-treatment period, a strategy very much in the spirit of the standard Machine Learning approach of minimizing a loss function between predictions and ground truth.

Econometric Thinking in Machine Learning

Interestingly, while more and more economists are experimenting with Machine Learning tools, there seems to be much less activity in the other direction. Theoretical computer scientists have long been concerned with causality (Pearl, 2009). However, these theoretical results primarily ask higher-level structural questions (such as identifiability), and thus have not found their way into Machine Learning practice to a significant degree (Varian, 2014).

There has been some interest in the Machine Learning community in discovering the causes of given effects based on relatively unstructured observational data (Guyon et al., 2008b,a). This goal seems starkly at odds with the maxim of “no causation without manipulation” coined by Holland (1986). However, even in econometrics, researchers have become willing to discuss statistical models for observational data in an explicitly causal framework (Angrist and Pischke, 2008).

There are also concrete methods from econometrics that have influenced algorithms developed by computer scientists. For example, the offline policy evaluation work by Strehl et al. (2010) and Li et al. (2012, 2011) essentially uses inverse propensity score weighting, initially developed in Rosenbaum and Rubin (1984, 1983), in order to estimate the reward of a policy used in a contextual bandit problem using data that has been generated under a different policy.

4.2.2 Causal Inference on Time Series Data

In many applications, the variables of interest are not a collection of single observations on each individual, but instead consist of a sequence of observations for each individual obtained over multiple time periods. In econometrics, such data is typically referred to as *longitudinal* or *panel data*. In these settings, the notion of both treatment and effect is potentially much

more ambiguous. In particular, a treatment effect in general may be considered a function of time (e.g. the decaying effect of informational intervention on behavior (Caplin, 2016), or time-varying treatment-response curves in medicine (Xu et al., 2016)).

Besides this more complicated interpretation of the treatment effect, causal inference on time series data is subject to a number of general challenges. Fundamentally, time series data generally is inherently serially correlated. Moreover, it often exhibits strong seasonality, in particular for long observation horizons. In addition, it is often much harder to justify basic unconfoundedness assumptions (Abadie, 2005).

A common experimental approach is a so-called Difference in Difference (DiD) design (see e.g. Ashenfelter and Card (1985) and Angrist and Pischke (2008)), which compares pre-post differences in the treatment and control group, respectively, assuming a common underlying trend. An additional challenge is caused by the serial correlation in the individual observations, which complicates statistical inference. For example, using conventional methods, which do not fully take the correlation structure into account, will result in overly optimistic standard errors (Hansen, 2007; Bertrand et al., 2004).

In the context of comparative studies, Abadie et al. (2010) suggest a method of “synthetic controls,” in which the outcome variable of the treated unit is compared against a convex combination of the outcome variables of the non-treated units, where the weights are determined by minimizing the error in the pre-treatment periods. They show that, if the outcome is generated by a factor model, then, under relatively mild conditions, the bias of the estimator goes to zero as the number of pretreatment periods increases. The authors use their method to study the effects of Proposition 99, a large-scale tobacco control program that was implemented in California in 1988. The synthetic control in this case is a weighted average across other states, with weights chosen such that the resulting synthetic state best reproduces the values of a set of predictors of cigarette consumption in California before the passage of Proposition 99. This approach is an example for across-subject matching.

There is also a relatively large body of work developed in the marketing literature. For instance, Brodersen et al. (2015) use Bayesian Structural Time-Series (BSTS) models to estimate the causal effect of discrete market interventions from times series data. Their approach is based on a state-space model that includes a local trend, seasonality, and arbitrary covariates (assumed unaffected by control). This state-space model explicitly models the counterfactual outcomes of the time series after the intervention. Inference is performed in a Bayesian fashion, using a spike-and-slab prior on the covariate coefficients. The basic motivation for this approach is to be able to use a potentially large number of time series as model inputs, and have the model selection be performed implicitly by the Bayesian inference. Because closed-form posteriors for the employed models in general do not exist, approximate inference is performed using Markov Chain Monte Carlo (MCMC) sampling. The authors test their approach both on synthetic data and on an online advertising campaign on search-related website visits.

While data sets used in empirical research on program evaluation (e.g. labor statistics, test scores, data from medical trials, etc.) have long included large numbers of subjects, the granularity and frequency of the data on each subject has traditionally been quite limited

(e.g. typically less than maybe a few dozen variables collected at a frequency of days, months or even years). In the last decades, the increasing availability of computational resources and the low cost of data storage has led to an explosion of automated data collection in more and more areas. It has become quite common today to have access to data with a time resolution of minutes (e.g. smart meter consumption data) or even seconds (e.g. personalized health devices such as fitness trackers) on the level of an individual. In the remainder of this chapter we develop a framework to perform causal inference on such kind of high-frequency data on individuals with repeated treatment exposure using within-subject matching.

4.3 Individual Treatment Effects Under Binary Treatments

Our goal is to estimate the causal effect of some intervention (the *treatment*) on the real-valued outcome variable of an individual subject i , based on time series data collected on the individual. Throughout this chapter we index time by $t \in \mathbb{N}_+$ and without loss of generality we assume that the observation times are $t = 1, \dots, N$. We start by focusing on binary treatments, and denote by $D_{it} \in \{0, 1\}$ the binary variable taking the value 1 if individual i received treatment at time t and 0 otherwise.⁴ We will extend our model to more general treatments types in the following sections.

Let $Y_{it}^0 \in \mathbb{R}$ be the outcome that would be observed if individual i received treatment 0 at time t , and let Y_{it}^1 be the outcome that would be observed if she received treatment 1 at time t . These variables are referred to as *potential outcomes* (Rubin, 1974). In addition, let $X_{it} \in \mathcal{X}$ be a vector of covariates at time t in the covariate space $\mathcal{X} \subset \mathbb{R}^{n_x}$. We denote by $\mathcal{D}_i = ((X_{i1}, Y_{i1}), (X_{i2}, Y_{i2}), \dots)$ the full data set collected for individual i .

The basic requirement of any method for causal inference is to be able to attribute any measured effect in the outcome to the treatment alone. One of the potential issues with comparing outcomes across time is that the underlying data generating process may change over time. To ensure that we can still obtain a causal estimate of individual treatment effects therefore requires an additional assumption on the stationarity of the conditional distribution of the outcome given the observed covariates.

Assumption 11 (Stationarity of Conditional Distribution of Potential Outcomes). *The conditional distribution of the potential outcomes given the covariates is stationary. That is,*

$$Y_{it}^0, Y_{it}^1 \perp\!\!\!\perp t \mid X_{it} \tag{4.1}$$

Note that Assumption 11 *does not* require the joint distribution of covariates and potential outcomes to be stationary. In fact, the marginal distribution of the covariates X_{it} may change over time, which under Assumption 11 requires the joint distribution of $(X_{it}, Y_{it}^0, Y_{it}^1)$

⁴For the purpose of this thesis, we will refer to $D_{it} = 0$ as “no treatment,” noting that this choice is somewhat arbitrary as $D_{it} = 0$ can denote any reference treatment.

to change as well, namely in such a way that (4.1) holds. Consider for example the case where the outcome is the power consumption of a residential household (which we will study in detail in Chapter 5). This outcome is clearly subject to seasonalities. Assumption 11 asserts that any seasonal patterns in the household’s consumption are captured by the covariates. This can be achieved by controlling for the time of the year, the ambient temperature, and possibly other covariates in X_{it} . If, however, the overall level of consumption shifts because the household invested in more energy efficient appliances (and this investment is unobserved), then Assumption 11 would not be satisfied anymore.

Our goal is to estimate the *Conditional Individual Treatment Effect* (CITE) and the *Individual Treatment Effect* (ITE):

Definition 2 (Conditional Individual Treatment Effect).

$$\mu_i(x) := \mu_i^1(x) - \mu_i^0(x) = \mathbb{E}[Y_{it}^1 \mid X_{it} = x] - \mathbb{E}[Y_{it}^0 \mid X_{it} = x] \quad (4.2)$$

Definition 3 (Individual Treatment Effect).

$$\mu_i := \mathbb{E}[\mu_i(x)] = \mathbb{E}[Y_{it}^1] - \mathbb{E}[Y_{it}^0] \quad (4.3)$$

In addition to the (full) CITE from Definition 2, we may also be interested in the ITE conditional on only a subset of covariates, in particular if X_{it} is high-dimensional. Specifically, if $\check{X}_{it} \in \mathbb{R}^{n_{\check{x}}}$ is a sub-vector of X_{it} (i.e., with $n_{\check{x}} < n_x$), then it is natural to consider

$$\mu_i(\check{x}) := \mu_i^1(\check{x}) - \mu_i^0(\check{x}) = \mathbb{E}[Y_{it}^1 \mid \check{X}_{it} = \check{x}] - \mathbb{E}[Y_{it}^0 \mid \check{X}_{it} = \check{x}] \quad (4.4)$$

as the individual treatment effect conditional on the covariates \check{X} . It is important to note such a restricted CITE is only meaningful with respect to the particular joint distribution of X_{it} under the underlying data generating process. That is, while the full CITE (4.2) is *externally valid* (for the same individual i), this is not necessarily the case for the restricted CITE (4.4).⁵

Remark 1 (Consistent response under repeated treatments). *By considering a treatment effect as in Definition 3, we implicitly assume that the effect of treatment on an individual is consistent under repeated treatment exposure. This is a simplification that in some situations may not be warranted, in particular if we have reason to believe that the effect varies significantly with the number of treatment exposures (see for example the work on modeling the decaying effect of behavioral treatments on a subject’s attention (Caplin, 2016)). Nevertheless, estimating the effect under this simplifying assumption will still yield a relevant estimate, namely the average effect of treatment across treatment periods. Moreover, to model such non-consistent treatment effects explicitly, we could directly include the number of past treatments into the more general treatment model developed in Section 4.4.*

⁵In principle, we could of course still compute a restricted CITE under a different covariate distribution if we had the full CITE, simply by marginalizing with respect to the desired distribution.

The fundamental problem of causal inference (Holland, 1986) is that, at any time t , at most one of the two potential outcomes can be observed. Denote by

$$Y_{it} = (1 - D_{it})Y_{it}^0 + D_{it}Y_{it}^1 \quad (4.5)$$

the observed outcome at time t , and let $\mathcal{T}_i := \{t \in \mathcal{T} : D_{it} = 1\}$ denote the set of treatment periods for individual i . We will make the following simplifying assumption:

Assumption 12 (No missing observations). *For each period t , exactly one of the two potential outcomes is observed. That is, (4.5) is well defined for all sequences $(D_{it})_t$.*

We emphasize that we allow for lagged outcome and covariate observations in the vector X_{it} , which will naturally result in the X_{it} being serially correlated. For example, an autoregressive process of order k can be modeled by including past outcome variables $(Y_{it-k}, \dots, Y_{it-1})$ into X_{it} . Letting \tilde{X}_{it} denote the vector of exogenous covariates, we have

$$X_{it} = (Y_{it-k}, \dots, Y_{it-1}, \tilde{X}_{it,1}, \dots, \tilde{X}_{it,n_{\tilde{x}}}) \quad (4.6)$$

and so $n_x = k + n_{\tilde{x}}$. We will make the standard assumption that the exogenous covariates \tilde{X}_{it} are not affected by either treatment. Furthermore, we consider the treatment periods as part of the data and thus as given.⁶

4.3.1 A Generic Outcome Model

Consider the following generic outcome model:

$$Y_{it} = f_i^1(X_{it}) + \epsilon_{it} = f_i^0(X_{it}) + D_{it}[\mu_i(X_{it}) + \gamma_{it}] + \epsilon_{it} \quad (4.7)$$

where f_i^0 is the (unknown) conditional mean function of the potential outcome Y_{it}^0 in the absence of treatment (i.e. if $D_{it} = 0$), and $(\epsilon_{it})_t$ and $(\gamma_{it})_t$ are zero-mean stationary random processes representing the idiosyncratic noise in the outcome and the treatment effect. We further assume that the processes $(\epsilon_{it})_t$ and $(\gamma_{it})_t$ are independent (in particular, $\epsilon_{it} \perp \gamma_{i\tau}$ for all t, τ). Note that this allows for including lagged values of the outcome variable Y_{it} as elements of the covariate vector X_{it} . In particular, the model includes general autoregressive models with exogenous inputs. Moreover, we do also allow for the noise processes to be serially correlated.

It is easy to check that (4.7) satisfies Assumption 11. Indeed, as $Y_{it}^0 = f_i^0(X_{it}) + \epsilon_{it}$ we have $Y_{it}^0 \perp t \mid X_{it}$ by stationarity of ϵ_{it} . Similarly, as $Y_{it}^1 = f_i^0(X_{it}) + \mu(X_{it}) + \gamma_{it} + \epsilon_{it}$ it follows that $Y_{it}^1 \perp t \mid X_{it}$ by stationarity and independence of ϵ_{it} and γ_{it} .

⁶The question of how to design treatments assignments that are, in some sense, optimal (e.g. that minimize the MSE of the estimate) is, while intriguing, outside the scope of this thesis.

Remark 2 (Persistency of treatment effects). *Our model (4.7) assumes that the direct effect of treatment on the outcome persists only for a single period. In the future, we may also be interested in models that explicitly allow treatment during period t to have a direct effect on the outcome in periods $\tau > t$. However, note that even in the basic model (4.7) treatment effects in general may persist over multiple periods due to the potential autoregressive components of the affected output included in the covariates.*

The basic idea of our approach is to estimate the conditional mean function f_i^0 by fitting some model \hat{f}_i^0 in a first stage, and then compare the observations in treatment periods with their respective counterfactuals estimated by the model \hat{f}_i^0 in a second stage. For this reason we will speak of the overall estimation strategy as a two-stage estimation process.

There is a wide variety of methods available for obtaining a function approximation \hat{f}_i^0 of f_i^0 , ranging from classic (regularized) Least Squares on linear parametric models over non-parametric Kernel regression methods to modern non-parametric Machine Learning algorithms such as K -Nearest Neighbors Regression, Decision Tree Regression, Random Forests, Support Vector Regression or Neural Networks (see Hastie et al. (2009) for a discussion of various such methods). The choice of which models is “best” depends on the properties of f_i^0 , the particular application at hand, and a large number of other factors, such as accuracy, robustness, simplicity, speed, etc.. We are interested in making general statements about our estimates for the treatment effect without restricting ourselves to a specific model or a specific class of models, and hence remain agnostic about the kind of model used.⁷ We will use the terminology of “training” a model and “predicting” an outcome as is standard in the Machine Learning literature.

Clearly, as we can only observe the outcome Y_{it} and the covariates X_{it} , we need to ensure that we fit our model \hat{f}_i^0 on a data set that does not include treatment periods, otherwise we will end up with a biased model. This is an important point and gets at the heart of the distinction between the standard use of Machine Learning methods for prediction and what we require for our purposes of estimation. To put it in the language of Machine Learning: in the context of performing causal inference some of the features are special and have to be treated differently. But the mere split in treatment and non-treatment periods is not our only concern: Since the noise processes may be serially correlated, we also need to ensure to minimize correlation between model errors and outcome observations. We will discuss these challenges in much more depth later in this chapter. For the purpose of this section we assume that we have a *training data set* $\mathcal{D}_{i,\text{tr}} \subset \mathcal{D}_i$ and an *estimation data set* $\mathcal{D}_{i,\text{est}} \subset \mathcal{D}_i$ such that $\mathcal{D}_{i,\text{tr}} \cap \mathcal{D}_{i,\text{est}} = \emptyset$. Figure 4.1 illustrates how this can be achieved in a setting where the serial correlation between observations decays rapidly: In such a situation it may be enough to exclude data around the treatment observations in order to ensure the data in $\mathcal{D}_{i,\text{tr}}$ and $\mathcal{D}_{i,\text{est}}$ are “approximately independent.”

⁷If we consider linear models for both stages, this approach in essence will do something very similar to a single regression model with dummy variables for the treatment that is common in analyzing panel data in econometrics. However, our two-stage approach is much more flexible in that it allows for a broad range of possible models for both stages.

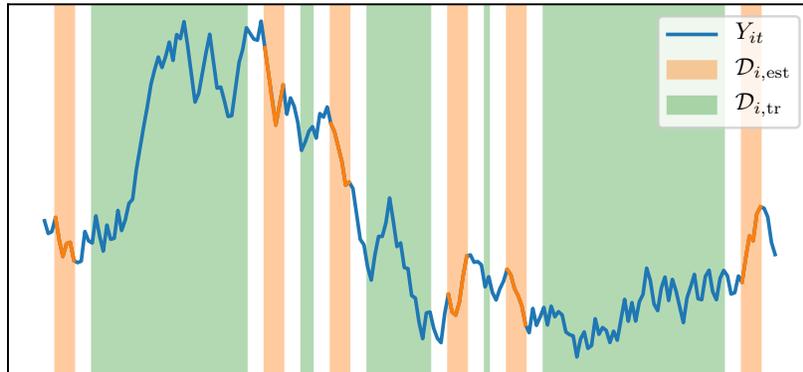


Figure 4.1: Splitting observations into training and estimation data sets (illustration)

We train the model \hat{f}_i^0 on the training data set $\mathcal{D}_{i,\text{tr}}$ by regressing the observations Y_{it}^0 onto the covariates X_{it} . Let

$$e_i(X_{it}) := \hat{f}_i^0(X_{it}) - f_i^0(X_{it})$$

denote the model error. Note that as the errors ϵ_{it} are unobserved, what the model really estimates is $\mathbb{E}[Y_{it}^0 | X_{it}] = f_i^0(X_{it}) + \mathbb{E}[\epsilon_{it} | X_{it}]$. In particular, if ϵ_{it} and X_{it} are correlated (which may be the case due to lagged output values in the covariate vector and serial correlation of the noise), then $e_i(X_{it}) = \mathbb{E}[\epsilon_{it} | X_{it}]$ for a perfect model \hat{f}_i^0 .

For each $t \in \mathcal{T}_i$ we construct the quantity

$$\delta_{it} := Y_{it} - \hat{f}_i^0(X_{it}) = f_i^0(X_{it}) - \hat{f}_i^0(X_{it}) + \mu_i(X_{it}) + \epsilon_{it} + \gamma_{it} \quad (4.8)$$

which can be viewed as a (noisy) one-sample estimate of the CITE of individual i at $x = X_{it}$. It is important to realize that the error

$$\kappa_{it} := \delta_{it} - \mu_i(X_{it}) = f_i^0(X_{it}) - \hat{f}_i^0(X_{it}) + \epsilon_{it} + \gamma_{it} \quad (4.9)$$

in general is heteroscedastic and serially correlated. Its mean is the model bias at the value of the covariate vector. In particular, for given X_{it} we have

$$\mathbb{E}[\kappa_{it} | X_{it}] = f_i^0(X_{it}) - \mathbb{E}_{\mathcal{D}_{i,\text{tr}}}[\hat{f}_i^0(X_{it})] = -\text{Bias } \hat{f}_i^0(X_{it})$$

For every fixed x , the model bias $\text{Bias } \hat{f}_i^0(x)$ is a random variable, whose variation is due to the randomness of the training data and the model fitting process.⁸ Hence $\text{Bias } \hat{f}_i^0$ is a random function on \mathbb{R}^{n_x} .

⁸Some regression methods, for example Random Forest Regression, are non-deterministic in nature and hence introduce additional variation.

We can also obtain an expression for the covariance between errors:

$$\begin{aligned}
 & \text{Cov}(\kappa_{it}, \kappa_{i\tau} \mid X_{it}, X_{i\tau}) \\
 &= \mathbb{E} \left[(f_i^0(X_{it}) - \hat{f}_i^0(X_{it}) + \epsilon_{it} + \gamma_{it})(f_i^0(X_{i\tau}) - \hat{f}_i^0(X_{i\tau}) + \epsilon_{i\tau} + \gamma_{i\tau}) \mid X_{it}, X_{i\tau} \right] \\
 &\quad - \mathbb{E}[f_i^0(X_{it}) - \hat{f}_i^0(X_{it}) + \epsilon_{it} + \gamma_{it} \mid X_{it}] \mathbb{E}[f_i^0(X_{i\tau}) - \hat{f}_i^0(X_{i\tau}) + \epsilon_{i\tau} + \gamma_{i\tau} \mid X_{i\tau}] \\
 &= \mathbb{E} \left[(\hat{f}_i^0(X_{it}) - \mathbb{E}[\hat{f}_i^0(X_{it}) \mid X_{it}] + \mathbb{E}[\hat{f}_i^0(X_{it}) \mid X_{it}] - f_i^0(X_{it})) \right. \\
 &\quad \left. (\hat{f}_i^0(X_{i\tau}) - \mathbb{E}[\hat{f}_i^0(X_{i\tau}) \mid X_{i\tau}] + \mathbb{E}[\hat{f}_i^0(X_{i\tau}) \mid X_{i\tau}] - f_i^0(X_{i\tau})) \mid X_{it}, X_{i\tau} \right] \\
 &\quad - \mathbb{E} \left[(\epsilon_{it} + \gamma_{it})(\hat{f}_i^0(X_{i\tau}) - f_i^0(X_{i\tau})) \mid X_{i\tau} \right] - \mathbb{E} \left[(\epsilon_{i\tau} + \gamma_{i\tau})(\hat{f}_i^0(X_{it}) - f_i^0(X_{it})) \mid X_{it} \right] \\
 &\quad + \mathbb{E}[(\epsilon_{it} + \gamma_{it})(\epsilon_{i\tau} + \gamma_{i\tau})] - \text{Bias } \hat{f}_i^0(X_{it}) \text{ Bias } \hat{f}_i^0(X_{i\tau}) \\
 &= \text{Cov } \hat{f}_i^0(X_{it}, X_{i\tau}) + \text{Cov}(\epsilon_{it}, \epsilon_{i\tau}) + \text{Cov}(\gamma_{it}, \gamma_{i\tau}) \\
 &\quad - \text{Cov}(\epsilon_{it} + \gamma_{it}, \text{Bias } \hat{f}_i^0(X_{i\tau}) \mid X_{i\tau}) - \text{Cov}(\epsilon_{i\tau} + \gamma_{i\tau}, \text{Bias } \hat{f}_i^0(X_{it}) \mid X_{it})
 \end{aligned}$$

where

$$\text{Cov } \hat{f}_i^0(X_{it}, X_{i\tau}) := \mathbb{E}_{\mathcal{D}_{i,\text{tr}}} \left[(\hat{f}_i^0(X_{it}) - \mathbb{E}_{\mathcal{D}_{i,\text{tr}}}[\hat{f}_i^0(X_{it})]) (\hat{f}_i^0(X_{i\tau}) - \mathbb{E}_{\mathcal{D}_{i,\text{tr}}}[\hat{f}_i^0(X_{i\tau})]) \right]$$

is the covariance between the model predictions $\hat{f}_i^0(X_{it})$ and $\hat{f}_i^0(X_{i\tau})$, which depends on the degree of correlation between observations in the training data.

The complexity of the above expressions, even for just a single error κ_{it} , illustrate the challenges in the theoretical analysis of our general model.

Example 1 (Linear Model). *If we use a linear model for \hat{f}_i^0 , that is, if $\hat{f}_i^0(X) = \hat{\eta}_i^\top X$, then it is straightforward to show that $\text{Cov } \hat{f}_i^0(X_{it}, X_{i\tau}) = X_{it} \text{Cov}(\hat{\eta}_i) X_{i\tau}^\top$, where $\text{Cov}(\hat{\eta}_i)$ is the covariance matrix of the estimate $\hat{\eta}_i$ and is given by $\text{Cov}(\hat{\eta}_i) = \mathbb{E}_{\mathcal{D}_{i,\text{tr}}}[(\hat{\eta}_i - \mathbb{E}_{\mathcal{D}_{i,\text{tr}}}[\hat{\eta}_i])(\hat{\eta}_i - \mathbb{E}_{\mathcal{D}_{i,\text{tr}}}[\hat{\eta}_i])^\top]$. If $f_i^0(X) = \eta_i^\top X$, i.e., if the conditional mean function f_i^0 is indeed linear in X_{it} , and we use OLS to compute $\hat{\eta}_i$, then $\text{Bias } \hat{f}_i^0(X_{it}) = 0$, hence δ_{it} is an unbiased estimator of $\mu_i(X_{it})$ for all values of X_{it} . Moreover, a straightforward calculation yields that $\text{Cov}(\hat{\eta}_i) = \mathbb{E}_{\mathcal{D}_{i,\text{tr}}}[(Z_i Z_i^\top)^{-1} Z_i \epsilon_i \epsilon_i^\top Z_i^\top (Z_i Z_i^\top)^{-1}]$, where $Z_i = (X_{i\theta})_{\theta \in \mathcal{D}_{i,\text{tr}}}$ is the matrix obtained by stacking the covariate vectors for the observations in the training set $\mathcal{D}_{i,\text{tr}}$, and $\epsilon_i = (\epsilon_{i\theta})_{\theta \in \mathcal{D}_{i,\text{tr}}}$. If we further assume that training and estimation data set are independent, then $\mathbb{E}[\kappa_{it}] = 0$ and*

$$\text{Cov}(\kappa_{it}, \kappa_{i\tau}) = \mathbb{E}_{\mathcal{D}_{i,\text{est}}} \left[X_{it} \mathbb{E}_{\mathcal{D}_{i,\text{tr}}} [(Z_i Z_i^\top)^{-1} Z_i \epsilon_i \epsilon_i^\top Z_i^\top (Z_i Z_i^\top)^{-1}] X_{i\tau}^\top + \epsilon_{it} \epsilon_{i\tau} + \gamma_{it} \gamma_{i\tau} \right]$$

At this point we should remind ourselves of our primary goal, which is to obtain an estimate $\hat{\mu}_i$ of the CITE μ_i . Heuristically speaking, if the errors κ_{it} were zero-mean independent random variables, then simply regressing the $\{\delta_{it}\}_{t \in \mathcal{T}_i}$ on $\{X_{it}\}_{t \in \mathcal{T}_i}$ would, in principle,⁹ yield

⁹That is, under a correctly specified parametric model, or non-parametric regression with appropriate bandwidth selection.

an estimator of the CITE that is consistent¹⁰ on the support of the distribution of the X_{it} . However, in our case the errors will generally be correlated (due to both the inherent serial correlation of the noise terms ϵ_{it} and γ_{it} in the output observations, as well as the correlation of the model predictions), and hence such an estimator may not be consistent.¹¹

The problem of inconsistency of standard errors due to serial correlation of observation errors in time series has been discussed by a number of authors, including for example Bertrand et al. (2004); Hansen (2007); Angrist and Pischke (2008). Cochrane and Orcutt (1949) investigate the efficiency loss caused by observations with autocorrelated errors and develop a procedure to regain the lost efficiency. Beach and MacKinnon (1978) develop a similar maximum likelihood procedure for regression with first-order autocorrelated errors, which essentially tries to estimate the correlation coefficient of the errors and transform the observations accordingly. Their method is based on alternating maximization to maximize the full likelihood function. Nickell (1981) examines biases in dynamic fixed effects models. Specifically, he derives analytical expressions of the bias of the OLS estimates in panel data for linear models including a first-order autoregressive term. However, he is primarily concerned with conventional economic panel, for which the number of observed time periods per individual is typically quite small. In such a setting, he states that “A typical set of panel data has a rather large number of individuals and a rather small number of time periods and it is in just these circumstances that the biases, which are essentially of the Hurwicz type, are most serious. The fact that they will go to zero when the number of time periods becomes infinite is scant consolation.” As we, in contrast, focus on high-frequency data sets, there is reason to hope that in our case the consolation offered by large numbers of observations per individual is not quite as scant. Intuitively, if the correlation between observations decays over time,¹² and if treatment periods are sufficiently spread out, then the bias may be small enough to yield a useful estimator. We will make this intuition precise in the following sections.

When talking about bias and variance of an estimator in our setting, it is essential to understand that the regression function \hat{f}_i^0 in general is not fixed a priori, but the result of the algorithm fitting the model on the training data during the first stage estimation process. The additional variation introduced through this step needs to be accounted for in the two-stage estimate. In the following, we will use the terms *ex-interim bias* and *ex-interim variance* to denote bias and variance of the estimator instantiated with the particular realization \hat{f}_i^0 of the model fitted on the training set. Similarly, we will refer to the bias and variance under the distributions of both training and estimation data as the *ex-ante bias* and *ex-ante variance*, respectively.

¹⁰Here consistency is with respect to the number $|\mathcal{T}_i|$ of treatment observations.

¹¹In fact, Angrist and Pischke (2008) point out that in some situations there may be not consistent estimator if the noise is serially correlated.

¹²More formally: If for each t the autocorrelation $R_{AC}(t, \tau)$ of the stochastic process goes to zero as $|\tau - t|$ goes to infinity.

4.3.2 A Sample Mean Estimator for the ITE

We start with estimating the ITE (4.3) from Definition 3. In the simplest case, in which the treatment effect is indeed homogeneous across covariates (i.e. $\mu_i(X_{it}) = \mu_i$), the estimate has a straightforward interpretation.¹³ In practice, however, a homogeneous ITE may often not be a very reasonable assumption. For example, in the case of estimating the change in electricity consumption of a DR participant in response to a DR event, one would assume the effect for a participant with an AC unit to generally be higher when the AC is running, which is usually during times with higher ambient temperature. So one would not expect a homogeneous effect across ambient temperature. Nevertheless, even if the underlying treatment effect is heterogeneous, estimating the marginal ITE still provides a useful and relevant summary statistic.¹⁴

Arguably, the most straightforward estimator for the ITE μ_i is given by the sample mean of the δ_{it} , i.e.,

$$\hat{\mu}_i = \frac{\sum_t D_{it} \delta_{it}}{\sum_t D_{it}} = \frac{1}{|\mathcal{T}_i|} \sum_{t \in \mathcal{T}_i} \delta_{it} \quad (4.10)$$

The estimation error of this estimator is

$$\eta_i := \hat{\mu}_i - \mu_i = \frac{1}{|\mathcal{T}_i|} \sum_{t \in \mathcal{T}_i} \kappa_{it} = -\frac{1}{|\mathcal{T}_i|} \sum_{t \in \mathcal{T}_i} (e_i(X_{it}) - \epsilon_{it} - \gamma_{it}) \quad (4.11)$$

For statistical inference, we are interested in understanding the distribution of this estimation error. Deriving finite sample properties of the error is complicated by the assumption of a generic model \hat{f}_i^0 as well as by the potential correlation of the covariates and noise between different treatment times. We will provide some general results on this in the following sections.

Basic Limiting Behavior

It is easy to see from the classical central limit theorem that in the special case when $e_i(X_{it})$, ϵ_{it} and γ_{it} are all i.i.d. with finite variance, then η_i is asymptotically normal with mean $-\mathbb{E}_{\mathcal{D}_{i,\text{est}}} [e_i(X_{it})]$. Unfortunately these conditions rarely hold in practice. In particular, not only the error processes, but also the covariate process will often be serially correlated. A useful generalization can be obtained if the involved random processes are weakly coupled. To this end, recall the following definition:

¹³Conceptually, it is straightforward to extend this estimator to problems with a finite set of covariates simply by stratifying across the different covariate values. This, however, may quickly result in unacceptably high variance if the covariate space has many elements.

¹⁴See also Remark 1 in the previous section.

Definition 4 (Strong Mixing). *Let W_1, W_2, \dots be a random process and let*

$$\alpha_n := \sup \left\{ \left| \mathbb{P}(A \cap B) - \mathbb{P}(A)\mathbb{P}(B) \right| : A \in \sigma(\{W_\tau\}_{\tau=1}^t), B \in \sigma(\{W_\tau\}_{\tau=t+n}^\infty), t \geq 1 \right\}$$

where \mathbb{P} is the probability measure on the σ -field generated by $(W_t)_{t \geq 1}$. The process W_1, W_2, \dots is said to be strong α -mixing if $\alpha_n \rightarrow 0$ as $n \rightarrow \infty$.

With this we can state the following result:

Proposition 2. *Let $Z_{it} = e_i(X_{it}) - \epsilon_{it} - \gamma_{it}$ and suppose that Z_{i1}, Z_{i2}, \dots is stationary and strong α -mixing with $\alpha_n = \mathcal{O}(n^{-5})$, and that $\mathbb{E}_{\mathcal{D}_{i,\text{est}}}[Z_{it}] = b$ and $\mathbb{E}_{\mathcal{D}_{i,\text{est}}}[Z_{it}^2] < \infty$. Then $\text{Var}(\eta) \rightarrow \frac{\sigma^2}{|\mathcal{T}_i|}$ with $\sigma^2 = \sum_{t=1}^\infty \mathbb{E}_{\mathcal{D}_{i,\text{est}}}[Z_{i1}Z_{it}]$. If $\sigma > 0$, then $\sqrt{|\mathcal{T}_i|} \eta_i \xrightarrow{d} \mathcal{N}(b, \sigma^2)$.*

Proof. Up to some small modifications, this is essentially a restatement of Theorem 27.4 in Billingsley (1995). \square

Proposition 2 is reassuring in that it states what we would hope for as a basic property of the estimate: If the involved processes are stationary and their coupling decays over time, then the estimator will be asymptotically normally distributed with a variance decaying at the standard rate of $1/|\mathcal{T}_i|$. Of course the challenge is to verify, for general models \hat{f}_i^0 , that the necessary assumptions are satisfied. Moreover, Proposition 2 also shows that any non-zero mean in the estimation error Z_{it} will determine the mean of the limiting distribution. A general challenge is that the different elements of Z_{it} may not be independent. In the special case that they are, we have the following:

Corollary 1. *Suppose that the processes $(X_{it})_{t \geq 1}$, $(\epsilon_{it})_{t \geq 1}$ and $(\gamma_{it})_{t \geq 1}$ are independent, stationary and α -mixing with $\alpha_n = \mathcal{O}(n^{-5})$. Suppose further that $\mathbb{E}_{\mathcal{D}_{i,\text{est}}}[e_i(X_{it})] = b$ and $\mathbb{E}_{\mathcal{D}_{i,\text{est}}}[e_i(X_{it})^2 + \epsilon_{it}^2 + \gamma_{it}^2] < \infty$. Then the conclusion of Proposition 2 holds with $\sigma^2 = \sum_{t=1}^\infty \mathbb{E}_{\mathcal{D}_{i,\text{est}}}[e_i(X_{i1})e_i(X_{it}) + \epsilon_{i1}\epsilon_{it} + \gamma_{i1}\gamma_{it}]$.*

Proof. Note that it follows immediately from Definition 4 that any Borel-measurable function of an α -mixing processes is also α -mixing with no greater mixing coefficients. This observation, in combination with the independence assumption, implies that the result follows Theorem 5.1 in Bradley (2005). \square

Unfortunately, for the independence assumption in Corollary 1 to hold, the covariate vector X_{it} may not contain lagged values of the outcome variable Y_{it} .¹⁵ At this level of generality, it appears hard to obtain more general results that would allow for this kind of interdependence without making additional assumptions on the conditional mean function. In the following, we instead look explicitly at how the properties of the first-stage model \hat{f}_i^0 appear in bias and variance of the estimate.

¹⁵In fact, the conditions of Corollary 1 may be weakened to the process $(e_i(X_{it}))_{t \geq 1}$, which might be independent of the error processes even if the original process $(X_{it})_{t \geq 1}$ is not.

Ex-Interim Bias and Variance

As a first step, suppose that we have trained a model \hat{f}_i^0 on the training data set. With respect to this model, the estimator has a certain bias and variance — we refer to them as *ex-interim* bias and variance. We can show the following:

Lemma 1 (Ex-Interim Bias and Variance of $\hat{\mu}_i$). *Fix the model \hat{f}_i^0 and the set of treatment periods \mathcal{T}_i . Conditional on this selection, we have*

$$\begin{aligned} \text{Bias}(\hat{\mu}_i | \hat{f}_i^0, \mathcal{T}_i) &= \mathbb{E}_{\mathcal{D}_{i,\text{est}}} [\hat{\mu}_i | \hat{f}_i^0, \mathcal{T}_i] - \mu_i \\ &= -\frac{1}{|\mathcal{T}_i|} \sum_{t \in \mathcal{T}_i} \left(\mathbb{E}_{\mathcal{D}_{i,\text{est}}} [e_i(X_{it}) | D_{it}=1] - \mathbb{E}[\epsilon_{it} | D_{it}=1] - \mathbb{E}[\gamma_{it} | D_{it}=1] \right) \end{aligned} \quad (4.12a)$$

$$\begin{aligned} \text{Var}(\hat{\mu}_i | \hat{f}_i^0, \mathcal{T}_i) &= \mathbb{E}_{\mathcal{D}_{i,\text{est}}} \left[\left(\hat{\mu}_i - \mathbb{E}_{\mathcal{D}_{i,\text{est}}} [\hat{\mu}_i | \hat{f}_i^0, \mathcal{T}_i] \right)^2 | \hat{f}_i^0, \mathcal{T}_i \right] \\ &= \frac{1}{|\mathcal{T}_i|^2} \sum_{t \in \mathcal{T}_i} \left(\text{Var}(e_i(X_{it}) | \hat{f}_i^0, D_{it}=1) + \text{Var}(\epsilon_{it} | D_{it}=1) + \text{Var}(\gamma_{it} | D_{it}=1) \right. \\ &\quad \left. + (\mathbb{E}[\epsilon_{it} | D_{it}=1] + \mathbb{E}[\gamma_{it} | D_{it}=1])^2 \right. \\ &\quad \left. - 2 \text{Cov}(\epsilon_{it} + \gamma_{it}, e_i(X_{it}) | \hat{f}_i^0, D_{it}=1) \right) \\ &\quad + \frac{2}{|\mathcal{T}_i|^2} \sum_{t \in \mathcal{T}_i} \sum_{\substack{\tau \in \mathcal{T}_i \\ \tau > t}} \left(\text{Cov}(\epsilon_{it}, \epsilon_{i\tau} | D_{it}=D_{i\tau}=1) + \text{Cov}(\gamma_{it}, \gamma_{i\tau} | D_{it}=D_{i\tau}=1) \right. \\ &\quad \left. + \text{Cov}(e_i(X_{it}), e_i(X_{i\tau}) | \hat{f}_i^0, D_{it}=D_{i\tau}=1) \right. \\ &\quad \left. - \text{Cov}(\epsilon_{it} + \gamma_{it}, e_i(X_{i\tau}) | \hat{f}_i^0, D_{it}=D_{i\tau}=1) \right. \\ &\quad \left. + \mathbb{E}[\epsilon_{it} + \gamma_{it} | D_{it}=1] \mathbb{E}[\epsilon_{i\tau} + \gamma_{i\tau} | D_{i\tau}=1] \right) \end{aligned} \quad (4.12b)$$

Proof. See Appendix C.2 □

From Lemma 1 we see that ex-interim bias and variance are driven by two causes: The model error and the idiosyncratic noise. The former is unavoidable, and can be mitigated by improving the fit of the model \hat{f}_i^0 (i.e. by using an appropriate model specification and a large training data set). The latter may be relatively benign and only increase the variance of the estimator. However, under *confounded* treatment assignment, i.e., when the treatment assignment is correlated with the noise, it may be much more problematic, since it would be adding systematic bias to the estimator.¹⁶ Under *unconfounded treatment assignment*, i.e. if treatment assignment and noise are independent, the situation simplifies:

¹⁶This situation is similar to the one encountered when estimating the average treatment effect in non-experimental settings, where treatment assignment may be correlated with individual-level covariates, which in turn may be correlated with the effect size. In that case, one typically makes a “selection on observables” assumption, which requires that treatment assignment and potential outcomes are conditionally independent given the observed covariates, which in effect allows to re-weight observations appropriately to obtain an unbiased estimate of the ATE.

Corollary 2. *Suppose that $\epsilon_{it}, \gamma_{it} \perp\!\!\!\perp D_{it}$. Then ex-interim bias and variance in Lemma 1 reduce to*

$$\text{Bias}(\hat{\mu}_i | \hat{f}_i^0, \mathcal{T}_i) = -\frac{1}{|\mathcal{T}_i|} \sum_{t \in \mathcal{T}_i} \mathbb{E}_{\mathcal{D}_{i,\text{est}}} [e_i(X_{it}) | D_{it}=1] \quad (4.13a)$$

$$\begin{aligned} \text{Var}(\hat{\mu}_i | \hat{f}_i^0, \mathcal{T}_i) &= \frac{1}{|\mathcal{T}_i|^2} \sum_{t \in \mathcal{T}_i} \left(\text{Var}(e_i(X_{it}) | \hat{f}_i^0, D_{it}=1) + \text{Var}(\epsilon_{it}) + \text{Var}(\gamma_{it}) \right. \\ &\quad \left. - 2 \text{Cov}(\epsilon_{it} + \gamma_{it}, e_i(X_{it}) | \hat{f}_i^0, D_{it}=1) \right) \\ &\quad + \frac{2}{|\mathcal{T}_i|^2} \sum_{t \in \mathcal{T}_i} \sum_{\substack{\tau \in \mathcal{T}_i \\ \tau > t}} \left(\text{Cov}(\epsilon_{it}, \epsilon_{i\tau}) + \text{Cov}(\gamma_{it}, \gamma_{i\tau}) \right. \\ &\quad \left. + \text{Cov}(e_i(X_{it}), e_i(X_{i\tau}) | \hat{f}_i^0, D_{it}=D_{i\tau}=1) \right. \\ &\quad \left. - \text{Cov}(\epsilon_{it} + \gamma_{it}, e_i(X_{i\tau}) | \hat{f}_i^0, D_{i\tau}=1) \right) \end{aligned} \quad (4.13b)$$

Proof. If $\epsilon_{it}, \gamma_{it} \perp\!\!\!\perp D_{it}$ we have $\mathbb{E}[\epsilon_{it} | D_{it}] = \mathbb{E}[\epsilon_{it}] = 0$ and $\mathbb{E}[\gamma_{it} | D_{it}] = \mathbb{E}[\gamma_{it}] = 0$. \square

Observe that while $\epsilon_{it}, \gamma_{it} \not\perp\!\!\!\perp X_{it}$ in general (this is because X_{it} may contain lagged values of the output Y_{it}), the condition that $\epsilon_{it}, \gamma_{it} \perp\!\!\!\perp D_{it}$ may still be satisfied if treatment assignment is conditional on a subset of the covariates that is independent from the error processes — for instance, the exogenous covariates \tilde{X}_{it} in (4.6) that do not include information derived from past output values, as stated in Lemma 2 below. A more general correlation between treatment assignment and the whole covariate vector X_{it} is problematic if the errors ϵ_{it} and γ_{it} are serially correlated. For example, suppose treatment D_{it} were assigned if the most recent output Y_{it-1} fell above a certain threshold, and that the ϵ_{it} were positively correlated. Then $\mathbb{E}[\epsilon_{it} | D_{it}=1] > 0$, which by (4.12a) will result in a positively biased estimate of the treatment effect.

Lemma 2. *Suppose that $\tilde{X}_{it} \in \mathbb{R}^{n_x}$ is a sub-vector of X_{it} such that $\epsilon_{it}, \gamma_{it} \perp\!\!\!\perp \tilde{X}_{i\tau}$ for all t, τ , and that $\epsilon_{it}, \gamma_{it} \perp\!\!\!\perp D_{it} | \tilde{X}_{it}$. Then $\epsilon_{it}, \gamma_{it} \perp\!\!\!\perp D_{it}$.*

Proof. This is straightforward and follows from the fact that $\epsilon_{it} \perp\!\!\!\perp D_{it} | \tilde{X}_{it}$ and $\epsilon_{it} \perp\!\!\!\perp \tilde{X}_{it}$ together imply that $\epsilon_{it} \perp\!\!\!\perp D_{it}, \tilde{X}_{it}$ in combination with the assumption that $\epsilon_{it} \perp\!\!\!\perp \gamma_{it}$. \square

Under the assumption that $\epsilon_{it}, \gamma_{it} \perp\!\!\!\perp D_{it}$, Corollary 2 shows that if the expected model error for the model \hat{f}_i^0 under the distribution of the data in the estimation data set is zero, then the estimator is ex-interim unbiased.¹⁷ This is not very surprising and in line with our intuition. A somewhat more interesting observation is that the ex-interim variance increases with increasing *variation* in the model error across the covariate space \mathcal{X} , but is unaffected by the magnitude of the model error. This is also rather intuitive, as adding any constant to the model \hat{f}_i^0 does not affect the ex-interim variance of the estimator. Furthermore, as can

¹⁷Note that the estimator may be ex-interim unbiased also for imperfect models \hat{f}_i^0 , for example if the expected error is zero for every $t \in \mathcal{T}_i$. In principle, it may also be ex-interim unbiased if non-zero expected errors happen to cancel out under the joint distribution of the estimation data over the treatment times.

be seen from the second to last term in (4.13b), the variation in the model error has a particularly strong effect in regions of \mathcal{X} that are “high-probability” regions¹⁸ under treatment assignment. Conversely, model error outside the support of the estimation data set does not affect the properties of the estimator.

Ex-Ante Bias and Variance

When considering the full estimation strategy, the model error for a given value of X_{it} is now itself a random variable, as the model \hat{f}_i^0 is fitted on the (random) training data. Moreover, the model fitting process itself is potentially non-deterministic (see footnote 8 on page 81). If the training and estimation data set are independent (we will discuss this condition in more detail later), we can still obtain expressions for the bias and variance of $\hat{\mu}_i$:

Lemma 3 (Ex-Ante Bias and Variance of $\hat{\mu}_i$). *Fix the set of treatment periods \mathcal{T}_i and suppose that $\mathcal{D}_{i,\text{est}} \perp\!\!\!\perp \mathcal{D}_{i,\text{tr}}$. Then*

$$\text{Bias}(\hat{\mu}_i | \mathcal{T}_i) = -\frac{1}{|\mathcal{T}_i|} \sum_{t \in \mathcal{T}_i} \left(\mathbb{E}_{\mathcal{D}_{i,\text{est}}} [\text{Bias } \hat{f}_i^0(X_{it}) | D_{it}=1] - \mathbb{E}[\epsilon_{it} | D_{it}=1] - \mathbb{E}[\gamma_{it} | D_{it}=1] \right) \quad (4.14a)$$

$$\begin{aligned} \text{Var}(\hat{\mu}_i | \mathcal{T}_i) &= \frac{1}{|\mathcal{T}_i|^2} \sum_{t \in \mathcal{T}_i} \left(\text{Var}(\epsilon_{it} | D_{it}=1) + \text{Var}(\gamma_{it} | D_{it}=1) + \mathbb{E}[\epsilon_{it} + \gamma_{it} | D_{it}=1]^2 \right. \\ &\quad + \mathbb{E}_{\mathcal{D}_{i,\text{est}}} [\text{Var } \hat{f}_i^0(X_{it}) | D_{it}=1] + \text{Var}_{\mathcal{D}_{i,\text{est}}} (\text{Bias } \hat{f}_i^0(X_{it}) | D_{it}=1) \\ &\quad \left. - 2 \text{Cov}_{\mathcal{D}_{i,\text{est}}} (\epsilon_{it} + \gamma_{it}, \text{Bias } \hat{f}_i^0(X_{it}) | D_{it}=1) \right) \\ &\quad + \frac{2}{|\mathcal{T}_i|^2} \sum_{t \in \mathcal{T}_i} \sum_{\substack{\tau \in \mathcal{T}_i \\ \tau > t}} \left(\text{Cov}(\epsilon_{it}, \epsilon_{i\tau} | D_{it}=D_{i\tau}=1) + \text{Cov}(\gamma_{it}, \gamma_{i\tau} | D_{it}=D_{i\tau}=1) \right. \\ &\quad + \mathbb{E}[\epsilon_{it} + \gamma_{it} | D_{it}=1] \mathbb{E}[\epsilon_{i\tau} + \gamma_{i\tau} | D_{i\tau}=1] \\ &\quad + \mathbb{E}_{\mathcal{D}_{i,\text{est}}} [\text{Cov } \hat{f}_i^0(X_{it}, X_{i\tau}) | D_{it}=D_{i\tau}=1] \\ &\quad + \text{Cov}_{\mathcal{D}_{i,\text{est}}} (\text{Bias } \hat{f}_i^0(X_{it}), \text{Bias } \hat{f}_i^0(X_{i\tau}) | D_{it}=D_{i\tau}=1) \\ &\quad \left. - \text{Cov}_{\mathcal{D}_{i,\text{est}}} (\epsilon_{it} + \gamma_{it}, \text{Bias } \hat{f}_i^0(X_{i\tau}) | D_{i\tau}=1) \right) \end{aligned} \quad (4.14b)$$

where

$$\begin{aligned} \text{Var } \hat{f}_i^0(X_{it}) &:= \mathbb{E}_{\mathcal{D}_{i,\text{tr}}} \left[(\hat{f}_i^0(X_{it}) - \mathbb{E}_{\mathcal{D}_{i,\text{tr}}} [\hat{f}_i^0(X_{it})])^2 \right] \\ \text{Cov } \hat{f}_i^0(X_{it}, X_{i\tau}) &:= \mathbb{E}_{\mathcal{D}_{i,\text{tr}}} \left[(\hat{f}_i^0(X_{it}) - \mathbb{E}_{\mathcal{D}_{i,\text{tr}}} [\hat{f}_i^0(X_{it})]) (\hat{f}_i^0(X_{i\tau}) - \mathbb{E}_{\mathcal{D}_{i,\text{tr}}} [\hat{f}_i^0(X_{i\tau})]) \right] \end{aligned}$$

¹⁸i.e. regions of \mathcal{X} with a lot of probability mass in the joint distribution of the estimation data under the treatment assignment mechanism.

Proof. See Appendix C.2. \square

As with the ex-interim bias and variance, things simplify if treatment assignment and errors are independent:

Corollary 3. *Suppose that $\epsilon_{it}, \gamma_{it} \perp\!\!\!\perp D_{it}$. Then ex-ante bias and variance are given by*

$$\text{Bias}(\hat{\mu}_i | \mathcal{T}_i) = -\frac{1}{|\mathcal{T}_i|} \sum_{t \in \mathcal{T}_i} \mathbb{E}_{\mathcal{D}_{i,\text{est}}} [\text{Bias } \hat{f}_i^0(X_{it}) | D_{it}=1] \quad (4.15a)$$

$$\begin{aligned} \text{Var}(\hat{\mu}_i | \mathcal{T}_i) = & \frac{1}{|\mathcal{T}_i|^2} \sum_{t \in \mathcal{T}_i} \left(\text{Var}(\epsilon_{it}) + \text{Var}(\gamma_{it}) + \mathbb{E}_{\mathcal{D}_{i,\text{est}}} [\text{Var } \hat{f}_i^0(X_{it}) | D_{it}=1] \right. \\ & + \text{Var}_{\mathcal{D}_{i,\text{est}}} (\text{Bias } \hat{f}_i^0(X_{it}) | D_{it}=1) \\ & \left. - 2 \text{Cov}_{\mathcal{D}_{i,\text{est}}} (\epsilon_{it} + \gamma_{it}, \text{Bias } \hat{f}_i^0(X_{it}) | D_{it}=1) \right) \\ & + \frac{2}{|\mathcal{T}_i|^2} \sum_{t \in \mathcal{T}_i} \sum_{\substack{\tau \in \mathcal{T}_i \\ \tau > t}} \left(\text{Cov}(\epsilon_{it}, \epsilon_{i\tau}) + \text{Cov}(\gamma_{it}, \gamma_{i\tau}) \right. \\ & + \mathbb{E}_{\mathcal{D}_{i,\text{est}}} [\text{Cov } \hat{f}_i^0(X_{it}, X_{i\tau}) | D_{it}=D_{i\tau}=1] \\ & + \text{Cov}_{\mathcal{D}_{i,\text{est}}} (\text{Bias } \hat{f}_i^0(X_{it}), \text{Bias } \hat{f}_i^0(X_{i\tau}) | D_{it}=D_{i\tau}=1) \\ & \left. - \text{Cov}_{\mathcal{D}_{i,\text{est}}} (\epsilon_{it} + \gamma_{it}, \text{Bias } \hat{f}_i^0(X_{i\tau}) | D_{i\tau}=1) \right) \end{aligned} \quad (4.15b)$$

Lemma 3 and Corollary 3 provide insight into how the bias-variance tradeoff of the model used in the first stage affects bias and variance of the full two-stage estimator $\hat{\mu}_i$. On the one hand, reducing model bias directly reduces bias of $\hat{\mu}_i$, as evident from (4.15a). On the other hand, decreasing model bias generally comes at the cost of increasing model variance, which directly increases the variance of the overall estimate according to the term $\mathbb{E}_{\mathcal{D}_{i,\text{est}}} [\text{Var } \hat{f}_i^0(X_{it}) | D_{it}=1]$ in (4.15b).

Remark 3. *Conditioning on D_{it} is important in the first terms of both (4.14a) and (4.15a), as treatment assignment may be correlated with the covariates, in which case the conditional expectation of the bias under the treatment assignment mechanism may be substantially different than its marginal expectation. This is true independent of whether noise processes and treatment assignment are independent or not.*

Example 2. *It is instructive to consider the hypothetical case where we have perfect knowledge of the conditional mean function, i.e., when $\hat{f}_i^0 = f_i^0$. Then $\text{Bias } \hat{f}_i^0 \equiv 0$, $\text{Var } \hat{f}_i^0 \equiv 0$*

and $\text{Cov } \hat{f}_i^0 \equiv 0$, and so (4.14a) and (4.14b) in Lemma 3 reduce to

$$\begin{aligned} \text{Bias}(\hat{\mu}_i | \mathcal{T}_i) &= \frac{1}{|\mathcal{T}_i|} \sum_{t \in \mathcal{T}_i} (\mathbb{E}[\epsilon_{it} | D_{it} = 1] + \mathbb{E}[\gamma_{it} | D_{it} = 1]) \\ \text{Var}(\hat{\mu}_i | \mathcal{T}_i) &= \frac{1}{|\mathcal{T}_i|^2} \sum_{t \in \mathcal{T}_i} \left(\text{Var}(\epsilon_{it} | D_{it} = 1) + \text{Var}(\gamma_{it} | D_{it} = 1) + (\mathbb{E}[\epsilon_{it} | D_{it} = 1] + \mathbb{E}[\gamma_{it} | D_{it} = 1])^2 \right) \\ &\quad + \frac{2}{|\mathcal{T}_i|^2} \sum_{t \in \mathcal{T}_i} \sum_{\substack{\tau \in \mathcal{T}_i \\ \tau > t}} \left(\text{Cov}(\epsilon_{it}, \epsilon_{i\tau} | D_{it} = D_{i\tau} = 1) + \text{Cov}(\gamma_{it}, \gamma_{i\tau} | D_{it} = D_{i\tau} = 1) \right. \\ &\quad \left. + \mathbb{E}[\epsilon_{it} + \gamma_{it} | D_{it} = 1] \mathbb{E}[\epsilon_{i\tau} + \gamma_{i\tau} | D_{i\tau} = 1] \right) \end{aligned}$$

If, in addition, $D_{it} \perp\!\!\!\perp \epsilon_{it}, \gamma_{it}$, then $\mathbb{E}[\epsilon_{it} | D_{it} = 1] = \mathbb{E}[\gamma_{it} | D_{it} = 1] = 0$, then $\text{Bias}(\hat{\mu}_i | \mathcal{T}_i) = 0$ and the above expressions reduce further to nothing more than those for bias and variance of the sample mean estimate of some scalar parameter θ_i with observations $\theta_{it} = \theta_i + \epsilon_{it} + \gamma_{it}$ under the same error correlation structure as in (4.7). In this case our prediction of the counterfactual is free of error, so the problem reduces to a simple estimation of the mean.

Unlike as in Example 2, we in general need to take into account the bias of the estimate $\hat{\mu}_i$ generated by the bias of the prediction model \hat{f}_i^0 , as well as the additional variance of the estimate caused by the interaction between the correlation of covariates and errors and the bias and variance of the model as a result of the model fitting process. Clearly, if we had access to a model \hat{f}_i^0 that is itself unbiased on the whole domain,¹⁹ then the estimator $\hat{\mu}_i$ would be ex-ante unbiased under the assumption that treatment assignment and error processes are independent.²⁰ This, however, does not mean that $\hat{\mu}_i$ will be ex-interim unbiased. In general, ex-interim bias persists for a given model and is not reduced by simply increasing the number of treatment periods in the estimation data set.

Conditions for Unbiasedness and Consistency

It is of course unrealistic to assume access to a model that is unbiased on the whole domain for a given training data set. A more natural condition is one that accounts for improvements in the model fit as the size of the training data set grows. To this end, let $(\mathcal{D}_{i,\text{tr}}^k)_{k \geq 1}$ be a sequence of training data sets, and denote by $\hat{f}_i^{0,k}$ the model fitted on data from $\mathcal{D}_{i,\text{tr}}^k$. Furthermore, let $\hat{\mu}_i^k$ denote the estimator obtained from (4.10) using the model $\hat{f}_i^{0,k}$.

Definition 5 (Asymptotic Ex-Ante Unbiasedness). *The sequence of estimators $(\hat{\mu}_i^k)_{k \geq 1}$ is asymptotically ex-ante unbiased if $\text{Bias}(\hat{\mu}_i^k | \mathcal{T}_i) \rightarrow 0$ as $k \rightarrow \infty$.*

We should point out that the condition in Definition 5 does not restrict the structure of the regression and estimation models to be fixed. In fact, a natural thing to do is to refine

¹⁹By this we mean that $\mathbb{E}_{\mathcal{D}_{i,\text{tr}}}[\hat{f}_i^0(x)] = f_i^0(x)$. i.e. that for every $x \in \mathcal{X}$ the fit is correct in expectation over the training data.

²⁰This relies on the independence of training and estimation data.

the complexity of the models \hat{f}_i^k and \hat{f}_i^k (e.g. by adding additional transformations of the covariates in a Generalized Linear Model) as the amount of available data grows.

The following Proposition gives a sufficient condition on the sequence of models $\hat{f}_i^{0,k}$ for asymptotic ex-ante unbiasedness.

Proposition 3 (Sufficient Condition for Asymptotic Ex-Ante Unbiasedness). *Suppose that $\epsilon_{it}, \gamma_{it} \perp\!\!\!\perp D_{it}$, and denote by ν_{it} the conditional distribution of the covariate vector X_{it} on \mathcal{X} under treatment assignment. Let ν be a probability measure on \mathcal{X} such that $\nu_{it} \ll \nu$ for all $t \in \mathcal{T}_i$ and that there exists $K < \infty$ such that $\frac{d\nu_{it}}{d\nu} \leq K$ for all $t \in \mathcal{T}_i$. Suppose further that $\text{Bias } \hat{f}_i^{0,k} \rightarrow 0$ in $L^1(\mathcal{X}, \nu)$. Then $\text{Bias}(\hat{\mu}_i^k | \mathcal{T}_i) \rightarrow 0$ as $k \rightarrow \infty$.*

Proof. This follows from a simple change of measure argument:

$$\left| \mathbb{E}_{\mathcal{D}_{i,\text{est}}} [\text{Bias } \hat{f}_i^{0,k}(X_{it}) | D_{it}=1] \right| \leq \int_{\mathcal{X}} |\text{Bias } \hat{f}_i^{0,k}(x)| \frac{d\nu_{it}}{d\nu} d\nu(x) \leq K \int_{\mathcal{X}} |\text{Bias } \hat{f}_i^{0,k}(x)| d\nu(x)$$

Now $\text{Bias } \hat{f}_i^{0,k} \rightarrow 0$ in $L^1(\mathbb{R}^{n_x}, \nu)$ by assumption, so

$$\left| \frac{1}{|\mathcal{T}_i|} \sum_{t \in \mathcal{T}_i} \mathbb{E}_{\mathcal{D}_{i,\text{est}}} [\text{Bias } \hat{f}_i^{0,k}(X_{it}) | D_{it}=1] \right| \leq K \int_{\mathcal{X}} |\text{Bias } \hat{f}_i^{0,k}(x)| d\nu(x) \rightarrow 0$$

The result then follows in conjunction with Corollary 3. □

Proposition 3 formalizes the rather intuitive fact that if the model bias vanishes asymptotically with the size of the training set, then so does the ex-ante bias of the two-stage estimator. The condition in Proposition 3 is not particularly strong. For example, using dominated convergence it is easily seen that it is satisfied if $\text{Bias } \hat{f}_i^{0,k}$ is uniformly bounded and $\text{Bias } \hat{f}_i^{0,k} \rightarrow 0$ point-wise. However, point-wise convergence on \mathcal{X} is not necessary for the condition in Proposition 3 to hold.

Suppose that the data-generating processes of training and estimation data sets are stationary, and that $\nu = \nu_i^{\text{tr}}$, the (marginal) distribution in the training set. Then the condition that $\nu_{it} \ll \nu$ in Proposition 3 requires the support of the distribution of covariates in the estimation set to be contained within the support of the distributions of the covariates in the training set. Put intuitively, the model can only be expected to predict well on data sufficiently similar to the data that has been used for fitting the model. In this sense, the condition is similar in spirit to the so-called ‘‘overlap’’ condition on the covariate distributions of treatment and control groups in program evaluation under unconfoundedness (Rosenbaum and Rubin, 1984).

Example 3 (Linear Model). *Say we use a linear model of the form $\hat{f}_i^{0,k}(X) = \hat{\eta}_i^k \cdot X$, and suppose that $f_i^0(X) = \eta_i \cdot X$, i.e., that the true conditional mean f_i^0 is indeed a linear function of the observed covariates. A sufficient condition for the condition from Proposition 3 to hold is that $\hat{\eta}_i^k \rightarrow \eta_i$ in L^2 and that each X_j has finite second moment. Indeed, in this case it is easy to see from Cauchy’s inequality that $|\hat{\eta}_i^k \cdot X - \eta_i \cdot X| \leq \|\hat{\eta}_i^k - \eta_i\|_2 (\sum_j \mathbb{E}[X_j^2])^{1/2} \rightarrow 0$. Now it is*

well known that $\hat{\eta}_i^k \rightarrow \eta_i$ in L^2 for OLS estimates $\hat{\eta}_i^k$ under standard conditions, and this remains valid also if there is serial correlation between covariates under additional but relatively mild conditions on the strength of the correlation (for example strong α -mixing (Billingsley, 1995)). The condition on the variances is a weak condition that should be satisfied in most applications in practice.

Unbiasedness is of course not the only desirable property of an estimator. What we would like is to establish is a notion of consistency — we want our estimate to converge in probability to the true value with increasing sample size. Since “sample size” in this context involves both training and estimation data, the conventional notion of consistency has to be modified as follows:

Definition 6 (Consistency). *Consider a sequence $(\mathcal{D}_{i,\text{tr}}^k, \mathcal{D}_{i,\text{est}}^k)_{k=1}^\infty$ of tuples with treatment periods $\mathcal{T}_i^k \subset \mathcal{D}_{i,\text{est}}^k$ such that $|\mathcal{D}_{i,\text{tr}}^k| \rightarrow \infty$ and $|\mathcal{T}_i^k| \rightarrow \infty$, and let $\hat{\mu}_i^k$ denote the estimate obtained by applying $\hat{\mu}_i$ to \mathcal{T}_i^k with the model $\hat{f}_i^{0,k}$ being trained on $\mathcal{D}_{i,\text{tr}}^k$. We say that $\hat{\mu}_i$ is consistent if $\hat{\mu}_i^k \xrightarrow{P} \mu_i$.*

Clearly, while asymptotic ex-ante unbiasedness is necessary for consistency, it is not sufficient. As we have seen earlier, one of the primary challenges is that our observations are serially correlated. Intuitively, if the strength of this correlation decays with temporal separation, and if the observations become sufficiently separated in time, then we may still achieve consistency. Theorem 1 below formalizes this.

Theorem 1 (Sufficient Conditions for Consistency). *Suppose that $\hat{\mu}_i$ is asymptotically ex-ante unbiased, that $|\mathcal{T}_i^k| \rightarrow \infty$ with $\min\{|t - \tau| : t, \tau \in \mathcal{T}_i^k, \tau \neq t\} \rightarrow \infty$, that $\mathcal{D}_{i,\text{est}}^k \perp\!\!\!\perp \mathcal{D}_{i,\text{tr}}^k$ for all k , and that $D_{it} \perp\!\!\!\perp \epsilon_{it}, \gamma_{it}$. Suppose further that*

- (i) $\exists \sigma_\epsilon, \sigma_\gamma < \infty$ s.t. $\text{Var}_{\mathcal{D}_{i,\text{est}}^k}(\epsilon_{it}) < \sigma_\epsilon^2$ and $\text{Var}_{\mathcal{D}_{i,\text{est}}^k}(\gamma_{it}) < \sigma_\gamma^2$
- (ii) $\exists K_1 < \infty$ and $k' < \infty$ s.t. $\mathbb{E}_{\mathcal{D}_{i,\text{est}}^k}[\text{Var} \hat{f}_i^{0,k}(X_{it}) | D_{it}] < K_1$ for all $k > k', t \in \mathcal{T}_i^k$
- (iii) $\exists K_2 < \infty$ and $k' < \infty$ s.t. $\text{Var}_{\mathcal{D}_{i,\text{est}}^k}(\text{Bias} \hat{f}_i^{0,k}(X_{it}) | D_{it}) < K_2$ for all $k > k', t \in \mathcal{T}_i^k$
- (iv) $\text{Cov}(\epsilon_{it}, \epsilon_{i\tau}) \rightarrow 0$ and $\text{Cov}(\gamma_{it}, \gamma_{i\tau}) \rightarrow 0$ as $|t - \tau| \rightarrow \infty$
- (v) $\mathbb{E}_{\mathcal{D}_{i,\text{est}}^k}[\text{Cov} \hat{f}_i^{0,k}(X_{it}, X_{i\tau}) | D_{it} = D_{i\tau} = 1] \rightarrow 0$ as $k \rightarrow \infty$ and $|t - \tau| \rightarrow \infty$
- (vi) $\text{Cov}_{\mathcal{D}_{i,\text{est}}^k}(\text{Bias} \hat{f}_i^{0,k}(X_{it}), \text{Bias} \hat{f}_i^{0,k}(X_{i\tau}) | D_{it} = D_{i\tau} = 1) \rightarrow 0$ as $k \rightarrow \infty$ and $|t - \tau| \rightarrow \infty$
- (vii) $\text{Cov}_{\mathcal{D}_{i,\text{est}}^k}(\epsilon_{it} + \gamma_{it}, \text{Bias} \hat{f}_i^{0,k}(X_{i\tau}) | D_{i\tau} = 1) \rightarrow 0$ as $k \rightarrow \infty$ and $|t - \tau| \rightarrow \infty$

Then the estimator (4.10) is consistent in the sense of Definition 6.

Proof. $\text{Bias}(\hat{\mu}_i | \mathcal{T}_i) \rightarrow 0$ as $k \rightarrow \infty$ due to the assumption of ex-ante unbiasedness. Further, with some book-keeping it is easy to see from (4.14b) that under assumptions (i) - (vii) of Theorem 1 we also have that $\text{Var}(\hat{\mu}_i | \mathcal{T}_i) \rightarrow 0$ as $k \rightarrow \infty$. Convergence of probability is then immediate from Markov's inequality. \square

Condition (i) in Theorem 1 is basic. Conditions (ii) and (iii) are relatively straightforward and essentially require asymptotic bounds on the bias and variance of the model. If model bias and variance of $\hat{f}_i^{0,k}$ are uniformly bounded on all of \mathcal{X} for k sufficiently large, then (ii) and (iii) are trivially satisfied. Condition (iv) requires the serial correlation of the noise processes $(\epsilon_{it})_{t \geq 1}$ and $(\gamma_{it})_{t \geq 1}$ to vanish with increasing temporal separation, and is a very natural condition if we want to have any hope for a consistent estimator, even if we had a perfect model \hat{f}_i^0 . This condition is for example satisfied if $(\epsilon_{it})_{t \geq 1}$ and $(\gamma_{it})_{t \geq 1}$ are strongly α -mixing. Conditions (v) and (vi) are slightly harder to interpret. If model predictions were uncorrelated then $\text{Cov} \hat{f}_i^{0,k} \equiv 0$. However, due to serial correlation in the training data, data in the neighborhood of some x is likely to come from observations made around the same time, so one would expect that using this data to train a model would yield prediction errors that are indeed spatially correlated. The same applies to the model bias. On the other hand, for estimation the same argument goes the other way: If covariates X_{it} and $X_{i\tau}$ in different treatment periods are only weakly correlated, then this will reduce the impact of the spatial correlation of the model predictions on the estimates. Finally, condition (vii) is necessary since X_{it} may contain lagged values of Y_{it} , hence if the errors ϵ_{it} and γ_{it} are serially correlated this will result in $X_{i\tau}$ being correlated with ϵ_{it} and γ_{it} . Observe also that if $\text{Bias} \hat{f}_i^0$ is uniformly bounded, then (iv) implies (vii).

We emphasize that the conditions (i) - (vii) are not independent from asymptotic ex-ante unbiasedness. In fact, it is one of the fundamental tenets of statistical learning that reducing bias generally comes at the cost of increasing variance. The assertion of Theorem 1 is simply that if model bias vanishes asymptotically with increasing amount training data *while model variance is controlled*, then the resulting two-stage estimator will be consistent.

Remark 4. *The conditions in Theorem 1 can be weakened: In particular, rather than requiring asymptotic bounds in the form of constants it is enough to have some of the individual components grow sublinearly. For instance, instead of condition (ii) it suffices to have that $\mathbb{E}_{\mathcal{D}_{i,\text{est}}^k} [\text{Var} \hat{f}_i^{0,k}(X_{it}) | D_{it} = 1] = \mathcal{O}(|\mathcal{T}_i^k|)$.*

From Theorem 1 we could now derive a wide range of simpler, more restrictive conditions that ensure consistency of our estimate (see our above discussion for some examples). While such abstract conditions are mathematically interesting, they become much more relevant and useful in the context of specific models. However, as we in this chapter focus on the general methodology, we refrain from further developing such conditions at this point.

The Assumption of Independent Training and Estimation Data Sets

Most of our results so far have assumed that the training and estimation data sets are independent. If this assumption does not hold true, then we cannot simply consider quantities like $\mathbb{E}_{\mathcal{D}_{i,\text{est}}} [\text{Bias } \hat{f}_i^0(X_{it})]$ as in Lemma 3, but instead have to take a single expectation $\mathbb{E}_{\mathcal{D}} [\hat{f}_i^0(X_{it}) - f_i^0(X_{it})]$ over the joint distribution of training and estimation data. Because our general framework allows for serial correlation of error and covariate processes, training and estimation sets will never be completely independent. However, if correlation is relatively local in time (e.g if the processes are α -mixing), then approximate independence may be achieved by ensuring sufficient temporally separation between training and estimation data sets.

If treatment assignment is independent of the realization of the covariate and output processes, then we can satisfy the independence assumption on $\mathcal{D}_{i,\text{tr}}$ and $\mathcal{D}_{i,\text{est}}$, modulo the caveat just mentioned above. This is the experimental ideal. However, we may also want to perform observational studies, in which we face a situation in which treatment has been administered in a potentially non-random fashion. If in such a situation treatment assignment is correlated with covariates, this may result in training and estimation data sets having different underlying distribution.

One simple case in which the condition $\mathcal{D}_{i,\text{tr}}$ and $\mathcal{D}_{i,\text{est}}$ can be (approximately) satisfied in an observational setting is one in which no treatment is administered in the beginning, and treatment assignment starts after a time \tilde{t} that is independent of the realization of the output and covariate processes. Then choosing $\mathcal{D}_{i,\text{tr}}$ as all observations prior to \tilde{t} and $\mathcal{D}_{i,\text{est}}$ as the treatment periods ensures approximate independence. In this case Assumption 11 (stationarity of the conditional distribution of the outcomes given the covariates) is crucial to be able to obtain unbiased estimators of the treatment effect.

4.4 Conditional Individual Treatment Effects Under General Treatments

Besides the (marginal) ITE, we would also like to estimate CITE, the conditional individual treatment effect. Moreover, rather than in the response to a simple binary treatment, we are often interested in an individual's response as a function of a treatment parameter, say the intensity of treatment. For instance, one goal in the context of evaluating Demand Response is estimate the demand curve, i.e., the effect of receiving a notification on consumption as a function of the stated incentive level. It turns out that the divide and conquer approach of our two-stage estimation process makes it, at least conceptually, quite easy to perform such estimation.

Let $R_{it} \in \mathcal{R} \cup \{\mathbf{0}\}$ with $\mathcal{R} \subset \mathbb{R}^{n_r}$ denote the vector of treatment parameters in period t . Here $R_{it} = \mathbf{0}$ means individual i receives no treatment,²¹ so that $\mathcal{T}_i = \{t : R_{it} \neq \mathbf{0}\}$. Consider

²¹We still have a situation in mind in which the number of non-treatment periods is much larger than the

the following generalization of the outcome model (4.7), where we allow the treatment effect μ_i to be a function of the covariate X_{it} and the treatment parameter R_{it} :²²

$$Y_{it} = f_i^0(X_{it}) + \mathbf{1}_{\{R_{it} \neq \mathbf{0}\}} [\mu_i(X_{it}, R_{it}) + \gamma_{it}] + \epsilon_{it} \quad (4.16)$$

We would like to obtain an estimate $\hat{\mu}_i : \mathcal{X} \times \mathcal{R} \rightarrow \mathbb{R}$ of the CITE μ_i . For a given model \hat{f}_i^0 we can compute the quantity δ_{it} for all treatment periods in much the same way as before, namely

$$\delta_{it} := Y_{it} - \hat{f}_i^0(X_{it}) = f_i^0(X_{it}) - \hat{f}_i^0(X_{it}) + \mu_i(X_{it}, R_{it}) + \epsilon_{it} + \gamma_{it} \quad (4.17)$$

which leaves us with a set $((X_{it}, R_{it}), \mu(X_{it}, R_{it}) + \kappa_{it})_{t \in \mathcal{T}_i}$ of pairs of observations and estimated outcomes, where $\kappa_{it} = f_i^0(X_{it}) - \hat{f}_i^0(X_{it}) + \epsilon_{it} + \gamma_{it}$ is, as before, the estimation error, and where (X_{it}, R_{it}) can be viewed as an augmented covariate associated to the estimated effect $\mu(X_{it}, R_{it}) + \kappa_{it}$. In order to estimate the CITE, we can now simply regress the predicted counterfactuals $(\delta_{it})_{t \in \mathcal{T}_i}$ on the covariate-parameter tuples $(X_{it}, R_{it})_{t \in \mathcal{T}_i}$, using a method of our choosing. This is summarized below in Algorithm 1.

Algorithm 1: Computing the CITE Estimate

Input: $\mathcal{D}_{i,\text{tr}}, \mathcal{D}_{i,\text{est}}$

Output: $\hat{\mu}_i$

Train model \hat{f}_i^0 on $\mathcal{D}_{i,\text{tr}}$ using cross-validation for hyperparameter optimization

Compute predicted counterfactuals $(\delta_{it})_{t \in \mathcal{T}_i}$ where $\delta_{it} := Y_{it} - \hat{f}_i^0(X_{it})$

Obtain $\hat{\mu}_i$ by regressing $(\delta_{it})_{t \in \mathcal{T}_i}$ on $((X_{it}, R_{it}))_{t \in \mathcal{T}_i}$

Note also that by regressing on a sub-vector \check{X}_{it} of X_{it} we can also obtain an estimate of (4.4), the ITE conditional on a subset of covariates.²³

We will require the following assumption:

Assumption 13 (Conditional Independence). *The errors ϵ_{it} and γ_{it} are conditionally independent of the treatment given the covariates, i.e., $\epsilon_{it}, \gamma_{it} \perp\!\!\!\perp R_{it} \mid X_{it}$.*

Let Y_{it}^r denote the outcome Y_{it} under treatment assignment $R_{it} = r$. In particular, $Y_{it}^{\mathbf{0}}$ denotes the outcome in the absence of treatment, just as in the case of binary treatment. It turns out that Assumption 13 implies what is essentially an analogue of the weak unconfoundedness condition from Hirano and Imbens (2004):²⁴

number of treatment periods.

²²Note that this formulation includes binary treatments as a special case, namely when $R_{it} \in \{\mathbf{0}, 1\}$.

²³We emphasize again that, unlike the full CITE, this restricted estimate is only meaningful under the joint distribution of X_{it} under the underlying data generating process.

²⁴While Hirano and Imbens (2004) are interested in the assignment of real-valued treatments to different individuals, we instead consider assignment of different treatments to the same individual in different periods.

Lemma 4 (Weak Unconfoundedness). *Under Assumption 13, $Y_{it}^r \perp\!\!\!\perp R_{it} \mid X_{it}$ for all $r \in \mathcal{R}$.*

Proof. Fix $r \in \mathcal{R}$. Then $\mathbb{E}[Y_{it}^r \mid X_{it}, R_{it}] = f_i^0(X_{it}) + \mu_i(X_{it}, r) + \mathbb{E}[\epsilon_{it} + \gamma_{it} \mid X_{it}, R_{it}]$. Now, under Assumption 13 we have that $\mathbb{E}[\epsilon_{it} + \gamma_{it} \mid X_{it}, R_{it}] = \mathbb{E}[\epsilon_{it} + \gamma_{it} \mid X_{it}]$, and therefore $\mathbb{E}[Y_{it}^r \mid X_{it}, R_{it}] = \mathbb{E}[Y_{it}^r \mid X_{it}]$. \square

For a given (imperfect) model \hat{f}_i^0 , it will generally be the case that the model error $e_i(X_{it}) = \hat{f}_i^0(X_{it}) - f_i^0(X_{it})$ is correlated with the covariates X_{it} . Fundamentally, this means that, even if the form of the regression model for $\hat{\mu}_i$ were correctly specified, the result $\hat{\mu}_i$ of the second stage regression will not necessarily converge to the true effect μ with increasing number of treatment observations.²⁵ This will still be the case even if those observations were independent. That is, the estimator will not be consistent in the conventional sense. Much like with the sample mean estimator for the ITE discussed in Section 4.3.2, the appropriate notion of consistency is with respect to both training and estimation data and will require joint conditions on the sequences $(\hat{f}_i^{0,k})_k$ and $(\hat{\mu}_i^k)_k$.

Definition 7 (Point-Wise Consistency for General Treatments). *Consider a sequence of tuples $(\mathcal{D}_{i,\text{tr}}^k, \mathcal{D}_{i,\text{est}}^k)_{k=1}^\infty$ with treatment periods $\mathcal{T}_i^k \subset \mathcal{D}_{i,\text{est}}^k$ such that $|\mathcal{D}_{i,\text{tr}}^k| \rightarrow \infty$ and $|\mathcal{T}_i^k| \rightarrow \infty$. Let $\hat{\mu}_i^k$ denote the estimate obtained by applying Algorithm 1 to $(\mathcal{D}_{i,\text{tr}}^k, \mathcal{D}_{i,\text{est}}^k)$. We say that the sequence $(\hat{\mu}_i^k)_k$ is point-wise consistent on $S \subset \mathcal{X} \times \mathcal{R}$ if $\hat{\mu}_i^k(x, r) \xrightarrow{P} \mu_i(x, r)$ for all $(x, r) \in S$.*

Point-wise consistency according to Definition 7 is with respect to some set $S \subset \mathcal{X} \times \mathcal{R}$. In many cases, it may be unnecessary or even unnatural to consider the treatment effect on all of $\mathcal{X} \times \mathcal{R}$. Unlike with the sample mean estimator for the ITE under binary treatment in Section 4.3, we cannot easily derive expressions for bias and variance for general second stage estimators. Instead, we have to settle for a more abstract result that links the asymptotic properties of the first and second stage models.

Proposition 4. *Suppose that $\hat{f}_i^{0,k}(x) \xrightarrow{P} \mathbb{E}[Y_{it}^0 \mid X_{it}=x]$ and $\hat{\mu}_i^k(x, r) \xrightarrow{P} \mathbb{E}[\delta_{it} \mid X_{it}=x, R_{it}=r]$ as $k \rightarrow \infty$ for all $(x, r) \in S$. Suppose further that Assumption 13 holds. Then $\hat{\mu}_i^k(x, r) \xrightarrow{P} \mu_i(x, r) + \mathbb{E}[\gamma_{it} \mid X_{it}=x]$ as $k \rightarrow \infty$ for all $(x, r) \in S$.*

Proof. Fix $(x, r) \in S$. Then

$$\begin{aligned} \hat{\mu}_i^k(x, r) &= \hat{\mu}_i^k(x, r) - \mathbb{E}[\delta_{it} \mid X_{it}=x, R_{it}=r] + \mathbb{E}[\delta_{it} \mid X_{it}=x, R_{it}=r] \\ &= \hat{\mu}_i^k(x, r) - \mathbb{E}[\delta_{it} \mid X_{it}=x, R_{it}=r] - \hat{f}_i^{0,k}(x, r) + f_i^0(x, r) \\ &\quad + \mu_i(x, r) + \mathbb{E}[\epsilon_{it} \mid X_{it}=x, R_{it}=r] + \mathbb{E}[\gamma_{it} \mid X_{it}=x, R_{it}=r] \\ &= \hat{\mu}_i^k(x, r) - \mathbb{E}[\delta_{it} \mid X_{it}=x, R_{it}=r] - \hat{f}_i^{0,k}(x, r) + \mathbb{E}[f_i^0(x, r) + \epsilon_{it} \mid X_{it}=x] \\ &\quad + \mu_i(x, r) + \mathbb{E}[\gamma_{it} \mid X_{it}=x] \end{aligned}$$

²⁵In order to convey the intuition we are being very informal here and disregard a number of technicalities, such as what type of convergence we actually mean. We will make this notion precise in the following.

where the last step follows from Assumption 13. But $\mathbb{E}[f_i^0 + \epsilon_{it} | X_{it} = x] = \mathbb{E}[Y_{it}^0 | X_{it} = x]$, and thus the result follows under our assumptions. \square

Remark 5. *An interesting observation from the proof of Proposition 4 is that, even though a model that is asymptotically correct in estimating $\mathbb{E}[Y_{it}^0 | X_{it} = x]$ may be incorrectly estimating f_i^0 , this bias cancels out in the regression in the second stage. For this to hold Assumption 13 is crucial, since otherwise $\mathbb{E}[\epsilon_{it} | X_{it} = x, R_{it} = r] \neq \mathbb{E}[\epsilon_{it} | X_{it} = x, R_{it} \neq 0]$.*

Proposition 4 essentially states that under a “point-wise asymptotically correct” sequence of models $\hat{\mu}_i^k$ the estimator will extract the conditional ITE as long as $\mathbb{E}[\gamma_{it} | X_{it} = x] = 0$. But since X_{it} may contain lagged output values, this condition will not generally hold for any fixed treatment assignment. However, even if the noise process $(\gamma_{it})_t$ is serially correlated, it may hold asymptotically if this correlation decays with increasing temporal separation, and if the separation between treatment periods increases. This is because the dependence between γ_{it} and X_{it} is due exclusively to the dependence between γ_{it} and $\gamma_{it'}$ for treatment periods $t' < t$.

Corollary 4. *Suppose that, in addition to the Assumptions of Proposition 4, we also have that $\mathbb{E}[\gamma_{it} | X_{it} = x] \rightarrow 0$ as $\inf\{|t - t'| : t, t' \in \mathcal{T}_i^k\} \rightarrow \infty$ for all $x \in \pi_{\mathcal{X}}S$. Then $\hat{\mu}_i^k(x, r) \xrightarrow{P} \mu_i(x, r)$ as $k \rightarrow \infty$ for all $(x, r) \in S$.*

In addition to the notion of point-wise consistency from Definition 7, we can also consider other relevant concepts of consistency, such as a notion of weak consistency.²⁶ However, we will refrain from indulging in such mathematical pleasures at this point and instead focus on aspects more relevant to finite sample analysis.

Controlling for Differences in Covariate Distributions

One issue with the generic estimation strategy discussed so far, both with ITE and CITE estimates, is that the empirical distribution of covariates may be quite different in the training and estimation data sets. As discussed earlier, there needs to at least be some degree of “overlap” between the distributions to have any hope of obtaining a consistent estimator. However, even if this is the case, if in the first stage the model \hat{f}_i^0 is trained without taking the different frequencies of the covariates into account, then the different distribution in the training data may lead to the resulting model having significant finite-sample bias exactly where it impacts the estimation in the second stage the most.

For example, if we fit the model \hat{f}_i^0 by minimizing a standard (unweighted) cumulative loss function, such as the MSE, then a small number of observations with covariates in regions of \mathcal{X} most relevant for estimation may be “washed out” by a very large number of observations in regions of \mathcal{X} with no covariates during treatment periods. As the other extreme, we could consider exclusively those observations in the training data that have corresponding

²⁶For instance, we could define a sequence $(\hat{\mu}_i^k)_k$ to be weakly consistent w.r.t. some function class Φ and reference measure λ if $\int_{\mathcal{X} \times \mathcal{R}} \hat{\mu}_i^k \varphi d\lambda(x, r) \xrightarrow{P} \int_{\mathcal{X} \times \mathcal{R}} \mu_i \varphi d\lambda(x, r)$ as $k \rightarrow \infty$ for all $\varphi \in \Phi$.

covariate vectors in the estimation data. While this would minimize bias, it would also lead to excessive variance in the estimator. Hence we need to find a tradeoff between the two. This tradeoff can be achieved by using a weighted loss function that weights observations in the training data by “relevance”. Selecting the weights can be achieved through different methods. For instance, one possibility would be to do Kernel weighting, an approach that is very similar to that of non-parametric matching estimators (Imbens, 2004). The respective hyperparameters can be optimized by using cross-validation on the training data set.

While Proposition 3 shows that the ITE estimator will be asymptotically ex-ante unbiased even if the discrepancy between the covariate distributions on treatment and estimation data is significant (as long as there is overlap), it says nothing about convergence rates. In practice, appropriate weighting may have a significant effect on the finite sample performance of our estimator. We observe this behavior in our synthetic experiments on electricity consumption data in Section 5.4 in the following chapter.

4.5 Average Treatment Effects

So far all of our results in this chapter have been developed for measuring treatment effects on the level of a single individual. For many policy questions, a natural quantity of interest is the average effect of a treatment across individuals in some population of interest \mathcal{P} . Commonly, it may not be possible to observe every individual in \mathcal{P} ,²⁷ and we need to draw our conclusions from some subset of \mathcal{P} .²⁸

Note that implicit in our outcome model (4.16) is the assumption that the treatment effect μ_i of individual i is unaffected by the particular assignment of treatment to other individuals. This is the standard Stable Unit Treatment Value Assumption (SUTVA) developed by Rubin (1980).

In the following we will state all results for general, non-binary treatments. The corresponding results for binary treatments are simple special cases. Within a given sample population \mathcal{P}_T exposed to treatment (the “treated”), we can define the following:

Definition 8 (EATT). *Let \mathcal{P}_T denote the set of individuals receiving treatment. The Empirical Average Treatment Effect on the Treated (EATT) is*

$$\mu_{\mathcal{P}_T}(x, r) := \frac{1}{|\mathcal{P}_T|} \sum_{i \in \mathcal{P}_T} \mu_i(x, r) \quad (4.18)$$

The individuals in \mathcal{P}_T have been selected according to some assignment mechanism \mathcal{M} . Denote $\mathcal{M}_i = \mathbf{1}_{\{i \in \mathcal{P}_T\}}$, i.e. $\mathcal{M}_i = 1$ if individual i is assigned to \mathcal{P}_T , and $\mathcal{M}_i = 0$ otherwise.

²⁷In fact, the population \mathcal{P} is often assumed infinite, primarily for mathematical convenience. See Abadie et al. (2014) for a discussion of settings in which this assumption is not justified.

²⁸This rather plain fact has been referred to as the “analogy principle” (Manski, 1988).

The *Average Treatment Effect on the Treated* (ATT) is the expected treatment effect of an individual in the underlying population \mathcal{P} selected by the assignment mechanism \mathcal{M} :

Definition 9 (ATT). *The Average Treatment Effect on the Treated (ATT) is*

$$\mu(x, r) := \mathbb{E}[\mu_i(x, r) \mid \mathcal{M}_i = 1] \quad (4.19)$$

The notion of an average treatment effect that is a function of some covariates may seem confusing to readers familiar with the classic literature on causal inference and program evaluation. The point is that averages are across individuals, so it makes sense to talk about an average effect across individuals that is heterogeneous across the covariate and treatment parameter spaces \mathcal{X} and \mathcal{R} , respectively. Fundamentally, the prerequisite for being able to estimate such a quantity is the possibility of characterizing treatment effect heterogeneity within a single individual. While this is typically infeasible in settings with a single treatment exposure per individual (which are most commonly considered in the literature), it may potentially be possible in our case due to multiple treatment exposures per individual.

Often we are also interested in the average effect of treatment across the entire underlying population of individuals:

Definition 10 (Average Treatment Effect). *The Average Treatment Effect (ATE) is*

$$\nu(x, r) := \mathbb{E}[\mu_i(x, r)] \quad (4.20)$$

Clearly, EATT, ATT and ATE are the same if the treatment effect is homogeneous across individuals, i.e. if $\mu_i \equiv \mu_j$ for any i, j in the underlying population \mathcal{P} . If, instead, the treatment effect is heterogeneous across individuals, then ATE and ATT may differ significantly, while we would expect EATT to approach ATT with increasing sample population size $|\mathcal{P}_T|$.

Furthermore, it is easy to see that if \mathcal{P}_T is a random sample from the underlying population,²⁹ then ATE and ATT coincide. This is what makes it relatively straightforward to estimate the ATE in randomized experiments. For general assignment mechanisms, however, one has to be worried about a potential correlation between the assignment to treatment and the treatment effect itself. If such correlation exists and is not controlled for (e.g. by performing propensity score matching), then comparing treated and non-treated subjects in general will result in selection bias. Avoiding selection bias is a paramount concern in settings with few (possibly a single) observation per individual, in which across-subject matching is performed to obtain estimates of counterfactuals from the control group. Our methods do not really use a control group in the conventional sense, as estimates of the counterfactual outcome(s) are constructed via within-subject matching across the different observation periods. Nevertheless, if we are interested in estimating the ATE, we still have to worry about the composition of the sample population \mathcal{P}_T being different than that of the underlying population. We address this point below, but focus first on estimating the ATT.

²⁹More formally, we require that assignment to \mathcal{P}_T is independent of the ITE, i.e. that $\mathcal{M}_i \perp \mu_i$.

Estimating EATT and ATT

Once we have obtained estimates for the ITEs, it is straightforward to construct an estimate of both the EATT³⁰ and ATT by marginalizing over individuals in the sample population \mathcal{P}_T . Indeed, for each $i \in \mathcal{P}_T$ we can compute an ITE estimate $\hat{\mu}_i : \mathcal{X} \times \mathcal{R} \mapsto \mathbb{R}$. The natural estimator for both EATT and ATT then is simply

$$\hat{\mu}_{\mathcal{P}_T}(x, r) = \hat{\mu}(x, r) = \frac{1}{|\mathcal{P}_T|} \sum_{i \in \mathcal{P}_T} \hat{\mu}_i(x, r) \quad (4.21)$$

It is straightforward to see that if the ITE estimators are (point-wise) consistent for each individual, then so is the EATT estimate:

Corollary 5. *Suppose that \mathcal{P}_T is finite and that $\hat{\mu}_i$ is point-wise consistent on S for each $i \in \mathcal{P}_T$. Then the EATT estimate $\hat{\mu}_{\mathcal{P}_T}$ in (4.21) is also point-wise consistent on S , that is, $\hat{\mu}_{\mathcal{P}_T}^k(x, r) \xrightarrow{P} \mu_{\mathcal{P}_T}(x, r)$ as $k \rightarrow \infty$ for all $(x, r) \in S$.*

Intuitively, the same should be true also for the ATT estimate if the sample population grows and consists of independent random draws from \mathcal{P} . To this end, let $\rho_{\mathcal{M}}$ denote the probability measure on \mathcal{P} characterizing a single draw from \mathcal{P} under the assignment mechanism \mathcal{M} . For our result we will require uniformity in the convergence of the individual estimates under selection by \mathcal{M} :

Assumption 14. *$\hat{\mu}_i$ is point-wise consistent on S uniformly in i $\rho_{\mathcal{M}}$ -almost everywhere on \mathcal{P} . That is, for any $\epsilon > 0$ there exists $k' < \infty$ such that $|\hat{\mu}_i^k(x, r) - \mu_i(x, r)| < \epsilon$ for all $k > k'$ for $\rho_{\mathcal{M}}$ -almost all $i \in \mathcal{P}$.*

We also need the following technical assumption requiring the variance of the treatment effect under the selection \mathcal{M} to be uniformly bounded:

Assumption 15. *For all $(x, r) \in S$ there exists $K < \infty$ such that $\text{Var}(\mu_i(x, r) | \mathcal{M}) < K$.*

With this we have the following result:

Proposition 5. *Suppose that Assumptions 14 and 15 hold. Suppose further that $\mathcal{P}_T^k \subset \mathcal{P}$ for each k , with $|\mathcal{P}_T^k| \rightarrow \infty$ and each \mathcal{P}_T^k consisting of i.i.d. draws from \mathcal{P} according to $\rho_{\mathcal{M}}$. Then $\hat{\mu}$ is point-wise consistent on S , that is, $\hat{\mu}^k(x, r) \xrightarrow{P} \mu(x, r)$ as $k \rightarrow \infty$ for all $(x, r) \in S$.*

³⁰Note that estimating the EATT, while quite straightforward in the case of within-subject matching, is a more delicate problem under across-subject matching approaches. In the latter case, for a given sample population one would obtain an estimate by exposing only a random subset of the population to treatment. If the sample population is small and there is strong heterogeneity in the effect, then the resulting estimate may differ substantially from the actual EATT.

Proof. Fix $(x, r) \in S$ and note that

$$\hat{\mu}^k(x, r) - \mu(x, r) = \frac{1}{|\mathcal{P}_T^k|} \sum_{i \in \mathcal{P}_T^k} \hat{\mu}_i^k(x, r) - \mathbb{E}[\mu_i(x, r) | \mathcal{M}] = \frac{1}{|\mathcal{P}_T^k|} \sum_{i \in \mathcal{P}_T^k} (\hat{\mu}_i^k(x, r) - \mu_i(x, r)) - \xi^k$$

where $\xi^k = \mathbb{E}[\mu_i(x, r) | \mathcal{M}] - \frac{1}{|\mathcal{P}_T^k|} \sum_{i \in \mathcal{P}_T^k} \mu_i(x, r)$. Now $\mathbb{E}[\xi^k] = 0$ and, using Assumption 15, $\text{Var}(\xi^k) = \frac{1}{|\mathcal{P}_T^k|} \text{Var}(\mu_i(x, r) | \mathcal{M}) < \frac{K}{|\mathcal{P}_T^k|} \rightarrow 0$, and so $\xi^k \xrightarrow{P} 0$ by Markov's inequality. Moreover, Assumption 14 implies that $\frac{1}{|\mathcal{P}_T^k|} \sum_{i \in \mathcal{P}_T^k} (\hat{\mu}_i^k(x, r) - \mu_i(x, r)) \xrightarrow{P} 0$. \square

Assumption 14, which we used in Proposition 5, is quite strong. Specifically, a common concern that one may have is that some of the ITE estimates may be biased.³¹ However, since the ATT estimate $\hat{\mu}$ in (4.21) is obtained by marginalizing over ITEs, it may be unbiased even if the ITE estimates $\hat{\mu}_i$ are not. In particular, so long as the expected bias of the $\hat{\mu}_i$ under the selection mechanism \mathcal{M} is zero, then $\hat{\mu}(x, r)$ will be unbiased. But this of course does not necessarily mean that the estimator will be consistent — if the estimation errors are correlated between individuals this will require additional assumptions.

Proposition 6. *In the setting of Proposition 5, suppose that instead of Assumption 14 we have that $\mathbb{E}[\hat{\mu}_i^k(x, r) - \mu_i(x, r) | \mathcal{M}] \rightarrow 0$ and $\text{Cov}(\hat{\mu}_i^k(x, r), \hat{\mu}_j^k(x, r) | \mathcal{M}) \rightarrow 0$ as $k \rightarrow \infty$ for all $(x, r) \in S$. Then $\hat{\mu}$ is point-wise consistent on S .*

Proof. Very similar to the proof of Proposition 5, after noting that our assumptions imply that $\frac{1}{|\mathcal{P}_T^k|} \sum_{i \in \mathcal{P}_T^k} (\hat{\mu}_i^k(x, r) - \mu_i(x, r)) \xrightarrow{P} 0$. \square

Remark 6. *In the number of treatment observations is different for different individuals (i.e. if we have an “unbalanced panel”) we need to be careful when comparing our ATT estimate with that of a classic across-subject matching regression model. Specifically, if the treatment effect is heterogeneous across individuals, and if the number of observations is correlated with the treatment effect, then the estimate obtained from a standard (unweighted) regression model will be biased. To obtain an unbiased estimate of the ATT (4.19), each observation associated to an individual must be weighted in the regression by the inverse of the total number of observations of that individual. This issue is not present in our estimate based on marginalizing the ITE estimates, as each observation automatically gets weighted correctly by first computing the ITE.³²*

³¹For example, this may be the case if Assumption 11, i.e., that the conditional distribution of the potential outcomes given the covariates is stationary, is violated. For instance, in estimating residential energy consumption this may be the case if an individual has upgraded or replaced their AC system or some of their electric appliances.

³²Of course it does not eliminate the problem of individuals with very few observations having a disproportionate effect on the variance of the resulting estimator.

Estimating the ATE

In the case in which assignment to \mathcal{P}_T is random (i.e. if $\rho_{\mathcal{M}}$ is the uniform distribution on \mathcal{P}), estimating the ATE reduces to estimating the ATT. In general, however, assignment to \mathcal{P}_T may be correlated with an individual's treatment effect, in which case ATE and ATT may differ significantly. If we have access to additional subject-level covariates representing certain characteristics of the individuals, then, under the assumption that assignment to \mathcal{P}_T is conditionally independent given those subject-level covariates (that is, if the selection mechanism is unconfounded), we may still estimate the ATE by using well-established methods such as propensity score matching (Rosenbaum and Rubin, 1983), which amounts to appropriately weighting the ITEs in (4.21).

As the focus of our contributions is on estimating (conditional) ITEs using within-subject matching across multiple periods, we will assume for the remainder of this chapter that the sample population \mathcal{P}_T is indeed a random sample from the underlying population \mathcal{P} , and hence it is enough to focus on EATT and ATT. Following our comment above, it is relatively straightforward to extend our approaches to unconfounded selection mechanisms.

Why use Within-Subject Matching to Estimate Average Effects?

The usefulness of estimating individual treatment effects based on within-subject matching is clear. But in some cases it may also be advantageous to use this strategy to estimate average effects instead of across-subject matching approaches.

Firstly, in many settings it may be infeasible or impractical to have a control group in the first place, which essentially rules out across-subject matching methods. Secondly, it may be infeasible or impractical to expose different individuals to treatment in a coordinated fashion, which in some situations may make it impossible to pool observations across the treated, resulting in a treatment group of effectively much smaller size and consequently an imprecise estimator. By design, our method utilizes all of the available data for each individual, something that is not always easy to do in conventional across-subject matching approaches that may often assume a so-called balanced panel, i.e., the same number of observations during common observation periods.

Finally, within-subject matching can mitigate an inherent issue with across-subject matching approaches, which is that it is generally impossible to verify that the compared groups are statistically the same across all relevant characteristics. This is a particularly important concern for observational studies, but relevant also for randomized experiments with small sample sizes, for which it is possible that treatment and control group are systematically different along some unobserved characteristic that is correlated with the treatment outcome. These kinds of concerns are not present when using observations from a single individual.

4.6 Bootstrapped Confidence Intervals

Because of our agnosticism to the model used in our estimation strategy, the general form of the treatment response, and the potentially quite complicated correlation structure of the observations, it is very challenging, if not impossible, to derive the limiting distribution of our estimate in a general setting without making strong additional assumptions. Moreover, the limiting distribution of certain models \hat{f}_i^0 , such as Random Forests, may be unknown even in the case of i.i.d. observations.³³ Finally, our ITE estimates are, naturally, for a single individual, and thus the number of treatment observations will be limited. Therefore it is unlikely that the asymptotic limiting distribution will be a good approximation of the finite sample distribution. In order to still obtain confidence intervals around our estimates, we therefore turn to the bootstrap (Efron and Tibshirani, 1994).

4.6.1 Confidence Intervals for ITE and CITE

Bootstrapping in our setting gets complicated by a number of issues. Firstly, randomness is introduced to the estimate through various sources, including the training data $\mathcal{D}_{i,\text{tr}}$ used to train the model \hat{f}_i^0 , the model generation process itself (if it is non-deterministic), and the estimation data $\mathcal{D}_{i,\text{est}}$, which contains the noisy observations of the treatment periods. Secondly, due to the serial correlation in the time series, the observations are not i.i.d. as is typically assumed in the basic bootstrap. Finally, possible computational requirements due to fitting the model gets amplified significantly by the fact that for computing confidence intervals we typically require a large number of bootstrap estimates (Efron, 1987).

We can address the first two of these challenges by ensuring that we draw bootstrap samples from both the training and estimation data sets, and by performing a moving block bootstrap (MBB) to account for correlation in the data (Kunsch, 1989). In the spirit of the bootstrap, we should repeat the full estimation procedure for each bootstrap sample. This however may quickly become computationally infeasible if training the model \hat{f}_i^0 is expensive. In particular, this is likely to be the case if the model fitting involves cross-validation for optimizing hyper-parameters. To remedy this, a first step is to perform hyper-parameter search once, fix the optimal hyper-parameters, and then for each bootstrap sample from $\mathcal{D}_{i,\text{tr}}$ re-fit the model for this selection of hyper-parameters. Intuitively, this approach will work well if the optimal hyper-parameters are insensitive to being fit on different moving-block bootstrap samples. Performing this initial hyper-parameter optimization will significantly reduce computation time by eliminating repeated cross-validation, but it still involves re-fitting a model for each bootstrap sample. In some situations, this may still be computationally prohibitive. Heuristically, if there is a large amount of training data available, one would expect the variation between models trained on bootstrap samples to be relatively small, so that a large part of the variation of the estimates is due to the noise in the treatment observations. Intuitively speaking, bootstrapping the treatment observations in such situations will be more

³³See Wager and Athey (2015) for some recent advances on limiting distributions of Random Forests.

important for capturing the variance of the estimate, relative to bootstrapping the training data. This suggests the following three-step bootstrapping procedure for each individual:

1. Perform an initial hyper-parameter optimization using cross-validation on the full training data set $\mathcal{D}_{i,\text{tr}}$.
2. Generate a family of models $\{\hat{f}_{i,j}^0\}_{j=1}^{B_m}$, where $\hat{f}_{i,j}^0$ is trained on the j -th moving block bootstrap sample from the training data $\mathcal{D}_{i,\text{tr}}$ using the optimal hyper-parameters determined in step 1.
3. For each model $\hat{f}_{i,j}^0$ from step 2, generate B_t bootstrap estimates by sampling from the estimation data.

This procedure will generate $B = B_m B_t$ bootstrap estimates. What we mean by a bootstrap estimate in the context of general treatments requires some clarification. For each trained model $\hat{f}_{i,j}^0$, we draw B_t samples from the estimation data. The k -th sample $(X_{it}, R_{it}, Y_{it})_{i \in \mathcal{T}_{i,j,k}}$ consists of the covariate, treatment parameter and outcome at $|\mathcal{T}_i|$ random draws $\mathcal{T}_{i,j,k}$ from the treatment periods \mathcal{T}_i . If the samples in the estimation data are (approximately) independent, then we can obtain the $\mathcal{T}_{i,j,k}$ by simply drawing uniformly at random (with replacement) from \mathcal{T}_i — otherwise we may perform an appropriate block bootstrapping procedure. We will focus on the former case here, assuming that the treatment periods are sufficiently separated in time. For each j and k we compute the one-sample estimates

$$\delta_{it,j,k} := f_i^0(X_{it}) - \hat{f}_{i,j}^0(X_{it}) + \mu_i(X_{it}, R_{it}) + \epsilon_{it} + \gamma_{it}$$

for $t \in \mathcal{T}_{i,j,k}$ and obtain an estimate $\hat{\mu}_{i,j,k} : \mathcal{X} \times \mathcal{R} \rightarrow \mathbb{R}$ of μ_i by regressing the $\delta_{it,j,k}$ on the $(X_{it}, R_{it})_{t \in \mathcal{T}_{i,j,k}}$ (for the sample mean estimator in case of binary treatments, $\hat{\mu}_{i,j,k}$ would simply be the sample mean of the δ_{it}). The bootstrap distribution $\mathcal{B}_i(\hat{\mu}_i(x, r))$ of the CITE estimate $\hat{\mu}_i(x, r)$ is then given by

$$\mathcal{B}_i(\hat{\mu}_i(x, r)) = \{\hat{\mu}_{i,j,k}(x, r) : 1 \leq j \leq B_m, 1 \leq k \leq B_t\}$$

Based on this bootstrap distribution we can compute confidence intervals for $\hat{\mu}_i(x, t)$.

There are many different bootstrapping techniques that can be used to compute confidence intervals in our context. One of the most basic ones is the so-called percentile method from Efron and Tibshirani (1986), which has various extensions.

Percentile Method Fix $x \in \mathcal{X}$, $r \in \mathcal{R}$, and let $\beta_i := \mu_i(x, r)$ and $\hat{\beta} := \hat{\mu}_i(x, r)$. Denote by $\hat{B}_i(\hat{\beta})$ the empirical cumulative distribution function (cdf) of the bootstrap distribution $\mathcal{B}_i(\hat{\beta})$ of the estimate $\hat{\beta}$. Then

$$\text{CI}_P(\alpha_l, \alpha_u) = [\hat{B}^{-1}(\alpha_l), \hat{B}^{-1}(1 - \alpha_u)] \tag{4.22}$$

is a $1 - \alpha_l - \alpha_u$ confidence interval for β with asymptotic lower and upper coverages of $1 - \alpha_l$ and $1 - \alpha_u$, respectively, provided $\hat{\phi} \sim \mathcal{N}(\phi, \tau^2)$ for some for some monotone transformation $\hat{\phi} = g(\hat{\beta})$, $\phi = g(\beta)$.

Bias-Corrected Percentile Method Efron and Tibshirani (1986) also suggest a so-called bias-corrected (BC) percentile method that can account for a particular kind of bias. Let $\hat{z}_0 = \Phi^{-1}(\hat{B}(\hat{\beta}))$ where Φ^{-1} is the quantile function of a standard normal random variable (i.e. the inverse of the cdf Φ). Further, let $z^{(\alpha)}$ denote the $100 \cdot \alpha$ percentile point of a standard normal variate. Then

$$\text{CI}_{\text{BC}}(\alpha_l, \alpha_u) = [\hat{B}^{-1}(\Phi(2\hat{z}_0 + z^{(\alpha_l)})), \hat{B}^{-1}(\Phi(2\hat{z}_0 + z^{(1-\alpha_u)}))] \quad (4.23)$$

is a $1 - \alpha_l - \alpha_u$ confidence interval for β with asymptotic lower and upper coverages of $1 - \alpha_l$ and $1 - \alpha_u$, respectively. Efron (1982) shows that the bias-corrected bootstrap interval for $\hat{\beta}$ is exactly correct if $\hat{\phi} \sim \mathcal{N}(\phi - z_0\tau, \tau^2)$ for some monotone transformation $\hat{\phi} = g(\hat{\beta})$, $\phi = g(\beta)$ and some constant z_0 .

Bias-Corrected Accelerated Percentile Method A further modification of the above methods is provided the bias-corrected and accelerated (BC_α) method (Efron and Tibshirani, 1994). This method is quite similar to the BC method, and considers confidence intervals of the form

$$\text{CI}_{\text{BC}_\alpha}(\alpha_l, \alpha_u) = [\hat{B}^{-1}(\Phi(\hat{z}_0 + \tilde{z}(\alpha_l, \hat{z}_0))), \hat{B}^{-1}(\Phi(\hat{z}_0 + \tilde{z}(\alpha_u, \hat{z}_0)))] \quad (4.24)$$

where \hat{z}_0 as before and where

$$\tilde{z}(\alpha, \hat{z}_0) = \frac{\hat{z}_0 + z^{(\alpha)}}{1 - \hat{a}(\hat{z}_0 + z^{(\alpha)})}$$

The acceleration \hat{a} can be computed using a jackknife estimator (see eq. (14.15) in Efron and Tibshirani (1994)). The main advantage of the BC_α method over the (bias-corrected) percentile method is that it is second-order rather than just first-order accurate.³⁴

The three methods discussed above are of increasing generality. Indeed, it is easy to see that (4.24) reduces to (4.23) if $\hat{a} = 0$. Similarly, (4.23) reduces to (4.22) if $\hat{z}_0 = 0$.

Note that in order to obtain confidence intervals at different values (x, r) it is not necessary to re-train the models $\hat{f}_{i,j}^0$ and $\hat{\mu}_{i,j,k}$. Instead, it is enough to determine the $\hat{\mu}_{i,j,k}$ once and then evaluate them at the different covariates and treatment parameters of interest. The full three-step procedure for computing confidence intervals at the elements of $\mathcal{E} := (x_m^{\text{ev}}, r_m^{\text{ev}})$ is given in Algorithm 2.³⁵

³⁴Meaning that $\mathbb{P}(\beta < \hat{\beta}) = \alpha_l + c_l/n$ and $\mathbb{P}(\beta > \hat{\beta}) = \alpha_u + c_u/n$ rather than $\mathbb{P}(\beta < \hat{\beta}) = \alpha_l + c'_l/\sqrt{n}$ and $\mathbb{P}(\beta > \hat{\beta}) = \alpha_u + c'_u/\sqrt{n}$ for some constants c_l, c_u and c'_l, c'_u , where n is the sample size.

³⁵Here we let $*$ denote one of the bootstrapping methods from the previous section, i.e. percentile, bias-corrected and bias-corrected and accelerated.

Algorithm 2: Three-step Bootstrapped Confidence Intervals for the CITE

Input: $\mathcal{D}_{i,\text{tr}}, \mathcal{D}_{i,\text{est}}, B_m, B_t, l, \mathcal{E}, \alpha_l, \alpha_u$

Output: $(1 - \alpha_l, 1 - \alpha_u)$ CIs around $\hat{\mu}_i(x_m^{\text{ev}}, r_m^{\text{ev}})$ for each $(x_m^{\text{ev}}, r_m^{\text{ev}}) \in \mathcal{E}$
 Perform hyper-parameter optimization of f_i^0 using cross-validation on $\mathcal{D}_{i,\text{tr}}$

Compute $\hat{\mu}_i$ according to Algorithm 1

for $j = 1, \dots, B_m$ **do**

Draw overlapping MBB sample $\mathcal{D}_{i,\text{tr},j}$ from $\mathcal{D}_{i,\text{tr}}$ of block length l

Train model $\hat{f}_{i,j}^0$ on $\mathcal{D}_{i,\text{tr},j}$ using previously determined hyper-parameters

for $k = 1, \dots, B_t$ **do**

Draw bootstrap sample $\mathcal{T}_{i,j,k}$ of size $|\mathcal{T}_i|$ from \mathcal{T}_i

For each $t \in \mathcal{T}_{i,j,k}$ compute $\delta_{it,j,k} = Y_{it} - \hat{f}_{i,j}^0(X_{it})$

Determine $\hat{\mu}_{i,j,k}$ by regressing the $(\delta_{it,j,k})_{t \in \mathcal{T}_{i,j,k}}$ on $(X_{it}, R_{it})_{t \in \mathcal{T}_{i,j,k}}$

for each $(x_m^{\text{ev}}, r_m^{\text{ev}}) \in \mathcal{E}$ **do**

Compute $\text{CI}_*^{i,m}(\alpha_l, \alpha_u)$ around $\hat{\beta}_{i,m} := \hat{\mu}_i(x_m^{\text{ev}}, r_m^{\text{ev}})$ according to CI_* , with

$\hat{B}_{i,m} = \text{cdf}\{\hat{\beta}_{i,j,k,m} : j = 1, \dots, B_m, k = 1, \dots, B_t\}$

Remark 7. While the confidence intervals computed by Algorithm 2 are individually asymptotically valid, they are not jointly valid. That is, $\mathbb{P}(\beta_{i,m} \in \text{CI}_*^{i,m}(\alpha_l, \alpha_u)) \rightarrow 1 - \alpha_l - \alpha_u$ for each l with increasing sample and bootstrap sample size under the conditions discussed by Efron (1982), but $\mathbb{P}((\beta_{i,m} \in \text{CI}_*^{i,m}(\alpha_l, \alpha_u)) \cap (\beta_{i,l'} \in \text{CI}_*^{i,m'}(\alpha_l, \alpha_u))) \not\rightarrow 1 - \alpha_l - \alpha_u$ in general.

By the *coverage* of a confidence interval we mean the empirical frequency with which the confidence interval includes the true value in a set of (independent) realizations of the underlying data generating process. We evaluate the coverage of the confidence intervals obtained by Algorithm 2 using synthetic experiments on real data in Section 5.4 of the following chapter.

4.6.2 Confidence Intervals for EATT and ATT

For the EATT we estimate the average treatment effect within a given population, and thus we do not need to be concerned about the randomness associated with how individuals were selected into the population. Therefore, computing bootstrap confidence intervals for the EATT estimate $\hat{\mu}$ can be done in a straightforward fashion based on the three-stage bootstrapping procedure in Algorithm 2. In fact, a bootstrap sample for $\hat{\mu}$ can be obtained by simply computing one bootstrap sample of the ITE estimate for each of the individuals in the sample population \mathcal{P}_T , and then averaging across those.³⁶ Algorithm 3 formalizes this.

³⁶If we have already computed bootstrap distributions for each individual, we may simply perform an i.i.d. draw from each of the individual bootstrap distributions and then average over these samples.

Algorithm 3: Bootstrapped Confidence Intervals for the EATT

Input: $\{\mathcal{D}_{i,\text{tr}}\}_{i \in \mathcal{P}_T}$, $\{\mathcal{D}_{i,\text{est}}\}_{i \in \mathcal{P}_T}$, B_m , B_l , \mathcal{E} , α_l , α_u

Output: $(1 - \alpha_l, 1 - \alpha_u)$ CIs around $\hat{\mu}_{\mathcal{P}_T}(x_m^{\text{ev}}, r_m^{\text{ev}})$ for each $(x_m^{\text{ev}}, r_m^{\text{ev}}) \in \mathcal{E}$

for $i \in \mathcal{P}_T$ **do**

Compute $\hat{\mu}_i$ according to Algorithm 1

Determine BS distribution $\hat{B}_{i,m} = \{\hat{\beta}_{i,j,k,m} : j = 1, \dots, B_m, k = 1, \dots, B_t\}$ as in Algorithm 2.

for each $(x_m^{\text{ev}}, r_m^{\text{ev}}) \in \mathcal{E}$ **do**

Compute $\text{CI}_*^m(\alpha_l, \alpha_u)$ around $\hat{\beta}_m := \hat{\mu}_{\mathcal{P}_T}(x_m^{\text{ev}}, r_m^{\text{ev}}) = \frac{1}{|\mathcal{P}_T|} \sum_{i \in \mathcal{P}_T} \hat{\mu}_i(x_m^{\text{ev}}, r_m^{\text{ev}})$

according to CI_* , with $\hat{B}_m = \text{cdf}\{\hat{\beta}_{j,k,l} : j = 1, \dots, B_m, k = 1, \dots, B_t\}$ where

$\hat{\beta}_{j,k,m} = \frac{1}{|\mathcal{P}_T|} \sum_{i \in \mathcal{P}_T} \hat{\mu}_{i,j,k}(X_m^{\text{ev}}, R_m^{\text{ev}})$

When considering the ATT estimate, it is important that, in addition to the statistical variation in the ITE estimates of the treated individuals, one also accounts for the variation in the estimate introduced by the randomness in the selection of the sample population. Assuming that the draws from the underlying population are i.i.d.,³⁷ we can layer an i.i.d. bootstrap across individuals on top of the bootstrapping procedure for the ITE estimates in Algorithm 2. Specifically, to generate a bootstrap sample of the ATT estimate $\hat{\mu}$, we draw a set of $|\mathcal{P}_T|$ “individuals” from \mathcal{P}_T uniformly at random (with replacement). For each of these bootstrapped individuals, we generate a bootstrap sample of the ITE estimate $\hat{\mu}_i$. The average of these across individuals then forms the bootstrap sample for $\hat{\mu}$. Algorithm 4 provides the details.

Remark 8. *If we are only interested in the bootstrap distribution of the ATT estimate $\hat{\mu}$, and not the bootstrap distributions of the different ITE estimates $\hat{\mu}_i$, we can potentially speed up computation by avoiding to run Algorithm 2 for each individual. Instead, we can generate the bootstrap estimates $\hat{\mu}_{i,j,k}$ online, caching the intermediate models for later use. Conversely, if we are interested in the bootstrap distributions of both the ITEs and the ATT, rather than generating ITE bootstrap samples online, we can simply draw from the individual bootstrap distributions already computed. That is, for each bootstrap population draw from \mathcal{P}_T , we for each individual i draw one sample uniformly at random from the respective bootstrap distribution \hat{B}_i of the ITE estimate $\hat{\mu}_i$.*

³⁷Note that for this we do not need to assume that the assignment to \mathcal{P}_T is uniform. Further, if draws were not independent, we could use an appropriate block bootstrapping procedure for this, provided we have some information about the possible correlation structure of ITEs across individuals.

Algorithm 4: Bootstrapped Confidence Intervals for the ATT

Input: $\{\mathcal{D}_{i,\text{tr}}\}_{i \in \mathcal{P}_T}$, $\{\mathcal{D}_{i,\text{est}}\}_{i \in \mathcal{P}_T}$, B_m , B_l , \mathcal{E} , α_l , α_u

Output: $(1 - \alpha_l, 1 - \alpha_u)$ CIs around $\hat{\mu}(x_m^{\text{ev}}, r_m^{\text{ev}})$ for each $(x_m^{\text{ev}}, r_m^{\text{ev}}) \in \mathcal{E}$

for $i \in \mathcal{P}_T$ **do**

- ┌ Compute $\hat{\mu}_i$ according to Algorithm 1
- ┌ Determine BS distribution $\hat{B}_{i,m} = \{\hat{\beta}_{i,j,k,m} : j = 1, \dots, B_m, k = 1, \dots, B_l\}$ as in Algorithm 2.

for $p = 1, \dots, B_p$ **do**

- ┌ Draw i.i.d. bootstrap sample \mathcal{P}_p of size $|\mathcal{P}_T|$ from \mathcal{P}_T

for each $(x_m^{\text{ev}}, r_m^{\text{ev}}) \in \mathcal{E}$ **do**

- ┌ Compute $\text{CI}_*^m(\alpha_l, \alpha_u)$ around $\hat{\beta}_m := \hat{\mu}(x_m^{\text{ev}}, r_m^{\text{ev}}) = \frac{1}{|\mathcal{P}_T|} \sum_{i \in \mathcal{P}_T} \hat{\mu}_i(x_m^{\text{ev}}, r_m^{\text{ev}})$ according to CI_* , where $\hat{B}_m = \text{cdf}\{\hat{\beta}_{p,m} : p = 1, \dots, B_p\}$ with $\hat{\beta}_{p,m} := \frac{1}{|\mathcal{P}_T|} \sum_{i \in \mathcal{P}_p} \hat{\beta}_{i,m}^p$ where each $\hat{\beta}_{i,m}^p$ is an i.i.d. random draw from $\hat{B}_{i,m}$

As for the CITE, we evaluate the coverage of the confidence intervals for the ATT obtained by Algorithm 4 using synthetic experiments on real data in Section 5.4 of the following chapter.

4.6.3 Computation of Bootstrap Confidence Intervals

The task of training models on moving block bootstrap samples of the treatment data and subsequently using them to generate bootstrap estimates of the treatment effect is trivially parallelizable. We make ample use of this fact in our implementation in order to speed up computation on multi-core machines. Our implementation is in python, employing the `numpy` package for fast numerical computations and the `scikit-learn` package as a general purpose Machine Learning library providing a broad range of standard options for the models \hat{f}_i^0 and $\hat{\mu}_i$. We plan to release our implementation as open source as soon once it has been appropriately refactored.³⁸

³⁸At this point the software package still contains confidential information from our industry collaborator, which will need to be removed before we can make the software made publicly available.

Chapter 5

Estimating Individual Treatment Effects in a Residential Demand Response Program

In this chapter, we apply the framework developed in the previous chapter to data from a Residential Demand Response program in California. The research presented here is the first step in a larger ongoing research project conducted in collaboration with a Demand Response provider in California.

We begin by reviewing relevant work in the extant literature on estimating treatment effects of dynamic electricity pricing among residential customers. We then lay out the scope and describe the core elements of the ongoing field experiment. The design of this experiment is based on the analysis of an initial, non-experimental data set of 5,000 of the DR provider's customers, to which we apply our two-stage ITE estimation strategy. We first study how our algorithms perform on semi-synthetic data, that is, data in which different simulated experiments are generated by adding synthetic treatment effects to customers' underlying (non-treatment) electricity consumption data. This allows us to benchmark our results against a known synthetic ground truth, while at the same capturing the true variation of the underlying data generating process for the nominal consumption. We analyze the properties of the estimation error under a number of different models, and show that the confidence intervals for both ITEs and ATT generated by our multi-stage bootstrap provide good empirical coverage. We also study how well our estimates are able to characterize heterogeneity in the treatment effect across different individuals, a task we refer to as *ranking*. Moreover, we investigate the potential for *targeting*, i.e., performing the DR dispatch conditional on characterizations of customer-level heterogeneity of the treatment effect.

We apply our algorithms to (non-experimental) data with actual treatment exposures, and estimate individual treatment effects of receiving a Demand Response notification on customers' energy consumption during the DR period. We find a statistically significant reduction in consumption during DR periods for the sample population, although at 6.6% our ATT estimate is on the lower end of the range of estimates reported by previous related

experiments (Faruqui and Sergici, 2010). However, there is significant heterogeneity in the treatment effect across different individuals, which we are able to characterize using our ITE estimates. Finally, we use these ITE estimates to characterize treatment effect heterogeneity across different customer-level covariates including weather, consumption level, and the presence of home automation technology.

5.1 Related Work

There have been a number of empirical studies on the effect of various dynamic pricing schemes on electricity consumption. We focus here on those studies that were conducted among residential customers within the setup of a randomized experiment.

Multiple experiments on dynamic electricity pricing took place in the 1970s and 1980s, focusing on simple seasonal and ToU rates (Faruqui and Malko, 1983; Caves et al., 1984). While these studies did find that customers responded to higher prices during peak pricing periods by curtailing and/or deferring consumption, they for various reasons did not have a major impact on policy and rate design.¹ Following an extended hiatus, a large number of dynamic pricing studies have since been conducted in the aftermath of the California Electricity Crisis (Borenstein, 2002). In the following, we discuss some of these studies.

Wolak (2006) estimates the effect of Critical Peak Pricing on electricity consumption in a CPP experiment conducted among residential electricity customers in Southern California in the summer of 2005. Within a small sample population of 123 customers, he estimates that CPP events on average led to a 12 percent reduction in consumption during peak hours of the event day. He is not able to detect any spillover effects, that is, in his data the impact of CPP events is confined to the peak periods of CPP days. He does find some evidence for a larger reduction in consumption during warm CPP days.

Herter and Wayland (2010) study a CPP program among residential customers in the California Statewide Pricing Pilot conducted between 2003 and 2004. They estimate the effect of CPP on the reduction of electricity consumption during peak periods to be 5.1%. They also find that an increase of the CPP rate from \$0.50/kWh to \$0.68/kWh did not elicit a larger response from the participants, and conclude from this that customer response to CPP comes mainly from cutting discretionary end-use loads.

Faruqui and Sergici (2010) collect and survey the results of 15 studies on dynamic electricity pricing. They conclude that households indeed respond to higher prices by lowering electricity consumption, and that the magnitude of this response depends on a number of different factors, including the level of the price incentive itself, the load type of the customer (in particular whether the household has central air conditioning), and whether the price incentive is accompanied by additional enabling technologies, such as home automation devices. If no additional enabling technologies are used, they find the reported effects to be

¹Faruqui and Sergici (2010) name the high cost of time-of-use metering, the overly broad peak periods in the ToU rates, and marketing failures on the side of the utilities as reasons for the limited impact of these studies.

between 3 - 6% for ToU, and between 13 - 20% for CPP tariffs. With enabling technologies such as Programmable Communicating Thermostats (PCTs) or In-Home Devices (IHDs), the effect for CPP tariffs increases to between 27 - 44%. The same authors also provide a similar meta-study about the effect of informational feedback through IHDs on energy consumption, which aggregates results of a dozen trials conducted primarily in North America (Faruqui et al., 2010). The focus of the interventions in these trials is more one of energy efficiency (i.e. reduction of total consumption over a longer horizon), rather than peak load reduction, though some of them also involve time-varying rates. The average of the reported estimates of the effect of deploying IHDs on energy consumption is 7% under regular tariffs, and 14% if the customer is on an electricity prepayment system.

Allcott (2011) uses a randomized field experiment to evaluate a RTP pilot program that took place in 2003 among approximately 700 residential households in Chicago. In this program the, prices for the different hours of the day were made available to consumers by 4pm the day before, and notifications about particularly high prices were communicated through special day-ahead alerts. Allcott (2011)'s primary estimates are the demand elasticity with respect to deviation from average seasonal price (17.4W and 21.8W/(cents/kWh) for summer and non-summer), and with respect to the average hourly price (12W/ (cents/kWh) on average across hours).

Jessoe and Rapson (2014) use a randomized controlled trial conducted among 1,152 households in New England in July and August 2011 to test the effect of high-frequency information about residential electricity usage on the price elasticity of demand. They find that household exposed only to price treatments reduced their usage by between 0% and 7% on average during pricing events, while households who were also provided with an IHD exhibited usage reductions of 8% to 22%. They point out that this difference in effect is not due to price salience, and suggest that IHDs facilitate consumer learning and thus improves decision making when confronted with high prices. Unlike Wolak (2006), Jessoe and Rapson (2014) do find that the treatment effect to some degree spills over into non-event hours, both within an event day and to non-event days.

Ito et al. (2015) estimates the effect of both economic incentives and behavioral treatments (or "moral suasion") on household electricity consumption in a self-selected RCT conducted among roughly 700 households in Japan in 2012 and 2013. He finds that moral suasion resulted in a usage reduction of 8%, but that this effect diminished after the first few treatments. Economic incentives provided through Variable Critical Peak Pricing, on the other hand, lead to persistent usage reductions between 14% and 17% for the lowest and highest critical peak price, respectively. Ito et al. (2015) also provides estimates for the welfare gains from the two policies.

Harding and Lamarche (2016) evaluate the effect of ToU pricing on peak electricity consumption in a field experiment among around 1,000 households in the southern U.S. from July through September 2011. They find effects on the order of 10% for customers with and without IHDs, and a much larger effect, up to 48%, for customers with programmable thermostats. They observe substantial off-peak load shifting to evening and night hours among the participants with programmable thermostats. Harding and Lamarche (2016) also

analyze heterogeneity of the effect across demographics, weather, and usage distribution.

As the discussion above indicates, most of the empirical results in the literature are on ToU and CPP programs. Relatively few results on CPR programs have been reported. One of these is a study conducted by the Ontario Energy Board (2007), which compares the effect of ToU, CPP, and CPR tariffs in the context of a pricing experiment that took place in Ontario from August 2006 to February 2007. The estimates for the load reductions during summer critical peak hours are 17.5% for the CPR and 25.4% for the CPP tariff, and during the entire peak period they are 8.5% and 11.9% for CPR and CPP tariffs, respectively.

In a setting very similar to that of Harding and Lamarche (2016),² Bollinger and Hartmann (2016) estimate the price elasticity of consumers with different automation technologies facing ToU and VPP (variable peak pricing) tariffs. They find that home automation technology results in more elastic demand with greater reductions at high prices and expanded demand when prices are low, while the treatment effect associated with the information technology alone is caused by a permanent shift in demand, and does not exhibit any demand elasticity. The authors also perform a welfare analysis, concluding that while consumer surplus gains from the automation technology barely offset the cost, the supply-side benefits from deferred capacity investments are significant.

To estimate the potential for the utility to specifically target those households which realize the largest effects, Bollinger and Hartmann (2016) estimate individual treatment effects by matching households based on their empirical distribution of consumption in the pre-treatment period. Specifically, they construct an estimate of the counterfactual consumption during a particular hour for individual j by performing non-parametric locally linear regression on the Kolmogorov-Smirnov distance between the empirical cdf of pre-treatment consumption of individual j and the cdfs of the members of the control group.³ This across-subject matching methodology is interesting in its own right, though the authors do not discuss any of its theoretical properties, or what the benefit of this method is over the more straightforward “synthetic control” approach introduced by Abadie et al. (2010).

Most closely related to the contributions in this thesis is the work by Jessoe et al. (2015), who use a randomized controlled trial of a CPP program as a benchmark to evaluate the potential of obtaining causal estimates of treatment effects in non-experimental settings using high-frequency data. Similar to us, they argue that non-treatment periods within treated subjects may provide a valid counterfactual, and in some ways may actually be preferred to the counterfactual provided by the experimental control group. Specifically, they find that estimates based on within-subject variation are quite close to the experimental results in level and superior in precision. They also describe a simple specification test, performed on the control group and aimed at identifying circumstances under which their estimation

²In fact, it appears that Harding and Lamarche (2016) for their work use a subset of the data from Bollinger and Hartmann (2016).

³That is, the estimated counterfactual for individual j is a weighted average of the observations in the control group, where the weight for individual i is the Kernel-weighted Kolmogorov-Smirnov distance between the empirical cdfs of j and i .

strategy may be poorly suited.⁴

Finally, we also point to some of our prior work that employs methods similar to those developed in this thesis in a much more ad-hoc fashion. In (Zhou et al., 2016b) we use generic Machine Learning techniques in combination with non-parametric statistical methods to perform customer segmentation in a DR program according to characteristic load shapes; moreover, we relate non-experimental estimates of the reduction of consumption during DR events to the variability of a user’s consumption. Our results suggest that customers with more variable consumption patterns are more likely to reduce their consumption compared to customers with a more regular consumption behavior. In (Zhou et al., 2016a) we incorporate latent variables representing behavioral archetypes of electricity consumers into the process of short-term load forecasting, thereby differentiating between varying levels of energy consumption. The latent variables are constructed by fitting conditional mixture models of linear regressions and Hidden Markov Models (HMMs) on smart meter data of a DR program. Using these latent variables as covariates for prediction leads to a notable increase in the accuracy of short-term load forecasts compared to the case without latent variables. We estimate reductions during DR events conditional on the latent variables, and discover a higher level of reduction among users with automated smart home devices compared to those without.

5.2 A Large-Scale Randomized Controlled Trial

In this section we describe some relevant aspects of the Demand Response program of our industry partner, in the context of which our field experiment is currently being conducted, and discuss the main research goals of the experiment.

The Experimental Environment

The field experiment is being conducted among more than 10,000 residential electricity customers distributed across all geographic areas of California. These customers are a random sample from the (self-selected) customer base of our industry partner. Within the experiment population, the program will differ from those of other demand-side management programs implemented by large IOUs (e.g. PG&E’s “Smart Rate” Critical Peak Pricing program) in a number of ways.

Firstly, the mechanism with which participants will be notified of DR periods (the “dispatch mechanism”) will be unique in that DR periods are announced in (almost) real-time through email, text message, and smartphone app without prior notification. Most other programs on the residential level instead rely on some sort of notification ahead of time (typically day-ahead). The lack of such ahead notification makes it potentially much simpler

⁴Their test is necessary but not sufficient: a rejection provides evidence that a model may be misspecified, but a failure to reject does not ensure that a model is correctly specified. In that, it is similar to the Placebo tests we conduct in Section 5.5.2.

to analyze the effect of DR notifications on consumption, as one does not have to worry about changes to consumer behavior in anticipation of the DR period (however, customers can still engage in “load shifting,” i.e., defer some of their consumption to later periods). On the flip side, participants will not be able to use ahead notification to prepare for the DR event (e.g. by moving some of their electricity consumption to hours before the event), which might result in a smaller effect size.

Secondly, the incentive for customers for changing their energy consumption level will not be one-sided, but also include a downside risk. That is, unless participants opt out of a particular event, they will in fact be charged more for consuming above their baseline than under their base tariff. Compared to the asymmetric one-sided incentive of most CPR programs, this makes it much simpler to interpret changes in behavior as a demand curve.⁵

Research Goals

We use the field experiment to answer the following core research questions:

1. What is the causal effect of adopting home automation technology on customers’ energy usage during DR periods?
2. What is the “demand curve” of DR?
3. How does the performance of our proposed two-stage ITE estimators compare against that of benchmark experimental estimators?
4. What is the causal effect of “adaptive targeting?”

We provide a brief discussion of these research questions below — as the experiment is currently ongoing we cannot report any empirical results at this point (we do, however, provide empirical results from the observational data in our initial data set in Section 5.5).

The Effect of Adopting Home Automation Technology

As evident from Section 5.1, the effect of adopting home automation technology (such as PCTs or “smart plugs”) on the effect of residential electricity consumption has been of interest to many researchers. Most if not all work studying this question has been conducted in the context ToU or CPP programs. Our experiment, in contrast, offers the opportunity to estimate the effect in a CPR program. The challenge in our case is that, within the context of our industry partner’s product, we cannot randomly deploy a specific technology to customers, or deny customers the adoption of this technology. Instead, we use a Randomized Encouragement Design (see e.g. Hirano et al. (2000) and Duflo and Saez (2003)), through

⁵Thus, on the rational level the experiment can be interpreted as a Real-Time Pricing experiment. However, we suspect that the psychologically element of “rewards vs. prices” may turn out to be very important.

which we provide a randomly selected subset of the experiment population (the “encouraged” group) with a high financial incentive for installing and registering a PCT with the DR provider (in fact, we provide the thermostat for free). The rest of the (non-control) experiment population (the “non-encouraged” group) does not receive any incentives for installing a PCT. By comparing the consumption of the encouraged and non-encouraged groups during DR periods, we can estimate the causal effect of adopting home automation devices on the consumption level during these periods.

Estimating the Demand Curve of DR

One of the most fundamental concepts in economics is that of a demand curve, i.e., the relationship between price and demand for a good.⁶ Demand curves are rather elusive, for a variety of reasons it is typically quite hard to estimate them in practice. In our field experiment we have the unique opportunity to introduce a substantial amount of random variation in the reward levels (in \$/kWh) offered to the participants, both across individuals, and across periods within each individual. In most studies that involve such price variation, the variation is typically not random but correlated with the potential outcomes. For example, in a RTP pilot the price for electricity charged to the customer will generally correlate strongly with the system load. System load in turn is correlated with temperature (driven primarily by AC loads), and so the effect of the price is conflated with the load-dependent potential of reducing consumption, which renders it difficult to estimate the actual demand curve. As we are able to randomize rewards completely independently, we are not subject to these concerns. The design of our experiment involves five different reward levels, rather than only two or three (the typical number in studies on CPP/VPP). This allows us to determine the actual shape of the demand curve, rather than its value at only two points.

Benchmarking Non-Experimental Estimators

The assumptions underlying our within-subject matching estimation strategy are often hard to verify in practice. While we can conduct placebo tests to ensure we do not measure any phantom effects (see Section 5.5.2), the only way to truly validate our estimators is by benchmarking them against experimental or quasi-experimental estimates. One core goal of our field experiment is to do just this: Estimate average treatment effects⁷ under random assignment to treatment and control groups using a standard fixed effects regression model, and compare the estimates to those of our two-stage estimators that do not utilize the random group assignment.

With this validation we follow the spirit of LaLonde (1986), who provides a classic comparison benchmarking various non-experimental estimators against experimental estimates

⁶In the context of CPR programs, where participants are offered a reward for a reduction with respect to some baseline, we are interested in the level of reduction as a function of the reward offered. In that sense, we should really speak of an “opportunity cost curve.”

⁷Recall that in case of fully random treatment assignment ATT and ATE coincide.

in the context of the National Supported Work Demonstration (NSW), a study conducted during the 1970s in the United States. His comparison shows that, in the context of this study, many of the considered non-experimental estimators do not replicate the experimentally determined results. His article is widely cited by many economists as a reminder that researchers should be aware of the potential for specification errors in non-experimental evaluations. Similar comparisons have since been made by various other authors, including Heckman and Hotz (1989); Dehejia and Wahba (1999) and Smith and Todd (2005).

In a related comparison, Chetty et al. (2014) ask whether teachers impacts on students test scores, quantified by the so called “value-added” (VA) measure, are a good measure of their quality. They test for bias in VA approaches using previously unobserved parent characteristics and a quasi-experimental design based on changes in teaching staff. Using school district and tax records for more than one million children, they find that VA models which control for a students prior test scores provide unbiased forecasts of teachers impacts on student achievement.

In the context of residential electricity usage, our comparison is essentially the same exercise as the one conducted by Jessoe et al. (2015), but with a number of important differences. Firstly, our sample is much larger (more than 10,000 vs. 442 households), and our design exposes participants to a higher number of treatments (25-50 vs. 6 events), two factors that will significantly improve precision. Secondly, we use a number of advanced non-parametric Machine Learning techniques in our first stage model, which we expect to result in a significant reduction in variance of the individual treatment effect estimates, compared to the basic linear regression models used by Jessoe et al. (2015). Finally, we have access to a significant amount of pre-treatment data (at least 1 year per participant), while the data set used by Jessoe et al. (2015) only encompasses two months including treatment periods.

Estimating the Effect of Adaptive Targeting

Our ITE estimation methodology allows us, in principle, to identify heterogeneity in the treatment effect of a single individual. One goal of our experiment is to show that this information can be used in order to improve efficiency of the DR dispatch in an on-line fashion. The implicit underlying assumption here is that dispatching residential DR means competing for a scarce resource, namely the DR participant’s attention.⁸ The idea is that by using *adaptive targeting*, i.e., performing DR dispatch based on within-subject treatment effect heterogeneity, one can increase the aggregate response level for a given amount of resources (such as customer attention, total incentive payouts, etc.), which will result in overall efficiency gains. In Section 5.4.5 we demonstrate the potential of this adaptive targeting in a realistic setting by performing synthetic experiments on real data. Within the setup of our experiment, by exposing randomly chosen subpopulations to targeted messages

⁸Not making this assumption would basically mean that sending DR notifications every single hour of the day were a reasonable thing to do. Clearly this is not the case, and existing programs generally limit the number of dispatches to a relatively small number. For instance, PG&E’s “Smart Rate” CPP program will call a maximum of 15 CPP events per year (Pacific Gas and Electric Company, 2016a).

while continuing basic messaging to the rest, we can obtain causal estimates of the effect of targeting on the consumption during DR periods.

5.3 The Pre-Experiment Data Set

The empirical results in this section are based on an initial, non-experimental dataset provided by our industry partner in preparation of the field experiment. This data set is a random sample⁹ of around 5000 customers that have enrolled with the DR provider between June and September 2016.

5.3.1 Variables and Descriptive Statistics

In this section we present some basic information about the contents of the initial data set that we will use in the following sections to perform our semi-synthetic experiments and empirical evaluation.

Utility Association Table 5.1 contains the number of customers per IOU in the sample.^{10,11} Comparing the customer shares in the sample with those state-wide, we observe that SDG&E is overrepresented in the sample. This is likely due to the fact that the DR provider is concentrating its customer acquisition efforts in geographic areas with high penetration of AC among residential customers, as is the case for SDG&E’s service territory.

IOU	customers	% (sample)	% (state-wide)
PG&E	1962	40.6	45.8
SCE	1739	36.0	42.4
SDG&E	1126	23.3	11.9

Table 5.1: Customers by electric utility

Location For each customer, the data set includes the ZIP code and the LMP pricing node associated to the customer. Figure 5.1 shows the geographic location of the customers on the ZIP code level,¹² colored according to the associated IOU. The intensity of each marker corresponds to the number of customers in the sample with that ZIP code. Figure 5.1 also

⁹Excluded from the sample are a number of ZIP code areas in the inner San Francisco Bay Area.

¹⁰PG&E: Pacific Gas and Electric, SCE: Southern California Edison, SDG&E: San Diego Gas and Electric.

¹¹State-wide data taken from current electric service account numbers published by the IOUs.

¹²The geographic location for each customer was taken as the location of the centroid of the associated ZIP code tabulation area provided by the U.S. Census Bureau (United States Census Bureau, 2015).

includes the locations of the CIMIS weather stations, from which we obtained the weather data for our study (see Section 5.3.2).

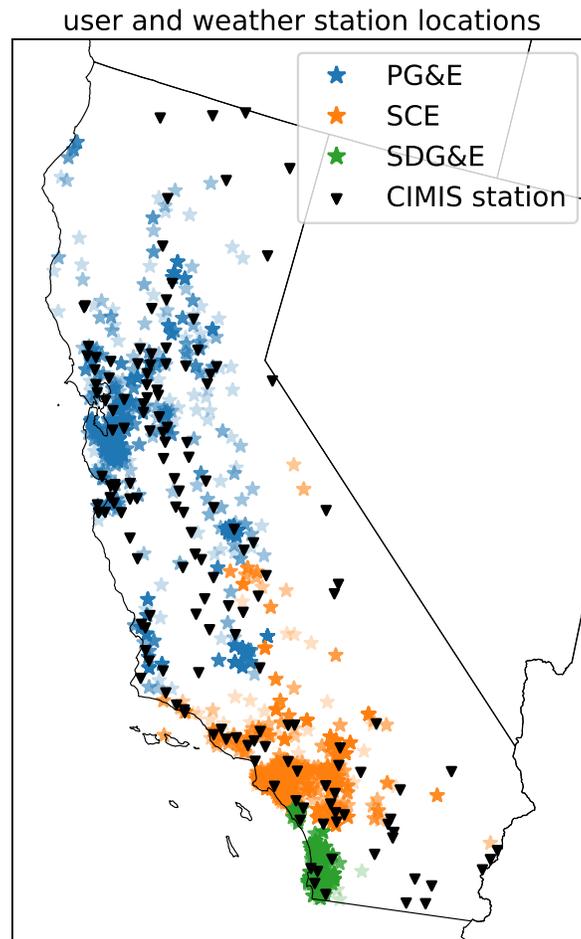


Figure 5.1: Geographic location of customers and weather stations

Smart Meter Electricity Consumption Data The core component of our data set are smart meter consumption readings for each customer on a 15min or hourly resolution. Since 15min data is available for only around 13% of the customers in the sample, we chose to restrict ourselves to hourly resolution and down-sampled the 15min observations by taking the total consumption within each hour of the day.

Figure 5.2 shows histograms of the length of the observation history per customer for each IOU. We observe that most customers have an observation history of less than 250 days, and that there are some clusters around longer availabilities. In particular, a group of SCE customers has around 500, and a group of SDG&E customer has around 850 days of historical data availability, while there are no PG&E customers with comparable long-term

availability. This is primarily due to ongoing challenges for third-party DR providers to obtain historical data from IOUs.¹³ More recent data from the ongoing field experiment shows that a consumption history of at least one year is available for most subjects in the experiment.

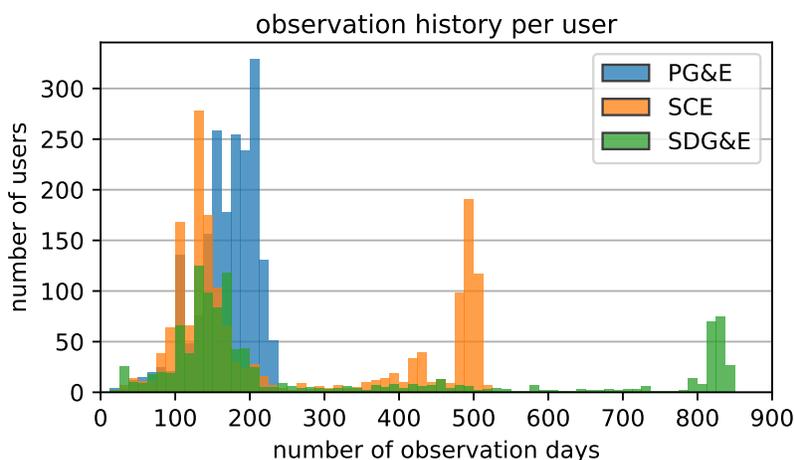


Figure 5.2: Histogram of length of observation histories

The data set also contains a number of customers that have photo-voltaic (PV) solar systems installed. This presents a challenge for estimating the response to DR notifications, as for many of these customers the meter reading is the net consumption¹⁴ from the point of view of the grid, i.e., actual consumption minus the power generated by the solar PV system. Unfortunately, at this point we do not have reliable information on which of the customers actually owns solar PV that is behind the smart meter and not metered separately. As a proxy, we exclude all customers with negative consumption readings, as well as all customers with an unusually high number of zero observations.

Demand Response Periods and Notifications The data set also contains all DR notifications sent to the customers in the sample. DR events are customer-specific, that is, not all customers are notified of every DR event. In this analysis we focus on DR events with a duration of one hour.¹⁵ Figure 5.3 shows the distribution of DR periods over the hours of

¹³ There are still a number technical challenges for third-party providers to interface directly with the IOUs' systems. The CPUC's decision 16-06-008 requires the utilities to hold working group meetings with stakeholders to develop a consensus proposal to streamline and simplify the direct participation enrollment process, including adding more automation, mitigating enrollment fatigue, and resolving any remaining electronic signature issues (California Public Utilities Commission, 2016).

¹⁴Under so-called net-energy metering (NEM), smart meters actually can report negative consumption values. We observe this for a small number of customers in our data set.

¹⁵This is the majority of DR events in the data set. The data set also contains a few events that are two hours long, and one event that is four hours long.

the day. This distribution looks largely like what one would expect, showing a concentration of events during the afternoon and evening hours when the system is under peak load and LMPs are more likely to spike to high levels.

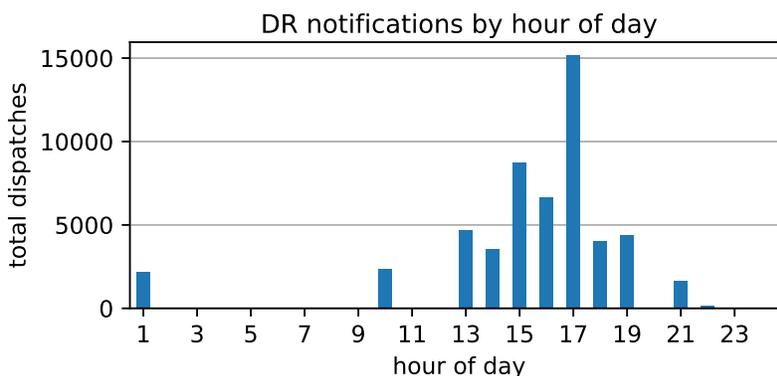


Figure 5.3: Frequency of DR notifications by hour of day

We point out that data in this initial data set was collected during periods in which customers were not fully salient of the reward incentive. That is, while reductions with respect to the DR baseline were in fact valued differently across different events ex-post, participants were not aware of the different marginal rewards at the time of making their consumption decision. This makes it quite hard to interpret the results – in particular, it makes it impossible to estimate price elasticities or a demand curve.¹⁶ Therefore the effect we estimate in this section is simply the average effect of receiving a DR notification on the consumption during the associated DR period.

Finally, we also observe through which channels a customer is notified about an event (email, SMS, smartphone app notification), but at this point we do not make use of this information.¹⁷

Data on Home Automation Technology DR participants have the option of registering a number of different home automation devices, such as programmable communicating thermostats (PCT), “smart plugs” and electrical vehicle (EV) charging stations, with the DR provider. Participants can choose to grant the DR provider authority to remotely control their device during customer-defined hours of the day, but have the option to opt out and override the remote decision at any time. Such an approach is typically referred to as “Direct Load Control” (DLC) in the literature (Albadi and El-Saadany, 2008; Palensky and Dietrich, 2011). The standard hypothesis is that customers with automated devices will exhibit larger

¹⁶In our field experiment, participants receive the reward level stated ex-ante at the time of the DR notification.

¹⁷For a more in-depth analysis it would be interesting to understand to what extent the response depends on the notification delivery method.

treatment effects, which has been confirmed by other studies conducted in similar settings (see our overview in Section 5.1). We will study this aspect in Section 5.5.3 of this chapter.

PCT	EV	other
151	15	78

Table 5.2: Penetration of home automation devices

Table 5.13 shows the penetration of active automated devices in the initial data set. The “other” category consists primarily of “smart plugs”, i.e., electrical outlets that can be remote controlled, typically via the internet. We point out that this data is not representative of the penetration of devices across the DR provider’s entire customer base. This is because the initial data set only includes customers that have signed up after June 1, 2016. Therefore many of these customers have likely not been exposed to information and incentives provided with the aim of increasing adoption and registration rates of home automation devices. For simplicity, we will refer to those customers with an active home automation device as “automated customers” and to those without as “non-automated customers.”

5.3.2 Weather and Locational Marginal Electricity Prices

In addition to the variables received directly by the DR provider, we obtained additional data on weather and LMPs for the time frame of interest. We use the weather data primarily as covariates to control for its effects on consumption on the customer level, but also employ it to uncover heterogeneity in the treatment effects between areas with different climates. We obtain the pricing data in order to perform analyses of the impact of the DR program on generation costs based on the data collected in the field experiment.

Weather We obtained hourly weather data¹⁸ for the full range of the consumption data for 154 active weather stations in the CIMIS system (California Department of Water Resources, 2015). The data was then spatially interpolated for each observation hour to the centroids of the ZIP code tabulation areas using barycentric interpolation. The accuracy of the resulting time series data on the ZIP code level is obviously limited, but as we use it primarily as a covariate for our prediction models, it is sufficient for our purposes since it is highly correlated with the actual local weather.¹⁹

¹⁸This data includes a variety of variables, including air temperature, precipitation, solar radiation, vapor pressure, relative humidity, dew point, wind speed and wind direction. At this point we only make use of the air temperature.

¹⁹Indeed, as we normalize the temperature in the covariates, any constant error in the temperature has no effect on the model. Furthermore, there is no reason to believe that systematic correlation between the *errors* in the temperature data and the electricity consumption of the customers should be a concern.

Locational Marginal Prices For each of the LMP pricing nodes associated to at least one customer in the data set, we scraped the Day-Ahead and Real-Time LMPs for the time frame from June 1, 2016 through September 30, 2016, using the CAISO API (California Independent System Operator Corporation, 2015b). We observe a relatively strong heterogeneity of LMPs both across the different DR periods and across different pricing nodes. Figure 5.4 shows the mean LMP across all DR periods for all pricing nodes in the sample. It is apparent how LMPs are much higher in Southern California, especially in the Los Angeles basin and San Diego areas. These areas have seen particularly high and volatile electricity prices in the last year, an effect driven primarily by two factors: The higher than usual temperatures and, importantly, the shortage of natural gas in the area in the aftermath of the Aliso Canyon gas leak that was discovered in October 2015.²⁰

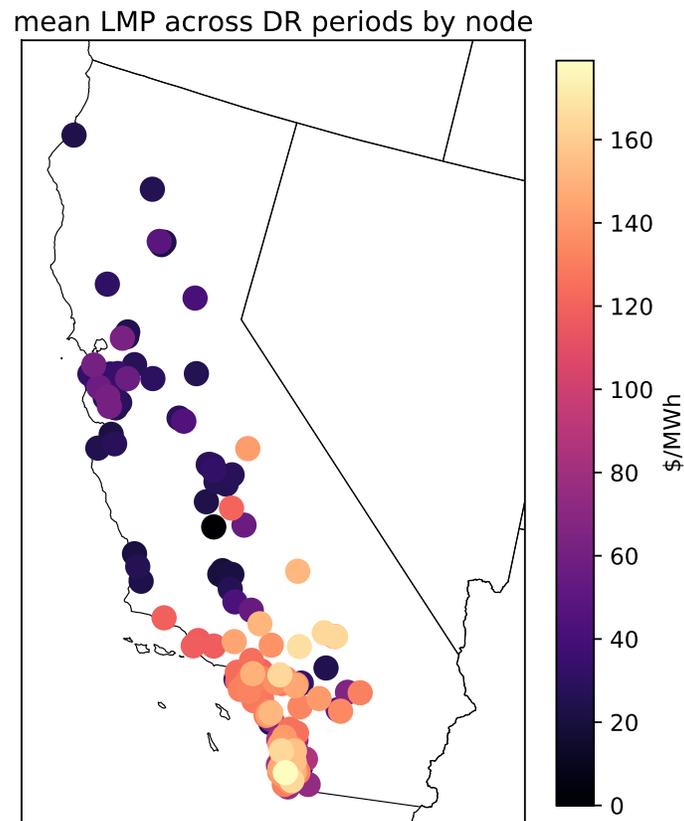


Figure 5.4: Heterogeneity of LMP across pricing nodes during DR periods

Figure 5.5 shows the mean LMP across all pricing nodes for the DR periods in the data

²⁰See California Independent System Operator Corporation (2016) for the effect of the Aliso Canyon leak on the electricity market.

set. While there are some very high price hours in June, most of the DR periods see an average LMP across pricing nodes of less than \$200/MWh. There is a high variance of prices across nodes within each DR period hour.²¹ This primarily driven by congestion on the tie line between NP-15 and SP-15, two main CAISO zones, which roughly correspond to Northern and Southern California.

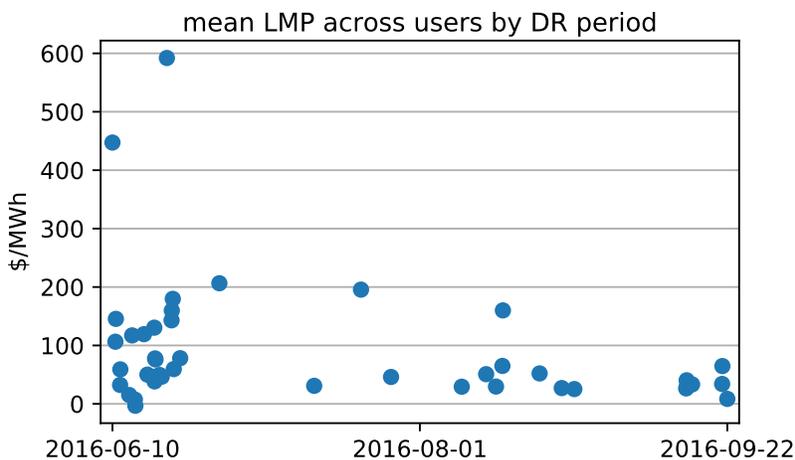


Figure 5.5: Heterogeneity of LMP across DR periods

5.3.3 Applying the Two-Stage ITE Estimation Framework

In the following sections we apply our estimation framework from Chapter 4 to the initial data provided by our industry partner. To do so, for a given customer we first need to split the data into training and estimation data sets. The estimation data set is relatively straightforward: For each participant these are the periods for which they received DR notifications, together with the associated covariates. We need to be more careful about the estimation data set. In particular, we need to be concerned about “load-shifting,” i.e., customers increasing their consumption in periods around the actual DR event. In the analysis of most DR programs, one needs to be concerned about customers shifting their consumption to hours both before and after the DR event. This is because participants usually receive a notification ahead of time (typically one day ahead). However, in our case customers do not receive such notifications ahead of time, which means we can regard the consumption up to and including the hour right before the DR event as unaffected by the event.²² Using these observations in our models f_i^0 means we need to predict consumption

²¹LMPs during DR periods range from $-\$3.24/\text{MWh}$ (11am on Jun 14, 2016 at node MIRALOMA_6_N001) to $\$990.81/\text{MWh}$ (7pm on Jun 21, 2016 at node ROADWAY_1_N001).

²²Here we are implicitly assuming that customers are not themselves trying to predict DR events. If subjects started building statistical models for the occurrence probability of DR events and basing their decisions on these models, then our assumptions would not be valid, and this problem would instead become

only one hour ahead, which results in much smaller variance of the prediction.²³

Figure 5.6 illustrates the basic estimation concept: From the consumption in the training periods, we generate the model \hat{f}_i^0 for the conditional mean function of the consumption in non-DR periods. For each DR period $t \in \mathcal{T}_i$ we then estimate the counterfactual consumption by predicting consumption using the model \hat{f}_i^0 from the associated covariates, which include lagged consumption values (here we make use of the fact that notifications are sent without prior warning, so the lagged consumption values are unaffected by treatment). The difference between the predicted and observed consumption values is our δ_{it} , the estimate of the demand reduction. As discussed in Section 4.4, we can then estimate the treatment response by regressing the δ_{it} on the variables of interest. In the most basic case of estimating the marginal ITE, we simply compute the sample mean.

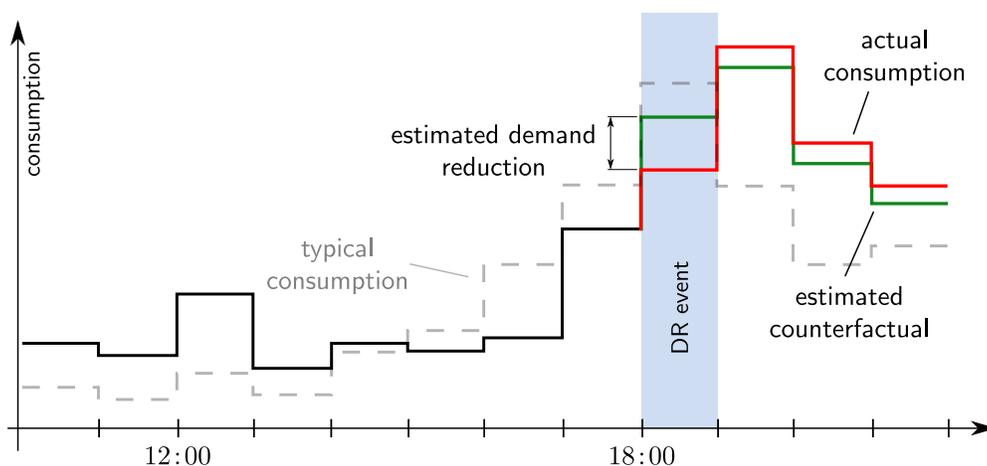


Figure 5.6: Predicting the counterfactual consumption during DR periods

5.4 Synthetic Experiments on Real Data

In order to validate our estimation strategy, we perform a number of synthetic experiments on real data. That is, we take the observed historical consumption time series, and for each customer add simulated responses for a number of simulated DR periods. We then use our two-stage estimation strategy to estimate both individual and average treatment effects.

The basic idea of these synthetic experiments is to capture the true underlying variation in consumption, while being able to benchmark the coverage of our confidence intervals using the known treatment effects. This provides us with a way to verify that we can perform correct statistical inference using our methods. Moreover, it helps us develop an

a rather interesting game-theoretic problem. However, it seems relatively safe to assume that the number of such customers is quite small.

²³As we will see in the following section, the variance of these estimates is still very significant.

understanding of the statistical power in this setting, which in combination with the non-experimental estimates of the treatment effects in Section 5.5 in turn informs the design of the treatment group assignment in our randomized field experiment.

5.4.1 Setup

We first need to determine the form of the treatment response. The most straightforward formulation is a *level* specification of the form

$$\text{kWh}_{it} = f_i^0(X_{it}) + D_{it}(\mu_i + \gamma_{it}) + \epsilon_{it} \quad (5.1)$$

where f_i^0 is the conditional mean function of the consumption kWh_{it} , and for which the treatment effect μ_i characterizes a simple absolute shift in consumption (in kWh) from that in the absence of treatment. Since electricity consumption is highly variable, and frequently very low during individual hours, such a simple additive effect may not be the most appropriate model. In particular, it is hard to compare different individuals with very different average consumption levels, as well as different DR periods with very different consumption levels within an individual.

Instead, it is common to assume a relative change in consumption with respect to the counterfactual consumption level. A standard model is the *semi-log* specification²⁴

$$\log(\text{kWh}_{it}) = f_i^0(X_{it}) + D_{it}(\mu_i + \gamma_{it}) + \epsilon_{it} \quad (5.2)$$

where f_i^0 is the conditional mean function of the logarithm of the consumption kWh_{it} . The interpretation as a relative change during treatment periods becomes clear when taking the exponential and writing its Taylor expansion around zero:

$$\begin{aligned} \text{kWh}_{it} &= \text{kWh}_{it}^{\text{nom}} \exp(\mu_i + \gamma_{it}) \\ &= \text{kWh}_{it}^{\text{nom}} (1 + (\mu_i + \gamma_{it}) + \mathcal{O}(\mu_i + \gamma_{it})) \end{aligned}$$

where $\text{kWh}_{it}^{\text{nom}} := \exp(f_i^0(X_{it}) + \epsilon_{it})$ is the counterfactual consumption in the absence of treatment. Hence for small μ_i the effect is very close to that of a relative change in consumption.²⁵ However, if μ_i is large, then this is not the case anymore, and one has to transform the estimate to obtain an estimate of the change in relative consumption.

From the continuous mapping theorem it follows that if $\hat{\mu}_i$ is a consistent estimator for μ_i in (5.2), then $g(\hat{\mu}_i) := \exp(\hat{\mu}_i)$ will be a consistent estimator for the relative consumption $\text{kWh}_{it}/\text{kWh}_{it}^{\text{nom}}$. However, the challenge with estimating ITEs is that we typically have relatively few observations and so the distribution of the estimate $\hat{\mu}_i$ has relatively high variance. From Jensen's inequality we have that $\mathbb{E}_{\mathcal{D}_{i,\text{est}}}[g(\hat{\mu}_i)] \geq g(\mathbb{E}_{\mathcal{D}_{i,\text{est}}}[\hat{\mu}_i])$. Thus we need to be cognizant of the fact that transforming the ITE estimate $\hat{\mu}_i$ may yield a negatively

²⁴In an effort to reduce notational burden, we use the same notation f_i^0 and μ_i for both the level and the semi-log specification. Which specification we are working with will be clear from the context.

²⁵Recall that the process γ_{it} is zero mean under our assumptions from Chapter 4.

biased estimate of the average relative change in consumption.²⁶ In the following, we will primarily use the semi-log specification for individual treatment response, keeping in mind that we cannot always simply interpret it as a percentage change in consumption.

We also need to be somewhat careful when marginalizing across different individuals to obtain estimates of the ATT. Again using Jensen’s inequality, we have that $g(\mathbb{E}_{\mathcal{P}}[\mu_i]) \leq \mathbb{E}_{\mathcal{P}}[g(\mu_i)]$. A classic across-subject matching regression model from the program evaluation literature estimates the ATT $\mathbb{E}_{\mathcal{P}}[\mu_i]$ in the semi-log specification (5.2), so in order to obtain the associated estimate from our methodology we first need to marginalize our ITE estimates $\hat{\mu}_i$ across subjects before transforming them to the relative change.²⁷

Note that if we have determined our confidence intervals via bootstrapping using the percentile method, we can also directly transform the endpoints of the confidence intervals. Indeed, since g is a monotonic transformation, the ordering in the transformed bootstrap distribution is the same (see Efron and Tibshirani (1994) for details on this *transformation respecting* property). In general this does not hold if we use other methods for determining the confidence intervals.²⁸

Specification of the First-Stage Model

Our first stage prediction models \hat{f}_i^0 make use of the following covariates:

- kWh_{*iτ*}: smart meter reading²⁹ of the total energy consumption in period τ
- $\vartheta_{i\tau}$: mean outside air temperature in period τ
- HoD_{*iτ*}: hour of the day
- MoY_{*iτ*}: month of the year
- BDay_{*iτ*}: indicator for whether the day is a business day

In order to make results comparable across customers also for prediction models that are not scale-invariant,³⁰ we standardize both outside air temperature and consumption data to have zero mean and unit variance. We use the lagged consumption values of the past five hours as well as the ambient air temperature in the current and two previous hours.

²⁶This subtlety is often overlooked in the empirical evaluation literature, where it usually is not a big concern, since there the distribution of the ATT estimate typically has much smaller variance.

²⁷Marginalizing the transformed estimates still yields an reasonable estimate, albeit one that is generally not considered in the literature.

²⁸For instance, if we were to assume that $\sqrt{n} \hat{\mu}_i \xrightarrow{d} \mathcal{N}(\mu_i, \sigma^2)$ and performed the associated normal approximation, then we could not transform the confidence intervals and would need to re-compute them appropriately. In particular, by the delta method we have $\sqrt{n} g(\hat{\mu}_i) \xrightarrow{d} \mathcal{N}(g(\mu_i), [g'(\mu_i)]^2 \sigma^2) = \mathcal{N}(g(\mu_i), \exp(\mu_i)^2 \sigma^2)$, and hence $g(\hat{\mu}_i \pm z^* \frac{\sigma}{\sqrt{n}}) \neq g(\hat{\mu}_i) \pm z^* \frac{\sigma}{\sqrt{n}} \exp(\mu_i)$, especially if $z^* \frac{\sigma}{\sqrt{n}}$ is large.

²⁹15-min readings are summed across each hour.

³⁰For example this includes Support Vector Regression (Smola and Schölkopf, 2004).

With this, the covariate (or feature) vector associated to the observation in period t for our non-parametric models is

$$X_{it} = [\text{kWh}_{it-5} \quad \dots \quad \text{kWh}_{it-1} \quad \vartheta_{it-2} \quad \vartheta_{it-1} \quad \vartheta_{it} \quad \text{HoD}_{it} \quad \text{MoY}_{it} \quad \text{BDay}_{it}] \quad (5.3)$$

We also consider parametric models, for which we explicitly specify HoD_{it} , MoY_{it} and BDay_{it} as categorical variables (we do not do this for our nonparametric models). For the linear models, we also consider the interaction between HoD_{it} and BDay_{it} in (5.3).³¹

This choice of features and specifications, which we have found to work quite well for a wide range of different customers, has been developed iteratively through a combination of domain knowledge, systematic exploration, and trial and error. We do not aim to give a detailed characterization of this process here — after all our framework is purposely model-agnostic, so it can benefit from whatever prediction model works “best” for the given application. For additional details on our model see Zhou et al. (2016b,a), where we also provide an analysis of the out of sample prediction performance.

Treatment Effect Models

We consider two treatment effect models. A simple *uniform* treatment effect $\mu_i^{\text{unif}} \in \mathbb{R}$, and a *linear* treatment effect of the form $\mu_i^{\text{lin}} = \mu_i^{\text{lin}}(R_{it}) = \beta_i R_{it}$, where R_{it} is the reward level (in \$/kWh) and $\beta_i \in \mathbb{R}$ is the price elasticity of the individual’s linear demand function. Note that the overall effect is always interpreted with respect to the specification: For the level specification the uniform effect characterizes a constant response, while for the semi-log specification it characterizes a relative response with constant rate.

We use the obvious second-stage models to estimate the parameters, namely the sample mean estimator from Section 4.3.2 for the uniform response level μ_i^{unif} , and a standard linear regression with response variable δ_{it} and covariate R_{it} for the price elasticity β_i .

5.4.2 Individual Treatment Effect Estimates

Setup

Our synthetic experiments take place from March 15 through June 1, 2016. We chose this time frame based on a number of different considerations: Since all customers in the data set signed up on or after June 1, this range does not include any DR periods, so we do not run the risk of contaminating the sample actual responses to DR notifications. With 10 weeks it is also of a duration that is on the order of what will be the timeframe of observation for a single customer in the actual field experiment (between 90-120 days). Finally, we chose this period as late as possible, as we are limited by the amount of historical consumption data that is available for a large subset of the sample (see also the discussion in Section 5.3.1).

This leaves us with a set of sample of 547 customers that have (i) no negative consumption values, (ii) historical consumption data available until at least June 1, 2016, and (iii) have at

³¹In the R formula mini language this would read $\text{HoD}_{it}:\text{BDay}_{it}$.

least 363 days of consecutive consumption readings before June 1, 2016 (the average number of days is 484, the maximum 1,062).

Within the synthetic study period, we randomly chose 25 synthetic DR periods according to the same joint distribution over business days and hour of the day as in the observed DR periods after June 1 (see Figure 5.3 in the previous section). We choose 25 events to be consistent with the design of our field experiment. For our synthetic experiments all DR events are shared by all participants.

We specify synthetic responses in the semi-log specification (5.2) for both the uniform and the linear treatment effect model introduced in the previous section. For each simulation, we randomly generate parameters by i.i.d. sampling $\mu_i^{\text{unif}} \sim U[-0.3, 0]$ and $\beta_i \sim U[-0.15, 0]$, respectively. The associated ATTs are $\mu^{\text{unif}} = \mathbb{E}[\mu_i^{\text{unif}}] = -0.15$ and $\beta = \mathbb{E}[\beta_i] = -0.075$, respectively.³²

According to the above synthetic response models, we generate synthetic consumption observations $\tilde{Y}_{it} = Y_{it} + D_{it}\mu_i^{\text{unif}}$ and $\tilde{Y}_{it} = Y_{it} + D_{it}\mu_i^{\text{lin}}(R_{it})$ for all periods t and individuals i , where the rewards R_{it} are i.i.d. draws from $U[0, 3]$. This range is informed by the RCT, in which we will expose participants to reward levels ranging from \$0.05/kWh (essentially zero) to \$3.00/kWh.³³ With this, the maximum possible reduction is $1 - \exp(-0.3) = 25.9\%$ under the uniform and $1 - \exp(-0.15 \cdot 3) = 36.2\%$ under the linear response model.

We use three different first-stage models: L_2 -regularized linear regression (Ridge), K-Nearest Neighbors Regression (KNN), and Random Forest Regression (RFR) using 200 individuals estimators. For each model, we optimize the associated hyperparameters³⁴ using 4-fold cross validation on the training set. As our approach is designed to be model-agnostic, we do not discuss these prediction methods in detail at this point; we point to Hastie et al. (2009) for a general overview and to Zhou et al. (2016a,b) and Mirowski et al. (2014) for a discussion in the context of short-term load forecasting of residential energy consumption.

Minimizing Conditional Bias

An important characteristic of the DR events is that they are not randomly distributed across the year, but that they occur primarily during afternoon and evening hours of the summer months. Therefore, a model that is trained by minimizing some measure of the average error across the whole training data set (e.g. the MSE) may potentially show conditional bias during those periods (see our general discussion of this issue in Section 4.4). This of course poses a potentially major problem for statistical inference.

A sufficiently expressive model class \hat{f}_i^0 should be able to estimate the correct conditional mean as the amount of relevant training data increases.³⁵ However, in our case the amount

³²For brevity we are being somewhat imprecise about the linear effect model here: As we would like to estimate the parameters β_i , we refer to the average value of β_i in the underlying population as the ATT, while the ATT in the strict sense is really $\mathbb{E}[\beta_i R_{it}]$ and thus depends on the distribution of rewards.

³³In the experiment rewards will not be randomized from a continuum, but draws from a discrete set.

³⁴For Ridge this is the regularization parameter α , for KNN it is the number of neighbors, and for RFR it is the maximum depth of each individual tree as well as the minimum sample size in a leaf node.

³⁵This implicitly assumes that the outcomes associated to this data are not or only weakly correlated.

of relevant training data is limited, due not least to the fact that we are unable to use observations during and after DR periods. Instead, we have to find other ways to ensure this conditional bias is minimized.

There are multiple ways to address this issue. The most basic one is to simply only consider training data from “similar” periods. The problem with this approach is that all the information contained in the rest of the training data is disregarded, which increases the variance of the estimate. An alternative approach is to keep the model fitting process untouched, but to correct the model prediction with an estimate of the conditional bias obtained on an independent subset of the training data. Such an estimate of the bias of course will itself be noisy, so that the overall variance of the estimate increases. Finally, as a “soft” version of the first approach, one can also weight the individual samples in the training data according to the distribution of the treatment periods (or some transformation thereof). Clearly, the variance of the resulting estimate will increase the more localized (i.e. less uniform) these weights are.

The above discussion illustrates the fundamental bias-variance tradeoff in any estimation problem. When considering individuals, the variance of an ITE estimate is typically quite large due to the limited number of treatment observations per individual, so allowing for some bias while reducing variance may significantly reduce the MSE of the ITE estimate. Conversely, when estimating the EATT or ATT of a large group of individuals, then averaging over individuals reduces the variance of the estimate by $1/N$ (assuming the ITE estimates are independent), but any systematic bias will persist. Thus, if we are interested in ATT estimates, we should be concerned to minimize the conditional bias within each individual, in particular if this bias is systematic across individuals.

	mean	std dev
basic	-0.01447	0.09956
bias-corrected	-0.01562	0.09952
hour-adjusted	-0.01180	0.10040
weighted	-0.00603	0.10542

Table 5.3: Estimation errors $\hat{\mu}_i - \mu_i$ under different methods

Table 5.3 shows the estimates³⁶ of the mean and standard deviation of the estimation errors $\hat{\mu}_i - \mu_i$ under the methods described above. Here “basic” is the standard estimate from the first-stage model \hat{f}_i^0 trained on the training set with uniform sample weights, “bias-corrected” uses the same first-stage model but corrects the estimate using an estimate of the conditional bias obtained on an independent subset of the training data, “hour-adjusted” restricts the training set to observations made during hours of the day which saw DR periods, and “weighted” uses a model with non-uniform sample weights.

³⁶These estimates are based on 4376 independent ITE estimations among 547 different customers.

For the weighted model, we use sample weights w_{it} given by

$$w_{it} \propto w_{it}^{\text{HoD}} w_{it}^{\text{MoY}} \quad (5.4)$$

where

$$w_{it}^{\text{HoD}} = \eta^{\text{HoD}} + \sum_{\tau \in \mathcal{T}_i} \mathbf{1}_{\{\text{HoD}_{it} = \text{HoD}_{i\tau}\}} \quad (5.5a)$$

$$w_{it}^{\text{MoY}} = \eta^{\text{MoY}} + \sum_{\tau \in \mathcal{T}_i} \mathbf{1}_{\{\text{MoY}_{it} = \text{MoY}_{i\tau}\}} \quad (5.5b)$$

with parameters $\eta^{\text{HoD}} > 0$ and $\eta^{\text{MoY}} > 0$. Intuitively, this method assigns the highest weight to hours that are most likely to be DR periods, according to the empirical distribution of DR events in the estimation data. The parameters η^{HoD} and η^{MoY} “mix in” a uniform distribution to ensure that all observations in the training data are used for fitting the model. By increasing η^{HoD} and η^{MoY} , the resulting distribution becomes more uniform, and the model increasingly “borrows strength” from adjacent hours and months. As $\eta^{\text{HoD}}, \eta^{\text{MoY}} \rightarrow \infty$, the weights become uniform and thus the weighted model approaches the basic model. As usual, the tradeoff is one between bias and variance: By increasing the parameters η^{HoD} and η^{MoY} we reduce the variance of the estimate at the cost of increasing the bias. The choice depends on the particular problem at hand — if the goal is to estimate the ATTE within a large group of individuals, one will be typically worried more about bias in the ITE estimates than about variance, while in a situation where the goal is to distinguish the ITE of different individuals, variance is typically the main concern (so long as there is reason to believe that model bias behaves similarly for different individuals).

The results in Table 5.3 were obtained with $\eta^{\text{HoD}} = 0.25 \frac{|\mathcal{T}_i|}{24}$ and $\eta^{\text{MoY}} = 0.25 \frac{|\mathcal{T}_i|}{12}$. We can observe that weighting the training data results in a significant reduction in bias. The effect of restricting the training data to the same hours has a positive, but smaller effect. The method of estimating and correcting the bias does not seem to provide any benefits in this setting. In general, the weights η^{HoD} and η^{MoY} can be obtained by using cross-validation on the training data.

Results

For each effect model, we estimate the ITE for each individual in the sample population using our two-stage estimation procedure under the different first-stage models discussed above, and compute associated bootstrapped confidence intervals. We use the weighting approach for bias reduction from the previous section for both the Ridge and RFR first-stage models, but not for the KNN first-stage model. We analyze the distribution of the estimation errors, which we can compute based on the “ground truth” ITEs μ_i , as well as the coverage of the confidence intervals.

Uniform Response Figure 5.7 plots 4,376 ITE estimates $\hat{\mu}_i$ against the associated true effects μ_i for two-stage estimation using different first-stage models, as well as a dummy-

coded single-stage linear regression estimator.³⁷ We observe that, while the all models are generally able to extract the underlying line with unit slope (i.e., they show little bias), there is significant variance in the estimates.

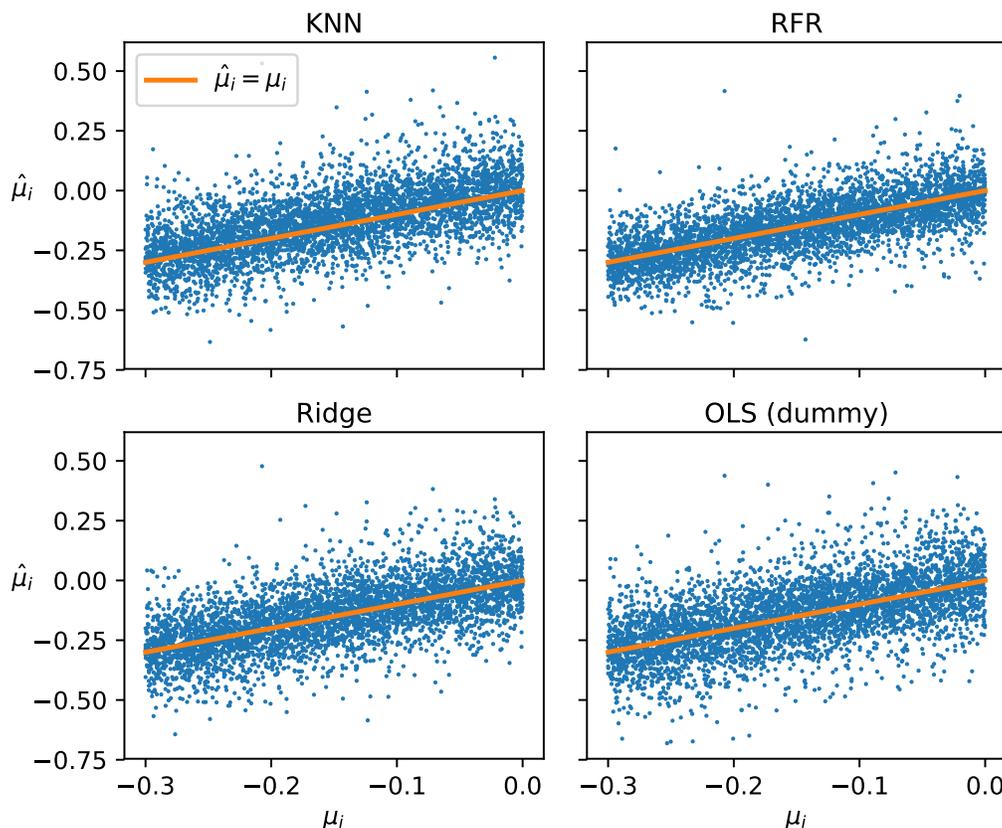


Figure 5.7: ITE estimates vs. true effects, uniform response

Figure 5.8 shows the empirical distribution of the estimation error. The associated means and standard deviations are given in Table 5.4. For the unweighted KNN model, we can observe a much larger bias than for the weighted models.

Table 5.5 contains the empirical coverage of the 95% confidence intervals computed using the multi-stage bootstrapping procedure for $B_m = 10$ and $B_t = 250$,³⁸ as well as the parametric confidence interval from the OLS dummy regression. Except for the percentile method used in combination with the first-stage RFR model, the empirical coverage is insufficient. The primary reason for this is that there are only 25 treatment observations per individual, and so the bootstrap distribution of the estimate is often not a very good estimate of its true underlying distribution. This is a fundamental challenge in settings with

³⁷Due to our data limitations, we “re-use” individuals from the data set of 547 individuals by assigning different randomized response functions.

³⁸Here error (lower) and error (upper) indicate the coverage error of the lower and the upper end of the confidence intervals, respectively.

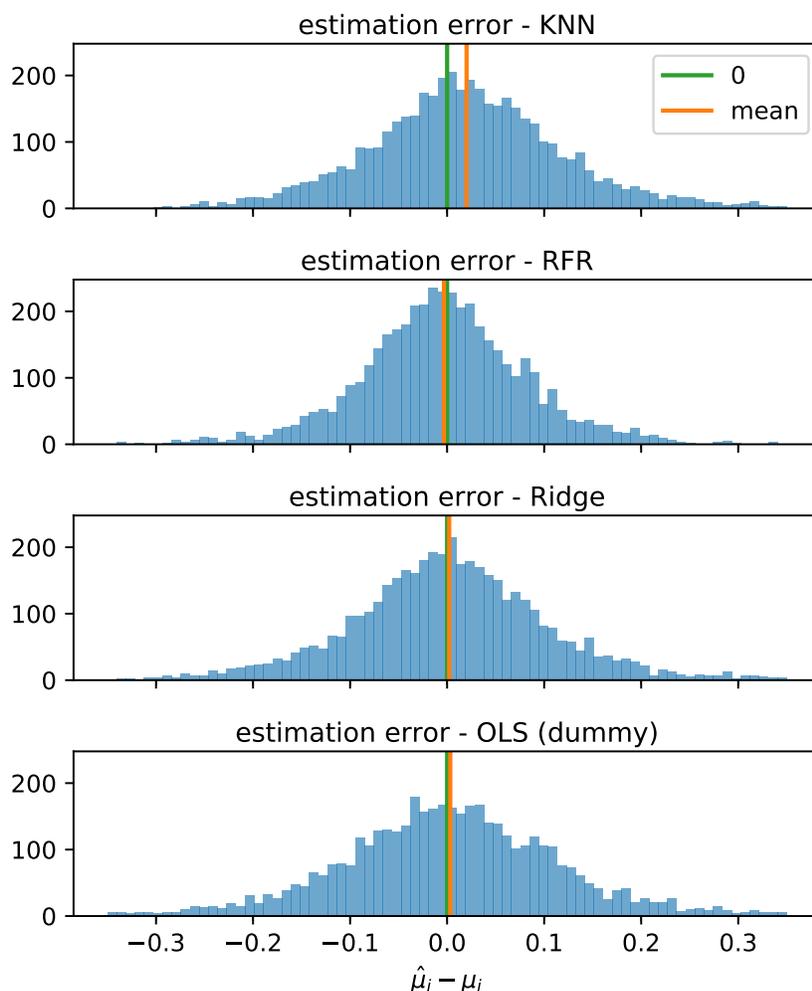


Figure 5.8: Empirical distribution of ITE estimation errors, uniform response

few treatment observations, and is likely also the reason why the bias-corrected (BC) and bias-corrected accelerated (BCa) methods perform rather poorly compared to the more basic percentile (perc) method.³⁹ Finally, the bias of the KNN estimate can also be observed in the asymmetry of the coverage errors at the lower and upper end of the confidence intervals.

Linear Response Figures B.1 and B.2 and Tables B.1 and B.2 in Appendix B provide the corresponding results for the linear response model based on 5470 ITE estimates. Note that in this case we estimate the parameter β_i of a linear response function, rather than the (marginal) individual treatment effect μ_i . The fact that the results look quite similar to those in the case of the uniform response is reassuring, and indicates that estimating

³⁹This shows that the second-order accuracy of the BC_α method discussed in Section 4.6.1 may not be all that beneficial in situations with small effective sample sizes.

model	mean	std dev
KNN (unweighted)	0.0201	0.1053
RFR	-0.0033	0.0890
Ridge	0.0022	0.1044
OLS (dummy)	0.0033	0.1175

Table 5.4: Properties of the ITE estimation error, uniform response

model	method	coverage	error	error (lower)	error (upper)
KNN (unweighted)	BC	0.8999	-0.0501	-0.0404	-0.0097
	BCa	0.9047	-0.0453	-0.0362	-0.0090
	perc	0.9097	-0.0403	-0.0353	-0.0049
RFR	BC	0.8560	-0.0940	-0.0390	-0.0550
	BCa	0.8576	-0.0924	-0.0388	-0.0536
	perc	0.9874	0.0374	0.0202	0.0172
Ridge	BC	0.8885	-0.0615	-0.0314	-0.0301
	BCa	0.8912	-0.0588	-0.0280	-0.0308
	perc	0.8887	-0.0613	-0.0310	-0.0303
OLS (dummy)		0.9027	-0.0473	-0.0271	-0.0202

Table 5.5: Empirical coverage of 95% ITE confidence intervals, uniform response

heterogeneity of treatment response within individuals is indeed possible using our methods.

Informed by the results in this section, we will in the following focus on Random Forest Regression as our first stage model, and on the percentile method as the method for obtaining bootstrap confidence intervals. This two-stage estimator combines small bias with relatively accurate (slightly conservative) confidence intervals.

5.4.3 Average Treatment Effect Estimates

It becomes somewhat challenging to correctly evaluate the coverage of the confidence intervals for the ATT estimate. In principle, we should run a very large number of independent, synthetic experiments, for each experiment compute the multi-stage bootstrapped confidence interval, and evaluate whether it covers the ground truth (i.e. the mean of the underlying distribution from which the $\hat{\mu}_i$ are drawn). However, this is rather challenging to do efficiently, since generating a full synthetic experiment and estimating the associated ATT requires a lot of computation. We can approximate this by generating a smaller number of synthetic experiments, randomly draw subgroups from the experiment population, and compute confidence intervals for the associated pseudo-experiment. The challenge with this

is that for large group sizes the errors will be correlated, so we have to be somewhat careful with the estimates. We generate 8 full synthetic experiments, from each of which we draw 200 randomly selected subgroups and compute the ATT estimate and associated confidence interval, for a total of 1,600 confidence intervals.

In this section we focus on the Random Forest Regression as our first stage model. We estimate the coverage of the bootstrap confidence intervals obtained using the percentile method for the ATT for sample population sizes of 10, 50, 100, 250 and 500 individuals. We also generate parametric confidence intervals by interpreting the ITE estimates of a given subgroup as a sample from a normal distribution with unknown mean and variance (i.e. we use t-based confidence intervals). We compute these parametric confidence intervals for the estimates obtained from the two-stage estimators as well as from the single-stage regression with dummy variables.

Uniform Response In order to assess how well a normal distribution approximates the distribution of the ITE estimates we make use of quantile-quantile plots. Figure 5.9 shows the quantile-quantile plots of the ITE estimate $\hat{\mu}_i$ against a $\mathcal{N}(\bar{\mu}_i, \hat{\sigma}_{\mu_i}^2)$ distribution, where $\bar{\mu}_i$ and $\hat{\sigma}_{\mu_i}$ are empirical mean and standard deviation of the respective distribution of the estimate. From the plots we see that while a normal is a relatively good approximation of the error distributions around zero, the error distributions are more dispersed and heavy-tailed, in particular for higher values.

Table 5.6 reports the empirical coverage of the ATT confidence intervals. We see that all three methods provide relatively accurate coverage, except for a group size of $|\mathcal{P}_T| = 500$. There are two primary reasons for this behavior for $|\mathcal{P}_T| = 500$: the first one is that the RFR model is slightly biased, which has little effect on the coverage if the variation of the mean is high due to the small group size, but which becomes important for large groups. The second reason has nothing to do with the confidence intervals themselves, but with the limitations of how we estimate the coverage. As indicated above, it is not feasible to run a full synthetic experiment for every single ATT estimation problem, so we repeatedly draw random subgroups for which we estimate the ATT. For $|\mathcal{P}_T| = 500$ the sample group size is almost as large as the underlying population, which results in the observations being correlated between the different groups. Hence the empirical coverage for $|\mathcal{P}_T| = 500$ are likely to underestimate the actual coverage.

$ \mathcal{P}_T $	10	50	100	250	500
RFR, bootstrap (perc)	1.0000	0.9525	0.9525	0.9762	0.9137
RFR, t-param	0.9456	0.9506	0.9442	0.9319	0.9116
OLS (dummy), t-param	0.9481	0.9534	0.9522	0.9431	0.9351

Table 5.6: Empirical coverage of 95% ATT confidence intervals, uniform response

If the ITE estimates are independent across individuals, and if individuals are selected

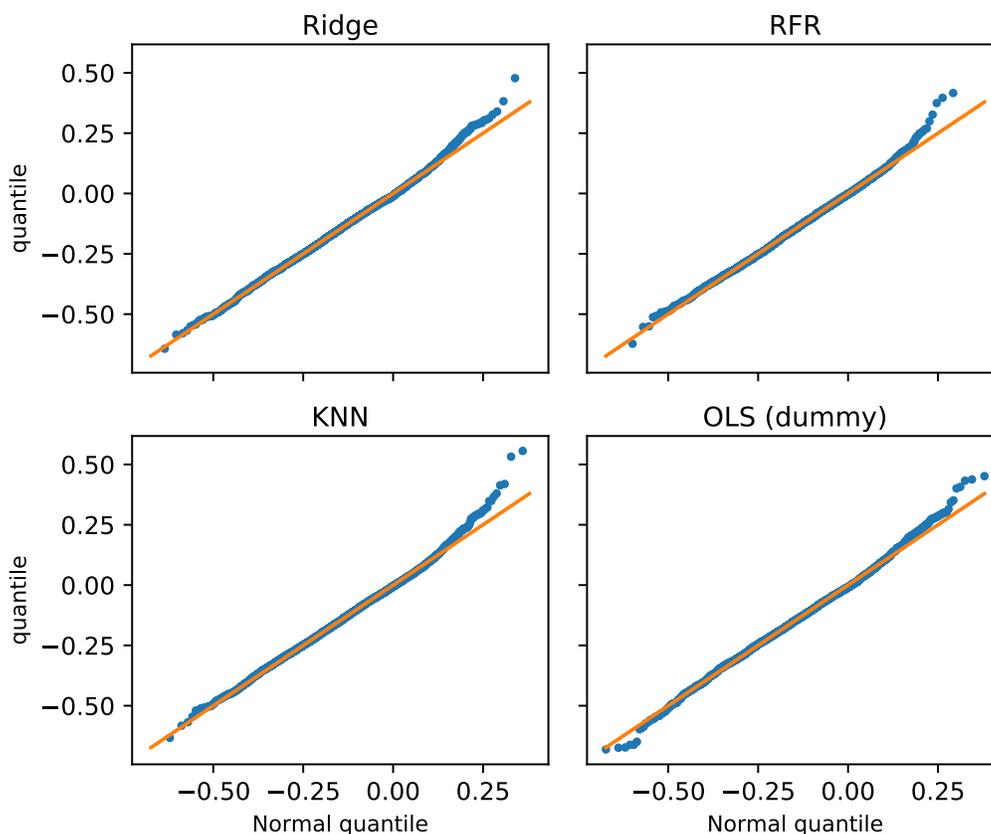


Figure 5.9: QQ-plot ITE estimates $\hat{\mu}_i$ against Normal, uniform response

randomly from the population, then by the central limit theorem the ATT estimate will be asymptotically normal as the group size increases. Therefore it is not surprising that the parametric confidence intervals based on the normal approximation perform well in our case. However, we caution against the blind use of this approximation; In many cases of practical interest the ITE estimates may be correlated, in which case the distribution of the ATT estimates may be quite far from normal. This correlation between the estimates may be due to a range of different reasons — for instance, smart meter readings could be systematically corrupted across customers. In such a case simply considering the ITE point estimates alone and not including the individual variation may result in significantly underestimating the variance of the ATT estimate, and hence in overly optimistic confidence intervals.

The average width of the confidence intervals used in determining the empirical coverage is given in Table 5.7. Unsurprisingly, the parametric confidence intervals are generally of smaller width, which illustrates the fact that parametric test typically provide higher statistical power than non-parametric ones.

For all of the methods in Table 5.7, we observe the expected decrease in width at a rate of $|\mathcal{P}_T|^{-1/2}$ (for independent estimates the variance decreases at rate $1/N$, hence the standard deviation at rate $1/\sqrt{N}$). We can use this to extrapolate and compute the sample sizes

$ \mathcal{P}_T $	10	50	100	250	500
RFR, bootstrap (perc)	0.1993	0.0949	0.0685	0.0428	0.0304
RFR, t-param	0.1903	0.0772	0.0541	0.0340	0.0240
OLS (dummy), t-param	0.1206	0.0487	0.0340	0.0213	0.0151

Table 5.7: Average width of 95% ATT confidence intervals, uniform response

necessary for measuring treatment effects of different magnitudes. In experimental design one typically refers to the *minimum detectable effect* (MDE), which is the smallest effect size under which the null hypothesis of no treatment effect can be rejected at a given significance level. In Figure 5.10 we show the data points from Table 5.7 together with the associated least-squares fit of the parameterization $N = a + b|\mu - \mu'|^{-2}$, which clearly illustrates the growth of the sample size required to estimate very small effects reliably.

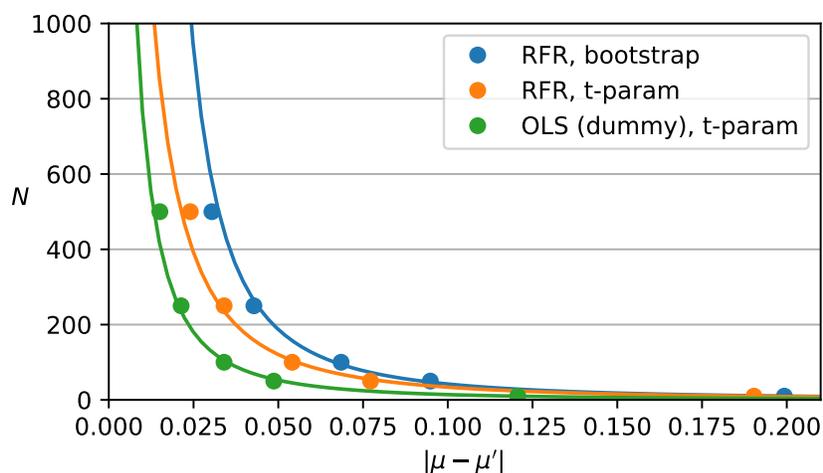


Figure 5.10: Sample size vs. MDE, uniform response

Figure 5.11 shows the distribution of the ATT estimate as well as the lower and upper end of the associated bootstrapped 95% confidence interval for different group sizes for the RFR first-stage model.

Linear Response As in the previous section, we see close similarities in the results for the linear response model, shown in Figure 5.12 and Tables 5.8 and 5.9. One observation that sticks out is that, for the RFR first stage model, the parametric confidence intervals have insufficient coverage, while the bootstrapped confidence intervals for the identical first-stage model have relatively good coverage. We therefore do not plot the associated sample size over MDE curve in Figure 5.13. Comparing Figures 5.10 and 5.13 it is clear that, for the same sample size, we can detect smaller effects on the parameter β than is possible for the ATT μ . Intuitively, the main driver for this is the reward range over which we randomize:

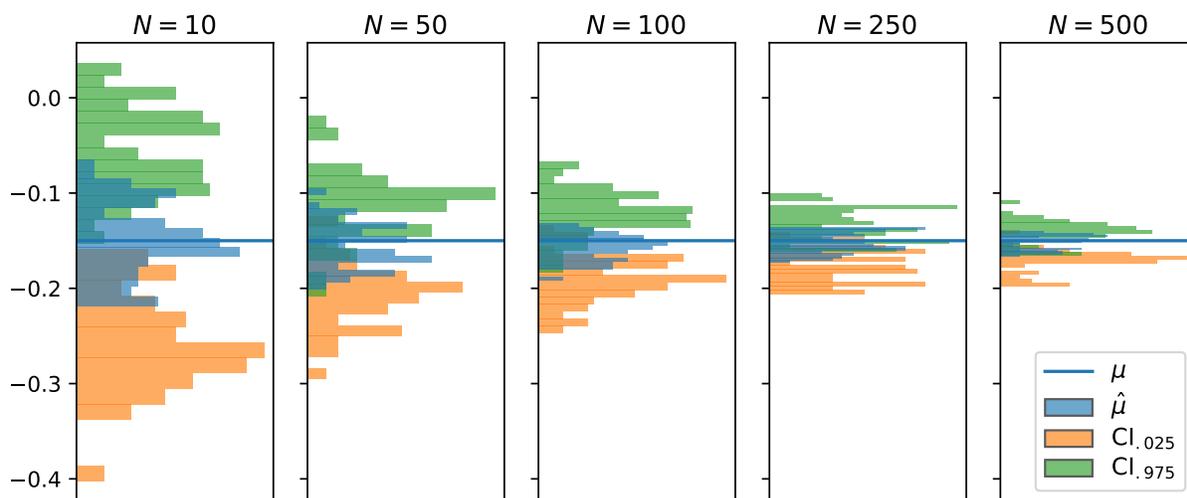


Figure 5.11: Estimates and CIs for different group sizes, RFR model, uniform response

$ \mathcal{P}_T $	10	50	100	250	500
RFR, bootstrap	0.9880	1.0000	0.9810	0.9480	0.9190
RFR, t-param	0.9429	0.9303	0.8892	0.8151	0.6915
OLS (dummy), t-param	0.9415	0.9481	0.9516	0.9417	0.9422

Table 5.8: Empirical coverage of 95% ATT confidence intervals, linear response

$ \mathcal{P}_T $	10	50	100	250	500
RFR, bootstrap	0.1167	0.0550	0.0391	0.0245	0.0175
RFR, t-param	0.1056	0.0432	0.0303	0.0191	0.0135
OLS (dummy), t-param	0.0607	0.0244	0.0171	0.0107	0.0076

Table 5.9: Average width of 95% ATT confidence intervals, linear response

The larger the reward differences are, the easier it is to estimate the slope β_i of the response function of an individual (for a given reward range $R_{it} \in [R_{\min}, R_{\max}]$, it is easy to see that statistical power is maximized for the discrete distribution that assigns probability 0.5 to both R_{\min} and R_{\max}).

Finally, we also show the distribution of ATT estimate and associated confidence interval bounds under the RFR first stage model for the linear response in Figure 5.14.

Summary In this section we have illustrated that marginalizing the ITEs obtained using our two-stage within-subject matching methods leads to relatively precise estimates of the

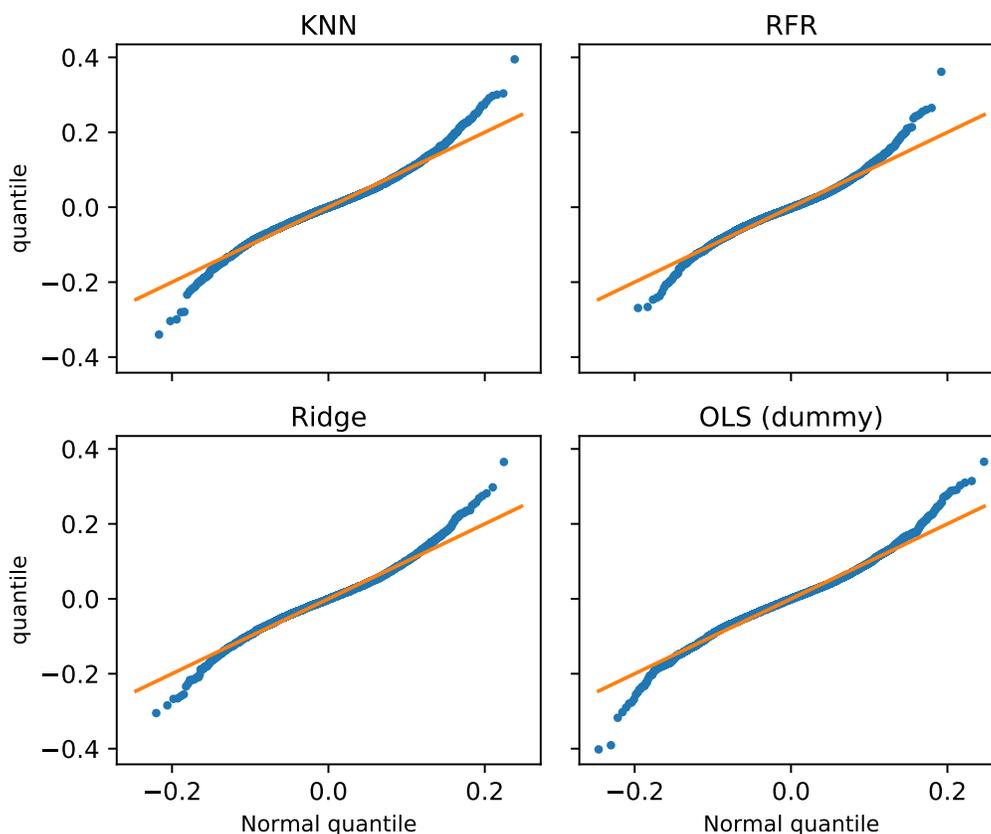


Figure 5.12: QQ-plot parameter estimates $\hat{\beta}_i$ against Normal, linear response

ATT,⁴⁰ and that our bootstrapping algorithms provide valid confidence intervals for these estimates. This is remarkable, and the importance of the fact that this estimation scheme does not require a control group in the conventional sense, or really any randomization of treatments *across individuals*, can hardly be overstated.⁴¹ In particular, it has the potential to provide the ability to perform causal inference in settings where conducting a randomized controlled trial may not be an option. To evaluate this potential, we will benchmark our method against the experimental estimates obtained in the randomized experiment currently being conducted.

⁴⁰By employing standard methods, such as propensity score matching, we may also use these estimates as the basis for estimating the ATE if the treatment assignment on the level of individuals is not randomized, but conditionally independent of the ITE given the customer-level covariates.

⁴¹However, it relies on some other assumptions that may not be necessary in a classic across-subject matching approach, primarily the one that the conditional distribution of the outcomes given the covariates is stationary, or at least does not change in a systematic way across different individuals.

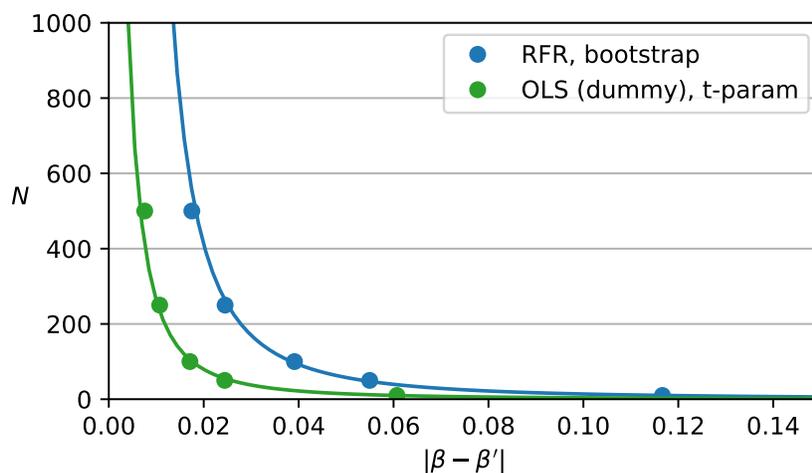


Figure 5.13: Sample size vs. MDE, linear response

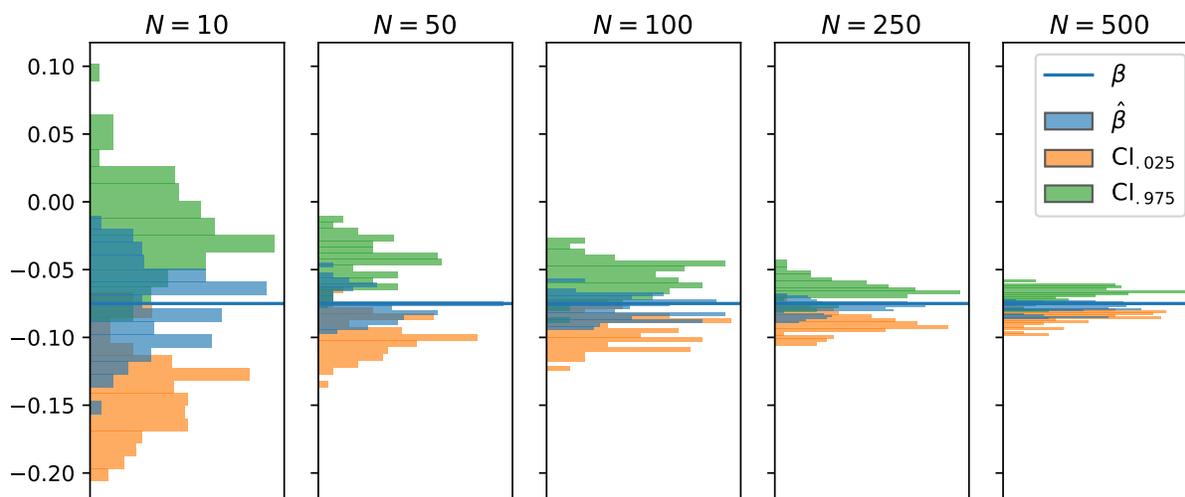


Figure 5.14: Estimates and CIs for different group sizes, RFR model, linear response

5.4.4 Ranking

While the usefulness of ITE estimates as a basis for obtaining average treatment effect estimates is clear, the primary interest in estimating ITEs is, however, the ability to characterize heterogeneity in the treatment effect, both between different individuals, and within each individual. This is valuable for a number of different reasons — in particular, for investigating whether there is a systematic cause of this heterogeneity. For example, in our application we may want to understand if the treatment effect size is correlated with some other customer-level covariate, such as geographic (e.g. climate zone) or socioeconomic (e.g. income level) parameters.

In conventional across-subject matching evaluation, one has to be careful with such an approach, as assignment to the different treatment or control groups may itself be correlated with the covariates of interest. This is especially true for non-experimental studies, but may also be an issue in randomized experiments with small population sizes. When estimating ITEs using within-subject variation these concerns do not apply, and one can perform this kind of ex-post analysis with relative impunity.⁴²

Understanding heterogeneity in the treatment effect can be extremely useful, both from operational and policy perspectives. For example, a DR provider may try to educate specifically those customers who do not respond much to DR notifications. Or they may concentrate their recruitment efforts on electricity customers with characteristics that are correlated with high response levels. This is also important for policymakers, who can use this kind of information to spend taxpayer money more efficiently.

One basic task is to rank study participants according to their individual treatment effect. In this section we explore how we can use our ITE estimators to achieve this task. We do this in a somewhat ad-hoc fashion, and point out that the task of ordering customers according to their treatment effects is quite different than that of obtaining a point estimate of the ITE. Thus one should expect that there are better, more principled ways of performing such ranking.

Our basic ranking method is straightforward: For each estimation model introduced in the previous section, we sort the individuals in the sample according to their ITE estimates in ascending order. In order to evaluate the result, we look at the different quantiles of this ordering. Specifically, for each quantile we determine the fraction of individuals in that quantile contained in the respective quantile of the ordering according to the true ITE. For each model and quantile q this gives us a the “fraction correctly identified,” a number normalized between zero and one, where one corresponds to a perfect identification of the individuals in the quantile, and zero means that no individual has been correctly identified. Random guessing should, in expectation, result in a fraction of correctly identified of q .

Uniform Response Figure 5.15 plots the fraction of correctly identified over the quantiles for the uniform response. We can verify that random guessing (“chance”) indeed achieves the expected identification ratio. All other methods perform relatively similar, with some small differences for $q < 0.5$. In particular, in the range $0.2 < q < 0.4$ the two-stage models achieve slightly better ratios than the single-stage regression with treatment dummies. In general, we observe that the fraction of correctly identified is relatively small for the lower quantiles, and that the biggest improvement over random guessing is achieved in for quantiles $0.2 < q < 0.3$. Decreasing the variance of the ITE estimators (while not affecting bias in a non-systematic fashion across individuals) will improve the fraction of correctly identified.

⁴²Of course one is not immune to “p-value hacking,” i.e. the fallacy of exploring the data long enough along a large number of dimensions so as to eventually finding something that one may falsely consider “statistically significant.”

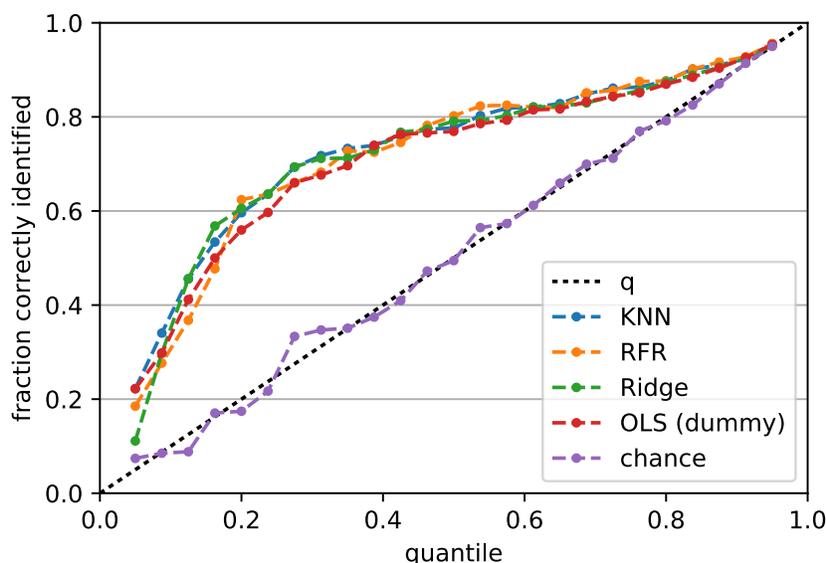


Figure 5.15: Identifying quantiles using ITE estimates, uniform response

The results on the fraction correctly identified are more or less in line with what we would expect. However, they do not provide the full picture, as what we are really interested in is the treatment effect in the identified subgroup. To this end, Figure 5.16 shows the conditional EATT among the group selected for the different quantiles. Here “omniscient” corresponds to the ordering according to the true ITE, which is the theoretical (but practically unachievable) benchmark. At the other extreme, random guessing in expectation achieves the marginal EATT, with the variance going to zero as $q \rightarrow 1$. Even though the identification rates in Figure 5.15 are not very high, especially for lower quantiles, it turns out that in terms of the conditional EATT, the ranking according to ITE estimates performs quite well in comparison to the omniscient benchmark. The reason for that is that it is easiest to discern individuals with large differences in ITE, and thus it is easiest to identify those individuals that matter most for reducing the conditional EATT.

In the metric in Figure 5.16 it also becomes more clear that the two-stage approaches do achieve a better performance consistently in the quantile range $0.1 < q < 0.7$.

Linear Response For the linear response model the results in Figures 5.17 and 5.18 are qualitatively very similar to those of the uniform response model. We can observe slightly more homogeneous results across the different models, which is primarily due to the fact that when estimating the parameter β_i using linear regression, the reduction in variance provided by the two-stage estimators does not have as much of a positive effect as it does when estimating the mean under the uniform response.

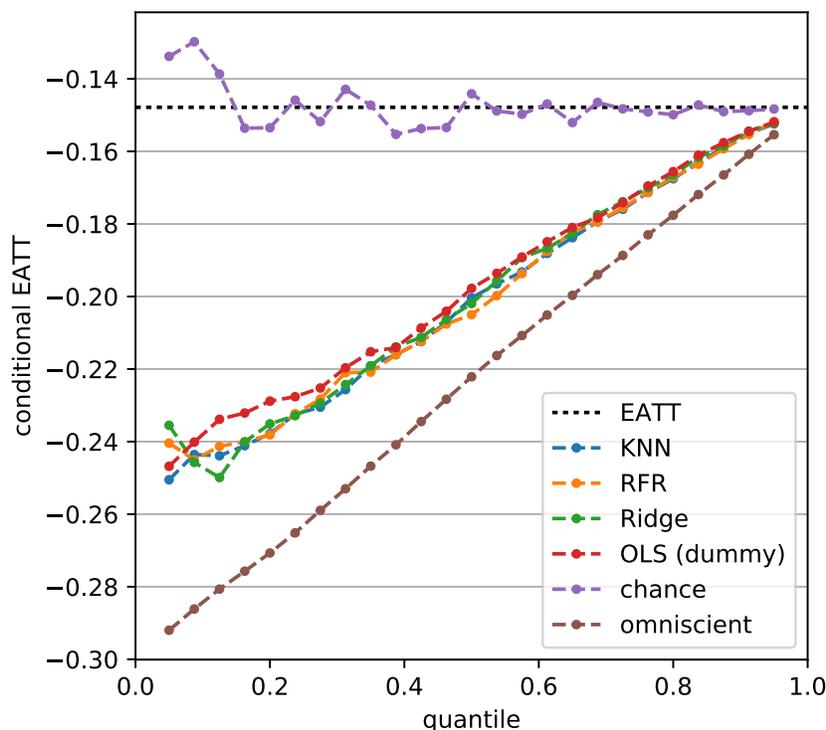


Figure 5.16: Effect of ranking on conditional EATT, uniform response

5.4.5 Adaptive Targeting

In the previous section we investigated the problem of ranking individuals according to their ITEs under a given, fixed treatment assignment mechanism. A more ambitious goal would be to improve efficiency of the DR dispatch in an “on-line” fashion by exploiting heterogeneity of treatment effect on the level of each individual. The implicit underlying assumption here is that dispatching residential DR means competing for a scarce resource, namely the attention of the customer (Simon, 1971).⁴³ The idea is that by using *adaptive targeting*, i.e., performing DR dispatch based on within-subject treatment effect heterogeneity, one can increase the aggregate response level for a given amount of resources (such as customer attention, total incentive payouts, etc.), which will result in efficiency gains.

One would expect the ITE to be heterogeneous across different dimensions within each customer, and this heterogeneity to be specific to each customer. For example, the effect of a DR notification on a participant’s consumption behavior clearly depends on whether the participant is actually able to change their behavior at that point in time, and it is natural to expect that whether this is the case depends on the hour of the day (as this will

⁴³Not making this assumption would basically mean that sending DR notifications every single hour of the day were a reasonable thing to do. Clearly this is not the case, and existing programs generally limit the number of dispatches to a relatively small number. For instance, PG&E’s “Smart Rate” CPP program will call a maximum of 15 CPP events per year (Pacific Gas and Electric Company, 2016a).

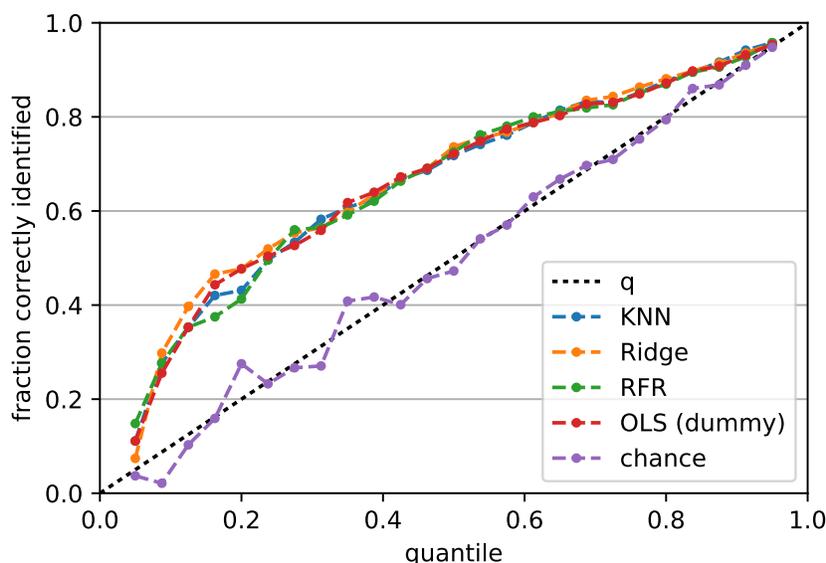


Figure 5.17: Identifying quantiles using ITE estimates, linear response

affect whether the customer is at home or at work). Since different people follow different schedules, the hours in which the CITE is large will be different between different customers. By targeting only those customers that, for the given hour of the day, are most responsive, one can arguably increase the total response for a given number of DR notifications.

Similarly, one would also expect that different individuals have different price elasticities, i.e., respond differently to different marginal monetary incentives.⁴⁴ Specifically, the response of some customers may not depend very much on whether they are offered a high or a low monetary incentive, while the response of others may be strongly affected by the incentive level. For simplicity we refer to customers of the former group as “low marginal reducers,” and to customers of the latter group as “high marginal reducers.”⁴⁵ If we interpret “supply” as the amount of reduction in consumption in response to the incentive, we would say that low and high marginal reducers have a small and large elasticities of supply, respectively.

This suggests that it should be possible to increase the total response per incentive offered by price-discriminating between high and low marginal reducers, that is, exposing low marginal responders to lower rewards than high marginal responders. Of course such price discrimination immediately raises questions of fairness. We think there are number of potential ways to address this, but this is a discussion that is somewhat orthogonal to the contributions of this thesis, so we will not embark on it at this point. However, we do

⁴⁴Reiss and White (2005) document considerable heterogeneity in households’ long-term price elasticities using empirical data from 1,300 households in California. Ito (2014) finds evidence that electricity consumers respond average prices rather than marginal or expected marginal price by comparing billing data of consumers in Southern California.

⁴⁵The assumption of two customer groups is for convenience of exposition, in general one would expect there to be a continuum of customer types.

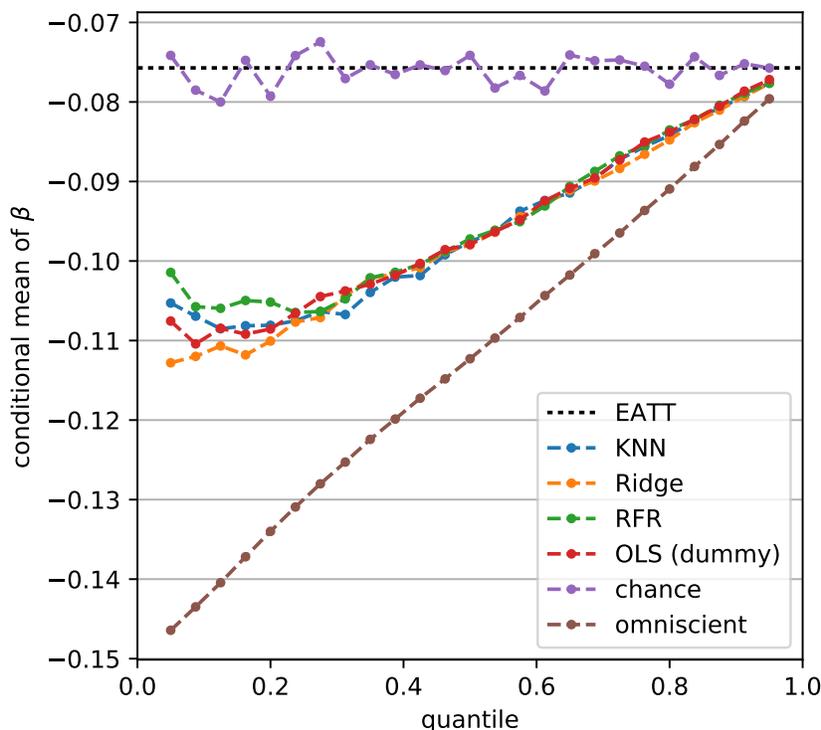


Figure 5.18: Effect of ranking on conditional EATT, linear response

point out that the problem of selecting DR reward levels based on customer responsiveness is closely related to Ramsey Pricing and the problem of optimal taxation of commodities (see e.g. Slemrod (1990) for an overview).

Evaluating the Potential of ITE-Based Price Discrimination

We now investigate the potential of adaptive targeting based on a relatively simple approach for price-discriminating between DR participants. Figure 5.19 illustrates the two prototype customer classes of low and high marginal reducers that we discussed above.

For the purpose of this section, we will assume for simplicity that each consumer has an affine response function to the DR reward. We specifically allow for participants with negative intercept,⁴⁶ which are those that will respond to the DR notification itself, even if it does not involve a financial incentive.⁴⁷ We consider distribution of rewards for each DR period as

⁴⁶Note that in Figure 5.19 we plot the *negative* treatment effect (i.e. the reduction) on the y -axis.

⁴⁷It has long been known that both intrinsic and extrinsic motivation (Bénabou and Tirole, 2003, 2006) play an important role in making economic decisions. This suggests that merely informing electricity consumers about the consequences of their consumption decisions may have a significant effect on consumption. Ferraro and Price (2013) found evidence of this when studying intrinsically motivated conservation using a RCT in the context of residential water demand. In a recent field experiment on electricity consumption, Ito et al. (2015) also find significant short-run effects of such informational treatments, but point out that

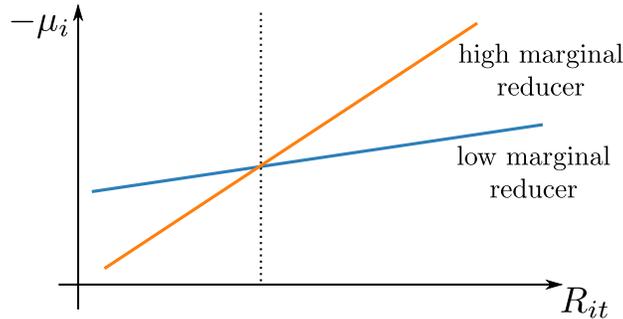


Figure 5.19: Illustration of the ITE for low and high marginal reducers

given exogenously, and hence the problem is essentially one of distributing these rewards to the different participants. The motivation for this comes from the fact that our experiment will feature random price variation, and in order to obtain a causal estimate of the effect of targeting we must have the distribution of rewards be the same as in the control group, which receives non-targeted notifications and incentives. In a more general sense, we can think of this problem as a building block of the larger problem of optimally selecting the reward levels.

We randomly assign the customers in the synthetic experiment population to a group of *low* and a group of *high* marginal reducers. The synthetic ITE is affine in the reward and given by $\mu_i(R_{it}) = a_i + b_i R_{it}$. The parameters a_i and b_i are i.i.d. random draws from the following distributions:

$$a_i \sim \begin{cases} U[-0.125, -0.075] & \text{if } i \in \text{low} \\ U[-0.01, 0.0] & \text{if } i \in \text{high} \end{cases} \quad b_i \sim \begin{cases} U[-0.025, 0.0] & \text{if } i \in \text{low} \\ U[-0.25, -0.1] & \text{if } i \in \text{high} \end{cases} \quad (5.6)$$

We generate a synthetic experiment using the ITEs as described above with the reward level R_{it} (\$/kWh) in each DR period drawn i.i.d. from $U[0, 2]$. For each participant, we estimate both intercept and slope of the response function using our two-stage ITE estimator. Based on our results from Section 5.4.2 we choose Random Forest regression with 200 individual decision trees as the first-stage model, with hyperparameters for each customer optimized on the non-treatment data, and a simple least squares estimator (with intercept) as the second stage model.

Figure 5.20 shows the slopes of the treatment effect plotted against the respective intercepts, both for the ground truth and for the estimates, where the colors indicate the group association. The estimates are extremely noisy, which was to be expected given the results in previous sections. Nevertheless, while precise estimates of ITEs are out of reach in this setting, we will see below that we can still achieve significant improvements in the aggregate response by using the estimates to target (i.e. price differentiate between) participants.

effects diminish quickly after repeated interventions. Faruqui et al. (2010) survey the results of various other studies in the area of electricity consumption.

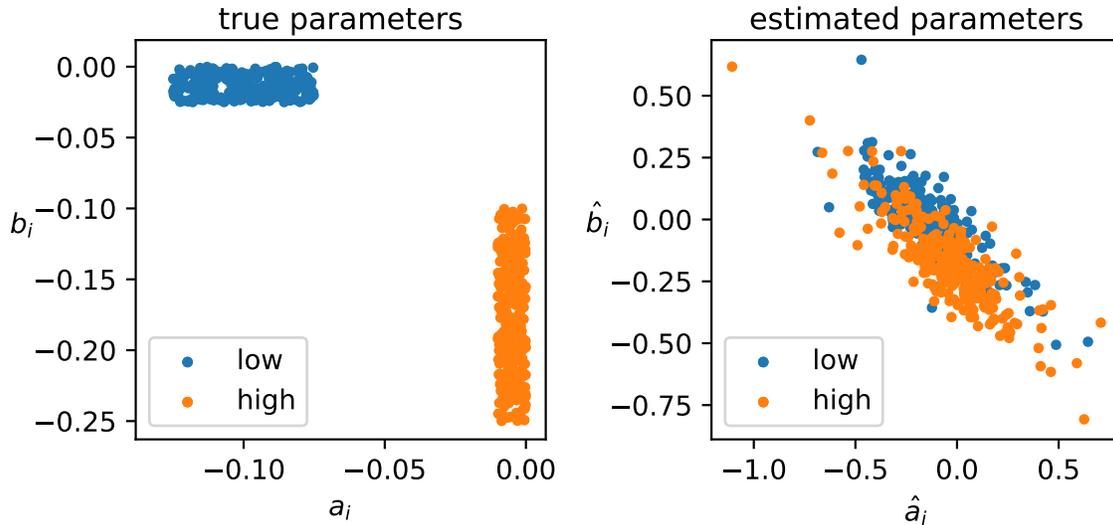


Figure 5.20: True and estimated parameters of the treatment effect

In order to split the sample into two groups based on our estimates, we use somewhat of an ad-hoc approach. Specifically, we assign a given customer i to the *high* group whenever

$$\hat{b}_i - \hat{a}_i < \frac{1}{|\mathcal{P}_T|} \sum_{i \in \mathcal{P}_T} (\hat{b}_i - \hat{a}_i) \quad (5.7)$$

and to the *low* group otherwise. Remembering that a reduction in consumption corresponds to a negative treatment effect, the intuition behind this choice is that low marginal reducers typically have less negative slopes b_i and more negative intercepts a_i .

We acknowledge that there are certainly better, more principled ways to infer the association to the different customer groups. For example, one could try to train a generic supervised Machine Learning classifier for this task, using data from synthetic experiments to provide the labels. Another option would be to use unsupervised clustering methods to identify the groups. However, our goal for this example is not so much to optimize classification performance of this particular classification problem, but rather to illustrate the kind of targeting mechanism that will be deployed in our field experiment. Thus we choose to not further pursue the question of how to optimally classify participants at this point.

Our targeting mechanism in this example is very simple:⁴⁸ We price differentiate between the (estimated) high and low groups by presenting low marginal reducers with lower and high marginal reducers with higher rewards. To ensure a fair comparison, we of course need to use the same *marginal* distribution of prices as in the “vanilla” (non-targeted) dispatch mechanism.

⁴⁸The challenge with using more complex targeting mechanisms is not so much with the mechanism itself, but rather with integrating them with the existing backend systems of the DR provider’s dispatch and notification mechanisms.

Estimating the Effect of Targeting An essential point in estimating the effect of targeting is that we must not use the same synthetic experiment from which we estimate the response parameters also for evaluating the potential effect of targeting. This is because the estimates are driven by the particular realization of the data generating process, and so using the same realization to evaluate the effect of targeting will likely yield a biased estimate overstating the effect. In order to ensure a fair comparison, we instead need to benchmark our estimated group assignments on a new synthetic experiment, under the original treatment effect models. Our methodology for this, including the initial estimation step described above, can be summarized as follows:

1. Conduct a synthetic experiment with random price variation and estimate the parameters of the specified response model for each customer.
2. Split the participants into a group of *low* and a group of *high* marginal reducers, according to a decision criterion such as (5.7). Denote by \hat{n}_{low} and \hat{n}_{high} the number of participants in the *low* and *high* groups, respectively.
3. Run a second, independent synthetic experiment. Compute the true EATT (using the known ground truth) and record the empirical distribution of the rewards.
4. For each DR period, randomly re-assign the top \hat{n}_{high} rewards observed in that period to the individuals assigned to the *high* group. Similarly, randomly re-assign the bottom \hat{n}_{low} rewards to the individuals in the *low* group.
5. Generate treatment responses according to the individual treatment effect models under the newly assigned prices.
6. Compute the EATT of the resulting synthetic experiment.

Under the simple ad-hoc split (5.7), we identify the correct group of 69% of customers in the sample. Table 5.10 shows the resulting EATT after targeting, as well as the relative improvement over the nominal case. It also includes the EATT under the omniscient split (i.e. when group assignment is according to the ground truth), as well as for a placebo test, in which customers are randomized into the two groups.

	nominal	targeted	omniscient	placebo
EATT	-0.1478	-0.1639	-0.1887	-0.1477
improvement	–	10.9%	27.7%	-0.1%

Table 5.10: Effect of targeting on the EATT in the synthetic experiment

We observe a significantly larger (i.e. more negative) treatment effect under the targeted price assignment.⁴⁹ In fact, this effect is larger than one might expect from a mere 69%

⁴⁹In this example, targeting realizes around 40% of the theoretically achievable improvement under the omniscient split.

correct identification rate. The point to note here is that we are more likely to correctly identify those customers with very high or very low marginal response, which are the same ones from which we see the largest change in effect. The result of the placebo test reassures us that our procedure for estimating the effect of targeting is not systematically biased.

Targeting and Online Experimentation One fundamental question that arises in adaptive targeting is how to estimate the heterogeneity of the response if the reward assignment itself is already targeted. This is the classic problem of exploration vs. exploitation: We would like to randomly sample as many different reward levels as possible in order to precisely estimate the customer’s full response function, but at the same time we also want to make use of the information we have gathered. This tension is at the heart of any online decision making process in which one has to estimate properties of the “environment,” including multi-armed bandit problems, Reinforcement Learning, and, as in our case, online experimentation.

In our field experiment we are quite limited in the complexity of our “exploration” scheme, due to the technical and operational limits of the DR provider’s dispatch mechanism. We will therefore be running simple two-phase adaptive targeting, in which we send participants randomized prices for approximately 25 DR notifications in a first phase, and targeted prices based on the resulting ITE estimates during a second phase. A random selection of half the customers will not be subject to adaptive targeting, and continue to receive the same non-targeted prices as in the first phase. This group will serve as the control group, and comparing the treatment effect in the targeted and control groups will allow us to obtain an experimental estimate of the causal effect of targeting on the ATE.

5.5 Empirical Results on Observational Data

In this section we use the initial data set to estimate individual and average treatment effects in the DR provider’s Residential DR program. The primary purpose of this observational study is to inform the final design of the field experiment, in particular as to what effect size to expect, and hence which minimum detectable effect (MDE) to design the treatment group assignments in the experiment for.

First and foremost, we should emphasize that this is non-experimental data, and that we do not have detailed insight into the mechanism under which DR events were called, and how customers were dispatched during those events. Thus one needs to be careful in interpreting the results in this section. That said, the experiment population is a random sample from the overall customer base, and the data quality is quite good.⁵⁰ We have validated our methodology extensively on synthetic experiments using real data in the previous section. Our ATT estimates are in the general range of what has been observed in randomized controlled trials in similar settings (Faruqui and Sergici, 2010), and our placebo tests measure

⁵⁰Except for the lack of historical data for some of the customers. The cause for this, however, seems to be random in nature and uncorrelated with treatment assignment.

no effect in non-DR periods. All these factors combined leave us confident that our estimates are credible.

The Data Set

We exclude all customers with solar PV on Net Energy Metering (NEM) tariffs, as well as those with large amounts of missing consumption values. Since all customers in this initial data set had signed up with the DR provider after June 1, 2016, many of them had received only a small number of DR notifications. For many of these individuals our ITE estimates will be excessively noisy and not very informative. Hence we restrict our attention in this initial evaluation to those customers in the sample who received at least 12 DR notifications.⁵¹ In addition, in order to have sufficient training data for each individual household, we consider only households with at least 220 observation days. This leaves us with a sample of 844 households. Table 5.11 contains summary statistics on the observations and DR periods, and Table 5.12 provides summary statistics on the consumption and the air temperature across the households in the sample (restricted to the hours between 7am and 10pm).

	min	mean	max
observation days	220	495	1178
pre-treatment days	113	404	1079
DR notifications	12	14.4	22

Table 5.11: Summary statistics on observations and DR notifications

	min	mean	max	std dev
mean hourly usage (kWh)	0.01	0.74	13.29	0.67
max hourly usage (kWh)	0.25	4.74	29.47	2.90
mean daily usage (kWh)	0.56	15.89	293.31	14.51
max daily usage (kWh)	1.78	47.43	473.23	34.99
mean air temp (°C)	15.05	20.45	26.94	1.50
max air temp (°C)	28.48	38.33	46.38	2.41

Table 5.12: Summary statistics on consumption and temperature (hours 7-21)

5.5.1 Treatment Effect Estimates

We use the same semi-log specification for the ITEs as in our synthetic experiments in the previous section. In addition, we also consider the level specification in order to analyze

⁵¹This number may seem arbitrary, but follows from the observation that while many customers received 12 notifications, there were much fewer who received 13 or more notifications.

absolute changes in consumption. Since households actually received notifications, and, as we will see later, adjusted their consumption behavior during the associated time periods, we have to be more careful in determining what we consider training data. In particular, we cannot use the hours directly after a DR event to train our models \hat{f}_i^0 for the counterfactual consumption, since participants are likely to defer some of their consumption from the DR period to those hours. Previous studies on CPP programs have observed that this “load-shifting” effect is indeed present, but that consumption (at least on the aggregate level of treatment groups) is statistically indistinguishable from regular consumption behavior within a few hours after a CPP event (see e.g. Jessoe and Rapson (2014) and Herter and Wayland (2010)). Therefore, erring on the side of caution, we choose to drop the 10 hours following each DR notification from our training data set.

Based on the results from our synthetic experiments, we choose a Random Forest Regressor with 200 individual decision trees as our first stage model \hat{f}_i^0 , and weight the training samples according to the weighting scheme (5.4) in order to minimize model bias during treatment periods. The hyperparameters of the first-stage model are optimized for each individual separately by performing four-fold cross-validation on the training data. As discussed earlier in this chapter, customers in this initial data set were not salient of the specific price incentives they faced in different periods, so we estimate the effect of receiving a DR notification on the consumption level during the DR period through uniform treatment effect model using a sample mean second stage estimator. We compute bootstrapped 95% confidence intervals using the multi-stage bootstrapping procedure with $B_m = 25$, $B_t = 250$ and $B_p = 2000$ based on the percentile method outlined in Section 4.6.⁵²

Relative Change in Consumption During DR Periods

Figure 5.21 shows the ITE estimates $\hat{\mu}_i$ under the semi-log specification ordered by value (increasing from bottom to top), with the blue circles representing the point estimates, and the thin horizontal lines representing the associated 95% confidence interval.⁵³

As was to be expected from the synthetic experiments, there is a very large uncertainty with respect to each individual estimate, and considerable heterogeneity across individuals, both in the value of the point estimate and the width of the associated confidence interval. On the actual data we analyze here this uncertainty is even more pronounced than in the synthetic experiments, since individuals on average only receive 14 DR events (compared to

⁵²These parameter choices mean that more than 2.1 billion individual bootstrap estimates are used to generate the confidence intervals for the ATT estimate.

⁵³In rare cases it may happen that the point estimate $\hat{\mu}_i$ lies outside the associated confidence interval. This is not an error, but a result of the fact that on the level of an individual the empirical distribution of the change in consumption during DR periods may be very far from unimodal. In particular, as the number of observations is small, outliers can have a significant effect on the mean of the distribution. If this is the case, then the bootstrapped confidence intervals may center around the mode of the empirical distribution, which may be quite different from the sample mean, which is strongly affected by the outlier. This is in stark contrast to confidence intervals obtained by assuming a normal limiting distribution, which are always symmetric around the sample mean.

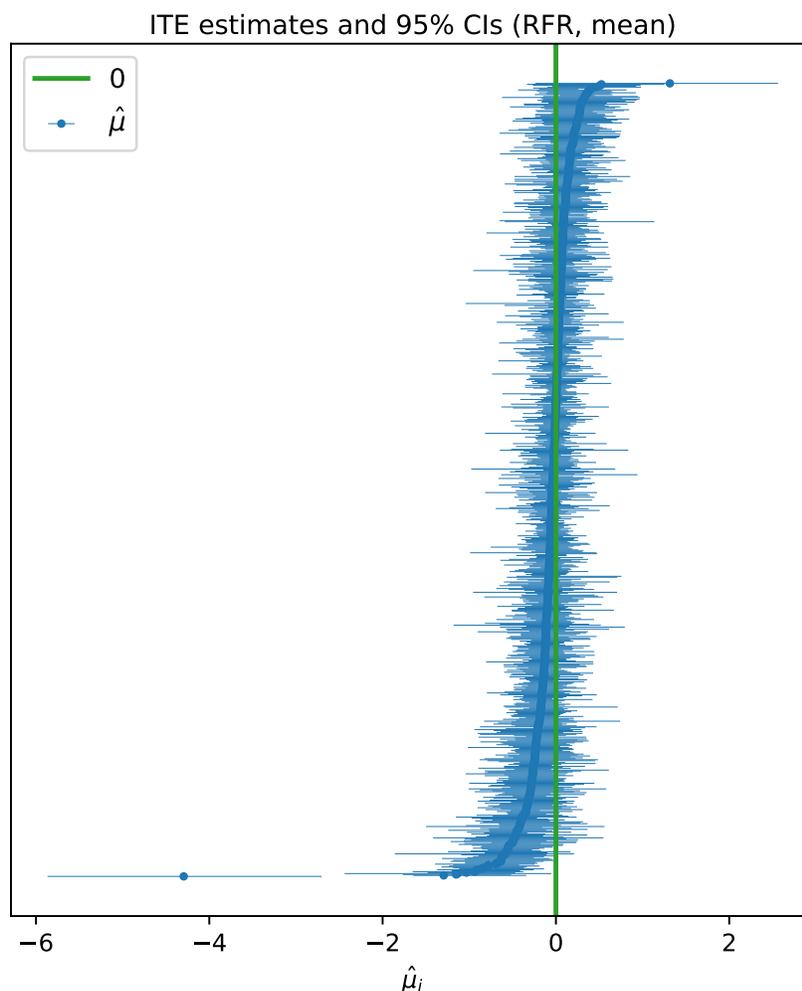


Figure 5.21: Empirical Results: ITE Estimates (semi-log)

25 in the synthetic experiment). Moreover, unlike as in the synthetic experiments, the actual treatment responses are presumably non-homogeneous across the different DR periods, which adds additional variance to the estimate.

What is most striking about Figure 5.21 is the outlier at the very bottom left of the plot. The point estimate for this customer is $\hat{\mu}_i = -4.2664$, which translates to an average reduction in consumption of 98.6% during DR events. Initially, we suspected this estimate to be the result of corrupt data, but it turns out that this customer in fact did reduce consumption to almost zero in almost all DR periods. This is shown in Figure 5.22, which displays the consumption time series around all DR events for this customer.⁵⁴ This showcases the

⁵⁴Note that the second event in the first plot in the third row of Figure 5.22 is not used in the estimation, as we do not observe a sufficient number of non-treatment consumption values prior to the event.

potential for using individual treatment effects to explore treatment heterogeneity ex-post.⁵⁵

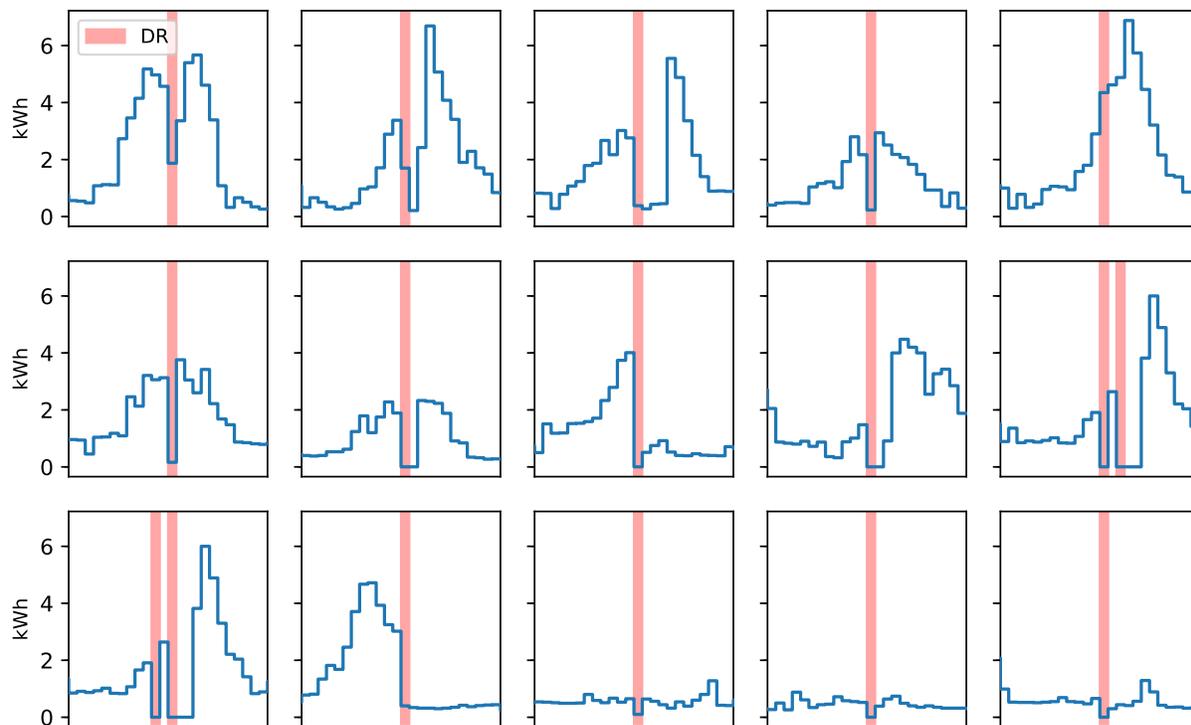


Figure 5.22: Consumption of the “outlier” around DR periods

Figure 5.23 shows the same data as Figure 5.21, but scales the x -axis. There are 59 DR participants in the sample of 844 for which we can reject the null hypothesis of a zero treatment effect at 95% significance level. For 55 of these participants we estimate a negative treatment effect (i.e., they on average reduce their consumption during DR periods), and for 4 of them we estimate a positive effect. If, as our synthetic experiments in the previous section suggest, the empirical coverage of our confidence intervals is correct, then under the null of a homogeneous zero treatment effect we would expect around $0.025 \cdot 844 = 21$ false positives on either side. The fact that within the group of 55 participants we observe more than twice as many with negative ITE estimate, but only 4 with positive ITE estimate, indicates that, on aggregate, customers do reduce consumption during DR periods.

This observation is confirmed by our ATT estimate of $\hat{\mu} = -0.0682$ and the associated 95% confidence interval $CI_{\hat{\mu}}^{.95} = (-0.0924, -0.0492)$ computed by Algorithm 4. Both the point estimate and the confidence interval are shown on top of the ITE estimates in Figure 5.23. The t -based 95% confidence interval for $\hat{\mu}$ is $(-0.0861, -0.0503)$, and thus quite close to the one obtained by our bootstrapping method. For either method we can easily

⁵⁵Of course one must still be mindful of avoiding the practice of what is often referred to as “p-value hacking,” i.e., looking at the data long and hard enough until finding something “statistically significant.”

reject the null of a zero ATT at 95% significance level. Moreover, the estimate is within the range of estimates obtained by other randomized experiments in similar settings (Ontario Energy Board, 2007; Faruqui and Sergici, 2010; Jessoe and Rapson, 2014).

As can be seen from Figure 5.23, individual treatment effects are highly heterogeneous across customers. While the ATT estimate $\hat{\mu}$ corresponds to a mere 6.6% reduction in consumption during DR periods across the sample population, the estimate of the conditional ATT for those customers with negative ITE estimate for which we can reject the null at 95% significance is $\hat{\mu}|_{CI_{\mu_i}^{95} < 0} = -0.6177$, which corresponds to a 46.1% reduction in consumption.⁵⁶ These insights into the level of heterogeneity are in line with the observations made by Bollinger and Hartmann (2016) and, on an aggregate level for participants with PCTs, with those of Harding and Lamarche (2016).

Absolute Change in Consumption During DR Periods

The results for the absolute change in consumption obtained using the level specification are qualitatively similar. Figure 5.24 shows the results corresponding to Figure 5.23 under the level specification. Unsurprisingly, compared to the semi-log specification the confidence intervals are somewhat narrower for those participants with an ITE estimate around zero. Many of these customers exhibit a generally low consumption level, and so there is little variation observed in their consumption in absolute terms (this illustrates some of the difficulties with comparing absolute change in consumption across different customers).

The ATT estimate under the level specification is $\hat{\mu} = -0.0691$, with associated 95% confidence interval $CI_{\hat{\mu}}^{95} = (-0.0918, -0.0454)$.⁵⁷ The corresponding t-based 95% confidence interval $(-0.0874, -0.0508)$ is again quite similar to the bootstrapped one. This point estimate of the absolute reduction roughly corresponds to the average customer switching off a 75W lightbulb for the duration of the DR event. We can reject the null hypothesis of a zero treatment effect at 95% significance level for 64 (lower consumption) and 4 (higher consumption) individuals, respectively.

5.5.2 Placebo Tests

In order to rule out systematic estimation errors in our approach, we conduct placebo tests. For these tests, we remove the DR periods from the data set, and for each individual randomly label periods in the experiment time frame as synthetic DR periods, according to the individual's empirical distribution of actual DR periods (over months, days and hours) during the observation time frame. This means that we can assume that participants' consumption behavior in the placebo hours is statistically the same as during DR hours, with

⁵⁶We have to be careful in interpreting this kind of result — conditioning on low values will of course always yield a more negative effect. The point here is less about the exact value of the conditional estimate (and whether it is unbiased), and more about the fact that there is significant heterogeneity across individuals.

⁵⁷This estimate being close in numbers to the estimate under the semi-log specification is a coincidence that is explained by the fact that the mean consumption level across customers and DR periods is of 1.14kWh.

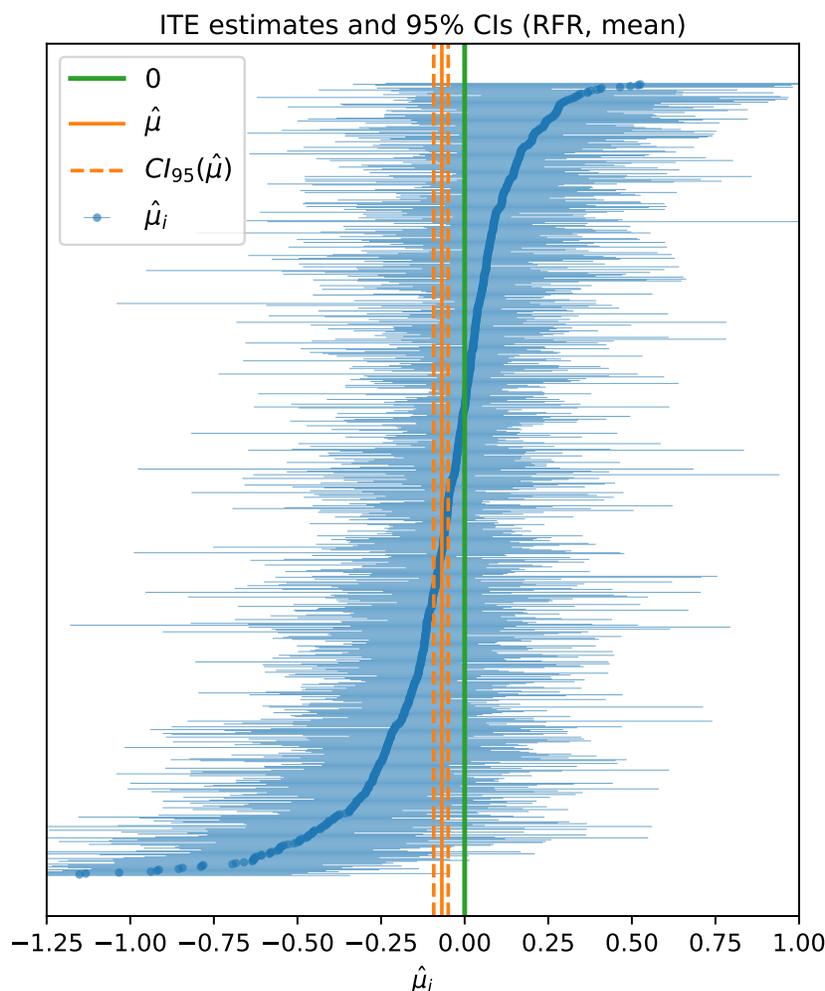


Figure 5.23: Empirical Results: ITE and ATT estimates (semi-log)

the only difference being the lack of exposure to DR notifications. We then estimate ITEs and ATT in exactly the same fashion as on the actual treatment data. If this yielded a significant non-zero average effect (vis-à-vis a significantly non-symmetric distribution of ITE estimates around zero), this would indicate a systematic error in our estimation approach. We emphasize that while such a placebo test can help identify systematic errors, it is not guaranteed to do so (that is, passing the placebo test is necessary but not sufficient).

Figure 5.25 shows the results of the placebo test under the semi-log specification. Here we color-coded the estimates according to whether the individual's actual (non-placebo) estimate belonged to the lower (yellow) or the upper (blue) half of all estimates. The fact that these two groups are well mixed in the placebo test indicates that the outcomes of our estimation strategy are not driven by some subject-level characteristics that are unrelated to treatment exposure. Instead, they are driven by randomness and, of course, the ITE.

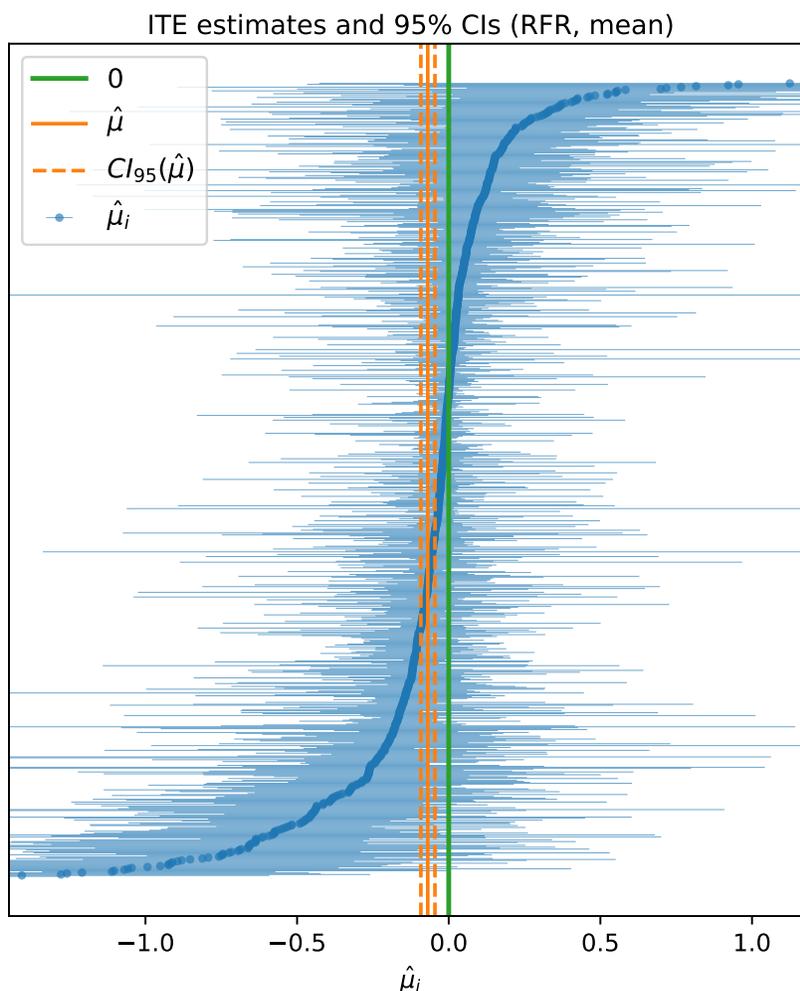


Figure 5.24: Empirical Results: ITE and ATT estimates (level)

The ATT estimate under the semi-log specification for the placebo test is $\hat{\mu} = 0.0090$, with associated 95% confidence interval $CI_{\hat{\mu}}^{95} = (-0.0059, 0.0254)$. The corresponding t-based confidence interval is $(-0.0022, 0.0201)$. Hence in either case we cannot reject the null of a zero ATT. Moreover, there are only 8 individuals for which we can reject the null of a zero ITE (2 negative, 6 positive), compared to the 59 (55 negative, 4 positive) in the actual treatment data.

The results of the corresponding placebo test for the level specification are shown in Figure 5.26, with an ATT of $\hat{\mu} = 0.0076$ and confidence interval $CI_{\hat{\mu}}^{95} = (-0.0105, 0.0236)$.⁵⁸ For the level specification we can reject a zero ITE for 8/4 (low/high) individuals, compared to 64/4 (low/high) in the actual treatment data.

The results of these placebo tests reassure us that the effects detected in the previous

⁵⁸The t-based confidence interval is $(-0.0034, 0.0187)$.

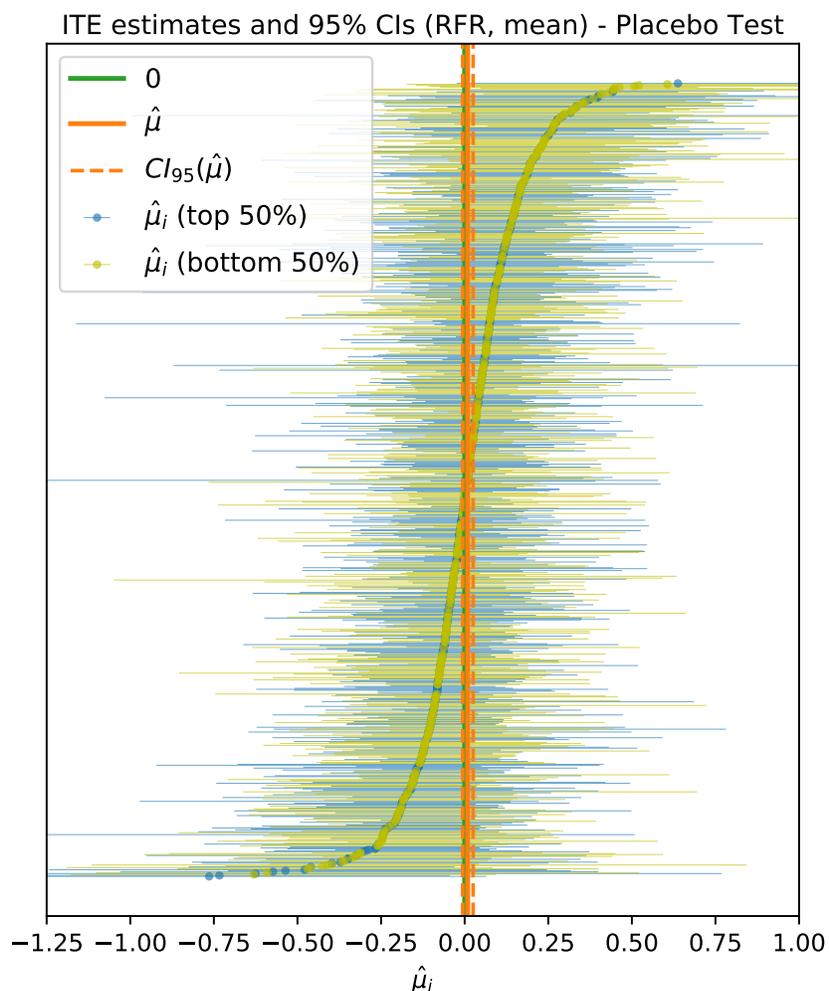


Figure 5.25: ITE Estimates (semi-log) - Placebo Test

section are indeed real effects, and that they can serve as the basis for informing the design of the treatment group assignment in our randomized controlled trial.

5.5.3 Heterogeneity in the ITE

In this section we investigate the heterogeneity in the ITE estimates along a number of dimensions. Given some of the data limitations we face with the initial data set, this should be regarded more as exploratory analysis that informs the design of the field experiment, rather than a complete and thorough study.

To start, we note that we do not find any statistically significant difference in the ATT conditional on the electric utility service territory, suggesting that, at least in terms of how they respond to DR events, customers of the different IOUs are statistically quite similar.

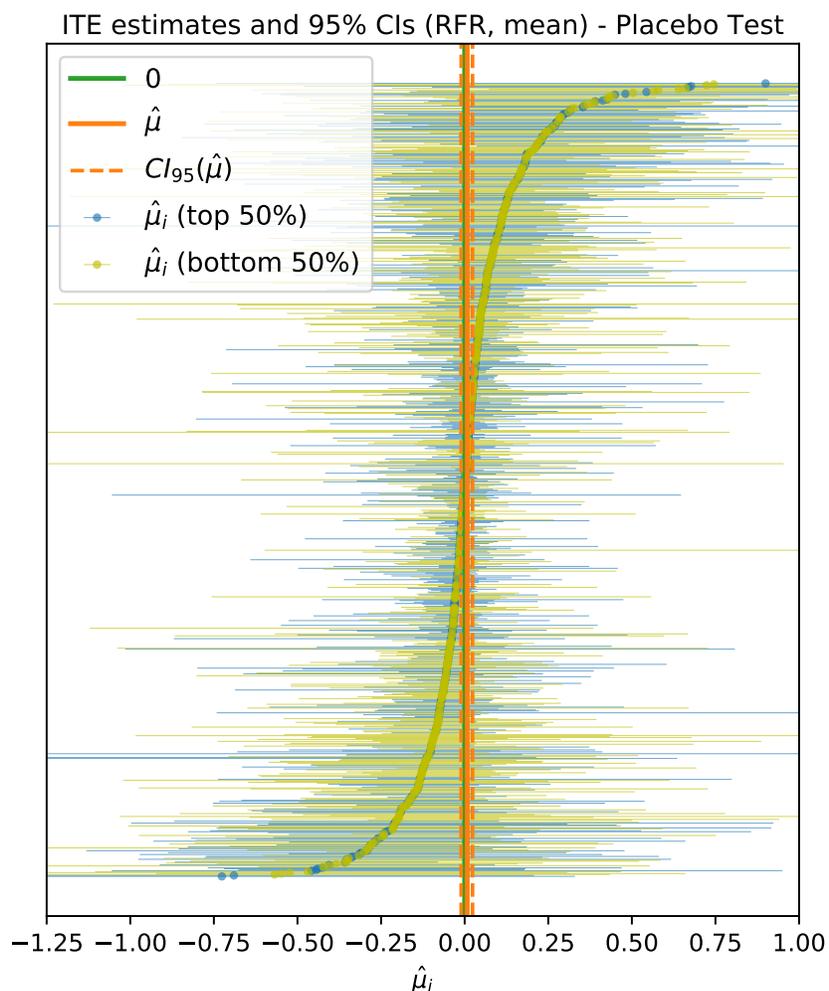


Figure 5.26: ITE Estimates (level) - Placebo Test

Temperature

Much of the peak electricity load in California is due to residential AC.⁵⁹ In combination with the fact that reducing or deferring AC load is a simple way for residential customers to reduce their peak consumption, one would expect a dependency of the treatment effect on the ambient temperature during DR periods.⁶⁰ Based on the semi-log specification, we do not find a statistically significant correlation between the ITE estimate and the mean air temperature across DR periods in our data. That is, we cannot claim that the *relative* change in consumption during DR periods is affected by the temperature.

However, using the level specification, we do find a dependency of *absolute* change in

⁵⁹Residential and commercial air conditioning together represent at least 30 percent of summer peak electricity loads in California (California Public Utilities Commission)

⁶⁰This has been documented in many studies, including (Wolak, 2006; Harding and Lamarche, 2016).

consumption. Specifically, Figure 5.27 shows the ITE estimates under this specification plotted against the mean air temperature during DR periods for the respective individual, together with the associated least-squares fit. While there is significant noise, regressing the the ITE estimate on the temperature yields a coefficient of -0.0088 with associated 95% confidence interval $(-0.014, -0.003)$.⁶¹ In other words, the average participant would decrease their consumption during DR periods by 0.1kWh on average following a 11.3°C increase in ambient temperature.

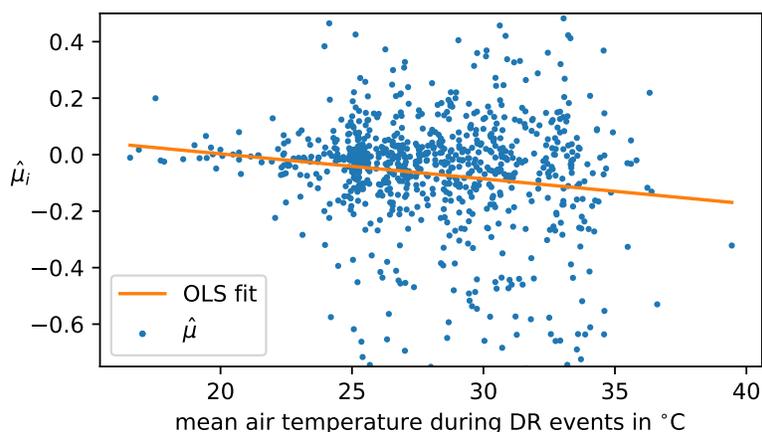


Figure 5.27: ITE estimates (level) vs. mean local air temperature during DR events

Consumption Level

The higher a customer’s electric load, the more likely it is that they have the opportunity to reduce consumption during DR periods (in particular, high peak load is often due to the use of AC, which is easily curtailed by changing the set point or overriding the AC unit). Figure 5.28 shows the ITE estimates under the level specification against the peak hourly consumption of the household. As one would expect, there is a negative correlation between the absolute change in consumption and the hourly peak load. However, with a coefficient of -0.0105 with 95% confidence interval $(-0.020, -0.001)$ this correlation is not very strong. Indeed, under this estimate an average additional 0.1kWh reduction in consumption during DR periods would correspond to an average increase of 9.5kWh in peak load, which is higher than the peak load of most households.

As was the case with the ambient temperature, we do not find a statistically significant correlation between the relative change in consumption and the customer’s hourly peak load using the semi-log specification. If there was such a correlation, this could be interpreted as an indication that some of the heterogeneity in reduction is due to the *composition* of the participants’ consumption.⁶²

⁶¹We use confidence intervals based on robust covariance estimates.

⁶²It is reasonable to assume that consumers have some level of base load that is generally harder to curtail

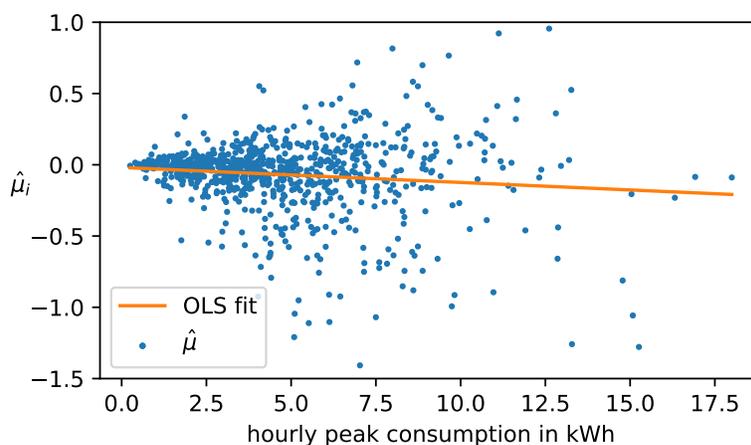


Figure 5.28: ITE estimates (level) vs. hourly peak consumption

Home Automation Technology

Another potentially important factor for driving responsiveness in CPP and CPR programs, as documented by many of the studies discussed in Section 5.1, is whether households use home automation technology such as “smart” thermostats. Table 5.13 shows the penetration of active automated devices in the full initial data set (see Section 5.3.1) as well as in the sample for which we estimated the ITE.

	population	sample
PCT	151	21
EV	15	0
other	78	10

Table 5.13: Penetration of home automation devices in estimation sample

In Figure 5.29 we present normalized histograms of the ITE estimates for non-automated (na) and automated (a) customers, together with the conditional ATTs (vertical lines).⁶³ The conditional ATTs are $\hat{\mu}^a = -0.3225$ and $\hat{\mu}^{na} = -0.0601$, which correspond to a 27.6% and 5.8% average reduction in consumption during DR periods among automated and non-automated customers.

Testing the null hypothesis of equal means using Welch’s t-test yields a p-value of 0.0009, which suggests that automated customers indeed show a statistically significant larger (more

or defer. Hence a higher peak load would indicate the consumer has additional loads that can be used more easily to reduce consumption.

⁶³Due to the small number of automated customers in the sample, the associated histogram is likely not a very good estimate of the actual distribution of the conditional ITE.

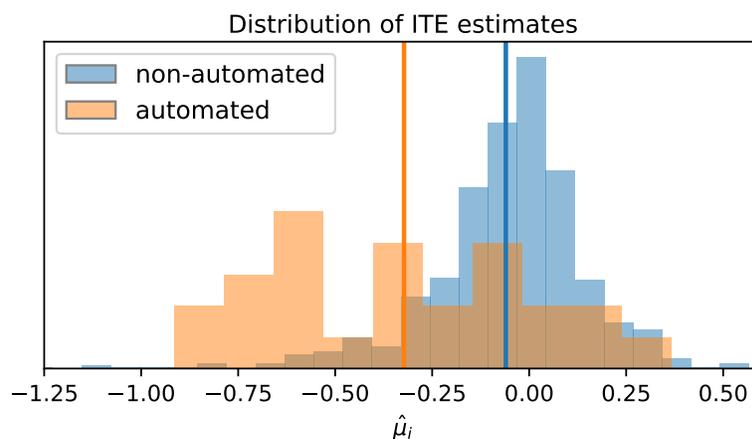


Figure 5.29: Heterogeneity in ITE estimates (semi-log) by automation

negative) treatment effect than non-automated customers.⁶⁴ These results are in line with the benefits that we would expect from DLC. With better data, it would be interesting to study how much of the benefits are attributable to larger reductions, and how much are attributable to more consistent reductions (e.g. by being able to turn off the AC remotely if the DR participant is not at home).

⁶⁴Welsh's t-test assumes that the ITE estimates in each group are normally distributed. If we want to avoid this assumption, we can perform a one-sided Mann-Whitney U test, which yields a p-value of $1.8 \cdot 10^{-5}$. Under the additional assumption that the distributions differ only in their location parameter, this p-value pertains to the likelihood of the observed sample under the null hypothesis of equal medians.

Appendix A

Supplementary Material for Chapter 2

A.1 Formulation of the Optimization Problem

Recall the utility function of a risk-neutral, utility-maximizing consumer from Section 2.3:

$$V(u, x, q, z, q^{BL}) := U(u, x, y, q) - \sum_{t=1}^T \left[p_t^R q_t - \mathbf{1}_{\{t \in \mathcal{E}\}} p_t^{DR} DR_t \right] - \text{FC} - \text{DC} \quad (\text{A.1})$$

A.1.1 Optimization Problem for a Fully Rational Consumer

A fully rational consumer faces the following optimization problem:

$$\max_{x, u, q, z, q^{BL}} V(x, u, q, z, q^{BL}) \quad (\text{A.2a})$$

$$\text{s.t. } z_t^{DR} \in \{0, 1\} \quad \forall t \quad (\text{A.2b})$$

$$DR_t = (q_t^{BL} - q_t) z_t^{DR} \quad \forall t \quad (\text{A.2c})$$

$$\text{consumption model constraints}_t(x_t, u_t) \quad \forall t \quad (\text{A.2d})$$

$$\text{baseline definition}(q^{BL}, q, z^{DR}) \quad (\text{A.2e})$$

$$\text{FC} = \text{fixed charge definition} \quad (\text{A.2f})$$

$$\text{DC} = \text{demand charge definition}(q) \quad (\text{A.2g})$$

Denote problem (A.2) by \mathcal{P} . We point out that, under our assumption of perfect foresight (which follows from Assumptions 1 and 2), participation in DR events, though voluntary, will take place automatically in (A.2) whenever it is ex-post beneficial to the consumer.

Before describe how to formulate the different elements of \mathcal{P} in the following sections, we briefly comment on the complexity of the optimization problem. Depending on the length of the simulation horizon, the number of DR events, and whether the tariff includes a demand charge, \mathcal{P} is a Mixed-Integer Program (MILP for the HVAC model and MIQP for the Quadratic Utility model) of considerable size. The largest of our simulations are

for the HVAC model and involve around 45,000 variables and 100,000 constraints. For the Quadratic Utility model our largest optimization problem has around 35,000 variables and 9,000 constraints. Despite the large size of the problem, the GUROBI solver (Gurobi Optimization, 2016) we use manages to solve even the largest of the problems in less than 100 seconds on a standard desktop computer.

Demand Charges

Let \mathcal{T}_{DC} denote the set of periods relevant for the demand charge in the horizon of interest¹. Then we have $DC = \sum_{p \in \mathcal{T}_{DC}} DC_p$, where, for each $p \in \mathcal{T}_{DC}$, $DC_p \geq q_t$ for all $t \in p$. From the objective function, it is easy to see that at the optimum, $DC_p = \max_{t \in p} q_t$ for each $p \in \mathcal{T}_{DC}$. In this paper, we focus on charges on the peak consumption during each month. However, the above formulation is rather general and can easily be modified to account for similar kinds of demand charges, as long as they can be reformulated as linear constraints.

Endogenous Definition of the DR Baseline

Various baselining methodologies have been proposed for Demand Response and are used by different ISOs. In the present analysis, we focus on the so-called “10 in 10” baseline defined by CAISO (California Independent System Operator Corporation, 2015a). In this methodology, the baseline for a particular hour is the average consumption in the previous n similar² non-event days (i.e. days without a DR event), where $n = 10$ for Business days and $n = 4$ for non-Business days. The CAISO methodology also allows for a so-called load-point adjustment (LPA), by which the raw customer baseline can be adjusted up or down by no more than 20%, depending on the consumption level during the morning of the event day. For simplicity, we will ignore the LPA in our analysis.

Let $q_{d,h}$ denote the total energy consumption of the DR participant in hour h of day d . Under some abuse of notation, let $q_{d,h}^{\text{BL}}$ denote the value of the CAISO baseline in the same period. In order to improve readability, we start with the simplest case and step by step build up a formulation that captures the complete baseline.

No other DR events: In the simplest case, there are no possible DR events in the past n similar days, and the baseline value can be written as

$$q_{d,h}^{\text{BL}} = \frac{1}{n} \sum_{d' \in \mathcal{D}} q_{d',h} \quad (\text{A.3})$$

where \mathcal{D} is the set of the n similar previous days³. In this case, it is easy to formulate the DR participants decision problem using the so-called “big-M” method (Williams, 2009). Abusing

¹E.g. the different months or the different billing periods within the simulation horizon $1, \dots, T$.

²Here “similar” just means the distinction between Business and non-Business days.

³For example, if d is a Business day, then \mathcal{D} contains the 10 previous Business days.

notation again, let $z_{d,h} \in \{0, 1\}$ be a binary variable indicating whether the participant reduces consumption w.r.t to baseline ($z_{d,h} = 1$) or not ($z_{d,h} = 0$) during hour h of day d . Moreover, let $r_{d,h} = (q_{d,h}^{\text{BL}} - q_{d,h})z_{d,h}$ denote the reduction w.r.t. the baseline value. The DR reward during this hour is $\text{LMP}_{d,h} \cdot r_{d,h}$, where $\text{LMP}_{d,h}$ is the Locational Marginal Price at the participant's pricing node. Thus, as the term $\text{LMP}_{d,h} \cdot r_{d,h}$ appears in the participant's objective function⁴ (which is to be maximized), then the constraints

$$r_{d,h} \geq 0 \tag{A.4a}$$

$$r_{d,h} \leq q_{d,h}^{\text{BL}} - q_{d,h} \tag{A.4b}$$

$$r_{d,h} \leq z_{d,h} \cdot M \tag{A.4c}$$

with M a suitably large constant⁵ fully encode the decision problem. In particular, if $z_{d,h} = 0$ then $r_{d,h}$ is forced to zero by constraint (A.4c). Otherwise, if M is large enough and $z_{d,h} = 1$, then (A.4c) is not binding, and by way of how it appears in the objective, $r_{d,h}$ at optimum will equal the baseline reduction $q_{d,h}^{\text{BL}} - q_{d,h}$ by constraint (A.4b). Observe that, importantly, all constraints (A.4) are linear inequality constraints.

However, if within the n previous similar days there is the possibility for participating in another DR event, then the previous decision of whether to reduce consumption and receive a DR payment will affect the computation of the baseline. Nevertheless, by introducing additional variables and constraints, it is still possible to formulate the baseline as a (possibly large) set of linear inequality constraints.

All DR events in same hour, one other possible event: For simplicity, assume first that DR events always occur during the same hour of the day. This allows us to simplify notation, drop the index h , and simply consider d as the period of interest. Further, suppose that there is exactly one day, say day $d_{[1]}$, in the previous n similar days during which the participant may choose to be rewarded for reducing consumption w.r.t. the baseline. Let $z_{d_{[1]}} \in \{0, 1\}$ denote the associated indicator variable. Again, with the term $\text{LMP}_d \cdot r_d$ in the objective function, the problem can be encoded via the following constraints:

$$r_d \geq 0 \tag{A.5a}$$

$$r_d \leq q_d^{\text{BL}} - q_d \tag{A.5b}$$

$$r_d \leq z_d \cdot M \tag{A.5c}$$

$$q_d^{\text{BL}} \leq \frac{1}{n} \left(\sum_{d' \in \mathcal{D}'} q_{d'} + q_{d_{[1]}} \right) + z_{d_{[1]}} \cdot M \tag{A.5d}$$

$$q_d^{\text{BL}} \leq \frac{1}{n} \left(\sum_{d' \in \mathcal{D}'} q_{d'} + q_{d^{(1)}} \right) + (1 - z_{d_{[1]}}) \cdot M \tag{A.5e}$$

where \mathcal{D}' is the set of $n-1$ similar previous days *excluding* day $d_{[1]}$, and $d^{(1)}$ is the first similar day prior to all days in \mathcal{D}' . In particular, if $z_{d_{[1]}} = 0$, the constraint (A.5e) will be inactive

⁴Note that the variable $r_{d,h}$ appears nowhere else in the objective or in other constraints.

⁵Choosing the constant M large enough but not too large is important for the optimization problem to be well-conditioned. In general this can be tricky, but in our case a straightforward and suitable choice is to set M to the maximum possible energy consumption or the participant in any given period.

and, because q_d^{BL} enters the objective through the reduction r_d , (A.5d) will be binding. In both cases, the baseline q_d^{BL} will be forced to the correct value.

All DR events in same hour, more than one event: We retain the simplifying assumption that all events happen during the same hour, but now consider the case when there are k possible event days that could affect the baseline in day d . Call those event days $d_{[1]}, \dots, d_{[k]}$ and denote by $z_{d_{[1]}}, \dots, z_{d_{[k]}} \in \{0, 1\}$ the associated indicator variables for participating in the respective event. Furthermore, let \mathcal{D}' be the set of $n - k$ previous similar days *excluding* days $d_{[1]}, \dots, d_{[k]}$. Finally, denote by $d^{(1)}, \dots, d^{(k)}$ the k first similar days prior do all days in \mathcal{D}' , and let $\mathcal{K} = \{1, \dots, k\}$. It turns out that we can encode the correct baseline by using 2^k big-M type constraints. This is easiest to see when $k = 2$, in which the constraints are:

$$r_d \geq 0 \tag{A.6a}$$

$$r_d \leq q_d^{\text{BL}} - q_d \tag{A.6b}$$

$$r_d \leq z_d \cdot M \tag{A.6c}$$

$$q_d^{\text{BL}} \leq \frac{1}{n} (\sum_{d' \in \mathcal{D}'} q_{d'} + q_{d_{[1]}} + q_{d_{[2]}}) + (z_{d_{[1]}} + z_{d_{[2]}}) \cdot M \tag{A.6d}$$

$$q_d^{\text{BL}} \leq \frac{1}{n} (\sum_{d' \in \mathcal{D}'} q_{d'} + q_{d^{(1)}} + q_{d_{[2]}}) + (1 - z_{d_{[1]}})M + z_{d_{[2]}} \cdot M \tag{A.6e}$$

$$q_d^{\text{BL}} \leq \frac{1}{n} (\sum_{d' \in \mathcal{D}'} q_{d'} + q_{d_{[1]}} + q_{d^{(1)}}) + z_{d_{[1]}}M + (1 - z_{d_{[2]}}) \cdot M \tag{A.6f}$$

$$q_d^{\text{BL}} \leq \frac{1}{n} (\sum_{d' \in \mathcal{D}'} q_{d'} + q_{d^{(1)}} + q_{d^{(2)}} + (2 - z_{d_{[1]}} - z_{d_{[2]}}) \cdot M \tag{A.6g}$$

It is rather straightforward to verify that for all possible combinations

$$(z_{d_{[1]}}, z_{d_{[2]}}) \in \{(0, 0), (0, 1), (1, 0), (1, 1)\}$$

the constraints (A.6) result in the correct definition of the baseline. For the general case, equations (A.6d)-(A.6g) can be replaced by the following 2^k constraints:

$$\forall K \in 2^{\mathcal{K}} \quad q_d^{\text{BL}} \leq \frac{1}{n} (\sum_{d' \in \mathcal{D}'} q_{d'} + \sum_{j \in K} q_{d_{[j]}} + \sum_{l \in \mathcal{K} \setminus K} q_{d^{(l)}}) + M \cdot \sum_{j \in K} z_{d_{[j]}} + M \cdot \sum_{l \in \mathcal{K} \setminus K} (1 - z_{d_{[l]}}) \tag{A.7}$$

where $2^{\mathcal{K}}$ denotes the power set (i.e., the set of all subsets) of \mathcal{K} .

DR events in different hours, more than one event: This is the most general case that covers the full problem. To encode whether there is one or more DR events during a given day, we use an auxiliary indicator variable for each day with at least one possible event. Specifically, suppose that during day d there are m hours during which the participant can choose to participate in DR. Call those hours h_1, \dots, h_m , and let $z_{d, h_i} \in \{0, 1\}$ for $i = 1, \dots, m$ the associated indicator variables for DR participation. Let $z_d \in \{0, 1\}$ be the binary variable

indicating whether the participant places at least one DR bid during day d . Then the variable z_d is fully determined by the following constraints:

$$z_d \geq z_{d,h_i} \quad \text{for } i = 1, \dots, m \quad (\text{A.8a})$$

$$z_d \leq \sum_{i=1}^m z_{d,h_i} \quad (\text{A.8b})$$

The general problem can then be formulated by defining auxiliary variables and associated sets of constraints as in (A.8) for each day with possibly multiple DR events, and by using those auxiliary variables z_d in the respective constraints (A.7) that determine the baseline value.

A.1.2 A Fixed-Point Algorithm for Computing a Baseline-Taking Equilibrium

A consumer that takes the baseline values q^{BL} as given exogenously, but otherwise behaves in a fully rational way, solves a slightly modified version of the optimization problem \mathcal{P} in (A.2). Specifically, q^{BL} does not play the role of an optimization variable anymore but instead is a fixed parameter, and the constraint (A.2e) is dropped from the problem. Let $\tilde{\mathcal{P}}$ denote this modified optimization problem.

We are interested in baseline-taking equilibria as formalized in Definition 1. A straightforward approach to finding such an equilibrium is a fixed-point iteration, as given by Algorithm 5. Here $\beta : q \mapsto q^{BL}$ is the function computing the baseline based on the consumption vector q (in our case, this is the CAISO 10 in 10 method). Furthermore, $d : q^{BL} \times q^{BL} \mapsto \mathbb{R}^+$ is a distance function between two baseline profiles, and $\epsilon > 0$ is a numerical tolerance parameter.⁶ At this point we do not have any theoretical convergence guarantees for Algorithm 5, but numerical simulations have shown it to converge reliably within a few (< 10) iterations for all our simulation scenarios.⁷ Algorithm 5 returns, up to numerical tolerances, a consumption vector q^* that is optimal with respect to the baseline $\beta(q^*)$, which is of course nothing but the a Baseline-taking Equilibrium according to Definition 1.

A.2 Dynamical System Models

While we investigate the two particular models of the Quadratic Utility with battery and the HVAC-equipped building in detail, we point out that our formulation also allows for general quadratic utility functions of the form

$$U(u, x, y, q) = \mathbf{w}^\top H \mathbf{w} + h^\top \mathbf{w} \quad (\text{A.9})$$

⁶In our simulations, we use the standard euclidean norm for the distance and a tolerance of $\epsilon = 10^{-2}$.

⁷Without making additional assumptions, deriving theoretical guarantees for convergence seems quite daunting, as in each iteration of the algorithm we are solving a full Mixed-Integer optimization problem.

Algorithm 5: Algorithm for Computing Baseline-taking Equilibrium

Data: ϵ
Result: Baseline-taking equilibrium (x^*, u^*, q^*, z^*)
Solve \mathcal{P} to obtain initial condition (x^0, u^0, z^0, q^0) ;
Let $q^{BL} = \beta(q^0)$;
for $k \in \mathbb{N}$ **do**
 Solve $\tilde{\mathcal{P}}$ for baseline q^{BL} to obtain $(\tilde{x}, \tilde{u}, \tilde{z}, \tilde{q})$;
 Let $\tilde{q}^{BL} = \beta(\tilde{q})$;
 if $d(q^{BL}, \tilde{q}^{BL}) < \epsilon$ **then**
 return $(x^*, u^*, q^*, z^*) := (\tilde{x}, \tilde{u}, \tilde{z}, \tilde{q})$;
 $q^{BL} \leftarrow \tilde{q}^{BL}$;

where $\mathbf{w} = [u^\top, x^\top, y^\top, q^\top]^\top$ and $H \preceq 0$. This is a rather general formulation that encompasses many different models of interest. As our pyDR package (Balandat et al., 2016a) is written in a modular fashion, it is straightforward to include other consumption models of the form (A.9).

A.2.1 Quadratic Utility Model with Battery

Calibration of the Consumption Utility Model

Recall from Section 2.3.2 that in the Quadratic Utility model a consumer who consumes quantity \tilde{q}_t in period t at price p_t^R derives stage utility

$$U_t(q_t) = a_t \tilde{q}_t - \frac{1}{2} b_t \tilde{q}_t^2 - p_t^R \tilde{q}_t \quad (\text{A.10})$$

In the absence of storage, optimal consumption without a budget constraint yields that

$$U'_t(q_t) = 0 \Leftrightarrow a_t - p_t^R = b_t \tilde{q}_t \Leftrightarrow \tilde{q} = \frac{a_t - p_t^R}{b_t} \quad (\text{A.11})$$

The parameters a_t and b_t are calibrated for each period based on observed consumption data and prices, by positing the (point) elasticity of demand, η . The elasticity is

$$\eta(p_t^R) \triangleq \frac{dq_t(p_t) p_t^R}{dp_t^R q_t} = -\frac{1}{b_t} \frac{p_t^R}{q_t} = -\frac{p_t^R}{b_t} \frac{b_t}{a_t - p_t^R} = -\frac{p_t^R}{a_t - p_t^R}$$

Solving for the parameters a_t and b_t yields

$$a_t = -\frac{p_t^R(1 - \eta)}{\eta} \quad b_t = -\frac{p_t^R}{q_t \eta} \quad (\text{A.12})$$

We calibrate the parameters a_t and b_t in each period using consumption data representative⁸ of customers with PG&E’s A1 tariff under the associated retail charges from the A1 tariff. To ensure comparability of the results, this calibration is the same for all simulations of the Quadratic Utility model.

Battery Parameters

The battery charge x_t (in kWh) is subject to the following constraints:

$$0 \leq x_t \leq \begin{cases} 0 \text{ kWh} & \text{for no battery} \\ 10 \text{ kWh} & \text{for medium battery} \\ 25 \text{ kWh} & \text{for large battery} \end{cases}$$

Charging and discharging⁹ are limited to $0 \leq u_{i,t} \leq u_i^{\max}$, where

$$u_1^{\max} = \begin{cases} 0 \text{ kW} & \text{for no battery} \\ 5 \text{ kW} & \text{for medium battery} \\ 25 \text{ kW} & \text{for large battery} \end{cases} \quad u_2^{\max} = \begin{cases} 0 \text{ kW} & \text{for no battery} \\ 7.5 \text{ kW} & \text{for medium battery} \\ 30 \text{ kW} & \text{for large battery} \end{cases}$$

We assume a leakage time constant of $T_{\text{leak}} = 96$ h and the same charging and discharging efficiency of $\eta_c = \eta_d = 0.95$. Discretizing (2.5) under zero-order hold sampling with a sampling time of 1h then yields the following matrices for the discrete-time system model (2.2):

$$A = 0.9974 \quad B = [0.95, -1.0526, 0] \quad E = D = 0 \quad c_q = [1, 0, 1]$$

A.2.2 Commercial Building HVAC System Model

We consider a simple Linear Time Invariant model for the HVAC system of a commercial building, with form and parameters from Gondhalekar et al. (2013). The continuous-time system dynamics are $\dot{x} = A_{ct}x + B_{ct}u + E_{ct}v$, where $x \in \mathbb{R}^3$, $u \in \mathbb{R}^2$, $v \in \mathbb{R}^3$ and

$$A_{ct} = \begin{bmatrix} -(k_1 + k_2 + k_3 + k_5)/c_1 & (k_1 + k_2)/c_1 & k_5/c_1 \\ (k_1 + k_2)/c_2 & -(k_1 + k_2)/c_2 & 0 \\ 5/c_3 & 0 & -(k_4 + k_5)/c_3 \end{bmatrix}$$

$$B_{ct} = \begin{bmatrix} 1/c_1 & -1/c_1 \\ 0 & 0 \\ 0 & 0 \end{bmatrix} \quad E_{ct} = \begin{bmatrix} k_3/c_1 & 1/c_1 & 1/c_1 \\ 0 & 1/c_2 & 0 \\ k_4/c_3 & 0 & 0 \end{bmatrix}$$

Here $x_1(t)$ represents the room temperature, and $x_2(t)$ and $x_3(t)$ represent interior-wall surface and exterior-wall core temperature at time t , respectively. The inputs $u_1(t)$ and $u_2(t)$

⁸We use the so-called ‘‘Dynamic Load Profile’’ (Pacific Gas and Electric Company, 2016a).

⁹In our implementation we also limit the direct consumption in order to simplify finding an ‘‘M’’ in the big-M formulation. However, the limit is so high that the constraints are never binding, and thus do not affect the solution of the optimization problem.

are heating and cooling power in period t , respectively. The disturbance vector $v(t)$ consists of outside air temperature ($v_1(t)$), solar radiation ($v_2(t)$) and internal heat gains ($v_3(t)$). All temperatures are in $^{\circ}\text{C}$, all other inputs are in kW. The parameters in the matrices are given in Table A.1.

c_1	c_2	c_3	k_1	k_2	k_3	k_4	k_5
$9.356 \cdot 10^5$	$2.97 \cdot 10^6$	$6.695 \cdot 10^5$	16.48	108.5	5.0	30.5	23.04

Table A.1: HVAC model parameters (Gondhalekar et al., 2013)

We discretize the continuous-time model using a zero-order hold scheme with 1 hour sampling time¹⁰ to obtain a discrete-time system of the form (2.2). The resulting matrices are

$$A = 0.1 \cdot \begin{bmatrix} 5.821 & 3.394 & 0.582 \\ 1.069 & 8.868 & 0.048 \\ 0.814 & 0.214 & 7.536 \end{bmatrix} \quad B = 10^{-3} \cdot \begin{bmatrix} 2.947 & -2.947 \\ 0.231 & -0.231 \\ 0.181 & -0.181 \end{bmatrix}$$

$$E = 10^{-3} \cdot \begin{bmatrix} 20.238 & 3.178 & 2.947 \\ 1.441 & 1.368 & 0.231 \\ 143.635 & 0.190 & 0.181 \end{bmatrix}$$

The total power consumption in this model is simply the sum of heating and cooling power, and so $c_q = [1, 1]$. The interior air temperature $x_{1,t}$ is required to satisfy ‘‘comfort constraints’’ of the form $x_{1,t}^{\min} \leq x_{1,t} \leq x_{1,t}^{\max}$, where

$$x_{1,t}^{\min} = \begin{cases} 21 & \text{if } 8\text{am} \leq t \leq 8\text{pm} \\ 19 & \text{otherwise} \end{cases} \quad x_{1,t}^{\max} = \begin{cases} 26 & \text{if } 8\text{am} \leq t \leq 8\text{pm} \\ 30 & \text{otherwise} \end{cases}$$

with the narrower band capturing the main work hours. Heating and cooling power consumption $u_{1,t}$ and $u_{2,t}$ satisfy actuator constraints of the form $0 \leq u_{1,t} \leq u_{1,t}^{\max}$ and $0 \leq u_{2,t} \leq u_{2,t}^{\max}$, respectively, where

$$u_{1,t}^{\max} = 500 \quad u_{2,t}^{\max} = \begin{cases} 150 & \text{for nodes PGEB, PGP2} \\ 200 & \text{for node PGCC} \\ 300 & \text{for nodes PGSA, PGF1} \end{cases}$$

Here we adjusted some of the constraints from Gondhalekar et al. (2013) upwards to account for the higher cooling requirements (hence larger HVAC systems) at the higher temperature pricing nodes PGCC (Central Coast), PGSA (Sacramento) and PGF1 (Fresno).

¹⁰We could also higher sampling frequencies, but those would result in much larger optimization models.

A.3 Economics Appendix

A.3.1 Deadweight Loss Bias-Variance Decomposition

We can compute the bias-variance decomposition of a tariff's deadweight loss (DWL), displayed in Table 2.1, under the assumption of a time-separable linear demand system, as follows. We recall the expression (2.6) for the DWL, under the assumption of time separability (i.e. cross-derivatives equal zero):

$$\text{DWL} = -\frac{1}{2} \sum_{j=1}^J e_j^2 \frac{\partial x_j}{\partial e_j}$$

We construct weights which are proportional to the demand derivatives, so that the coefficients on the tariff errors sum to one:

$$w_j \triangleq -\frac{1}{2} \frac{\partial x_j}{\partial e_j} / \left(-\frac{1}{2} \sum_{i=1}^J \frac{\partial x_i}{\partial e_i} \right) = \frac{\partial x_j}{\partial e_j} / \left(\sum_{i=1}^J \frac{\partial x_i}{\partial e_i} \right).$$

We treat the vector of weights \vec{w} as a notional probability mass function. All expectations are with respect to \vec{w} . Then our proxy loss function is the expected squared tariff error, under \vec{w} , and is proportional to the DWL:

$$L \triangleq \sum_{j=1}^J w_j e_j^2 = \frac{\sum_j e_j^2 \frac{\partial x_j}{\partial e_j}}{\sum_j \frac{\partial x_j}{\partial e_j}} = -2 \cdot \text{DWL} / \left(\sum_j \frac{\partial x_j}{\partial e_j} \right)$$

We define the tariff's *bias* as the expected difference between the retail price and the social marginal cost (SMC):

$$\text{Bias} \triangleq \sum_j w_j (p_j^R - \text{SMC}_j) = \sum_j w_j e_j$$

The *variance* is the expected squared difference of the tariff error from the bias:

$$\text{Var} \triangleq \sum_j w_j (p_i^R - \text{SMC}_i - \text{Bias})^2 = \sum_j w_j (e_j - \text{Bias})^2$$

The tariff proxy loss is the variance plus the square of the bias:

$$L = \text{Bias}^2 + \text{Var}.$$

Since $\text{DWL} = -(\sum_j \frac{\partial x_j}{\partial e_j})L/2$, The portion of DWL due to bias is $-(\sum_j \frac{\partial x_j}{\partial e_j}) \text{Bias}^2 / 2$, and the portion due to variance is $-(\sum_j \frac{\partial x_j}{\partial e_j}) \text{Var} / 2$. In Table 2.1, as in our simulations, we determine demand derivatives by assuming constant demand elasticity, and backing out demand derivatives from historical price and load levels for A-1 customers.

A.3.2 Social Marginal Cost and the Social Cost of Carbon

To calculate the social marginal cost (SMC) RTP tariff, we assume that the social marginal cost of generation is the private marginal generation cost, plus the externalized cost of pollution. For simplicity, we consider pollution costs to be entirely attributable to GHG emissions. We do not include estimated capacity costs in SMC tariffs, because we consider the quality of the data to be too low, and the calculation method too arbitrary, to merit inclusion in the tariffs that form the basis for our repository of simulation data. However, when we calculate welfare metrics in Sections 2.5.3 and 2.6.2, we also include estimated capacity costs, whose calculation we describe in Appendix A.3.3. The resulting inconsistency means that the SMC RTP benchmark tariff could be improved on.

To determine the social cost of GHG emissions, we start by assigning a social cost of carbon (SCC), of \$40 per metric tonne CO₂ (Jacobsen et al., 2016; Interagency Working Group on Social Cost of Carbon, 2013). From this, we subtract an estimate of the cost of carbon that was reflected in the price of carbon in the California cap and trade market, and thus internalized into wholesale prices. We determine the latter subtrahend to be \$0 in 2012, and \$12 in 2013 and 2014 (Hsia-Kiung et al., 2014).

In order to obtain the carbon cost per MWh, we multiply these carbon costs per tonne by the marginal operating emissions rate (MOER) of the CAISO grid, in tonnes per MWh. To obtain these MOERs, we use a dataset from the company WattTime, which gives hourly marginal operating emissions rates (MOERs) for the CAISO market, for the year 2015 (see also Callaway et al. (2015)). We do not have hourly data on emissions rates for the CAISO grid for 2012-2014, but we will see below that the hourly variation in MOERs in 2015 is small enough that it would have a very small impact on our welfare calculations. We then assume that the composition of marginal power plants did not dramatically change between 2012 and 2015, and use the 2015 mean MOER from the WattTime dataset.

Then the time-average marginal external carbon cost is

$$\underbrace{914.83}_{\text{MOER}} \frac{\text{lbs}}{\text{MWh}} \cdot \underbrace{\frac{1 \text{ tonne}}{2204.62 \text{ lb}}}_{=1} \cdot \underbrace{\frac{\$40}{\text{tonne}}}_{\text{SCC}} = \frac{\$16.60}{\text{MWh}} \quad \text{for 2012}$$

and

$$\underbrace{914.83}_{\text{MOER}} \frac{\text{lbs}}{\text{MWh}} \cdot \underbrace{\frac{1 \text{ tonne}}{2204.62 \text{ lb}}}_{=1} \cdot \underbrace{\frac{\$(40 - 12)}{\text{tonne}}}_{\text{SCC}} = \frac{\$11.62}{\text{MWh}} \quad \text{for 2013 and 2014}$$

To show that using the mean MOER in place of the hourly results does not result in excessive error, we rely on our observation that the coefficient of variation for CAISO MOERs in WattTime dataset is 6.9%. Since variation in external costs is due entirely to variation in the MOER, this results in a standard deviation of \$1.16 / MWh in 2012, and \$0.86 / MWh

in 2013 and 2014. Compared to the pricing errors we observe, this is extremely small.¹¹ We also tried fitting several models predicting MOERs from LMPs, and found LMPs to be quite poor predictors. The R-squared coefficient for a linear regression of MOER on LMP for 2015 is 0.03. Lowess locally linear regression showed a similarly poor fit.

Therefore, in our scenario we can view carbon costs essentially as a fixed adder, which “cancels out” some portion of the welfare loss caused by high volumetric markups, by bringing the social marginal cost up toward the retail price.

A.3.3 Calculation of Capacity Costs

In California, due to both price caps and limited price-responsiveness of demand, the energy markets are seen as inadequate to the task of ensuring sufficient capacity to meet peak demand. To address this, the CPUC requires that LSEs procure sufficient capacity to meet their estimated contribution to system peak load, with a 15% reserve margin, in a bilateral capacity market (Gannon et al., 2015). In the California capacity market, LSEs procure capacity at the monthly level.

To calculate marginal contributions to system capacity cost, we follow Boomhower and Davis (2016), whose primary concern is to evaluate the benefits of energy efficiency investments, which deliver time-varying reductions in consumption. They rely on the 2013-2014 CPUC Resource Adequacy Report (Gannon et al., 2015), which presents data on a survey of resource adequacy (RA) contracts (covering generally around 10% of all RA contracts), including the mean, weighted average, maximum, minimum, and 85th percentile of contract costs per kW-year. It presents these figures by month-of-year for all of CAISO (Gannon et al., 2015, Table 13), and regionally, aggregated over all months (Gannon et al., 2015, Table 12).

Boomhower and Davis (2016, Appendix B) take the 85th percentile of contract prices, in \$/kW-month, for all of California, disaggregated by month, to represent the marginal cost of adding or maintaining capacity. Then they consider several methods of allocating percentages of contribution to peak capacity needs across hours. The result of such an allocation procedure is to arrive at a \$/kWh capacity charge for each hour of the month, in proportion to their contribution to peak, such that the charges add up to the original \$/kW-month quantity. We follow one of their three methods, which assigns one third of the peak capacity cost to each of the peak three hours of system load. We obtain publicly available system load data from LCG Consulting (2016).

This method has several shortcomings, but it seems to be the best achievable with publicly available data. Firstly, the CPUC survey covers only 31% of such contracts in CAISO, and we have no assurance that this subset of contracts is representative of the population. Further, we can see in (Gannon et al., 2015, Table 12) that capacity contracts reported for capacity local to the Bay Area and other PG&E areas settle at lower prices than in other regions,

¹¹We can draw a similar conclusion with publicly available data. In CAISO, the time-sensitive estimate of the marginal cost of carbon differs from its average by less than 2.9%, or \$0.463/MWh, more than 95% of the time (Callaway et al., 2015, p. 19).

but we do not have monthly prices available for our geographical regions of interest – only for the CAISO territory as a whole.¹²

Pfeifenberger et al. (2012) provide additional background on this topic. They argue that because California’s long-run Resource Adequacy requirement is currently met by bilateral contracting, it is difficult to estimate the value of marginal capacity in California. They report that as of 2012, the CPUC cost-effectiveness test assumed that peak reductions from that DR resources provided savings of \$136. kW-year. In contrast, because of excess capacity, they argue that the capacity could be acquired for as little as \$18-38/kW-year (Pfeifenberger et al., 2012, p. 2).

A.4 Data

Our simulations use different historical data as inputs, including time series of CAISO LMPs, weather, and representative historical consumption and data on the various tariffs offered by PG&E. In this section we describe the sources for this data and how the raw data has been processed. All data used in our simulations is available for download (Balandat et al., 2016b) so it can be used with our python package `pyDR` (Balandat et al., 2016a) to reproduce the results reported in this paper.

Wholesale Electricity Prices CAISO defines a total of 23 so-called Sub-Load-Aggregation Points (SLAP). Among other roles, these SLAPs are the pricing points on which compensation for Demand Response resources registered as Proxy Demand Resources (PDR) is based (California Independent System Operator Corporation, 2013). We used real-time market (RTM) data for the years 2012-2014 for the following SLAPs: PGCC (Central Coast), PGEB (SF East Bay), PGF1 (Fresno), PGP2 (SF Peninsula) and PGSA (Sacramento). The data was scraped on 15 minute resolution from the CAISO OASIS API (California Independent System Operator Corporation, 2015b). For the hourly resolution of our model, we used the average of the the real-time LMP within each hour.

Tariffs The schedules for the PG&E commercial electricity tariffs used in our study are provided by PG&E in form of a spreadsheet (Pacific Gas and Electric Company, 2016b). The tariffs used for the simulation in this paper are also included in our python package `pyDR` (Balandat et al., 2016a).

Weather Our HVAC model relies on a number of external inputs. Historical outside temperature and Global Horizontal incident (GHI) solar radiation data for each of the geographic locations was obtained from the publicly available CIMIS data set (California Department

¹²In fact, the CPUC capacity (“Long Run Adequacy”) regulations are more complex than we have indicated, because in addition to the total capacity requirement, LSEs are also required to ensure that an administratively-determined portion of their capacity is in their local area, so that sufficient capacity is still available during grid congestion. Capacity prices vary considerably by local area.

of Water Resources, 2015). For each SLAP node, we chose the center of the largest city associated with the node as the representative location of our prototype consumer. We obtained localized estimates of the different weather variables by performing a barycentric interpolation over the CIMIS weather stations. The four different components of the solar radiation used in the HVAC building model simulations were computed using the open source python library pvlib (pvlib, 2016).

PDP events For peak day pricing, we use data on the historical occurrence of “Smart-Days,” which are days on which the PG&E residential “SmartRate” critical peak prices are charged. This data is available from Pacific Gas and Electric Company (2016e), and we use the data from 2012 to 2014, contained in Table A.2 below, in our simulations.

2012	7/9	7/10	7/11	7/23	9/4	9/13	9/14	10/1	10/2	10/3		
2013	6/7	6/28	7/1	7/2	7/19	8/19	9/9	9/10				
2014	5/14	6/9	6/30	7/1	7/7	7/14	7/25	7/28	7/29	7/31	9/11	9/12

Table A.2: PG&E “Smart Days”

Demand Profiles We calibrate the parameters for the quadratic utility model (see Section A.2.1 for details) on the demand from the “dynamic profiles” provided by PG&E (Pacific Gas and Electric Company, 2016a).

A.5 Tables of Results

In Tables A.3 - A.14, we display the simulation results for the QU consumer, for each tariff, DR type (including PDP as a DR type), battery size, and elasticity. Results are averaged across years (2012, 2013, 2014) and region (PGCC, PGEB, PGF1, PGP2, PGSA).

In the tables “CS” is consumer surplus (consumption utility minus consumer expenditure), “RES” is “retail energy surplus” (consumer expenditure minus LMP-weighted consumption and carbon externalized carbon costs), and “Cap” is capacity costs. “SS” is social surplus (here reported as $SS = CS + RES - Cap$), and “Gen” is marginal generation cost (LMP-weighted consumption, plus externalized carbon cost this implies that consumer expenditure can be calculated as $RES + Gen$). “VSEAR” is virtual social energy arbitrage revenue (the revenue that would be generated if all charge-discharge cycles were purchases and sales of energy, at the LMP plus externalized carbon cost; see footnote 62 in Section 2.5.2) and “VPEAR” is private energy arbitrage revenue (the savings in individual expenditure due to battery usage, holding control decisions at their actual values, including PDP revenues, but excluding baseline-dependent DR revenues).

Tariff	DR Type	CS	RES	Cap	SS	Gen	VSEAR	VPEAR
A1	None	37501	3002	198	40305	1011	0	0
A1	CAISO	37500	3004	194	40309	1004	0	0
A1	LMP-G	37498	3005	194	40309	1005	0	0
A1	BLT	37503	3000	197	40306	1008	0	0
A1	LMP-G BLT	37502	3001	198	40306	1009	0	0
A1TOU	None	37457	3045	197	40304	1010	0	0
A1TOU	CAISO	37455	3046	193	40308	1003	0	0
A1TOU	LMP-G	37454	3047	193	40308	1004	0	0
A1TOU	BLT	37459	3042	197	40305	1007	0	0
A1TOU	LMP-G BLT	37458	3044	197	40305	1008	0	0
A1TOU	PDP	37557	2948	197	40308	1011	0	0
A6TOU	None	36506	3915	186	40235	997	0	0
A6TOU	CAISO	36505	3916	183	40238	990	0	0
A6TOU	LMP-G	36504	3917	183	40238	991	0	0
A6TOU	BLT	36508	3912	186	40234	994	0	0
A6TOU	LMP-G BLT	36507	3914	186	40235	995	0	0
A6TOU	PDP	36930	3532	190	40272	1003	0	0
Opt Flat	None	40600	-46	205	40349	1048	0	0
Opt Flat	CAISO	40598	-44	201	40353	1041	0	0
Opt Flat	LMP-G	40598	-44	201	40353	1041	0	0
Opt Flat	BLT	40603	-47	205	40350	1045	0	0
Opt Flat	LMP-G BLT	40602	-47	205	40350	1045	0	0
A1 RTP	None	38429	2106	197	40338	1015	0	0
SMC RTP	None	40558	0	202	40356	1039	0	0

Table A.3: QU results; $E_d = -0.05$, Battery size = None

Tariff	DR Type	CS	RES	Cap	SS	Gen	VSEAR	VPEAR
A1	None	37501	3002	198	40305	1011	0	0
A1	CAISO	37612	2924	192	40344	981	31	-43
A1	LMP-G	37574	2967	188	40352	975	36	-30
A1	BLT	37555	2991	184	40362	968	43	-5
A1	LMP-G BLT	37538	3008	183	40363	968	43	-5
A1TOU	None	37459	3039	197	40301	1014	-3	3
A1TOU	CAISO	37572	2962	184	40350	980	31	-36
A1TOU	LMP-G	37532	3006	182	40356	976	35	-23
A1TOU	BLT	37514	3033	171	40376	965	45	0
A1TOU	LMP-G BLT	37495	3051	172	40375	965	45	0
A1TOU	PDP	37561	2939	197	40302	1017	-6	4
A6TOU	None	37142	3265	185	40222	1013	-16	636
A6TOU	CAISO	37235	3212	165	40283	977	21	607
A6TOU	LMP-G	37196	3248	172	40272	978	19	619
A6TOU	BLT	37191	3257	162	40286	971	26	633
A6TOU	LMP-G BLT	37170	3274	170	40274	976	21	634
A6TOU	PDP	37389	3058	189	40258	1019	-16	459
Opt Flat	None	40600	-46	205	40349	1048	0	0
Opt Flat	CAISO	40757	-179	211	40368	1024	25	-13
Opt Flat	LMP-G	40736	-157	211	40368	1024	25	-13
Opt Flat	BLT	40660	-61	191	40407	1003	45	-1
Opt Flat	LMP-G BLT	40652	-53	191	40407	1004	44	-1
A1 RTP	None	38546	2144	149	40541	866	155	119
SMC RTP	None	40733	0	122	40611	868	178	189

Table A.4: QU results; $E_d = -0.05$, Battery size = Medium

Tariff	DR Type	CS	RES	Cap	SS	Gen	VSEAR	VPEAR
A1	None	37502	3001	198	40305	1011	0	1
A1	CAISO	37725	2767	227	40265	1027	-15	-143
A1	LMP-G	37645	2863	220	40288	1009	3	-100
A1	BLT	37556	2989	185	40360	969	42	-4
A1	LMP-G BLT	37538	3007	183	40361	969	41	-4
A1TOU	None	37461	3038	197	40301	1014	-3	4
A1TOU	CAISO	37682	2815	205	40292	1020	-8	-123
A1TOU	LMP-G	37601	2908	192	40317	1006	5	-83
A1TOU	BLT	37513	3032	171	40374	966	44	1
A1TOU	LMP-G BLT	37496	3049	172	40373	966	44	1
A1TOU	PDP	37563	2938	197	40303	1016	-5	6
A6TOU	None	37842	2632	73	40401	1017	-10	1360
A6TOU	CAISO	38008	2476	78	40407	1011	-3	1286
A6TOU	LMP-G	37950	2539	76	40412	1005	3	1309
A6TOU	BLT	37878	2624	73	40429	988	18	1358
A6TOU	LMP-G BLT	37866	2636	73	40429	989	18	1358
A6TOU	PDP	37903	2584	81	40406	1020	-10	984
Opt Flat	None	40600	-46	205	40349	1048	0	0
Opt Flat	CAISO	40983	-498	308	40177	1118	-68	-51
Opt Flat	LMP-G	40929	-443	303	40184	1116	-67	-51
Opt Flat	BLT	40660	-62	192	40406	1004	44	-1
Opt Flat	LMP-G BLT	40652	-55	192	40406	1004	44	-1
A1 RTP	None	38608	2180	106	40683	769	254	182
SMC RTP	None	40852	0	61	40791	751	297	320

Table A.5: QU results; $E_d = -0.05$, Battery size = Large

Tariff	DR Type	CS	RES	Cap	SS	Gen	VSEAR	VPEAR
A1	None	18631	3002	198	21435	1011	0	0
A1	CAISO	18634	3002	194	21443	999	0	0
A1	LMP-G	18632	3005	194	21443	1001	0	0
A1	BLT	18635	2999	197	21437	1001	0	0
A1	LMP-G BLT	18633	3001	197	21437	1004	0	0
A1TOU	None	18587	3043	196	21433	1009	0	0
A1TOU	CAISO	18591	3043	192	21441	997	0	0
A1TOU	LMP-G	18588	3045	192	21441	999	0	0
A1TOU	BLT	18591	3039	195	21435	1000	0	0
A1TOU	LMP-G BLT	18590	3041	196	21435	1002	0	0
A1TOU	PDP	18688	2950	197	21440	1012	0	0
A6TOU	None	17684	3785	175	21294	983	0	0
A6TOU	CAISO	17688	3784	171	21301	971	0	0
A6TOU	LMP-G	17685	3787	172	21300	974	0	0
A6TOU	BLT	17688	3779	174	21293	973	0	0
A6TOU	LMP-G BLT	17686	3784	174	21295	977	0	0
A6TOU	PDP	18083	3467	182	21369	995	0	0
Opt Flat	None	21782	-48	213	21522	1085	0	0
Opt Flat	CAISO	21786	-47	209	21530	1073	0	0
Opt Flat	LMP-G	21785	-46	209	21530	1074	0	0
Opt Flat	BLT	21787	-49	212	21526	1076	0	0
Opt Flat	LMP-G BLT	21786	-48	212	21526	1077	0	0
A1 RTP	None	19569	2129	197	21501	1018	0	0
SMC RTP	None	21743	0	206	21537	1068	0	0

Table A.6: QU results; $E_d = -0.1$, Battery size = None

Tariff	DR Type	CS	RES	Cap	SS	Gen	VSEAR	VPEAR
A1	None	18631	3002	198	21435	1011	0	0
A1	CAISO	18743	2925	192	21475	982	31	-43
A1	LMP-G	18704	2968	189	21483	976	36	-30
A1	BLT	18685	2991	184	21492	968	43	-5
A1	LMP-G BLT	18668	3008	183	21493	968	42	-4
A1TOU	None	18590	3037	196	21430	1013	-3	3
A1TOU	CAISO	18703	2962	184	21480	980	30	-36
A1TOU	LMP-G	18662	3005	182	21485	976	35	-23
A1TOU	BLT	18644	3031	170	21504	964	45	0
A1TOU	LMP-G BLT	18626	3049	171	21503	964	45	0
A1TOU	PDP	18691	2940	197	21434	1017	-6	4
A6TOU	None	18318	3140	174	21284	999	-15	635
A6TOU	CAISO	18410	3091	154	21347	965	19	606
A6TOU	LMP-G	18372	3123	161	21334	966	18	618
A6TOU	BLT	18366	3132	152	21345	958	24	632
A6TOU	LMP-G BLT	18346	3147	160	21334	962	20	632
A6TOU	PDP	18542	2997	181	21357	1011	-15	458
Opt Flat	None	21782	-48	213	21522	1085	0	0
Opt Flat	CAISO	21942	-182	217	21544	1060	27	-13
Opt Flat	LMP-G	21921	-160	217	21544	1060	27	-13
Opt Flat	BLT	21844	-63	198	21583	1038	46	-1
Opt Flat	LMP-G BLT	21837	-55	198	21583	1039	46	-1
A1 RTP	None	19684	2168	149	21703	875	155	119
SMC RTP	None	21919	0	126	21792	900	181	192

Table A.7: QU results; $E_d = -0.1$, Battery size = Medium

Tariff	DR Type	CS	RES	Cap	SS	Gen	VSEAR	VPEAR
A1	None	18632	3001	198	21435	1011	0	1
A1	CAISO	18856	2769	227	21398	1029	-15	-143
A1	LMP-G	18776	2865	221	21420	1010	3	-100
A1	BLT	18686	2989	184	21490	969	41	-4
A1	LMP-G BLT	18668	3006	183	21491	969	41	-4
A1TOU	None	18591	3036	196	21430	1013	-3	4
A1TOU	CAISO	18813	2816	205	21424	1021	-8	-123
A1TOU	LMP-G	18731	2908	191	21448	1007	5	-84
A1TOU	BLT	18644	3030	170	21503	965	44	1
A1TOU	LMP-G BLT	18626	3047	171	21502	966	43	1
A1TOU	PDP	18694	2939	197	21436	1016	-5	6
A6TOU	None	18984	2613	53	21543	1015	-10	1351
A6TOU	CAISO	19151	2458	63	21546	1010	-3	1278
A6TOU	LMP-G	19092	2520	57	21555	1003	3	1301
A6TOU	BLT	19020	2605	53	21572	986	18	1349
A6TOU	LMP-G BLT	19009	2616	53	21572	987	18	1349
A6TOU	PDP	19040	2574	66	21549	1019	-10	980
Opt Flat	None	21782	-48	213	21522	1085	0	0
Opt Flat	CAISO	22171	-504	317	21350	1155	-66	-51
Opt Flat	LMP-G	22116	-448	311	21357	1153	-65	-51
Opt Flat	BLT	21845	-65	199	21581	1040	45	-1
Opt Flat	LMP-G BLT	21837	-57	199	21581	1040	45	-1
A1 RTP	None	19746	2206	107	21845	779	255	184
SMC RTP	None	22039	0	66	21973	783	303	326

Table A.8: QU results; $E_d = -0.1$, Battery size = Large

Tariff	DR Type	CS	RES	Cap	SS	Gen	VSEAR	VPEAR
A1	None	9196	3002	198	12000	1011	0	0
A1	CAISO	9207	2999	191	12015	988	0	0
A1	LMP-G	9202	3004	192	12014	993	0	0
A1	BLT	9205	2994	195	12005	987	0	0
A1	LMP-G BLT	9202	3000	196	12006	993	0	0
A1TOU	None	9153	3038	195	11996	1007	0	0
A1TOU	CAISO	9164	3035	188	12011	984	0	0
A1TOU	LMP-G	9159	3041	190	12010	989	0	0
A1TOU	BLT	9162	3030	192	12001	983	0	0
A1TOU	LMP-G BLT	9158	3036	193	12002	990	0	0
A1TOU	PDP	9253	2953	196	12010	1012	0	0
A6TOU	None	8345	3524	152	11717	955	0	0
A6TOU	CAISO	8356	3519	145	11731	932	0	0
A6TOU	LMP-G	8349	3527	148	11728	939	0	0
A6TOU	BLT	8355	3509	145	11718	931	0	0
A6TOU	LMP-G BLT	8349	3522	151	11720	941	0	0
A6TOU	PDP	8695	3339	167	11867	978	0	0
Opt Flat	None	12452	-51	228	12173	1159	0	0
Opt Flat	CAISO	12463	-53	221	12189	1136	0	0
Opt Flat	LMP-G	12461	-51	222	12189	1138	0	0
Opt Flat	BLT	12462	-51	225	12185	1136	0	0
Opt Flat	LMP-G BLT	12460	-50	226	12185	1139	0	0
A1 RTP	None	10154	2173	196	12132	1025	0	0
SMC RTP	None	12418	0	215	12203	1125	0	0

Table A.9: QU results; $E_d = -0.2$, Battery size = None

Tariff	DR Type	CS	RES	Cap	SS	Gen	VSEAR	VPEAR
A1	None	9196	3002	198	12000	1011	0	0
A1	CAISO	9309	2927	193	12043	983	30	-43
A1	LMP-G	9270	2970	189	12050	977	36	-30
A1	BLT	9250	2990	184	12056	967	42	-4
A1	LMP-G BLT	9233	3008	183	12058	968	42	-4
A1TOU	None	9155	3032	195	11993	1011	-3	3
A1TOU	CAISO	9269	2960	183	12046	981	30	-36
A1TOU	LMP-G	9228	3003	181	12051	975	34	-23
A1TOU	BLT	9210	3026	169	12067	962	44	0
A1TOU	LMP-G BLT	9191	3044	170	12066	962	44	0
A1TOU	PDP	9257	2944	196	12005	1018	-6	4
A6TOU	None	8976	2895	149	11721	970	-13	632
A6TOU	CAISO	9064	2856	132	11788	942	17	604
A6TOU	LMP-G	9028	2881	139	11770	941	17	615
A6TOU	BLT	9019	2886	132	11773	933	22	629
A6TOU	LMP-G BLT	9002	2901	138	11765	936	19	630
A6TOU	PDP	9151	2875	165	11861	994	-15	457
Opt Flat	None	12452	-51	228	12173	1159	0	0
Opt Flat	CAISO	12618	-188	229	12201	1132	31	-13
Opt Flat	LMP-G	12595	-165	229	12201	1131	31	-13
Opt Flat	BLT	12519	-67	212	12239	1109	50	-1
Opt Flat	LMP-G BLT	12510	-59	212	12239	1109	50	-1
A1 RTP	None	10265	2217	150	12332	894	156	120
SMC RTP	None	12594	0	136	12459	965	186	198

Table A.10: QU results; $E_d = -0.2$, Battery size = Medium

Tariff	DR Type	CS	RES	Cap	SS	Gen	VSEAR	VPEAR
A1	None	9197	3001	198	12000	1011	0	1
A1	CAISO	9422	2775	228	11969	1032	-15	-143
A1	LMP-G	9342	2869	221	11989	1012	3	-100
A1	BLT	9251	2989	184	12055	969	41	-4
A1	LMP-G BLT	9233	3006	183	12056	969	41	-4
A1TOU	None	9157	3032	195	11993	1011	-3	4
A1TOU	CAISO	9380	2818	204	11994	1023	-8	-123
A1TOU	LMP-G	9297	2909	190	12016	1007	5	-84
A1TOU	BLT	9209	3026	169	12066	964	43	1
A1TOU	LMP-G BLT	9192	3043	170	12065	964	43	1
A1TOU	PDP	9259	2943	196	12006	1017	-5	6
A6TOU	None	9564	2602	46	12120	1016	-11	1335
A6TOU	CAISO	9732	2449	57	12125	1012	-3	1262
A6TOU	LMP-G	9673	2511	50	12135	1005	2	1285
A6TOU	BLT	9600	2593	45	12148	986	18	1332
A6TOU	LMP-G BLT	9589	2605	46	12148	987	17	1333
A6TOU	PDP	9618	2571	46	12142	1021	-10	974
Opt Flat	None	12452	-51	228	12173	1159	0	0
Opt Flat	CAISO	12851	-515	333	12003	1228	-62	-51
Opt Flat	LMP-G	12795	-458	327	12010	1226	-61	-51
Opt Flat	BLT	12519	-69	213	12237	1111	48	-1
Opt Flat	LMP-G BLT	12511	-61	213	12237	1111	48	-1
A1 RTP	None	10327	2256	110	12473	800	258	186
SMC RTP	None	12718	0	76	12641	847	314	337

Table A.11: QU results; $E_d = -0.2$, Battery size = Large

Tariff	DR Type	CS	RES	Cap	SS	Gen	VSEAR	VPEAR
A1	None	6051	3002	198	8855	1011	0	0
A1	CAISO	6069	2996	187	8877	978	0	0
A1	LMP-G	6062	3004	190	8876	985	0	0
A1	BLT	6066	2989	188	8868	975	0	0
A1	LMP-G BLT	6060	2999	194	8865	983	0	0
A1TOU	None	6009	3034	193	8849	1006	0	0
A1TOU	CAISO	6027	3027	183	8871	973	0	0
A1TOU	LMP-G	6019	3036	186	8869	980	0	0
A1TOU	BLT	6024	3020	183	8861	969	0	0
A1TOU	LMP-G BLT	6018	3031	190	8858	979	0	0
A1TOU	PDP	6109	2956	195	8870	1013	0	0
A6TOU	None	5296	3263	128	8430	926	0	0
A6TOU	CAISO	5314	3257	119	8452	897	0	0
A6TOU	LMP-G	5303	3267	125	8446	906	0	0
A6TOU	BLT	5311	3240	117	8434	892	0	0
A6TOU	LMP-G BLT	5302	3260	126	8436	905	0	0
A6TOU	PDP	5596	3210	151	8656	962	0	0
Opt Flat	None	9412	-54	243	9115	1233	0	0
Opt Flat	CAISO	9430	-60	233	9138	1199	0	0
Opt Flat	LMP-G	9427	-56	234	9138	1202	0	0
Opt Flat	BLT	9427	-54	235	9138	1197	0	0
Opt Flat	LMP-G BLT	9425	-52	238	9135	1200	0	0
A1 RTP	None	7029	2217	194	9052	1033	0	0
SMC RTP	None	9382	0	223	9159	1182	0	0

Table A.12: QU results; $E_d = -0.3$, Battery size = None

Tariff	DR Type	CS	RES	Cap	SS	Gen	VSEAR	VPEAR
A1	None	6051	3002	198	8855	1011	0	0
A1	CAISO	6165	2930	193	8901	985	30	-43
A1	LMP-G	6125	2971	189	8907	978	35	-30
A1	BLT	6105	2990	184	8911	967	41	-4
A1	LMP-G BLT	6088	3007	183	8912	967	41	-4
A1TOU	None	6011	3028	193	8846	1009	-3	3
A1TOU	CAISO	6126	2959	182	8903	981	30	-36
A1TOU	LMP-G	6085	3001	180	8906	975	34	-23
A1TOU	BLT	6065	3022	168	8919	960	44	0
A1TOU	LMP-G BLT	6047	3040	169	8918	961	43	0
A1TOU	PDP	6113	2947	195	8865	1019	-6	4
A6TOU	None	5920	2672	118	8474	942	-11	628
A6TOU	CAISO	6004	2643	107	8540	921	15	602
A6TOU	LMP-G	5971	2660	113	8518	917	16	612
A6TOU	BLT	5959	2663	109	8513	908	20	625
A6TOU	LMP-G BLT	5945	2677	112	8509	911	19	626
A6TOU	PDP	6051	2757	149	8659	977	-13	455
Opt Flat	None	9412	-54	243	9115	1233	0	0
Opt Flat	CAISO	9584	-194	242	9147	1203	34	-13
Opt Flat	LMP-G	9560	-170	242	9148	1203	34	-13
Opt Flat	BLT	9483	-71	226	9186	1179	53	-1
Opt Flat	LMP-G BLT	9474	-63	226	9185	1179	53	-1
A1 RTP	None	7136	2265	153	9249	914	157	121
SMC RTP	None	9560	0	146	9414	1029	192	204

Table A.13: QU results; $E_d = -0.3$, Battery size = Medium

Tariff	DR Type	CS	RES	Cap	SS	Gen	VSEAR	VPEAR
A1	None	6052	3001	198	8855	1011	0	1
A1	CAISO	6278	2780	228	8830	1035	-15	-143
A1	LMP-G	6197	2872	221	8848	1014	3	-100
A1	BLT	6106	2988	184	8910	968	40	-4
A1	LMP-G BLT	6088	3006	183	8911	969	40	-4
A1TOU	None	6012	3027	193	8847	1009	-3	4
A1TOU	CAISO	6236	2820	203	8853	1025	-9	-123
A1TOU	LMP-G	6154	2909	189	8873	1008	4	-84
A1TOU	BLT	6065	3021	168	8919	962	43	1
A1TOU	LMP-G BLT	6047	3039	168	8917	962	42	1
A1TOU	PDP	6115	2947	195	8867	1018	-5	6
A6TOU	None	6430	2592	45	8977	1016	-11	1318
A6TOU	CAISO	6599	2443	56	8986	1015	-3	1245
A6TOU	LMP-G	6540	2504	49	8995	1007	2	1268
A6TOU	BLT	6467	2583	45	9005	986	17	1315
A6TOU	LMP-G BLT	6455	2595	45	9005	987	17	1316
A6TOU	PDP	6482	2575	46	9011	1025	-11	968
Opt Flat	None	9412	-54	243	9115	1233	0	0
Opt Flat	CAISO	9821	-526	346	8948	1301	-58	-51
Opt Flat	LMP-G	9763	-467	341	8955	1299	-58	-51
Opt Flat	BLT	9483	-74	226	9183	1181	51	-1
Opt Flat	LMP-G BLT	9474	-65	226	9183	1181	51	-1
A1 RTP	None	7199	2307	113	9392	821	261	188
SMC RTP	None	9686	0	87	9600	910	324	348

Table A.14: QU results; $E_d = -0.3$, Battery size = Large

Tariff	DR Type	Indiv	Gen	Cap	-SS
A1	None	146795	47275	20825	68100
A1	BLT	144649	44312	18211	62523
A1TOU	None	148514	47122	20724	67846
A1TOU	CAISO	144595	44006	17477	61483
A1TOU	BLT	146211	43926	17284	61210
A1TOU	LMP-G	146239	44011	17312	61323
A1TOU	LMP-G BLT	146998	44424	18188	62612
A1TOU	PDP	145289	47116	20715	67831
A10	None	149471	47936	11455	59391
A10	BLT	147712	46135	10904	57039
A10TOU	None	147242	47618	11184	58802
A10TOU	BLT	145564	45892	10702	56594
A10TOU	PDP	144832	47629	11150	58779
A6TOU	None	136553	45695	2867	48562
A6TOU	BLT	134969	44162	2778	46940
A6TOU	PDP	135830	45616	3515	49131
E19TOU	None	131046	49598	3361	52959
E19TOU	BLT	129784	48440	3234	51674
E19TOU	PDP	131046	49598	3361	52959
OptFlat	None	39261	47272	20825	68097
OptFlat	CAISO	27985	51660	19106	70766
OptFlat	BLT	36334	43510	17037	60547
OptFlat	LMP-G	29903	50291	19207	69498
OptFlat	LMP-G BLT	36729	43557	17037	60594
A1 RTP	None	112629	37073	15051	52124
SMC RTP	None	33304	33304	6500	39804

Table A.15: HVAC results

The HVAC simulation results, displayed in Table A.15, are also averaged across years and regions. “Indiv” is individual expenditures, “Gen” is generation cost, and “Cap” is capacity cost. If we normalize consumption utility to zero, then consumer surplus is -Indiv, retailer surplus is Indiv - Gen (or Indiv - Gen - Cap if we account for capacity costs), and social surplus is -Gen - Cap. We report -SS, i.e. Gen + Cap in the last column.

Appendix B

Supplementary Material for Chapter 5

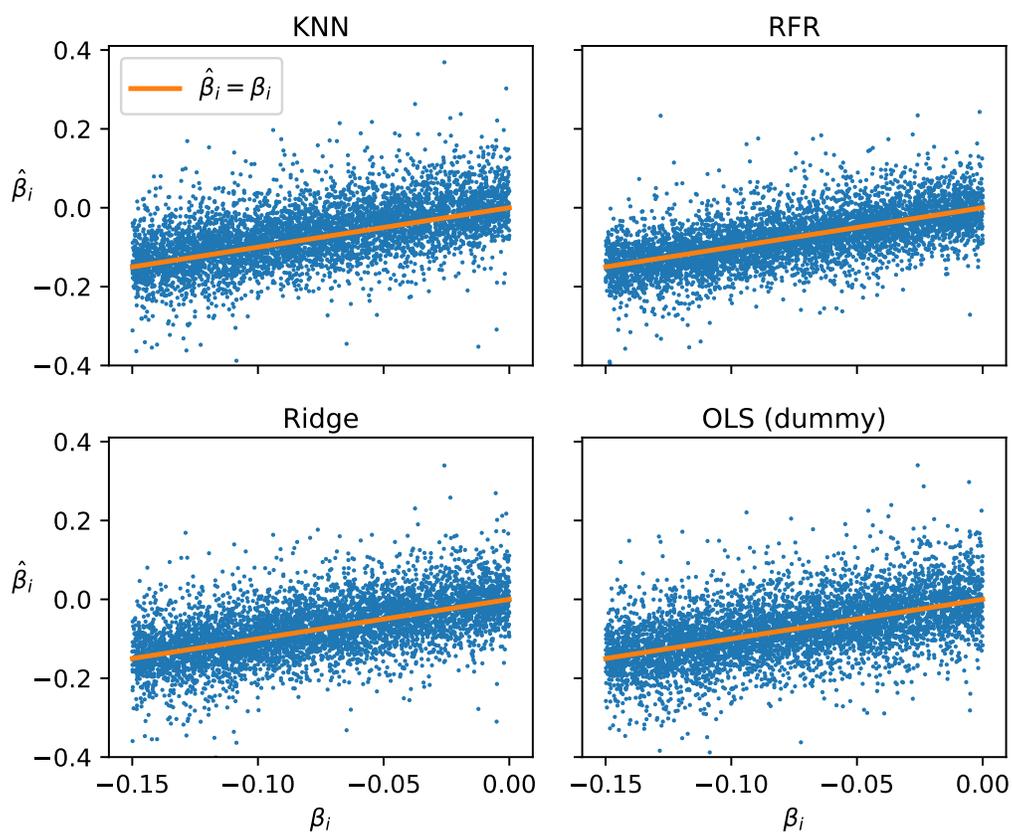


Figure B.1: Parameter estimates vs. true parameters, linear response

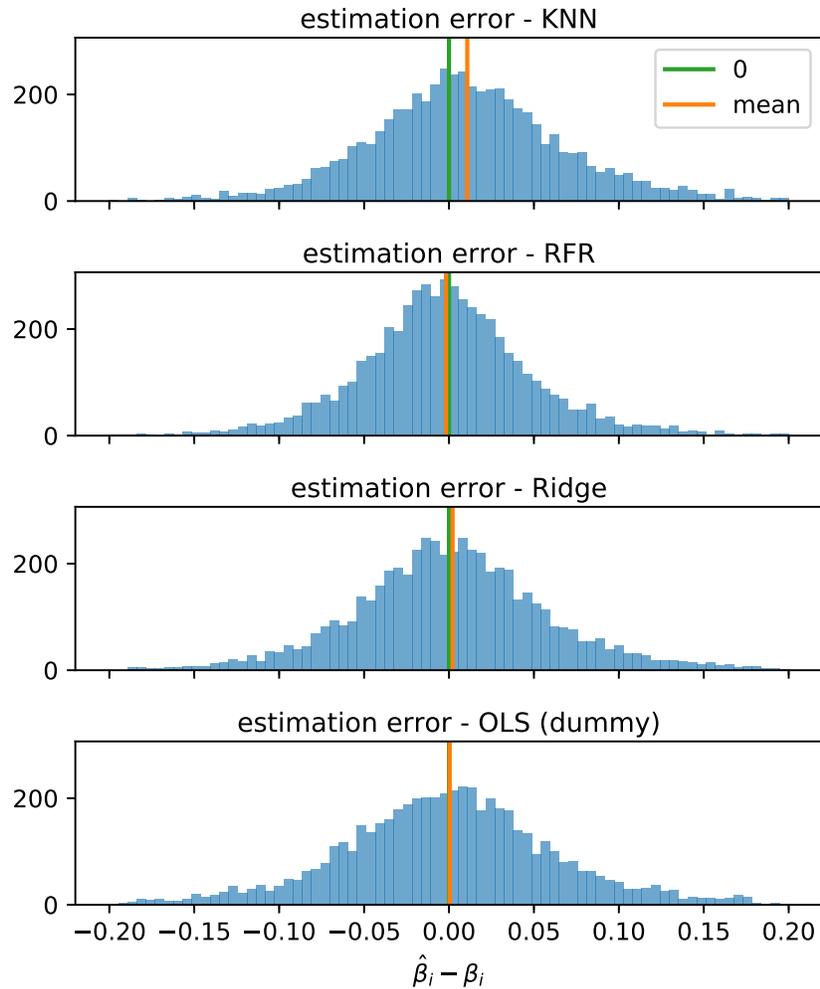


Figure B.2: Empirical distribution of parameter estimation errors, linear response

model	mean	std dev
KNN (unweighted)	0.0107	0.0622
RFR	-0.0018	0.0530
Ridge	0.0021	0.0608
OLS (dummy)	0.0004	0.0674

Table B.1: Properties of the ITE estimation error, linear response

model	method	coverage	error	error (lower)	error (upper)
KNN	BC	0.8967	-0.0533	-0.0395	-0.0138
	BCa	0.8951	-0.0549	-0.0395	-0.0154
	perc	0.9027	-0.0473	-0.0366	-0.0106
RFR	BC	0.8554	-0.0946	-0.0426	-0.0520
	BCa	0.8569	-0.0931	-0.0415	-0.0516
	perc	0.9834	0.0334	0.0202	0.0131
Ridge	BC	0.8823	-0.0677	-0.0337	-0.0340
	BCa	0.8793	-0.0707	-0.0344	-0.0362
	perc	0.8909	-0.0591	-0.0306	-0.0286
OLS (dummy)		0.9093	-0.0407	-0.0220	-0.0187

Table B.2: Empirical coverage of 95% ITE confidence intervals, linear response

Appendix C

Proofs

C.1 Proofs from Chapter 3

Proof of Proposition 1. The state of (3.3) in interval Kt is given by

$$\tilde{x}_{Kt} = \tilde{A}^{Kt}x_0 + \sum_{\theta < Kt} \tilde{A}^{Kt-1-\theta} \tilde{B}(\tilde{u}_\theta + \tilde{w}_\theta) + D\tilde{u}_{Kt}$$

while the state of (3.4) in interval t is given by

$$x_t = A^t x_0 + \sum_{\tau < t} A^{t-1-\tau} B(u_\tau + w_\tau) + Du_t$$

Now $\tilde{u}_{tK+j} = u_t$ for all $j = 0, 1, \dots, K-1$ and all t by assumption. That is, “move-blocking” is applied so that the input \tilde{u} is the same as the input u . Then,

$$\begin{aligned} \tilde{x}_{Kt} &= \tilde{A}^{Kt}x_0 + \sum_{\tau < t} \left(\sum_{j < K} \tilde{A}^{Kt-1-(K\tau+j)} \tilde{B} \right) u_\tau + \sum_{\tau < Kt} \tilde{A}^{Kt-1-\tau} \tilde{B} \tilde{w}_\tau + D\tilde{u}_{Kt} \\ &= \tilde{A}^{Kt}x_0 + \sum_{\tau < t} \tilde{A}^{K(t-1-\tau)} \left(\sum_{j < K} \tilde{A}^{K-1-j} \tilde{B} \right) u_\tau + \sum_{\tau < Kt} \tilde{A}^{Kt-1-\tau} \tilde{B} \tilde{w}_\tau + D\tilde{u}_{Kt} \end{aligned}$$

Observe now that $A = e^{AT_S} = e^{AKT_R} = \tilde{A}^K$ and

$$\begin{aligned} B &= \int_0^{T_S} e^{A(T_S-\tau)} B d\tau = \sum_{j=0}^{K-1} \int_{jT_R}^{(j+1)T_R} e^{A(KT_R-\tau)} B d\tau \\ &= \sum_{j=0}^{K-1} \int_0^{T_R} e^{A(K-1-j)T_R} e^{A(T_R-\tau)} B d\tau \\ &= \sum_{j=0}^{K-1} \tilde{A}^{K-1-j} \tilde{B} \end{aligned}$$

With this we have that

$$\tilde{x}_{Kt} = A^t x_0 + \sum_{\tau < t} A^{(t-1-\tau)} B u_\tau + D\tilde{u}_{Kt} + \sum_{\tau < Kt} \tilde{A}^{Kt-1-\tau} \tilde{B} \tilde{w}_\tau$$

and hence

$$y_t - \tilde{y}_{Kt} = \sum_{\tau < t} C A^{t-1-\tau} B w_\tau - \sum_{\tau < Kt} C \tilde{A}^{Kt-1-\tau} \tilde{B} \tilde{w}_\tau$$

To show that the robustified schedule determined for the system (3.4) ensures robust constraint satisfaction for the system (3.3) it suffices to show that¹, for all $i = 1, \dots, n_y$,

$$\max_{w_0, \dots, w_{t-1}} \sum_{\tau < t} (C)_i: A^{t-1-\tau} B w_\tau = \max_{\tilde{w}_0, \dots, \tilde{w}_{Kt-1}} \sum_{\tau < Kt} (C)_i: \tilde{A}^{Kt-1-\tau} \tilde{B} \tilde{w}_\tau$$

where $-r_\tau^\uparrow \leq w_\tau \leq r_\tau^\downarrow$ for all $\tau < t$ and $-r_\tau^\uparrow \leq \tilde{w}_{K\tau+j} \leq r_\tau^\downarrow$ for all $\tau < t$ and all $j < K$. Since in terms of worst-case analysis there is no coupling between the different w_τ (and similarly between the different \tilde{w}_τ) it suffices to show that the above holds for $t = 1$. Again using $B = \sum_{j=0}^K \tilde{A}^{K-1-j} \tilde{B}$ we must show that

$$\begin{aligned} \max_{-r^\uparrow \leq w \leq r^\downarrow} y_{t,i} &= \max_{-r^\uparrow \leq w \leq r^\downarrow} \left(\sum_{j < K} (C)_i: \tilde{A}^{K-1-j} \tilde{B} \right) w \\ &= \max_{\substack{-r^\uparrow \leq \tilde{w}_j \leq r^\downarrow \\ j=0, \dots, K-1}} \sum_{j < K} (C)_i: \tilde{A}^{K-1-j} \tilde{B} \tilde{w}_j \\ &= \max_{\substack{-r^\uparrow \leq \tilde{w}_j \leq r^\downarrow \\ j=0, \dots, K-1}} \tilde{y}_{Kt,i} \end{aligned}$$

By choosing $\tilde{w}_j = w$ for all j it is clear that

$$\max_{-r^\uparrow \leq w \leq r^\downarrow} y_{t,i} \leq \max_{\substack{-r^\uparrow \leq \tilde{w}_j \leq r^\downarrow \\ j=0, \dots, K-1}} \tilde{y}_{Kt,i} \quad (\text{C.1})$$

Now write

$$\begin{aligned} \left(\sum_{j < K} (C)_i: \tilde{A}^{K-1-j} \tilde{B} \right) w &= \sum_{j < K} \sum_{k=1}^{n_u} \eta_{ijk} w_k \\ \sum_{j < K} (C)_i: \tilde{A}^{K-1-j} \tilde{B} \tilde{w}_j &= \sum_{j < K} \sum_{k=1}^{n_u} \eta_{ijk} \tilde{w}_{j,k} \end{aligned}$$

and observe that if the condition in Proposition 1 is satisfied then (C.1) holds with equality. \square

¹The argument for the minimum is identical.

C.2 Proofs from Chapter 4

Proof of Lemma 1. The bias (4.12a) is straightforward and follows directly from the definition of the estimation error (4.11). For the variance we have that

$$\begin{aligned}
& \mathbb{E}_{\mathcal{D}_{i,\text{est}}} \left[(\hat{\mu}_i - \mathbb{E}_{\mathcal{D}_{i,\text{est}}} [\hat{\mu}_i \mid \hat{f}_i^0, \mathcal{T}_i])^2 \mid \hat{f}_i^0, \mathcal{T}_i \right] \\
&= \frac{1}{|\mathcal{T}_i|^2} \mathbb{E}_{\mathcal{D}_{i,\text{est}}} \left[\left(\sum_{t \in \mathcal{T}_i} e_i(X_{it}) - \mathbb{E}_{\mathcal{D}_{i,\text{est}}} [e_i(X_{it})] - \epsilon_{it} - \gamma_{it} \right)^2 \mid \mathcal{T}_i \right] \\
&= \frac{1}{|\mathcal{T}_i|^2} \sum_{t \in \mathcal{T}_i} \mathbb{E}_{\mathcal{D}_{i,\text{est}}} \left[(e_i(X_{it}) - \mathbb{E}_{\mathcal{D}_{i,\text{est}}} [e_i(X_{it})])^2 + (\epsilon_{it} + \gamma_{it})^2 \right. \\
&\quad \left. - 2(\epsilon_{it} + \gamma_{it})(e_i(X_{it}) - \mathbb{E}_{\mathcal{D}_{i,\text{est}}} [e_i(X_{it})]) \mid D_{it}=1 \right] \\
&\quad + \frac{1}{|\mathcal{T}_i|^2} \sum_{t \in \mathcal{T}_i} \sum_{\substack{\tau \in \mathcal{T}_i \\ \tau \neq t}} \mathbb{E}_{\mathcal{D}_{i,\text{est}}} \left[(e_i(X_{it}) - \mathbb{E}_{\mathcal{D}_{i,\text{est}}} [e_i(X_{it})]) (e_i(X_{i\tau}) - \mathbb{E}_{\mathcal{D}_{i,\text{est}}} [e_i(X_{i\tau})]) \right. \\
&\quad \left. - (\epsilon_{it} + \gamma_{it})(e_i(X_{i\tau}) - \mathbb{E}_{\mathcal{D}_{i,\text{est}}} [e_i(X_{i\tau})]) \right. \\
&\quad \left. - (\epsilon_{i\tau} + \gamma_{i\tau})(e_i(X_{it}) - \mathbb{E}_{\mathcal{D}_{i,\text{est}}} [e_i(X_{it})]) \right. \\
&\quad \left. + (\epsilon_{it} + \gamma_{it})(\epsilon_{i\tau} + \gamma_{i\tau}) \mid D_{it}=D_{i\tau}=1 \right] \\
&= \frac{1}{|\mathcal{T}_i|^2} \sum_{t \in \mathcal{T}_i} \mathbb{E}_{\mathcal{D}_{i,\text{est}}} \left[(e_i(X_{it}) - \mathbb{E}_{\mathcal{D}_{i,\text{est}}} [e_i(X_{it})])^2 + (\epsilon_{it} - \mathbb{E}[\epsilon_{it} \mid D_{it}=1])^2 + (\gamma_{it} - \mathbb{E}[\gamma_{it} \mid D_{it}=1])^2 \right. \\
&\quad \left. + 2(\epsilon_{it} - \mathbb{E}[\epsilon_{it} \mid D_{it}=1])(\gamma_{it} - \mathbb{E}[\gamma_{it} \mid D_{it}=1]) + (\mathbb{E}[\epsilon_{it} \mid D_{it}=1] + \mathbb{E}[\gamma_{it} \mid D_{it}=1])^2 \right. \\
&\quad \left. + 2(\epsilon_{it} - \mathbb{E}[\epsilon_{it} \mid D_{it}=1] + \gamma_{it} - \mathbb{E}[\gamma_{it} \mid D_{it}=1])(\mathbb{E}[\epsilon_{it} \mid D_{it}=1] + \mathbb{E}[\gamma_{it} \mid D_{it}=1]) \right. \\
&\quad \left. - 2(\epsilon_{it} - \mathbb{E}[\epsilon_{it} \mid D_{it}=1] + \gamma_{it} - \mathbb{E}[\gamma_{it} \mid D_{it}=1])(e_i(X_{it}) - \mathbb{E}_{\mathcal{D}_{i,\text{est}}} [e_i(X_{it})]) \right. \\
&\quad \left. - 2(\mathbb{E}[\epsilon_{it} \mid D_{it}=1] + \mathbb{E}[\gamma_{it} \mid D_{it}=1])(e_i(X_{it}) - \mathbb{E}_{\mathcal{D}_{i,\text{est}}} [e_i(X_{it})]) \mid D_{it}=1 \right]
\end{aligned}$$

$$\begin{aligned}
& + \frac{1}{|\mathcal{T}_i|^2} \sum_{t \in \mathcal{T}_i} \sum_{\substack{\tau \in \mathcal{T}_i \\ \tau \neq t}} \mathbb{E}_{\mathcal{D}_{i,\text{est}}} \left[(e_i(X_{it}) - \mathbb{E}_{\mathcal{D}_{i,\text{est}}} [e_i(X_{it})]) (e_i(X_{i\tau}) - \mathbb{E}_{\mathcal{D}_{i,\text{est}}} [e_i(X_{i\tau})]) \right. \\
& \quad - (\epsilon_{it} - \mathbb{E}[\epsilon_{it} | D_{it} = 1] + \gamma_{it} - \mathbb{E}[\gamma_{it} | D_{it} = 1]) (e_i(X_{i\tau}) - \mathbb{E}_{\mathcal{D}_{i,\text{est}}} [e_i(X_{i\tau})]) \\
& \quad - (\mathbb{E}[\epsilon_{it} | D_{it} = 1] + \mathbb{E}[\gamma_{it} | D_{it} = 1]) (e_i(X_{i\tau}) - \mathbb{E}_{\mathcal{D}_{i,\text{est}}} [e_i(X_{i\tau})]) \\
& \quad - (\epsilon_{i\tau} - \mathbb{E}[\epsilon_{i\tau} | D_{i\tau} = 1] + \gamma_{i\tau} - \mathbb{E}[\gamma_{i\tau} | D_{i\tau} = 1]) (e_i(X_{it}) - \mathbb{E}_{\mathcal{D}_{i,\text{est}}} [e_i(X_{it})]) \\
& \quad - (\mathbb{E}[\epsilon_{i\tau} | D_{i\tau} = 1] + \mathbb{E}[\gamma_{i\tau} | D_{i\tau} = 1]) (e_i(X_{it}) - \mathbb{E}_{\mathcal{D}_{i,\text{est}}} [e_i(X_{it})]) \\
& \quad + (\epsilon_{it} - \mathbb{E}[\epsilon_{it} | D_{it} = 1]) (\epsilon_{i\tau} - \mathbb{E}[\epsilon_{i\tau} | D_{i\tau} = 1]) \\
& \quad + (\gamma_{it} - \mathbb{E}[\gamma_{it} | D_{it} = 1]) (\gamma_{i\tau} - \mathbb{E}[\gamma_{i\tau} | D_{i\tau} = 1]) \\
& \quad + (\epsilon_{it} - \mathbb{E}[\epsilon_{it} | D_{it} = 1]) (\gamma_{i\tau} - \mathbb{E}[\gamma_{i\tau} | D_{i\tau} = 1]) \\
& \quad + (\epsilon_{i\tau} - \mathbb{E}[\epsilon_{i\tau} | D_{i\tau} = 1]) (\gamma_{it} - \mathbb{E}[\gamma_{it} | D_{it} = 1]) \\
& \quad + (\epsilon_{it} - \mathbb{E}[\epsilon_{it} | D_{it} = 1] + \gamma_{it} - \mathbb{E}[\gamma_{it} | D_{it} = 1]) \cdot \\
& \quad \quad (\mathbb{E}[\epsilon_{i\tau} | D_{i\tau} = 1] + \mathbb{E}[\gamma_{i\tau} | D_{i\tau} = 1]) \\
& \quad + (\epsilon_{i\tau} - \mathbb{E}[\epsilon_{i\tau} | D_{i\tau} = 1] + \gamma_{i\tau} - \mathbb{E}[\gamma_{i\tau} | D_{i\tau} = 1]) \cdot \\
& \quad \quad (\mathbb{E}[\epsilon_{it} | D_{it} = 1] + \mathbb{E}[\gamma_{it} | D_{it} = 1]) \\
& \quad + (\mathbb{E}[\epsilon_{i\tau} | D_{i\tau} = 1] + \mathbb{E}[\gamma_{i\tau} | D_{i\tau} = 1]) (\mathbb{E}[\epsilon_{it} | D_{it} = 1] + \mathbb{E}[\gamma_{it} | D_{it} = 1]) \\
& \quad \quad \left. | D_{it} = D_{i\tau} = 1 \right]
\end{aligned}$$

Now (4.12b) follows by exploiting symmetry and the independence of ϵ_{it} and γ_{it} . \square

Proof of Lemma 3. Fix treatment times and parameters as well as the covariates $X_i := \{X_{it}\}_{t \in \mathcal{T}_i}$ during treatment periods. Conditional on this selection,

$$\mathbb{E}_{\mathcal{D}_{i,\text{tr}}} [\hat{\mu}_i | \mathcal{T}_i, X_i] - \mu_i = -\frac{1}{|\mathcal{T}_i|} \sum_{t \in \mathcal{T}_i} (\text{Bias } \hat{f}_i^0(X_{it}) - \epsilon_{it} - \gamma_{it})$$

Taking the expectation w.r.t. $\mathcal{D}_{i,\text{est}}$ (here we use that $\mathcal{D}_{i,\text{est}} \perp \mathcal{D}_{i,\text{tr}}$) yields (4.14a). Moreover,

$$\begin{aligned}
\text{Var}(\hat{\mu}_i | \mathcal{T}_i, X_i) &= \mathbb{E} \left[(\hat{\mu}_i - \mathbb{E}[\hat{\mu}_i | \mathcal{T}_i, X_i])^2 | \mathcal{T}_i, X_i \right] \\
&= \frac{1}{|\mathcal{T}_i|^2} \sum_{t \in \mathcal{T}_i} \sum_{\tau \in \mathcal{T}_i} \mathbb{E} \left[(\mathbb{E}_{\mathcal{D}_{i,\text{tr}}} [\hat{f}_i^0(X_{it})] - \hat{f}_i^0(X_{it})) (\mathbb{E}_{\mathcal{D}_{i,\text{tr}}} [\hat{f}_i^0(X_{i\tau})] - \hat{f}_i^0(X_{i\tau})) | X_i \right] \\
&= \frac{1}{|\mathcal{T}_i|^2} \sum_{t \in \mathcal{T}_i} \left(\text{Var } \hat{f}_i^0(X_{it}) + 2 \sum_{\substack{\tau \in \mathcal{T}_i \\ \tau > t}} \text{Cov } \hat{f}_i^0(X_{it}, X_{i\tau}) \right)
\end{aligned}$$

From the law of total variance, we have

$$\text{Var}(\hat{\mu}_i | \mathcal{T}_i) = \mathbb{E} \left[\text{Var}(\hat{\mu}_i | \mathcal{T}_i, X_i) | \mathcal{T}_i \right] + \text{Var} \left(\mathbb{E}[\hat{\mu}_i | \mathcal{T}_i, X_i] | \mathcal{T}_i \right)$$

Now, again using that $\mathcal{D}_{i,\text{est}} \perp\!\!\!\perp \mathcal{D}_{i,\text{tr}}$,

$$\begin{aligned} \mathbb{E}\left[\text{Var}(\hat{\mu}_i | \mathcal{T}_i, X_i) | \mathcal{T}_i\right] &= \frac{1}{|\mathcal{T}_i|^2} \sum_{t \in \mathcal{T}_i} \mathbb{E}_{\mathcal{D}_{i,\text{est}}} \left[\text{Var} \hat{f}_i^0(X_{it}) | D_{it}=1 \right] \\ &\quad + \frac{2}{|\mathcal{T}_i|^2} \sum_{t \in \mathcal{T}_i} \sum_{\substack{\tau \in \mathcal{T}_i \\ \tau > t}} \mathbb{E}_{\mathcal{D}_{i,\text{est}}} \left[\text{Cov} \hat{f}_i^0(X_{it}, X_{i\tau}) | D_{it}=D_{i\tau}=1 \right] \end{aligned}$$

and

$$\begin{aligned} &\text{Var}\left(\mathbb{E}[\hat{\mu}_i | \mathcal{T}_i, X_i] | \mathcal{T}_i\right) \\ &= \frac{1}{|\mathcal{T}_i|^2} \mathbb{E}\left[\left(\sum_{t \in \mathcal{T}_i} \text{Bias} \hat{f}_i^0(X_{it}) - \mathbb{E}_{\mathcal{D}_{i,\text{est}}}[\text{Bias} \hat{f}_i^0(X_{it})] - \epsilon_{it} - \gamma_{it}\right)^2 | \mathcal{T}_i\right] \\ &= \frac{1}{|\mathcal{T}_i|^2} \sum_{t \in \mathcal{T}_i} \mathbb{E}_{\mathcal{D}_{i,\text{est}}} \left[(\text{Bias} \hat{f}_i^0(X_{it}) - \mathbb{E}_{\mathcal{D}_{i,\text{est}}}[\text{Bias} \hat{f}_i^0(X_{it})])^2 + (\epsilon_{it} + \gamma_{it})^2 \right. \\ &\quad \left. - 2(\epsilon_{it} + \gamma_{it})(\text{Bias} \hat{f}_i^0(X_{it}) - \mathbb{E}_{\mathcal{D}_{i,\text{est}}}[\text{Bias} \hat{f}_i^0(X_{it})]) | D_{it}=1 \right] \\ &\quad + \frac{1}{|\mathcal{T}_i|^2} \sum_{t \in \mathcal{T}_i} \sum_{\substack{\tau \in \mathcal{T}_i \\ \tau > t}} \mathbb{E}_{\mathcal{D}_{i,\text{est}}} \left[(\text{Bias} \hat{f}_i^0(X_{it}) - \mathbb{E}_{\mathcal{D}_{i,\text{est}}}[\text{Bias} \hat{f}_i^0(X_{it})]) \cdot \right. \\ &\quad \quad (\text{Bias} \hat{f}_i^0(X_{i\tau}) - \mathbb{E}_{\mathcal{D}_{i,\text{est}}}[\text{Bias} \hat{f}_i^0(X_{i\tau})]) \\ &\quad \quad - (\epsilon_{it} + \gamma_{it})(\text{Bias} \hat{f}_i^0(X_{i\tau}) - \mathbb{E}_{\mathcal{D}_{i,\text{est}}}[\text{Bias} \hat{f}_i^0(X_{i\tau})]) \\ &\quad \quad - (\epsilon_{i\tau} + \gamma_{i\tau})(\text{Bias} \hat{f}_i^0(X_{it}) - \mathbb{E}_{\mathcal{D}_{i,\text{est}}}[\text{Bias} \hat{f}_i^0(X_{it})]) \\ &\quad \quad \left. + (\epsilon_{it} + \gamma_{it})(\epsilon_{i\tau} + \gamma_{i\tau}) | D_{it}=D_{i\tau}=1 \right] \\ &= \frac{1}{|\mathcal{T}_i|^2} \sum_{t \in \mathcal{T}_i} \mathbb{E}_{\mathcal{D}_{i,\text{est}}} \left[(\text{Bias} \hat{f}_i^0(X_{i\tau}) - \mathbb{E}_{\mathcal{D}_{i,\text{est}}}[\text{Bias} \hat{f}_i^0(X_{i\tau})])^2 + (\epsilon_{it} - \mathbb{E}[\epsilon_{it} | D_{it}=1])^2 \right. \\ &\quad \quad + (\gamma_{it} - \mathbb{E}[\gamma_{it} | D_{it}=1])^2 + 2(\epsilon_{it} - \mathbb{E}[\epsilon_{it} | D_{it}=1])(\gamma_{it} - \mathbb{E}[\gamma_{it} | D_{it}=1]) \\ &\quad \quad + (\mathbb{E}[\epsilon_{it} | D_{it}=1] + \mathbb{E}[\gamma_{it} | D_{it}=1])^2 \\ &\quad \quad + 2(\epsilon_{it} - \mathbb{E}[\epsilon_{it} | D_{it}=1] + \gamma_{it} - \mathbb{E}[\gamma_{it} | D_{it}=1])(\mathbb{E}[\epsilon_{it} | D_{it}=1] + \mathbb{E}[\gamma_{it} | D_{it}=1]) \\ &\quad \quad - 2(\epsilon_{it} - \mathbb{E}[\epsilon_{it} | D_{it}=1] + \gamma_{it} - \mathbb{E}[\gamma_{it} | D_{it}=1])(\text{Bias} \hat{f}_i^0(X_{i\tau}) - \mathbb{E}_{\mathcal{D}_{i,\text{est}}}[\text{Bias} \hat{f}_i^0(X_{i\tau})]) \\ &\quad \quad \left. - 2(\mathbb{E}[\epsilon_{it} | D_{it}=1] + \mathbb{E}[\gamma_{it} | D_{it}=1])(\text{Bias} \hat{f}_i^0(X_{i\tau}) - \mathbb{E}_{\mathcal{D}_{i,\text{est}}}[\text{Bias} \hat{f}_i^0(X_{i\tau})]) | D_{it}=1 \right] \end{aligned}$$

$$\begin{aligned}
& + \frac{1}{|\mathcal{T}_i|^2} \sum_{t \in \mathcal{T}_i} \sum_{\substack{\tau \in \mathcal{T}_i \\ \tau \neq t}} \mathbb{E}_{\mathcal{D}_{i,\text{est}}} \left[(\text{Bias } \hat{f}_i^0(X_{it}) - \mathbb{E}_{\mathcal{D}_{i,\text{est}}} [\text{Bias } \hat{f}_i^0(X_{it})]) (\text{Bias } \hat{f}_i^0(X_{i\tau}) - \mathbb{E}_{\mathcal{D}_{i,\text{est}}} [\text{Bias } \hat{f}_i^0(X_{i\tau})]) \right. \\
& \quad - (\epsilon_{it} - \mathbb{E}[\epsilon_{it} | D_{it} = 1] + \gamma_{it} - \mathbb{E}[\gamma_{it} | D_{it} = 1]) \cdot \\
& \quad \quad (\text{Bias } \hat{f}_i^0(X_{i\tau}) - \mathbb{E}_{\mathcal{D}_{i,\text{est}}} [\text{Bias } \hat{f}_i^0(X_{i\tau})]) \\
& \quad - (\mathbb{E}[\epsilon_{it} | D_{it} = 1] + \mathbb{E}[\gamma_{it} | D_{it} = 1]) (\text{Bias } \hat{f}_i^0(X_{i\tau}) - \mathbb{E}_{\mathcal{D}_{i,\text{est}}} [\text{Bias } \hat{f}_i^0(X_{i\tau})]) \\
& \quad - (\epsilon_{i\tau} - \mathbb{E}[\epsilon_{i\tau} | D_{i\tau} = 1] + \gamma_{i\tau} - \mathbb{E}[\gamma_{i\tau} | D_{i\tau} = 1]) \cdot \\
& \quad \quad (\text{Bias } \hat{f}_i^0(X_{it}) - \mathbb{E}_{\mathcal{D}_{i,\text{est}}} [\text{Bias } \hat{f}_i^0(X_{it})]) \\
& \quad - (\mathbb{E}[\epsilon_{i\tau} | D_{i\tau} = 1] + \mathbb{E}[\gamma_{i\tau} | D_{i\tau} = 1]) (\text{Bias } \hat{f}_i^0(X_{it}) - \mathbb{E}_{\mathcal{D}_{i,\text{est}}} [\text{Bias } \hat{f}_i^0(X_{it})]) \\
& \quad + (\epsilon_{it} - \mathbb{E}[\epsilon_{it} | D_{it} = 1]) (\epsilon_{i\tau} - \mathbb{E}[\epsilon_{i\tau} | D_{i\tau} = 1]) \\
& \quad + (\gamma_{it} - \mathbb{E}[\gamma_{it} | D_{it} = 1]) (\gamma_{i\tau} - \mathbb{E}[\gamma_{i\tau} | D_{i\tau} = 1]) \\
& \quad + (\epsilon_{it} - \mathbb{E}[\epsilon_{it} | D_{it} = 1]) (\gamma_{it} - \mathbb{E}[\gamma_{it} | D_{it} = 1]) \\
& \quad + (\epsilon_{i\tau} - \mathbb{E}[\epsilon_{i\tau} | D_{i\tau} = 1]) (\gamma_{i\tau} - \mathbb{E}[\gamma_{i\tau} | D_{i\tau} = 1]) \\
& \quad + (\epsilon_{it} - \mathbb{E}[\epsilon_{it} | D_{it} = 1] + \gamma_{it} - \mathbb{E}[\gamma_{it} | D_{it} = 1]) \cdot \\
& \quad \quad (\mathbb{E}[\epsilon_{i\tau} | D_{i\tau} = 1] + \mathbb{E}[\gamma_{i\tau} | D_{i\tau} = 1]) \\
& \quad + (\epsilon_{i\tau} - \mathbb{E}[\epsilon_{i\tau} | D_{i\tau} = 1] + \gamma_{i\tau} - \mathbb{E}[\gamma_{i\tau} | D_{i\tau} = 1]) \cdot \\
& \quad \quad (\mathbb{E}[\epsilon_{it} | D_{it} = 1] + \mathbb{E}[\gamma_{it} | D_{it} = 1]) \\
& \quad + (\mathbb{E}[\epsilon_{i\tau} | D_{i\tau} = 1] + \mathbb{E}[\gamma_{i\tau} | D_{i\tau} = 1]) (\mathbb{E}[\epsilon_{it} | D_{it} = 1] + \mathbb{E}[\gamma_{it} | D_{it} = 1]) \\
& \quad \quad \left. | D_{it} = D_{i\tau} = 1 \right]
\end{aligned}$$

Again exploiting symmetry and independence of ϵ_{it} and γ_{it} , we obtain (4.14b) after collecting terms. \square

Bibliography

- Alberto Abadie. Semiparametric difference-in-differences estimators. *The Review of Economic Studies*, 72(1):1–19, 2005.
- Alberto Abadie, Alexis Diamond, and Jens Hainmueller. Synthetic control methods for comparative case studies: Estimating the effect of california’s tobacco control program. *Journal of the American Statistical Association*, 105(490):493–505, 2010.
- Alberto Abadie, Susan Athey, Guido W. Imbens, and Jeffrey M. Wooldridge. Finite population causal standard errors. Working Paper 20325, National Bureau of Economic Research, July 2014.
- Mohamed H. Albadi and Ehab F. El-Saadany. A summary of demand response in electricity markets. *Electric Power Systems Research*, 78(11):1989–1996, 2008.
- Barbara R. Alexander. Dynamic pricing? Not so fast! A residential consumer perspective. *Electricity Journal*, 23(6):39–49, 2010.
- Hunt Allcott. Rethinking real-time electricity pricing. *Resource and Energy Economics*, 33(4):820 – 842, 2011. Special section: Sustainable Resource Use and Economic Dynamics.
- Torben G. Andersen, Tim Bollerslev, Francis X. Diebold, and Paul Labys. Modeling and forecasting realized volatility. *Econometrica*, 71(2):579–625, 2003.
- Joshua D. Angrist and Jörn-Steffen Pischke. *Mostly Harmless Econometrics: An Empiricist’s Companion*. Princeton University Press, 2008.
- Kelman Anthony and Francesco Borrelli. Bilinear model predictive control of a HVAC system using sequential quadratic programming. In *Proceedings of the 18th IFAC World Congress*, pages 9869–9874, August 2011.
- Orley Ashenfelter and David Card. Using the longitudinal structure of earnings to estimate the effect of training programs. *The Review of Economics and Statistics*, 67(4):648–660, 1985.
- Anil Aswani, Neal Master, Jay Taneja, Andrew Krioukov, David Culler, and Claire J. Tomlin. Energy-efficient building HVAC control using hybrid system LBMPC. In *Proceedings of the*

- 4th IFAC Nonlinear Model Predictive Control Conference*, pages 496–501, Leeuwenhorst, NL, 2012a.
- Anil Aswani, Neal Master, Jay Taneja, Virginia Smith, Andrew Krioukov, David Culler, and Claire J. Tomlin. Identifying models of HVAC systems using semiparametric regression. In *American Control Conference (ACC)*, pages 3675–3680, 2012b.
- Susan Athey. Machine learning and causal inference for policy evaluation. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 5–6, Sydney, 2015.
- Susan Athey and Guido W. Imbens. Recursive partitioning for heterogeneous causal effects. *ArXiv e-prints*, April 2015.
- Maximilian Balandat, Frauke Oldewurtel, Mo Chen, and Claire J. Tomlin. Contract design for frequency regulation by aggregations of commercial buildings. In *52nd Annual Allerton Conference on Communication, Control, and Computing*, pages 38–45, 2014.
- Maximilian Balandat, Clay Campaigne, and Lillian Ratliff. The pyDR python package, September 2016a. <https://github.com/Balandat/pyDR>.
- Maximilian Balandat, Clay Campaigne, and Lillian Ratliff. Data for pyDR simulations, August 2016b. https://www.ocf.berkeley.edu/~balandat/pyDR_data.zip.
- Charles M. Beach and James G. MacKinnon. A maximum likelihood procedure for regression with autocorrelated errors. *Econometrica*, 46(1):51–58, 1978.
- Alexandre Belloni, Daniel Chen, Victor Chernozhukov, and Christian Hansen. Sparse models and methods for optimal instruments with an application to eminent domain. *Econometrica*, 80(6):2369–2429, 2012.
- Roland Bénabou and Jean Tirole. Intrinsic and extrinsic motivation. *The Review of Economic Studies*, 70(3):489–520, July 2003.
- Roland Bénabou and Jean Tirole. Incentives and prosocial behavior. *American Economic Review*, 96(5):1652–1678, December 2006.
- Marianne Bertrand, Esther Duflo, and Sendhil Mullainathan. How much should we trust differences-in-differences estimates? *The Quarterly Journal of Economics*, 119(1):249–275, 2004.
- Patrick Billingsley. *Probability and Measure*. Wiley, 3rd edition, 1995.
- Bryan Bollinger and Wesley R. Hartmann. Welfare effects of home automation technology with dynamic pricing. Working Paper, May 2016.

- Judson Boomhower and Lucas W. Davis. Do energy efficiency investments deliver at the right time? Working Paper 271, Energy Institute at Haas, UC Berkeley, 2016.
- Severin Borenstein. The trouble with electricity markets: Understanding california's restructuring disaster. *Journal of Economic Perspectives*, 16(1):191–211, March 2002.
- Severin Borenstein. Time-varying retail electricity prices: Theory and practice. In *Electricity Deregulation: Choices and Challenges*. University of Chicago Press, 2005.
- Severin Borenstein. Peak-Time Rebates: Money for Nothing?, 2014. <http://www.greentechmedia.com/articles/read/Peak-Time-Rebates-Money-for-Nothing>.
- Severin Borenstein. The economics of fixed cost recovery by utilities. Working Paper 272, Energy Institute at Haas, UC Berkeley, 2016.
- Severin Borenstein and Stephen P. Holland. On the efficiency of competitive electricity markets with time-invariant retail prices. *The Rand Journal of Economics*, 36(3):469–493, 2005.
- Severin Borenstein, Michael Jaske, and Arthur Rosenfeld. Dynamic Pricing, Advanced Metering and Demand Response in Electricity Markets. Working Paper 105, Center for the Study of Energy Markets, 2002.
- Robert L. Borlick. Pricing Negawatts. *Public Utilities Fortnightly*, 14:148, 2010.
- Robert L. Borlick, Joseph Bowring, James Bushnell, Paul A. Centolella, Hung-Po Chao, Ahmad Faruqui, Michael Giberson, Dinos Gonatas, Scott Harvey, Benjamin F. Hobbs, William W. Hogan, Joseph P. Kalt, Robert J. Michaels, Shmuel S Oren, David B. Patton, Craig Pirrong, Susan Liese Pope, Larry E. Ruff, Richard Schmalensee, Roy J. Shanker, Vernon L. Smith, and Richard D. Tabors. Brief of Robert L. Borlick, Joseph Bowring, James Bushnell, and 18 Other Leading Economists As Amici Curiae In Support Of Petitioners, 2012.
- Richard C. Bradley. Basic properties of strong mixing conditions. A survey and some open questionsa survey and some open questions. *Probability Surveys*, 2:107–144, 2005.
- Leo Breiman. Random forests. *Machine Learning*, 45(1):5–32, 2001.
- Kay H. Brodersen, Fabian Gallusser, Jim Koehler, Nicolas Remy, and Steven L. Scott. Inferring causal impact using bayesian structural time-series models. *The Annals of Applied Statistics*, 9(1):247–274, March 2015.
- David B. Brown, James E. Smith, and Peng Sun. Information Relaxations and Duality in Stochastic Dynamic Programs. *Operations Research*, 58(4-Part-1):785–801, 2010.
- M. Kate Bundorf, Jonathan Levin, and Neale Mahoney. Pricing and welfare in health plan choice. *American Economic Review*, 102(7):3214–48, December 2012.

- James Bushnell, Benjamin F. Hobbs, and Frank A. Wolak. When it comes to Demand Response, is FERC its own worst enemy? *Electricity Journal*, 22(8):9–18, 2009.
- California Department of Water Resources. California Irrigation Management System, 2015. <http://www.cimis.water.ca.gov/>.
- California Independent System Operator Corporation. Load granularity refinements issue paper and pricing analysis study. Presentation, August 2013. https://www.aiso.com/Documents/Presentation-LoadGranularityRefinementsIssuePaperAug21_2013.pdf.
- California Independent System Operator Corporation. Fifth replacement FERC electric tariff, July 2015a. https://www.aiso.com/Documents/ConformedTariff_Jun12_2015.pdf.
- California Independent System Operator Corporation. *Interface specification for OASIS*, 4.2.6 edition, Mar 2015b. http://www.aiso.com/Documents/OASIS-InterfaceSpecification_v4_2_6Clean_Independent2015Release.pdf.
- California Independent System Operator Corporation. Aliso canyon gas-electric coordination, April 2016. https://www.aiso.com/Documents/DraftFinalProposal_AlisoCanyonGas_ElectricCoordination.pdf.
- California Public Utilities Commission. How high is california’s electricity demand, and where does the power come from? <http://www.cpuc.ca.gov/cfaqs/howhighiscaliforniaselectricitydemandandwheredoesthepowercomefrom.htm>.
- California Public Utilities Commission. Decision 16-06-008, Jun 2016.
- California State Senate. Senate Bill No. 695, October 2011.
- California State Senate. Senate Bill No. 1122, September 2012.
- Duncan Callaway, Meredith Fowlie, and Gavin McCormick. Location, location, location: The variable value of renewable energy and demand-side efficiency resources. Working Paper 264, Energy Institute at Haas, UC Berkeley, 2015.
- Clay Campaigne, Maximilian Balandat, and Lillian Ratliff. Welfare effects of dynamic electricity pricing. Working Paper, October 2016.
- Andrew Caplin. Measuring and modeling attention. *Annual Review of Economics*, 8(1), 2016.
- Douglas W. Caves, Laurits R. Christensen, and Joseph A. Herriges. Consistency of residential customer response in time-of-use electricity pricing experiments. *Journal of Econometrics*, 26(1):179–203, 1984.

- Hung-Po Chao and Mario DePillis. Incentive effects of paying demand response in wholesale electricity markets. *Journal of Regulatory Economics*, 43(3):265–283, 2013.
- Zhi Chen, Lei Wu, and Yong Fu. Real-time price-based demand response management for residential appliances via stochastic optimization and robust optimization. *IEEE Transactions on Smart Grid*, 3(4):1822–1831, December 2012.
- Raj Chetty, John N. Friedman, and Jonah E. Rockoff. Measuring the impacts of teachers I: Evaluating bias in teacher value-added estimates. *American Economic Review*, 104(9):2593–2632, September 2014.
- David Cochrane and Guy H. Orcutt. Application of least squares regression to relationships containing auto-correlated error terms. *Journal of the American Statistical Association*, 44(245):32–61, 1949.
- Benoît Colson, Patrice Marcotte, and Gilles Savard. An overview of bilevel optimization. 153(1):235–256, 2007.
- Katie Coughlin, Mary Ann Piette, Charles Goldman, and Sila Kiliccote. Estimating Demand Response load impacts: Evaluation of baseline load models for non-residential buildings in California. Technical Report 63728, Lawrence Berkeley National Laboratory, January 2008.
- Michael A. Crew, Chitru S. Fernando, and Paul R. Kleindorfer. The theory of peak-load pricing: A survey. *Journal of Regulatory Economics*, 8(3):215–248, 1995.
- Richard K. Crump, V. Joseph Hotz, Guido W. Imbens, and Oscar A. Mitnik. Nonparametric tests for treatment effect heterogeneity. *Review of Economics and Statistics*, 90(3):389–405, June 2008.
- Joseph A. Cruz and David S. Wishart. Applications of machine learning in cancer prediction and prognosis. *Cancer Informatics*, 2, 2006.
- Eric Cutter, Lakshmi Alagappan, and Snuller Price. Impact of market rules on energy storage economics. In *32nd IAAEE International Conference*, volume 24, June 2009.
- Boris Defourny, Damien Ernst, and Louis Wehenkel. Multistage stochastic programming. In *Decision Theory Models for Applications in Artificial Intelligence*, pages 97–143. IGI Global, 2011.
- Rajeev H. Dehejia and Sadek Wahba. Causal effects in nonexperimental studies: Reevaluating the evaluation of training programs. *Journal of the American Statistical Association*, 94(448):1053–1062, 1999.
- Kjell Doksum. Empirical probability plots and statistical inference for nonlinear models in the two-sample case. *The Annals of Statistics*, 2(2):267–277, 1974.

- Qifen Dong, Li Yu, Wen-Zhan Song, Lang Tong, and Shaojie Tang. Distributed demand and response algorithm for optimizing social-welfare in smart grid. In *26th IEEE International Parallel Distributed Processing Symposium (IPDPS)*, pages 1228–1239, May 2012.
- Esther Duflo and Emmanuel Saez. The role of information and social interactions in retirement plan decisions: Evidence from a randomized experiment. *The Quarterly Journal of Economics*, 118(3):815–842, 2003.
- Bradley Efron. *The Jackknife, the Bootstrap and Other Resampling Plans*. CBMS-NSF Regional Conference Series in Applied Mathematics. Society for Industrial and Applied Mathematics, 1982.
- Bradley Efron. Better bootstrap confidence intervals. *Journal of the American Statistical Association*, 82(397):171–185, 1987.
- Bradley Efron and Robert Tibshirani. Bootstrap methods for standard errors, confidence intervals, and other measures of statistical accuracy. *Statistical Science*, 1(1):54–75, 1986.
- Bradley Efron and Robert Tibshirani. *An Introduction to the Bootstrap*. Chapman & Hall/CRC Monographs on Statistics & Applied Probability. Taylor & Francis, 1994.
- Liran Einav and Jonathan D. Levin. The data revolution and economic analysis. Working Paper 19035, National Bureau of Economic Research, May 2013.
- Liran Einav, Amy Finkelstein, Stephen P. Ryan, Paul Schrimpf, and Mark R. Cullen. Selection on moral hazard in health insurance. *American Economic Review*, 103(1):178–219, February 2013.
- Gary Fagilde. HVAC by the Numbers: How to Save by Cutting Waste. blog post. <https://www.pge.com/en/mybusiness/save/smbblog/article/HVACnumbers.page>.
- Ahmad Faruqui. The global movement toward cost-reflective tariffs (presentation), 2015.
- Ahmad Faruqui and J.Robert Malko. The residential demand for electricity by time-of-use: A survey of twelve experiments with peak load pricing. *Energy*, 8(10):781–795, 1983.
- Ahmad Faruqui and Sanem Sergici. Household response to dynamic pricing of electricity: a survey of 15 experiments. *Journal of Regulatory Economics*, 38(2):193–225, 2010.
- Ahmad Faruqui, Sanem Sergici, and Ahmed Sharif. The impact of informational feedback on energy consumption – a survey of the experimental evidence. *Energy*, 35(4):1598 – 1608, 2010.
- Ahmad Faruqui, Ryan Hledik, and Jennifer Palmer. Time-Varying and Dynamic Rate Design. In *Global Power Best Practice Series*, pages 1–52. Regulatory Assistance Project, July 2012.

- FERC. Order No. 755: Frequency Regulation Compensation in the Organized Wholesale Power Markets, October 2011a.
- FERC. Order No. 745, Demand Response Compensation in Organized Wholesale Energy Markets, March 2011b.
- Paul J. Ferraro and Michael K. Price. Using nonpecuniary strategies to influence behavior: Evidence from a large-scale field experiment. *Review of Economics and Statistics*, 95(1): 64–73, March 2013.
- Sergio Firpo. Efficient semiparametric estimation of quantile treatment effects. *Econometrica*, 75(1):259–276, 2007.
- Jaime Gannon, Donald Brooks, Lily Chow, and Simone Brant. The 2013 – 2014 Resource Adequacy Report. Technical report, The California Public Utilities Commission, August 2015.
- Ravi Gondhalekar, Frauke Oldewurtel, and Colin N. Jones. Least-restrictive robust periodic model predictive control applied to room temperature regulation. *Automatica*, 49(9): 2760–2766, 2013.
- Gurobi Optimization. Gurobi Optimizer Reference Manual, 2016.
- Isabelle Guyon, Constantin Aliferis, Greg Cooper, André Elisseeff, Jean-Philippe Pellet, Peter Spirtes, and Alexander Statnikov. Causation and prediction challenge at WCCI 2008. In *JMLR Workshop and Conference Proceedings*, volume 3, Hong Kong, June 2008a.
- Isabelle Guyon, Dominik Janzing, and Bernhard Schölkopf. Causality: Objectives and assessment (NIPS 2008 workshop). In *JMLR Workshop and Conference Proceedings*, Whistler, Canada, December 2008b.
- Douglas A. Halamay, Ted K. A. Brekken, Asher Simmons, and Shaun McArthur. Reserve requirement impacts of large-scale integration of wind, solar, and ocean wave power generation. *Sustainable Energy, IEEE Transactions on*, 2(3):321–328, July 2011.
- Christian B. Hansen. Generalized least squares inference in panel and multilevel models with serial correlation and fixed effects. *Journal of Econometrics*, 140(2):670 – 694, 2007.
- He Hao, Anupama Kowli, Yashen Lin, Prabir Barooah, and Sean Meyn. Ancillary service for the grid via control of commercial building HVAC systems. In *American Control Conference (ACC)*, pages 467–472, June 2013.
- Arnold C. Harberger. The Measurement of Waste. *The American Economic Review*, 54(3): 58–76, 1964.

- Matthew Harding and Carlos Lamarche. Empowering consumers through data and smart technology: Experimental evidence on the consequences of time-of-use electricity pricing policies. *Journal of Policy Analysis and Management*, 35(4):906–931, 2016.
- Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *The Elements of Statistical Learning*. Springer Series in Statistics. Springer, 2009.
- Tatsuo Hatta and John Haltiwanger. Tax reform and strong substitutes. *International Economic Review*, 27(2):303–315, 1986.
- James J. Heckman and V. Joseph Hotz. Choosing among alternative nonexperimental methods for estimating the impact of social programs: The case of manpower training. *Journal of the American Statistical Association*, 84(408):862–874, 1989.
- Gregor P. Henze, Doreen E. Kalz, Simeng Liu, and Clemens Felsmann. Experimental analysis of model-based predictive optimal control for active and passive building thermal storage inventory. *HVAC&R Research*, 11(2):189–213, 2005.
- Karen Herter and Seth Wayland. Residential response to critical-peak pricing of electricity: California evidence. *Energy*, 35(4):1561–1567, 2010.
- Keisuke Hirano and Guido W. Imbens. The propensity score with continuous treatments. In *Applied Bayesian Modeling and Causal Inference from Incomplete-Data Perspectives*, pages 73–84. Wiley, 2004.
- Keisuke Hirano, Guido W. Imbens, Donald B. Rubin, and Xiao-Hua Zhou. Assessing the effect of an influenza vaccine in an encouragement design. *Biostatistics*, 1(1):69–88, 2000.
- William W Hogan. Demand response pricing in organized wholesale markets. Comments on demand response compensation in organized wholesale energy markets, ISO/RTO Council, 2010.
- William W Hogan. Time-of-Use Rates and Real-Time Prices. Working paper, John F. Kennedy School of Government, Harvard University, 2014.
- Paul W. Holland. Statistics and causal inference. *Journal of the American Statistical Association*, 81(396):945–960, 1986.
- Katherine Hsia-Kiung, Emily Reyna, and Timothy O’Connor. Carbon Market California: a comprehensive analysis of the Golden State’s cap-and-trade program: Year One: 2012–2013. Technical report, Environmental Defense Fund, 2014.
- Guido W. Imbens. Nonparametric estimation of average treatment effects under exogeneity: A review. *Review of Economics and Statistics*, 86(1):4–29, October 2004.

- Interagency Working Group on Social Cost of Carbon. Technical Update of the Social Cost of Carbon for Regulatory Impact Analysis Under Executive Order 12866. Technical Report May, Interagency Working Group on Social Cost of Carbon, US Government, 2013.
- Koichiro Ito. Do consumers respond to marginal or average price? evidence from nonlinear electricity pricing. *American Economic Review*, 104(2):537–563, February 2014.
- Koichiro Ito, Takanori Ida, and Makoto Tanaka. The persistence of moral suasion and economic incentives: Field experimental evidence from energy demand. Working Paper 20910, National Bureau of Economic Research, January 2015.
- Mark R. Jacobsen, Christopher R. Knittel, James M. Sallee, and Arthur A. van Benthem. Sufficient Statistics for Imperfect Externality-Correcting Policies. Working Paper 22063, National Bureau of Economic Research, March 2016.
- Katrina Jessoe and David Rapson. Knowledge is (less) power: Experimental evidence from residential energy use. *American Economic Review*, 104(4):1417–38, April 2014.
- Katrina K. Jessoe, Douglas L. Miller, and David S. Rapson. Can high-frequency data and non-experimental research designs recover causal effects? Validation using an electricity usage experiment. Working Paper, May 2015.
- Paul Joskow and Jean Tirole. Retail Electricity Competition. *RAND Journal of Economics*, 37(4):799–815, 2006.
- Roger Koenker and Gilbert Bassett. Regression quantiles. *Econometrica*, 46(1):33–50, 1978.
- Hans R. Kunsch. The jackknife and the bootstrap for general stationary observations. *The Annals of Statistics*, 17(3):1217–1241, 1989.
- Jean-Jacques Laffont and David Martimort. *The Theory of Incentives: The Principal-Agent Model*. Princeton University Press, 2002.
- Robert J. LaLonde. Evaluating the econometric evaluations of training programs with experimental data. *The American Economic Review*, 76(4):604–620, 1986.
- LCG Consulting. energyonline. website, 2016. <http://www.energyonline.com>.
- Brian K. Lee, Justin Lessler, and Elizabeth A. Stuart. Improving propensity score weighting using machine learning. *Statistics in medicine*, 29(3):337–346, 2010.
- Lihong Li, Wei Chu, John Langford, and Xuanhui Wang. Unbiased offline evaluation of contextual-bandit-based news article recommendation algorithms. In *Proceedings of the Fourth ACM International Conference on Web Search and Data Mining*, pages 297–306, Hong Kong, 2011.

- Lihong Li, Wei Chu, John Langford, Taesup Moon, and Xuanhui Wang. An unbiased offline evaluation of contextual bandit algorithms with generalized linear models. In *JMLR: Workshop and Conference Proceedings*, volume 26, 2012.
- Yudong Ma, Francesco Borrelli, Brandon Hancey, Brian Coffey, Sorin Bengea, and Philip Haves. Model predictive control for the operation of building cooling systems. *IEEE Transactions on Control Systems Technology*, 20(3):796–803, May 2012.
- Mehdi Maasoumy, Catherine Rosenberg, Alberto Sangiovanni-Vincentelli, and Duncan S. Callaway. Model Predictive Control Approach to Online Computation of Demand-Side Flexibility of Commercial Buildings HVAC Systems for Supply Following. In *American Control Conference (ACC)*, Portland, 2014.
- Yuri V. Makarov, Clyde Loutan, Jian Ma, and Phillip de Mello. Operational impacts of wind generation on california power systems. *Power Systems, IEEE Transactions on*, 24(2):1039–1050, May 2009.
- Christopher D Manning and Hinrich Schütze. *Foundations of statistical natural language processing*. MIT Press, 1999.
- Charles F. Manski. *Analog estimation methods in econometrics*. Chapman & Hall, 1988.
- Ali Mesbah. Stochastic model predictive control: An overview and perspectives for future research. *IEEE Control Systems Magazine*, 2016. In Press.
- Janine Migden-Ostrander, Betty Watson, Dave Lamont, and Richard Sedano. Decoupling case studies: Revenue regulation implementation in six states. Technical report, The Regulatory Assistance Project, 2014.
- Piotr Mirowski, Sining Chen, Tin Kam Ho, and Chun-Nam Yu. Demand Forecasting in Smart Grids. *Bell Labs Technical Journal*, 18(4), 2014.
- Jeremy Neubauer and Mike Simpson. Deployment of behind-the-meter energy storage for demand charge reduction. Technical Report NREL/TP-5400-63162, National Renewable Energy Laboratory, January 2015.
- Whitney K. Newey. Semiparametric efficiency bounds. *Journal of Applied Econometrics*, 5(2):99–135, 1990.
- Stephen Nickell. Biases in dynamic models with fixed effects. *Econometrica*, 49(6):1417–1426, 1981.
- Frauke Oldewurtel. *Stochastic Model Predictive Control for Energy Efficient Building Climate Control*. PhD thesis, ETH Zurich, 2011.

- Frauke Oldewurtel, Alessandra Parisio, Colin N. Jones, Dimitrios Gyalistras, Markus Gw-
erder, Vanessa Stauch, Beat Lehmann, and Manfred Morari. Use of model predictive
control and weather forecasts for energy efficient building climate control. *Energy and
Buildings*, 45(0):15 – 27, 2012.
- Frauke Oldewurtel, Theodor Borsche, Matthias Bucher, Philipp Fortenbacher, Ma-
rina González Vayá, Tobias Haring, Johanna L. Mathieu, Olivier Mégel, Evangelos Vrettos,
and Göran Andersson. A framework for and assessment of demand response and energy
storage in power systems. In *Proceedings of IREP Symposium: Bulk Power System Dy-
namics and Control - IX Optimization, Security and Control of the Emerging Power Grid*,
2013.
- Ontario Energy Board. Ontario Energy Board smart price pilot final report, July 2007.
- Pacific Gas and Electric Company. Non-residential rates, a.
- Pacific Gas and Electric Company. PG&E Tariff Book, b. [https://www.pge.com/tariffs/
ERS.SHTML](https://www.pge.com/tariffs/ERS.SHTML).
- Pacific Gas and Electric Company. Peak Day Pricing Fact Sheet: Dual Par-
ticipation for Peak Day Pricing, 2011. [http://www.pge.com/includes/docs/
pdfs/mybusiness/energysavingsrebates/demandresponse/peakdaypricing/
PDPDualParticipationCustFactSheet.pdf](http://www.pge.com/includes/docs/pdfs/mybusiness/energysavingsrebates/demandresponse/peakdaypricing/PDPDualParticipationCustFactSheet.pdf).
- Pacific Gas and Electric Company. Understanding energy use and prices, July 2016a. [http:
//www.pge.com/tariffs/energy_use_prices.shtml](http://www.pge.com/tariffs/energy_use_prices.shtml).
- Pacific Gas and Electric Company. Electric Rates for Commercial/General Service (A-1, A-
6, A-10, E-19), January 2016b. [https://www.pge.com/tariffs/Comm1_160101-160229.
xls](https://www.pge.com/tariffs/Comm1_160101-160229.xls).
- Pacific Gas and Electric Company. SmartRate FAQ, 2016c. [https://www.pge.
com/en_US/residential/rate-plans/rate-plan-options/smart-rate-add-on/
discover-smart-rate/discover-smart-rate.page?](https://www.pge.com/en_US/residential/rate-plans/rate-plan-options/smart-rate-add-on/discover-smart-rate/discover-smart-rate.page?)
- Pacific Gas and Electric Company. Renewable feed-in tariffs, 2016d. [https://www.pge.com/
en_US/for-our-business-partners/floating-pages/renewable-feed-in-tariffs/
renewable-feed-in-tariffs.page](https://www.pge.com/en_US/for-our-business-partners/floating-pages/renewable-feed-in-tariffs/renewable-feed-in-tariffs.page).
- Pacific Gas and Electric Company. SmartDay history. website, 2016e. [https://www.pge.com/en_US/residential/rate-plans/rate-plan-options/
smart-rate-add-on/smart-day-history/smart-day-history.page](https://www.pge.com/en_US/residential/rate-plans/rate-plan-options/smart-rate-add-on/smart-day-history/smart-day-history.page).
- Peter Palensky and Dietmar Dietrich. Demand side management: Demand response, intel-
ligent energy systems, and smart loads. *IEEE Transactions on Industrial Informatics*, 7
(3):381–388, August 2011.

- Judea Pearl. *Causality*. Cambridge University Press, 2009.
- Johannes P. Pfeifenberger, Kathleen Spees, and Samuel A. Newell. Resource Adequacy in California. Technical Report October, The Brattle Group, 2012.
- Warren B. Powell. *Approximate Dynamic Programming: Solving the curses of dimensionality*. John Wiley & Sons, 2nd edition, 2011.
- pvlib. pvlib python library, 2016. <https://github.com/pvlib/pvlib-python>.
- Benjamin Recht, Maryam Fazel, and Pablo A. Parrilo. Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM Review*, 52(3):471–501, 2010.
- Peter C. Reiss and Matthew W. White. Household electricity demand, revisited. *The Review of Economic Studies*, 72(3):853–883, 2005.
- Francesco Ricci, Lior Rokach, and Bracha Shapira. *Introduction to recommender systems handbook*. Springer, 2011.
- Paul R. Rosenbaum and Donald B. Rubin. The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70(1):41–55, 1983.
- Paul R. Rosenbaum and Donald B. Rubin. Reducing bias in observational studies using subclassification on the propensity score. *Journal of the American Statistical Association*, 79(387):516–524, 1984.
- Donald B. Rubin. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, 66(5):688–701, 1974.
- Donald B. Rubin. Randomization analysis of experimental data: The fisher randomization test comment. *Journal of the American Statistical Association*, 75(371):591–593, 1980.
- Pedram Samadi, Hamed Mohsenian-Rad, Robert Schober, and Vincent W. S. Wong. Advanced demand side management for the future smart grid using mechanism design. *IEEE Transactions on Smart Grid*, 3(3):1170–1180, September 2012.
- Roy J. Shanker. Comments of Roy J. Shanker Ph.D. regarding the energy market demand response nopr on behalf of the PJM power providers (P3). Technical Session Docket No. RM10-17-000, Federal Energy Regulatory Commission, 2010.
- Wenbo Shi and Vincent W. S. Wong. Real-time vehicle-to-grid control algorithm under price uncertainty. In *IEEE International Conference on Smart Grid Communications (SmartGridComm)*, pages 261–266, October 2011.
- Herbert A. Simon. *Designing Organizations for an Information-Rich World*, pages 37–72. Computers, communications, and the public interest. 1971.

- Kanwardeep Singh, Narayana P. Padhy, and Jaydev Sharma. Influence of price responsive demand shifting bidding on congestion and Imp in pool-based day-ahead electricity markets. *IEEE Transactions on Power Systems*, 26(2):886–896, May 2011.
- Joel Slemrod. Optimal taxation and optimal tax systems. *Journal of Economic Perspectives*, 4(1):157–178, March 1990.
- Jeffrey A. Smith and Petra E. Todd. Does matching overcome LaLonde’s critique of nonexperimental estimators? *Journal of Econometrics*, 125(1-2):305–353, 2005.
- Alex J. Smola and Bernhard Schölkopf. A tutorial on support vector regression. *Statistics and Computing*, 14(3):199–222, 2004.
- Alexander L. Strehl, John Langford, Lihong Li, and Sham Kakade. Learning from logged implicit exploration data. In *Advances in Neural Information Processing Systems 23 (NIPS)*, pages 2217–2225, Vancouver, December 2010.
- Richard Szeliski. *Computer vision: algorithms and applications*. Springer Science & Business Media, 2010.
- Inc. Tesla Motors. Tesla powerwall technical specifications. website, Dec 2016. <https://www.tesla.com/powerwall>.
- John N. Tsitsiklis and Yunjian Xu. Pricing of fluctuations in electricity markets. *European Journal of Operational Research*, 246(1):199–208, 2015.
- Kai M. Tsui and Shing C. Chan. Demand response optimization for smart home scheduling under real-time pricing. *IEEE Transactions on Smart Grid*, 3(4):1812–1821, December 2012.
- United States Census Bureau. 2015 U.S. Gazetteer Files, 2015. <http://www.census.gov/geo/maps-data/data/gazetteer2015.html>.
- U.S. Department of Energy. Buildings Energy Data Book: Building Sector Energy Consumption, March 2012. <http://buildingsdatabook.eren.doe.gov/TableView.aspx?table=1.1.9>.
- U.S. Energy Information Administration. State Energy Data System. <http://www.eia.gov/state/seds/#>.
- Pedro A. Valdés-Sosa, Jose Miguel Bornot-Sánchez, Mayrim Vega-Hernández, Lester Melie-García, Agustin Lage-Castellanos, and Erick Canales-Rodríguez. *Granger Causality on Spatial Manifolds: Applications to Neuroimaging*, pages 461–491. Handbook of Time Series Analysis. Wiley, 2006.

- John S. Vardakas, Nizar Zorba, and Christos V. Verikoukis. A survey on demand response programs in smart grids: Pricing methods and optimization algorithms. *IEEE Communications Surveys Tutorials*, 17(1):152–178, 2015.
- Hal R. Varian. Big data: New tricks for econometrics. *Journal of Economic Perspectives*, 28(2):3–28, May 2014.
- Evangelos Vrettos, Frauke Oldewurtel, Fengtian Zhu, and Göran Andersson. Robust Provision of Frequency Reserves by Office Building Aggregations. In *Proceedings of the 19th IFAC World Congress*, August 2014.
- Stefan Wager and Susan Athey. Estimation and inference of heterogeneous treatment effects using random forests. *ArXiv e-prints*, October 2015.
- H. Paul Williams. *Logic and Integer Programming*. Operations Research & Management Science. Springer, 2009.
- Frank A. Wolak. Residential customer response to real-time pricing: The anaheim critical peak pricing experiment. *CSEM Working Paper*, 2006.
- Frank A Wolak, James Bushnell, and Benjamin F Hobbs. The California ISO’s Proxy Demand Resource (PDR) Proposal. Technical report, 2009.
- Arthur Wright. Maximum Demand Indicator, 1902. US Patent.
- Lin Xu. Demand response net benefits test. Technical report, California Independent System Operator Corporation, July 2011.
- Yanbo Xu, Yanxun Xu, and Suchi Saria. A bayesian nonparametric approach for estimating individualized treatment-response curves. *ArXiv e-prints*, 2016.
- Peng Yang, Gongguo Tang, and Arye Nehorai. A game-theoretic approach for optimal time-of-use electricity pricing. *IEEE Transactions on Power Systems*, 28(2):884–892, May 2013.
- Victor M. Zavala. Real-time optimization strategies for building systems. *Industrial & Engineering Chemistry Research*, 52(9):3137–3150, 2013.
- Datong P. Zhou, Maximilian Balandat, and Claire J. Tomlin. A bayesian perspective on residential demand response using smart meter data. In *Proceedings of the 54th Annual Allerton Conference on Communication, Control, and Computing*, October 2016a.
- Datong P. Zhou, Maximilian Balandat, and Claire J. Tomlin. Residential demand response targeting using machine learning with observational data. In *Proceedings of the 55th IEEE Conference on Decision and Control*, Las Vegas, USA, December 2016b.

Marco Zugno, Juan Miguel Morales, Pierre Pinson, and Henrik Madsen. A bilevel model for electricity retailers' participation in a demand response market environment. *Energy Economics*, 36:182–197, March 2013.