Pricing and Incentive Design in Applications of Green Technology Subsidies and Revenue Management

by

Ruben Lobel

B.Sc., Electrical Engineering
Pontifícia Universidade Católica do Rio de Janeiro (2005)

Submitted to the Sloan School of Management
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Operations Research

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2012

Author ....
Sloan School of Management
May 18, 2012

Certified by ....
Georgia Perakis
William F. Pounds Professor of Management
Professor of Operations Research and Operations Management
Thesis Supervisor

Accepted by ....
Dimitris Bertsimas
Boeing Professor of Operations Research
Co-director, Operations Research Center

Pricing and Incentive Design in Applications of Green Technology Subsidies and Revenue Management

by

Ruben Lobel

Submitted to the Sloan School of Management
on May 18, 2012, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Operations Research

Abstract

This thesis addresses three issues faced by firms and policy-makers when deciding how to price products and properly incentivize consumers. In the first part of the thesis, we focus on a firm attempting to dynamically adjust prices to maximize profits when facing uncertain demand, as for example airlines selling flights or hotels booking rooms. In particular, we develop a robust sampling-based optimization framework that minimizes the worst-case regret and dynamically adjusts the price according to the realization of demand. We propose a tractable optimization model that uses direct demand samples, where the confidence level of this solution can be obtained from the number of samples used. We further demonstrate the applicability of this approach with a series of numerical experiments and a case study using airline ticketing data.

In the second part of the thesis, we propose a model for the adoption of solar photovoltaic technology by residential consumers. Using this model, we develop a framework for policy makers to find optimal subsidy levels in order to achieve a desired adoption target. The technology adoption process follows a discrete choice model, which is reinforced by network effects such as information spread and learning-by-doing. We validate the model through an empirical study of the German solar market, where we estimate the model parameters, generate adoption forecasts and demonstrate how to solve the policy design problem. We use this framework to show that the current policies in Germany could be improved by higher subsidies in the near future and a faster phase-out of the subsidy program.

In the third part of the thesis, we model the interaction between a government and an industry player in a two-period game setting under uncertain demand. We show how the timing of decisions will affect the production levels and the cost of the subsidy program. In particular, we show that when the government commits to a fixed policy, it signals to the supplier to produce more in the beginning of the horizon. Consequently, a flexible policy is on average more expensive for the government than a committed policy.

Thesis Supervisor: Georgia Perakis
Title: William F. Pounds Professor of Management
Professor of Operations Research and Operations Management

Acknowledgments

As I think about all the people that helped me and made my life so great through this journey at MIT, I cannot imagine how I can do justice to their invaluable contribution. I will try my best to express how grateful I am for all their support.

First and foremost, I must thank my advisor, Georgia Perakis, for guiding me through this process. I am very grateful for our countless hours of meetings, where she shared her knowledge with me, patiently pushing me to better myself and getting things done. Perhaps the most unique trait of Georgia is how much she cares about her students, putting them before anything else, going above and beyond for us. For that, I will be eternally grateful.

From the beginning of my PhD program, Georgia has been there every step of the way, supporting and encouraging me. Together we learned and explored new research directions, shared many successes and frustrations. She helped me develop as a researcher. I hope that I can be a mentor to others like she was to me.

I would like to specially thank Steve Graves and Ignacio Perez-Arriaga for participating in my thesis committee. Their comments and feedback have made a great impact on this work, keeping it rigorous and grounded in the problems of the real world.

I am also very thankful to Vivek Farias, for his numerous comments on my research, talk rehearsals and personal career advice. I cannot thank him enough for participating in my academic development.

In the development of the research in this thesis, I am thankful for the feedback and comments from a diverse group of faculty, including Gabriel Bitran, Retsef Levi, Marco Campi, Antonio Conejo, Charles Fine, Marcelo Olivares, and Ariel Pakes.

This thesis also benefited from the research discussions with my brother Ilan Lobel and fellow students Nikos Trichakis, Kostas Bimpikis, Rajiv Menjoge, Eduardo Azevedo, Ozan Candogan, and Yuan Zhong. I am also thankful for the discussions with the students at the ESRG group at MIT.

I am specially thankful for the collaboration of Maxime Cohen, who actively participated in the research presented here in Chapter 4. It has been a great pleasure working with Maxime for the last two years.

I would like to thank all the ORC faculty for providing me with this amazing PhD program, for teaching the courses, attending the talks I presented, discussing research with me and providing career guidance. I would like to specially thank Dimitris Bertsimas for his guidance as the Co-

Director of the ORC, from whom I personally learned a lot. I am also thankful to the ORC staff, Laura Rose and Andrew Carvalho, for all their help and timely support.

I would also like to thank my former colleagues at PSR, among them Mario Veiga Pereira and Sergio Granville. I learned a lot from them and I would have never made it to MIT without their help and encouragement.

I would like to give special thanks to my friends in Cambridge. Their company is the main reason why this journey through MIT has been such a pleasant experience. I am thankful to Pallav Chhaochharia, Rajiv Menjoge, Karima Nigmatulina, Doug Fearing, Juliane Dunkel, Dan Iancu, Nikos Trichakis, Gerry Tsoukalas, Gerasimos Skaltsas, Yannis Simaiakis and Ozan Candogan for being great friends and colleagues at the MIT. To Alex Rikun, Jonathan Kluberg, Gareth Williams, I am also thankful for their joyful distractions needed in times of stress while sharing the office. I am also thankful for having the support of Georgia's research family, including Guillaume Roels, Sun Wei, Phil Keller, Joline Uichanco, and Maxime Cohen, among many others.

For my roommates in Harvard Street, Kostas Bimpikis and Nikolas Pyrgiotis, I am thankful for the great times we've had sharing a house together. Thanks to them and my extended greek family in Boston, I have learned to appreciate good olive oil.

I am immensely grateful to my girlfriend, Ilke Kalcioglu, who cared for me through the hardest part of my PhD. She provided me with all the love and support that I needed throughout this journey. I am very lucky to have her by my side.

Finally, I dedicate this thesis to my family. I could not imagine getting this far without the help of my brother, Ilan, and his wife, Inna. For being a best friend and an inspiration, I am very thankful to my brother. From my parents, Robert and Lucia, I am very thankful for their unconditional love. Everything I am today and all I have achieved, I owe it to my family. To them I dedicate this thesis.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

"How much should a company charge for its product?" This pricing question is critical to managers across many industries and the answer is almost never trivial. In some cases this is an internal concern for the firm. Sometimes the government is the one asking this question, as it tries to promote emerging technologies using price incentives for consumers. The overarching goal of this thesis is to address issues faced by firms and policy-makers when deciding how much customers should pay for a product. We develop answers to this core question of *pricing* in three different applications.

In the first part of the thesis, Chapter 2, we focus on the pricing problem of the firm manager when facing uncertain demand. In this setting, a firm offers multiple products to be sold over a fixed time horizon. We develop an approach to dynamically price these products that combines ideas from data-driven and robust optimization to address the uncertain and dynamic aspects of the problem. Each product sold consumes one or more resources, possibly sharing the same resources among different products. The firm is given a fixed initial inventory of these resources and cannot replenish this inventory during the selling season. We assume there is uncertainty about the demand seen by the firm for each product. The firm wants to determine a robust and dynamic pricing strategy that maximizes revenue over the time horizon. While the traditional robust optimization models are tractable, they give rise to static policies and are often too conservative. At the same time, stochastic optimization models suffer from tractability and data estimation issues. Our approach tries to establish a compromise between these conflicting modeling objectives.

The main contribution of Chapter 2 is the exploration of closed-loop pricing policies for different robust objectives, such as MaxMin, MinMax Regret and MaxMin Ratio. We introduce a sampling based optimization approach that can solve this problem in a tractable way, with a confidence level and a robustness level based on the number of samples used. We will show how this methodology can be used for data-driven pricing or adapted for a random sampling optimization approach when limited information is known about the demand uncertainty. Finally, we compare the revenue performance of the different models using numerical simulations and a case study using real airline data collected from a flight booking website. We then explore the behavior of each model under different sample sizes and sampling distributions.

In the second part, Chapter 3, we explore the role of governments who design price incentives for consumers to adopt new technologies, in this case for solar photovoltaic technology. The first goal is to establish a framework for policy makers to design long-term subsidy policies while considering the dynamics of the technological adoption process. The second goal is to perform an empirical study using this framework. We collected real market data about the sales of solar panels in Germany and fitted the data to the model we developed, obtaining market forecasts and potential policy recommendations.

Unlike traditional sources of electricity generation, a rooftop solar panel is a clean generation source that can be installed in small scales directly at the consumers' household.This lowers transmission costs and losses, as well as offsetts the electricity demand during peak hours. On the other hand, these panels are expensive and will require subsidizing consumers until the technology itself and its distribution channels are mature. This can only happen once a large consumer base has already been established. With the objective of jump-starting this consumer base, governments propose a target adoption level to be reached within a given time frame. Current approaches to subsidy design can be short-sighted, as they do not take into consideration the dynamics of cost improvements and information spread that is inherit to the adoption of a new technology. In this chapter, we develop a modeling framework for policy-makers to find the optimal subsidy levels that will achieve the given target while incorporating these network effects.

Furthermore, we validate our model through an empirical study of the German solar market, where model parameters are estimated. We then use this model to generate adoption forecasts and demonstrate how to solve the policy design problem. Through this study we are able to show that the current policies in Germany are not efficient. In particular, it suggests that their subsidies should be higher in the near future and the gradual phase-out of the subsidies should occur faster.

In the third part of the thesis, Chapter 4, we look at how governments and firms interact with each other to determine prices and production levels for new products. In particular, the goal is to understand the effect of the government's inability to commit to certain policies and how this can lead to under-production from the private sector.

In order to establish this comparison, we first model a two-period system where the government chooses a set of rebates for this time horizon and commits to it. Industry then responds by building up production capacity. In the other model, we assume government has the ability to change the rebate level during the second period after the first demand is realized. On the other hand, the industry player is already committed to a capacity investment at the beginning of the time horizon. The resulting effect is that the industry player, forecasting a potential decrease in subsidy levels, will under-invest in the previous stage. This leads to a worse outcome for the entire system.

The remainder of this thesis is structured as follows. In Chapter 2, we cover the dynamic pricing problem. In Chapter 3, we develop the subsidy design problem applied to the German solar market. In Chapter 4, we develop the subsidy problem with strategic industry response, where we address the issue of policy commitment versus flexibility. Each of these chapters contains an appendix section with notation used in the chapter and the proofs developed within the chapter. In Chapter 5, we summarize some of the key conclusions of this thesis.

Chapter 2

Dynamic Pricing

2.1 Introduction

In many industries, managers are faced with the challenge of selling a fixed amount of inventory within a specific time horizon. Moreover, the firm offers an array of different products which will consume a shared set of resources. Examples include the case of airlines selling flight tickets, hotels trying to book rooms and retail stores selling products for the current season. All these cases share a common trait: a fixed initial inventory that cannot be replenished within the selling horizon. The firm's goal is to set prices at each stage of the given selling horizon that will maximize revenues while facing an uncertain demand.

As in most real-life applications, the intrinsic randomness of the firm's demand is an important factor that must be taken into account. Decisions based on deterministic forecasts of demand can expose the firm to severe revenue losses or fail to capture potential revenues. In the pricing problem, demand is usually modeled as a function of the prices set by the firm. To model uncertainty in demand, we represent some parameters of this function as random. These random parameters can be modeled assuming a fixed distribution, as in the stochastic optimization literature, or using the distribution-free framework of robust optimization. In the stochastic optimization approach, demand uncertainty is given by a certain probability distribution and the firm's objective is to find prices that maximize its expected revenue. A disadvantage of this approach is that it requires full knowledge of the demand distribution, often requiring the uncertainty to be independent across time periods. Note that the dynamic programming formulation used to solve the stochastic optimization model will have a serious dimensionality problem if there are many types of resources or if demands are correlated. On the other hand, in the robust optimization approach, one does not assume a known distribution for the demand uncertainty, but assumes only that it lies within a bounded uncertainty set. The goal in this case is to find a pricing policy that robustly maximizes the revenue within this uncertainty set, without further assumptions about the distribution or correlation of demands across time periods. Nevertheless, as a drawback, the robust solution is often regarded as too conservative and the robust optimization literature has mainly focused on static problems (open-loop policy). The dynamic models that search for closed-loop policies, i.e. which account for previously realized uncertainty and adjust the pricing decisions, can easily become intractable. Our goal in this chapter is to come up with an approach that can tackle these two issues, by developing a framework for finding non-conservative and adjustable robust pricing policies.

In order to achieve this goal, we are going to develop pricing optimization models that are based on sampled scenarios of demand. At first, using the assumption that we have access to historical samples of demand, we will develop a data-driven robust approach to pricing. Furthermore, in the case these samples are not available, we will develop a framework to generate random scenario samples using mild assumptions on the demand distribution and still obtain a sampling based robust solution. We will also provide theoretical bounds on the number of samples required to obtain a given confidence level about our solutions and develop numerical examples to demonstrate how to apply this methodology in a practical setting and to obtain managerial insights about these pricing models.

2.1.1 Literature Review

A good introduction to the field of revenue management and dynamic pricing would include the overview papers of (Elmaghraby and Keskinocak 2003), (Bitran and Caldentey 2003) and the book by (Talluri and van Ryzin 2004). The revenue management literature in general assumes a fixed set of pre-determined prices and tries to allocate capacity for different price levels. This problem can be traced back to (Littlewood 1972), who first tackled the network capacity control problem. The complexity of solving these large network models led to the development of numerous heuristics. (Belobaba 1987; Belobaba 1992) later developed heuristics such as EMSR-a and EMSR-b which performed very well in practice for solving the capacity control problem. (Talluri and van Ryzin 1998) further developed the concept of bid-price controls. (Bertsimas and Popescu 2003) solved this problem by approximating the revenue-to-go function and (Bertsimas and de Boer 2005) introduced an algorithm based on stochastic gradient. (Akan and Ata 2009) recently developed a martingale characterization of bid-price controls. On the other hand, the dynamic pricing literature attempts to solve a similar problem, but does not assume fixed prices, rather it assumes that we can choose different prices at each period. In that case, we also assume that the aggregate demand seen by the firm is a function of the price.

Most of the revenue management and pricing literature uses the stochastic optimization framework, which makes distributional assumptions on the demand model and often doesn't capture correlation of the demand uncertainty across time periods. Avoiding such problems, there are two different approaches prominent in the literature that will be most relevant for our research: data-driven and robust optimization. So far, these approaches have been studied separately and the modeling choice usually depends on the type of information provided to the firm about the demand uncertainty.

The operations management literature has explored sampling based optimization as a form of data-driven nonparametric approach to solving stochastic optimization problems with unknown distributions. In this case, we use past historical data, which are sample evaluations coming from the true demand distribution. In revenue management, this approach has been pioneered by (van Ryzin and McGill 2000) who introduced an adaptive algorithm for booking limits. (Bertsimas and de Boer 2005) and (van Ryzin and Vulcano 2008) developed a stochastic gradient algorithm to solve a revenue management problem using the scenario samples. Another typical form of data-driven approach is known as Sample Average Approximation (SAA), when the scenario evaluations are averaged to approximate the expectation of the objective function. (Kleywegt, Shapiro, and Homem-de Mello 2001) deal with a discrete stochastic optimization model and show that the SAA solution converges almost surely to the optimal solution of the original problem when the number of samples goes to infinity and derive a bound on the number of samples required to obtain at most a certain difference between the SAA solution and the optimal value, under some confidence level. (Zhan and Shen 2005) apply the SAA framework for the single period price-setting newsvendor problem. (Levi, Roundy, and Shmoys 2007) also apply the SAA framework to the newsvendor problem (single and multi-period) and establish bounds on the number of samples required to guarantee with some probability that the expected cost of the sample-based policies approximates the expected optimal cost and (Levi, Perakis, and Uichanco 2010), using the assumption that the demand distribution is log-concave, develop a better bound on the number of samples to obtain a similar guarantee as in (Levi, Roundy, and Shmoys 2007). More specifically in the dynamic pricing literature, the data-driven approach has been used by (Rusmevichientong, Van Roy, and Glynn 2006), to develop a non-parametric data-driven approach to pricing, and also more recently by

(Eren and Maglaras 2009). (Besbes and Zeevi 2008) assume there is no prior data and develop an approach that requires an exploration phase to obtain data and an exploitation phase to generate revenue. Other relevant work on the subject of revenue management with demand learning is by (Araman and Caldentey 2009) and by (Farias and Van Roy 2009).

The fairly recent field of robust optimization proposes distribution-free modeling ideas for making decision models under uncertainty. This area was initiated by (Soyster 1973) and it was further developed by (Ben-Tal and Nemirovski 1998; Ben-Tal and Nemirovski 1999; Ben-Tal and Nemirovski 2000), (Goldfarb and Iyengar 2003) and (Bertsimas and Sim 2004). A robust policy can be defined in different ways. In this chapter we will explore three different types of robust models: the MaxMin, the MinMax Regret (or alternatively MinMax Absolute Regret) and the MaxMin Ratio (or alternatively MaxMin Relative Regret or MaxMin Competitive Ratio). In inventory management, the MaxMin robust approach can be seen in (Scarf 1958), (Gallego and Moon 1993), (Ben-Tal, Goryashko, Guslitzer, and Nemirovski 2004), (Bertsimas and Thiele 2006). The following papers by (Adida and Perakis 2005; Nguyen and Perakis 2005; Perakis and Sood 2006; Thiele 2006; Birbil, Frenk, Gromicho, and Zhang 2006) are examples of the MaxMin robust approach applied to the dynamic pricing problem. This approach is usually appropriate for risk-averse managers, but it can give quite conservative solutions. For this reason we will explore the regret based models, which were originally proposed by (Savage 1951). (Lim and Shanthikumar 2007) and (Lim, Shanthikumar, and Watewai 2008) approach this problem from a different angle, where the pricing policies are protected against a family of distributions bounded by a relative entropy measure. In the broader operations management literature, (Yue, Chen, and M.-C. 2006) and (Perakis and Roels 2008) use the MinMax Absolute Regret for the newsvendor problem. A comparison of MaxMin and MinMax Absolute Regret for revenue management can be found in (Roels and Perakis 2007). An alternative approach is the relative regret measure, also known as the competitive ratio. In revenue management and pricing, (Ball and Queyranne 2006) and (Lan, Gao, Ball, and Karaesmen 2006) use this MaxMin Ratio approach.

Irrespective of the demand uncertainty models described above, multi-period decision models can also be categorized between (i) closed-loop policies, where policies use feedback from the actual state of the system at each stage, and (ii) open-loop policies, where the entire policy is defined statically at the beginning of the time horizon. The initial robust framework discussed in the papers above does not allow for adaptability in the optimal policy. This open-loop robust framework has been applied to the dynamic pricing problem in (Perakis and Sood 2006), (Nguyen and Perakis

2005), (Adida and Perakis 2005) and (Thiele 2006). Moving towards closed-loop solutions, (Ben-Tal, Goryashko, Guslitzer, and Nemirovski 2004) first introduced adaptability to robust optimization problems. (Ben-Tal, Golany, Nemirovski, and Vial 2005) propose an application of adjustable robust optimization in a supply chain problem. More specifically, they advocate for the use of affinely adjustable policies. Recent work by (Bertsimas, Iancu, and Parrilo 2009) was able to show that optimality can actually be achieved by affine policies for a particular class of one-dimensional multistage robust problems. Unfortunately this is not our case, therefore we must admit that the affine policies we are using will only achieve an approximation to the fully closed-loop policy. (Zhang 2006) develops a numerical study of the affinely adjustable robust model for the pricing problem using an MIP formulation. In this chapter, we present a model that introduces an affinely adjustable approach to the dynamic pricing problem and uses sampling based approach to solve the robust problem.

As mentioned by (Caramanis 2006), the sampling approach to the adaptable robust problem puts aside the non-convexities created by the influence of the realized uncertainties in the policy decisions. The natural question that arises is how many scenarios do we need to have any confidence guarantees on our model's solution. To answer this question, (Calafiore and Campi 2005; Calafiore and Campi 2006) define the concept of an ϵ-level robust solution and provide a theoretical bound on the sample size necessary to obtain this solution. The bound was later improved in (Campi and Garatti 2008) and (Calafiore 2009), which is provably the tightest possible bound for a class of robust problems defined as "fully-supported" problems. Recently (Pagnoncelli, Ahmed, and Shapiro 2009) suggested how to use this framework to solve chance constrained problems.

As for the numerical experiments, we will compare the out-of-sample revenues of our sampling-based pricing models by using Monte Carlo-based performance measures. (Chiralaksanakul and Morton 2004) and (Bayraksan and Morton 2006) proposed Monte Carlo-based procedures to access the quality of stochastic programs, building sampling-based confidence intervals for the optimality gap of the solution. To compare our pricing models, we will conduct paired tests and use similar concepts of sampling based confidence intervals around the average and worst-case difference in revenues.

2.1.2 Contributions

The main goal of this chapter is to solve the multi-item dynamic pricing problem in a practical way. In particular, we develop a methodology that is easy to implement, uses directly available data in the price optimization model, and at the same time is justified from a theoretical point of view. We accomplish this by using a sampling based optimization approach, which allows for a wide range of modeling complexity such as adjustable pricing policies, nonlinear demand functions, network effects, overbooking and salvage value. The solution concept that motivates our model is the robust pricing problem. Nevertheless this problem becomes intractable as we introduce all the modeling components mentioned before. We show in this chapter how this new combined robust sampling based framework is tractable, performs very well in numerical experiments and is theoretically grounded as a good approximation to the original robust problem, depending on the number of samples used.

More specifically, most robust approaches to the pricing problem use an open-loop model, where the information obtained by observing past demand is not used in the pricing decision of future periods. A key challenge, when introducing adjustability to dynamic pricing, is that the model easily turns into a non-convex problem, which is intractable to solve even when using traditional robust optimization techniques. Instead we propose a sampling based approach where we solve the problem using only a given set of demand scenarios. The new problem becomes a convex program and can be efficiently solved. The question that arises from this approach is how many samples do we need in order to have a performance guarantee on our sampling based solution. To answer this question, we define a notion of ϵ-robustness and show the sample size needed to achieve an ϵ-robust solution with some confidence level. Nevertheless, this type of data-driven approach uses the assumption that we have a large set of historical data available, which comes from the true underlying distribution of the uncertainty. This assumption can be quite restrictive in many real applications, for example when releasing a new product. For these cases, we introduce a new concept of random scenario sampling, where we use an artificial distribution over the uncertainty set to generate random sample points and apply the sampling based optimization framework. We are able to show bounds on the sample size required for the random sampling solution using a relatively mild assumption on the unknown underlying distribution, which we refer to as a bounded likelihood ratio assumption.

Although fairly general, the bounded likelihood ratio assumption can be seen as too abstract and hard to quantify in a practical setting. For this reason we further developed a method to obtain this likelihood ratio measure using more commonly used assumptions, which can be more easily verified by the firm. More specifically, using the assumption that the uncertain parameter of the demand follows a log-concave distribution, independent across products and time periods and with a known standard deviation, we obtain the likelihood ratio bound on the true unknown distribution, relative to a uniform random sampling distribution. Therefore we can use the bound developed before to determine the sample size that needs to be generated to obtain any level of robustness for the pricing problem.

The random scenario sampling framework we introduce is a rather powerful concept, given that we are now able to use a data-driven methodology to solve a robust problem without any actual historical data. Data-driven and robust optimization have been generally considered two separate fields, since they use very different initial assumptions of information about the problem. The bridge we develop between data-driven optimization and robust optimization has not been widely explored. We start with an intractable robust model and illustrate how to solve a close approximation to the robust solution using only randomly generated scenarios.

Another contribution of this work comes from the numerical experiments, where we study the simulated revenue performance of our dynamic pricing models. Our first experiment will provide empirical evidence that the sample size bound we present in this chapter is in fact tight for our dynamic pricing problem. In the other experiments, we will compare the performance of the robust pricing models, using three different types of robust objectives (MaxMin, MinMax Regret, MaxMin Ratio) and a Sample Average Approximation (SAA) model, which is more common in the stochastic optimization literature. When compared to the SAA benchmark, the robust models will usually obtain a smaller average revenue but will do significantly better in the lower revenue cases. On the other hand, when the number of samples provided are small, the robust models have shown to perform just as well as the SAA on average, while still maintaining a less volatile revenue outcome. Comparing between the robust models, we show that the traditional robust (MaxMin) is usually dominated by the regret based models, both in average and worst-case revenue. The two regret-based models (MinMax Regret and MaxMin Ratio) have a relatively similar performance, but the first one tends to be more robust and conservative than the second.

We also show in this chapter how to apply our methodology in practice, using a case study with actual airline data. Without considering competition effects, we will show how our robust pricing models might achieve significant improvements in the revenues per flight.

The remainder of the chapter is structured as follows. In Section 2, we introduce the modeling approach that we will consider and discuss the implementation issues. In Section 3, we show the simulated performance of the proposed models and interpret the numerical results. In Section 4, we develop an airline case study. Finally, in Section 5, we conclude with a summary of the discussions mentioned before and give possible directions for future work.

The appendices provide supplementary content to the reader that were omitted from the chapter for conciseness. Appendix A provides a summary of the notation used. Appendices B-F display the proofs of our theoretical results in Section 2. Appendices G-J provide information about the numerical experiments displayed Section 3, as well as additional experiments.

2.2 The Model

Before introducing the general model we propose in this chapter, we motivate the problem with the following example. Suppose a firm sells only one product over a two period horizon, with a limited inventory of C. Moreover, suppose the firm has a set of N historical data samples of demand and prices for each period of the sales horizon. We will assume for this example the demand is a linear function of the price plus δ, which is a random noise component: $\text{Demand}_t = a_t - b_t\text{Price}_t + \delta_t$. After estimating the demand function parameters (a_1, a_2, b_1, b_2) using the N data points, we are left with a set of estimation errors $\delta_t^{(1)}, ..., \delta_t^{(N)}$ for each time period $t = 1, 2$. A typical robust pricing approach would define an uncertainty set from which these errors are coming from and choose prices that maximize the worst case revenue scenarios within that set. It is not clear how one should define this uncertainty set given a pool of uncertainty samples and the resulting problem can also become too hard to solve, as we will show later in this section. The direct use of the uncertainty samples $\delta_t^{(i)}$ in the price optimization is what characterizes a sampling based optimization model, which we advocate for in this chapter. Our goal, as seen in the following model, is to find a pricing strategy that robustly maximizes the firm's revenue with respect to the N given observations of demand uncertainty:

$$\begin{aligned} \max_{p_1,p_2\geq 0} \quad \min_{i=1,...,N} \quad & p_1(a_1 - b_1p_1 + \delta_1^{(i)}) + p_2(a_2 - b_2p_2 + \delta_2^{(i)}) \\ s.t. \quad & (a_1 - b_1p_1 + \delta_1^{(i)}) + (a_2 - b_2p_2 + \delta_2^{(i)}) \leq C, \quad \forall i = 1, ..., N \end{aligned}$$

Note that the given uncertainty samples $\delta_t^{(i)}$ will approximate the uncertainty set from the traditional robust optimization approach. The major theoretical challenge now is to determine how many samples are needed for the sampling based model to approximate the original robust problem.

On the modeling aspect of the problem, one of the main problems with the solution concept presented above is that this MaxMin robust approach can often be too conservative. For this reason, we will propose other types of robust modeling. Another problem in this example is that the second period price p_2 does not depend on the uncertainty realized on the first period δ_1, which is what we call an open-loop model. Ideally, the second period pricing policy should be a function of the new information obtained in the first period, $p_2(\delta_1)$, which is known as a closed-loop model.

After the motivating example illustrated above, we proceed to generalize the problem to include components such as network effects, salvage/overbooking inventory and nonlinear demand models. Throughout this section we develop a solution approach that mitigates the modeling issues described before. Further on, in Section 2.1, we will address the theoretical issue of the number of samples required. We refer the reader to Appendix A for a summary of the notation that will be used in this chapter.

To address the network component of pricing problems, we consider the firm to have n products (eg. itineraries) and m resources (eg. flight legs). The $m \times n$ incidence matrix M determines which resources are consumed by each product sold. Also let T be the length of the time horizon. Define $p_{j,t}$ as the price of product j at time t and the demand function as $d_{j,t}$. We admit that the demand $d_{j,t}$ can be a function of the prices of all the products at any time period. For the sake of clarity, we will use demand functions that are affected only by the current product price $d_{j,t}(p_{j,t})$. This means that a certain product's demand will not be affected by the prices of the other products or by the past prices. It is important to note that the modeling techniques and the results presented in this chapter can be easily implemented with cross product/time effects (which also allow us to use demand models with reference prices). We are only restricting ourselves to effects of the current price to avoid complicating the notation. The firm's goal is to set a pricing policy for each product that robustly maximizes the revenue of the firm, $\sum_{t=1}^{T} \sum_{j=1}^{n} p_{j,t} d_{j,t}(p_{j,t})$. For technical reasons, which are discussed later, we require that the demand function satisfies the following convexity/concavity assumption.

Assumption 2.2.1 *Let $d_{j,t}(p_{j,t})$ be the nominal demand as a non-increasing function of the price $p_{j,t}$ for a given set of demand parameters. We further assume that $d_{j,t}(p_{j,t})$ is convex in $p_{j,t}$ and $\sum_{t=1}^{T} \sum_{j=1}^{n} p_{j,t} d_{j,t}(p_{j,t})$ is strictly concave in p.*

For example, suppose we assume a linear demand function, $d_{j,t}(p_{j,t}) = a_{j,t} - b_{j,t} p_{j,t}$, where $a_{j,t}$ and $b_{j,t}$ are scalars. The intuition behind Assumption 1 is that we want the space of pricing strategies to be a closed convex set and the objective function to be strictly concave, giving rise to a unique optimal solution. Examples of demand functions that satisfy Assumption 2.2.1, besides the linear demand function, which are common in the revenue management literature are the iso-elastic demand $d_{j,t}(p_{j,t}) = a(p_{j,t})^{-b}$ for $b \in [0,1]$, and the logarithmic demand $d_{j,t}(p_{j,t}) = -a\log(p_{j,t}/b)$. In the numerical experiment of Appendix H, we will illustrate our framework using the logarithmic demand model.

Without considering demand uncertainty, the objective is to find the optimal set of prices that maximizes the sum of the revenues for the entire horizon. The prices are nonnegative and the total demand seen by the firm for a given resource l should be less than its total capacity C_l or else the firm will pay an overbooking fee o_l for every unit of resource sold above capacity of resource l. For every unit of capacity not sold, the firm will get a salvage value of g_l. We require that $o_l > g_l$ for all resources $l = 1, ..., m$ to guarantee the concavity of the objective function. In most practical applications, the salvage value is small and the overbooking fee is large, which makes this assumption very easy to justify. Define w_l as the difference between the number of units sold and the capacity in each resource l, while $w_l^+ = \max(w_l, 0)$ and $w_l^- = \min(w_l, 0)$.

To capture the variability in demand, given the nominal demand function $d_{j,t}(p_{j,t})$, define $\tilde{d}_{j,t}(p_{j,t}, \delta_{j,t})$ as the actual demand for product j at time t, which is realized for some uncertain parameter $\delta_{j,t}$. For example, suppose we have a linear nominal demand function given by $d_{j,t} = a_{j,t} - b_{j,t}p_{j,t}$. Then introducing an additive uncertainty we would have: $\tilde{d}_{j,t}(p_{j,t}, \delta_{j,t}) = a_{j,t} - b_{j,t}p_{j,t} + \delta_{j,t}$. In the literature, it is mostly common to use additive or multiplicative uncertainty. In our framework, we admit any sort of dependence of the demand function on the uncertain parameters.

In general, $\delta = (\delta_{1,1}, ...\delta_{n,T})$ is a random vector with one component $\delta_{j,t}$ for each product j and each period t (it can be easily generalized for multiple uncertain components in each product or time period). We assume that δ is drawn from an unknown probability distribution $\mathcal{Q}$, with support on the set U, which we call the uncertainty set. We do not make any assumptions about the independence of δ across time or products, as opposed to most stochastic optimization approaches. Although not required in our general framework, we will use an independence assumption when giving an example of how to apply the random sampling approach (see Corollary 2.2.1).

In our framework, we assume there is a finite set of pricing decision variables s, which lie within the strategy space S and define the prices $p_{j,t}(s, \delta)$ for each product at each period. We assume that S is a finite dimensional and compact set. In the case of static pricing (open-loop policies), s is a vector of fixed prices decided before hand, independent of the realizations of demand. When using adjustable policies (closed-loop), the actual price at time t must naturally be a function only of the uncertainty up to time $t-1$. For conciseness, we will express the actual realized prices as $p_{j,t}(s, \delta)$

for both cases, while $p_{j,t}(s,\delta)$ is actually independent of δ in open-loop policies and in closed-loop policies a function of $\delta_{j,1},...,\delta_{j,t-1}$, including all products j. Also to avoid redundancy in notation, define $\tilde{d}_{j,t}(s,\delta) = \tilde{d}_{j,t}(p_{j,t}(s,\delta),\delta_{j,t})$. In other words, the policy s and the uncertainty δ determine the prices of all products at time t, therefore also determining the realized demand. Define the net leftover inventory of resource l as a function of the pricing policy and the realized uncertainty: $\tilde{w}_l(s,\delta) = \sum_{t=1}^{T}\sum_{j=1}^{n} M_{l,j}\tilde{d}_{j,t}(s,\delta) - C_l$.

As stated before, our goal is to find a pricing policy that will give a robust performance for all possible realizations of demand. One can think of the robust pricing problem defined as a game played between the firm and nature. The firm chooses a pricing policy s and nature chooses the deviations $\delta \in U$ that will minimize the firm's revenue. The firm seeks to find the best robust policy under the constraints that the pricing policy yields nonnegative prices and that the total demand must be less than or equal the capacity (although relaxed by overbooking fees and salvage value). To express the different types of robust objectives explored in this chapter, define $h^{obj}(s,\delta)$ as the objective function realization for a given pricing strategy s and uncertainty δ. The index obj can be replaced with one of three types of robust objectives that we consider in this work: the **MaxMin**, the **MinMax Regret** and the **MaxMin Ratio**.

$$\begin{aligned}
h^{MaxMin}(s,\delta) &= \sum_{t=1}^{T}\sum_{j=1}^{n} p_{j,t}(s,\delta)\tilde{d}_{j,t}(s,\delta) - \sum_{l=1}^{m} o_l\tilde{w}_l(s,\delta)^+ + g_l\tilde{w}_l(s,\delta)^- \\
h^{Regret}(s,\delta) &= \sum_{t=1}^{T}\sum_{j=1}^{n} p_{j,t}(s,\delta)\tilde{d}_{j,t}(s,\delta) - \sum_{l=1}^{m} o_l\tilde{w}_l(s,\delta)^+ + g_l\tilde{w}_l(s,\delta)^- \\
&\quad - \max_{y\geq 0}\left\{\sum_{t=1}^{T}\sum_{j=1}^{n} y_{j,t}\tilde{d}_{j,t}(y,\delta_t) - \sum_{l=1}^{m} o_l\tilde{w}_l(y,\delta)^+ + g_l\tilde{w}_l(y,\delta)^-\right\} \\
h^{Ratio}(s,\delta) &= \frac{\sum_{t=1}^{T}\sum_{j=1}^{n} p_{j,t}(s,\delta)\tilde{d}_{j,t}(s,\delta) - \sum_{l=1}^{m} o_l\tilde{w}_l(s,\delta)^+ + g_l\tilde{w}_l(s,\delta)^-}{\max_{y\geq 0}\left\{\sum_{t=1}^{T}\sum_{j=1}^{n} y_{j,t}\tilde{d}_{j,t}(y,\delta_t) - \sum_{l=1}^{m} o_l\tilde{w}_l(y,\delta)^+ + g_l\tilde{w}_l(y,\delta)^-\right\}}
\end{aligned}$$

In the first robust concept, **MaxMin**, we try to find a pricing policy that gives us the best revenue for all possible realizations of the uncertainty δ in the set U. More specifically, the actual revenue realized for some policy s and deviation δ is given by $\Pi(s,\delta) = \sum_{t=1}^{T}\sum_{j=1}^{n} p_{j,t}(s,\delta)\tilde{d}_{j,t}(s,\delta) - \sum_{l=1}^{m} o_l\tilde{w}_l(s,\delta)^+ + g_l\tilde{w}_l(s,\delta)^-$. Then for the MaxMin case, the objective function is simply given by $h^{MaxMin}(s,\delta) = \Pi(s,\delta)$. As a drawback, the MaxMin approach often finds conservative pricing policies. To avoid this issue, we also explore a robust approach called the **MinMax Regret**, also known as absolute regret. In this case, the firm wants to minimize the regret it will have from using a certain policy relative to the best possible revenue in hindsight, i.e. after observing the realization of demand. In other words, define the optimal hindsight revenue $\Pi^*(\delta)$ as the optimal revenue the firm could achieve, if it knew the demand uncertainty beforehand:

$$\begin{aligned} \Pi^*(\delta) = \quad & \max_{y,w} \quad \textstyle\sum_{t=1}^{T}\sum_{j=1}^{n} y_{j,t}\tilde{d}_{j,t}(y,\delta) - \displaystyle\sum_{l=1}^{m} o_l w_l^+ + g_l w_l^- \\ & s.t. \quad \textstyle\sum_{t=1}^{T}\sum_{j=1}^{n} M_{l,j}\tilde{d}_{j,t}(y,\delta) \leq C_l + w_l, \qquad \forall l = 1,...,m \\ & \qquad\;\; y \geq 0 \end{aligned} \tag{2.1}$$

The model above for the hindsight revenue $\Pi^*(\delta)$ is a deterministic convex optimization problem, which can be efficiently computed for any given δ (either in closed form or using an optimization package). Note that the capacity constraint can always be satisfied with equality in the optimal solution, since left-over capacity can be shifted to w_l while improving the objective function because of the salvage value, $g_l \geq 0$. Therefore, at optimality, the left-over capacity variable from the hindsight model above will have $w_l = \tilde{w}_l(y,\delta)$.

Define the absolute regret as the difference between the hindsight revenue and the actual revenue: $\Pi^*(\delta) - \Pi(s,\delta)$. To be consistent with the MaxMin formulation, we can define the objective function for the MinMax Regret as $h^{Regret}(s,\delta) = \Pi(s,\delta) - \Pi^*(\delta)$, which is the negative of the regret. We will continue calling this MinMax Regret, since it is a more common term in the literature (although this would be more precisely named MaxMin Negative Regret). Finally, the **MaxMin Ratio**, which is also known as the relative regret or competitive ratio, tries to bound the ratio between the actual revenue and the hindsight revenue. The objective function for the MaxMin Ratio can be concisely written $h^{Ratio}(s,\delta) = \frac{\Pi(s,\delta)}{\Pi^*(\delta)}$.

Consider the following model which we call the robust pricing model. We will go into details of the different modeling techniques later in this section, but they all fall within the following general framework:

$$\begin{aligned} & \max_{s\in S,z} \quad z \\ & s.t. \quad \left\{ \begin{array}{l} z \leq h^{obj}(s,\delta) \\ p(s,\delta) \geq 0 \end{array} \right\} \forall \delta \in U \end{aligned} \tag{2.2}$$

The inputs to the general robust model are: the structure of the pricing policy $p_t(s,\delta)$, parameterized by the decision variables s; the demand functions $\tilde{d}_t(p_t, \delta_t)$ for any given price p_t and uncertainty δ_t; the objective function $h^{obj}(s, \delta)$ and the uncertainty set U (which we will later replace with uncertainty samples from U, as we can see in Section 2.1). The outputs are the set of pricing decisions s and the variable z, which is a dummy variable used to capture the robust objective value of the pricing policy.

To introduce adjustability in our model, we must find a pricing policy that is a function of the previous realized uncertainties. On the other hand, searching over all possible functions is a rather hard task. An approach used in the literature to introduce adaptability to the robust model is the affinely adjustable robust optimization, where the search is limited to policies that are linear on the control variables. For example, define the price p using the variables u as the deterministic component of the price and v as the effect that past demand uncertainty should have on current prices. In this case $s = (u, v)$. Then a typical pricing policy can be defined as:

$$p_{j,t}(u, v, \delta) = u_{j,t} + \sum_{l=1}^{m} v_{j,l,t} \sum_{j'=1}^{n} \sum_{t'=1}^{t-1} M_{l,j'} \delta_{j',t'} \qquad (2.3)$$

This policy (2.3) will be used in our numerical experiments in Section 3. In this policy example above, we have an adjustable variable for each resource, where the summation $\sum_{j'=1}^{n} \sum_{t'=1}^{t-1} M_{l,j'} \delta_{j',t'}$ contains a lot of the inventory information necessary for our adjustable policy. As an intuition for choosing this type of policy, observe that when using additive uncertainty the sum of past uncertainties $\sum_{j'=1}^{n} \sum_{t'=1}^{t-1} M_{l,j'} \delta_{j',t'}$ is the amount of sales realized above or below the deterministic level expected up to time $t-1$ for the resource l. If the firm has sold more than expected, we should have an incentive to raise prices, given the reduced amount of inventory left. On the other hand, if uncertainties are negatively correlated across time, a series of low demands could induce a higher demand in the future, possibly giving the incentive to raise prices. Notice how static policies can be overly conservative because they cannot adjust the prices according to the different levels of the inventory, while the adjustable policies are more flexible. In the example above, we chose a pricing policy that is linear on δ. In some situations, we might rather capture the effect of past deviations in other forms, for instance δ^2 or $|\delta|$, depending on the structure of the demand function and the uncertainty set. In the framework we propose, pricing policies can be any function of previous realizations of δ, as long as the function is linear in the pricing decision variables for any given δ.

Assumption 2.2.2 *Let $p_t(s, \delta_1, ..., \delta_{t-1})$ be the pricing policy at time t for a given set of pricing* *decision variables s and uncertainties realized up to time t. We will restrict ourselves to policies that are linear on s for any given δ. Nevertheless the dependence on δ can be nonlinear. We also assume that s is restricted by the strategy space S, which is a finite dimensional, compact and convex set.*

The reason for Assumption 2 can be easily illustrated in the pricing policy in (2.3). When introducing this policy as an equality in the optimization model (2.2), we must guarantee that it defines a convex set in p, u and v for any given δ. This convexity will hold if the pricing policy is linear on the decision variables (u, v), as stated in Assumption 2.

Furthermore, the set S defined above captures only the constraints in the strategy space that do not depend on the uncertainty. In general, we expect S to be a box set, with upper and lower bounds for the pricing variables. The finite dimensionality and compactness of set S are introduced to guarantee the existence of a solution to the optimization problem.

Note that using adjustable policies, the inequality $z \leq h^{obj}(p(u, v, \delta), \delta)$, which defines the objective function, is neither concave nor convex with respect to the deviations δ. In Appendix B, we illustrate this with a simple instance of the adjustable MaxMin model. Because of this lack of convexity or concavity, the traditional robust optimization methods (i.e. solve the exact robust problem using duality arguments or simply searching over the boundary of uncertainty set) will be intractable. Note the example in Appendix B uses a simple linear demand model and simple objective function and it is not necessary to show that more complicated modeling will generally not resolve the convexity issue. In the next section we introduce the sampling based approach that we advocate for solving the robust pricing problem.

2.2.1 Sampling based optimization

Ideally, we would like to solve the exact robust pricing problem, but as we have seen in the previous section this can easily become intractable. Instead, assume that we are given N possible uncertainty scenarios $\delta^{(1)}, ..., \delta^{(N)}$, where each realization $\delta^{(i)}$ is a vector containing a value for each time period and product. We use the given sampled scenarios to approximate the uncertainty set, replacing the continuum of constraints in (2.2) by a finite number of constraints. It is only natural to question how good is the solution to the sampling-based problem relative to the original robust problem that was proposed in the previous section. In order to discuss this issue, we will explore in this section the theoretical framework needed to analyze the robustness of a sampling-based solution. For clarity purposes, we again refer the user to Appendix A for a summary of the notation used in this section to formulate the theoretical results. Define the following model as the sampling based counterpart of the robust pricing model:

$$\begin{aligned} \max_{s \in S, z} \quad & z \\ s.t. \quad & \left. \begin{cases} z \leq h^{obj}(s, \delta^{(i)}) \\ p(s, \delta^{(i)}) \geq 0 \end{cases} \right\} \forall i = 1 \ldots N \end{aligned} \tag{2.4}$$

The idea of using sampled uncertainty scenarios, or data points, in stochastic optimization models is often called sampling or scenario based optimization. Note that from Assumptions 1 and 2, the constraints in (2.4) define a convex set in (s, z) for any fixed vector δ. Therefore the scenario based problem is a convex optimization problem and can be solved by any nonlinear optimization solver. It is easy to argue that the exact robust problem that we initially stated in (2.2) has an optimal solution (for a proof under more concise notation, see Appendix C). Also, from Assumption 1, it follows that for any fixed uncertainty δ, there exist a unique optimal solution to the exact robust pricing problem. It remains to show how good is the approximation of the scenario based model relative to the exact robust model. Depending on what type of demand information and historical data that is initially provided to the firm, we propose two solution approaches: the Data-Driven and the Random Scenario Sampling. The first approach, Data-Driven, assumes we are given a large pool of uncertainty data drawn from the true underlying distribution, for example, from historical data. The second approach, Random Scenario Sampling, assumes that we don't have enough, if any, data points from the true distribution, but instead we have a sampling distribution which can be used to generate random data points.

Suppose we are given a large sample set of historical data of prices and demands. There are many ways to estimate the parameters of the demand function using such data. As an example, in the numerical study of Section 4, we used linear regression on the price-demand data to estimate a linear demand model. To be consistent with our distribution-free approach, we should not rely on hypothesis tests that use a normality assumption about the error. If such assumption can be made about the estimation errors in the first place, then it could potentially be explored in the optimization model. In this case, we use the linear regression simply as a tool to obtain estimates of the demand parameters we need, even if we cannot validate the fit with statistical testing. The estimation error obtained is the sample deviations δ in our model. Under the assumption that the data points come from the true underlying distribution, we will show a performance guarantee on how robust is the sampling based solution relative to the exact solution.

Define the set of decision variables $x = (s, z)$, such that it lies within the domain $x \in X$, where $X = S \times \Re$. Define c with the same dimension as x such that $c = (0, ..., 0, 1)$. Also define the equivalent constraint function as a scalar valued f such that: $f(x, \delta) = \max\left\{z - h^{obj}(s, \delta), -p(s, \delta)\right\}$.

Since each constraint in the definition of $f(x, \delta)$ above is convex in x for any given δ, the maximum between them is still convex in x. Moreover, $f(x, \delta) \leq 0$ is equivalent to the constraint set defined before in (2.2). Then problem (2.2) can be concisely defined as the following model (2.5):

$$\max_{x \in X} c'x, \qquad \text{s.t.} f(x, \delta) \leq 0, \forall \delta \in U \tag{2.5}$$

Since we cannot solve the problem with a continuum of constraints, we solve the problem for a finite sample of deviations $\delta^{(i)}$ from the uncertainty set U, where $i = 1, ..., N$. Then the concise version of the sampled problem (2.4) can be defined as the following model (2.6):

$$\max_{x \in X} c'x, \qquad \text{s.t.} f(x, \delta^{(i)}) \leq 0, \forall i = 1...N \tag{2.6}$$

The following definition is required to develop the concept of ϵ-robustness which we will use throughout the chapter.

Definition 2.2.1 *For a given pricing policy x and a distribution Q of the uncertainty δ, define the probability of violation $V_Q(x)$ as:*

$$V_Q(x) = \mathbf{P}_Q\{\delta : f(x, \delta) > 0\}.$$

Note that the probability of violation corresponds to a measure on the actual uncertainty realization δ, which has an underlying unknown distribution Q. In other words, for the MaxMin case, given the pricing policy $x = (s, z)$, $V_Q(x)$ is the probability that the actual realization of demand gives the firm a revenue lower than z, which we computed as the worst-case revenue, or that it violates non-negativity constraints. In reality, the constraints are "naturally" enforced. The firm won't set a negative price, so if such a deviation occurs, the constraints will be enforced at a cost to the firm's revenue. Therefore, it is easy to understand any violation as an unexpected loss in revenue. We can now define the concept of ϵ-robust feasibility.

Definition 2.2.2 *We say x is ϵ-level robustly feasible (or simply ϵ-robust) if $V_Q(x) \leq \epsilon$.*

Note that the given set of scenario samples is itself a random object and it comes from the probability space of all possible sampling outcomes of size N. For a given level $\epsilon \in (0,1)$, a "good" sample is one such that the solution $x^N = (s^N, z^N)$ to the sampling based optimization model will give us an ϵ-robust solution, *i.e.*, the probability of nature giving the firm some revenue below our estimated z^N is smaller than ϵ. Define the confidence level $(1-\beta)$ as the probability of sampling a "good" set of scenario samples. Alternatively, β is known as the "risk of failure", which is the probability of drawing a "bad" sample. Our goal is to determine the relationship between the confidence level $(1-\beta)$, the robust level ϵ and the number of samples used N.

Before we introduce the main result, there is one last concept that needs to be explained. Suppose that we do not have samples obtained from the true distribution $\mathcal{Q}$, i.e. we do not have enough historical data. Instead we are given the nominal demand parameters and the uncertainty set U. We would like to be able to draw samples from another chosen distribution $\mathcal{P}$ and run the sampling based pricing model (2.6). In order to make a statement about the confidence level of the solution and the sample size, we must make an assumption about how close $\mathcal{P}$ is to the true distribution $\mathcal{Q}$.

Definition 2.2.3 ***Bounded Likelihood Ratio:*** *We say that the distribution $\mathcal{Q}$ is bounded by $\mathcal{P}$ with factor k if for every subset A of the sample space:* $\mathbf{P}_{\mathcal{Q}}(A) \leq k\mathbf{P}_{\mathcal{P}}(A)$.

In other words, the true unknown distribution $\mathcal{Q}$ does not have concentrations of mass that are unpredicted by the distribution $\mathcal{P}$ from which we draw the samples. If $k=1$ then the two distributions are the same, except for a set of probability 0, and therefore the scenario samples come from the true distribution (which is usually the case in data-driven problems). Note that the assumption above will be satisfied under a more restrictive, but perhaps more common, assumption for continuous distributions of Bounded Likelihood Ratio $\frac{d\mathbf{P}_{\mathcal{Q}}(x)}{d\mathbf{P}_{\mathcal{P}}(x)} \leq k$.

At first glance, it seems hard for a manager to pick a bound k on the likelihood ratio that would work for his uncertainty set and sampling distribution without any knowledge of the true underlying distribution. In the following Theorem 2.2.1, we show how one can derive this ratio k from the standard deviation σ, when the true uncertainty distribution belongs to the family of log-concave distributions with bounded support and is independent across products and time periods. In this case, the standard deviation of the demand, which is a very familiar statistic to most managers, should be somehow obtained by the firm. Similar theorems could also be obtained by using other statistics about the volatility of the demand. Also note that the family of log-concave distributions, as defined by distributions where the log of the density function is concave, is a rather extensive family, which includes uniform, normal, logistic, extreme-value, chi-square, chi, exponential, laplace, among others. For further reference, a deeper understanding of log-concave distributions and their properties can be seen in (Bagnoli and Bergstrom 1989) and (Bobkov 1999).

Theorem 2.2.1 *Assume that the deviations $\delta_{j,t}$ for each product $j = 1, ..., n$ and time period $t = 1, ..., T$ are independent and bounded in a box-shaped uncertainty set: $\delta_{j,t} \in [lb, ub]$. We do not know the true distribution of $\delta_{j,t}$, but we assume it belongs in the family of log-concave distributions. Furthermore, the standard deviation of $\delta_{j,t}$ is known to be σ. We randomly sample points from a uniform distribution supported over the uncertainty set, with density $\frac{1}{ub-lb}$ for each $\delta_{j,t}$.*

The true distribution has a likelihood ratio bound $k = \left[(ub - lb)\frac{e\sqrt{6}}{\sigma\sqrt{2}}\right]^{nT}$.

If the distribution of δ is symmetric, the bound will be $k = \left[(ub - lb)\frac{1}{\sigma\sqrt{2}}\right]^{nT}$.

For a proof of Theorem 2.2.1, see Appendix D. The following theorem develops a bound on the probability that the solution of the sampled problem is not ϵ-robust, i.e. probability of drawing a "bad" sample.

Theorem 2.2.2 *Given the following scenarios of uncertainty $\delta^{(1)}, ..., \delta^{(N)}$, drawn from an artificial distribution $\mathcal{P}$, which bounds the true uncertainty distribution $\mathcal{Q}$ by a factor of k (see Definition 3 and Theorem 1). Let n_x be the dimension of the strategy space X, x^N be the solution of (2.6) using the N sample points, and ϵ be the robust level parameter. Then define a "risk of failure" parameter $\beta(N, \epsilon)$ as:*

$$\beta(N, \epsilon) \doteq \binom{N}{n_x}(1 - \epsilon/k)^{N-n_x}$$

Then with probability greater than $(1 - \beta(N, \epsilon))$, the solution found using this method is ϵ-level robustly feasible, i.e.,

$$\mathbf{P}_{\mathcal{P}}((\delta^{(1)}, ..., \delta^{(N)}) : V_{\mathcal{Q}}(x^N) \leq \epsilon) \geq (1 - \beta(N, \epsilon))$$

In other words, the level $\beta(N, \epsilon)$ is a bound on the probability of getting a "bad" sample of size N for a robust level ϵ. Then $1 - \beta(N, \epsilon)$ is the confidence level that our solution is ϵ-robust. As a corollary of Theorem 2.2.2, we can obtain a direct sample size bound for a desired confidence level β and robust level ϵ.

Corollary 2.2.1 *If the sample size N follows:*

$$N \geq N(\epsilon, \beta) \doteq (2k/\epsilon)\ln(1/\beta) + 2n_x + (2n_x k/\epsilon)\ln(2k/\epsilon)$$

Then with probability greater than $1 - \beta$ the solution to the sampled problem will be ϵ-level robustly feasible.

The bound in Corollary 2.2.1 is not necessarily the tightest value of N that will satisfy Theorem 2.2.2 for a given β and ϵ. Numerically solving for N the equation $\beta = \binom{N}{n_x}(1 - \epsilon/k)^{N-n_x}$ might give a smaller sample size requirement. On the other hand, Corollary 2.2.1 offers a direct calculation and insight about the relationship between the sample size and the confidence/robust level of the sampled solution. Note that $N(\epsilon, \beta)$ goes to infinity as either ϵ or β go to zero, which is rather intuitive. On the other hand the dependence on β is of the form $\ln(1/\beta)$ which means that the confidence parameter β can be pushed down towards zero without significant impact on the number of samples required. For implementation purposes, it allows us to keep a good level of confidence β and design N based on the ϵ-level of robustness desired.

For proof of Theorem 2.2.2, see Appendix E. The proof of this theorem is similar to the proof of (Calafiore and Campi 2006) (Theorem 1), but the latter could not be directly applied in our case since we generalized the input data to allow for the random sampling approach. Instead, we had to carefully introduce the bounded likelihood ratio assumption to connect the true distribution to the sampling distribution. Corollary 2.2.1 can be derived using algebraic manipulations by solving for N the expression $\beta \leq \binom{N}{n_x}(1 - \epsilon/k)^{N-n_x}$.

Note that this result can be quite useful in practice, specially if there is a small amount of data or no data samples at all. It is common for managers to assume the demand to lie within a certain range. If they can further make an assessment of the volatility of demand (standard deviation), we can apply this result. From Theorem 2.2.1, we obtain a likelihood ratio bound k, which directly implies the number of samples required to be drawn for a given confidence and robustness level, according to Theorem 2.2.2. To the best of our knowledge, this is the first result that gives a robust sampling size bound when using an artificial sampling procedure and such limited knowledge of the true distribution (log-concavity, range and standard deviation).

In more recent literature, (Campi and Garatti 2008) developed a better bound for the case where the samples come from the true distribution, which we call the data-driven case ($k = 1$). In fact they prove that this new bound is the tightest possible bound, since it holds with equality for the special case of fully-supported problems. They also show how the new bound will always provide a smaller sample size requirement than the one developed in Theorem 2.2.2. On the other hand, we have not been able to generalize this result for the case of random scenario sampling with artificial distributions ($k > 1$), which is a much more complicated problem. In Appendix F, we present the result from Campi and Garatti (2008) and in the subsection Appendix F.1 we demonstrate numerically that this bound is actually tight for an instance of the dynamic pricing problem. As shown in (Campi and Garatti 2008), the sample size bound in Corollary 2.2.1, when $k = 1$, requires significantly more samples for the same β and ϵ levels than Corollary F.1 in the appendix.

In order to apply these confidence bounds to our pricing problem, we only need to know the likelihood bound k, which is either data provided by the firm or calculated in a similar manner as in Theorem 2.2.1, and the dimension of the strategy space n_x, in other words, the number of decision variables in the optimization model. The latter will depend on the functional form of the pricing policy chosen. For example, in the case of static pricing, the number of decision variables is the number of time periods T times the number of products n (one fixed price per product/period) plus one for the robust objective variable: $n_x = Tn + 1$. For the pricing policy suggested in (2.3), which we use in Section 3, $n_x = 1 + Tn + (T-1)nm$. For instance, in the airline case study of Section 4, we deal with a 3-period single product adjustable model, therefore $n_x = 6$.

2.3 Numerical Results

In this section we present simulation results using the models described before. All optimization problems discussed below were solved using AMPL modeling language. The experiments were made on a server using a single core of 2.8GHz Intel Xeon processor with 32GB of memory.

The first experiment, in Section 3.1, is designed to compare the out-of-sample performance of the multiple pricing models discussed in Section 2 in order to obtain managerial insights on the effects of the models on the resulting pricing policies. In Section 3.2, we observe the performance of the different models when the sampling distribution is not the same as the true underlying distribution.

More specifically, we will compare the following models: MaxMin, MinMax Regret, MaxMin Ratio, and Sample Average Approximation. A summary of all the optimization models used in this section is available in Appendix G. Note that we are introducing here the Sample Average Approximation (SAA) approach only as a benchmark to our robust models and we will explain it in detail later in this section and present the optimization model in Appendix G.

In Appendix H, we display the out-of-sample revenues between the static and the adjustable policies. This experiment illustrates why we have not dedicated much time discussing the static models in this chapter, as it is clearly dominated by the adjustable counterpart. For this reason, all other experiments in this chapter use adjustable policies. In Appendix I we demonstrate an application of the model using a network model with multiple products and resources and in Appendix J we demonstrate the application of a nonlinear demand model. The qualitative conclusions of these two cases (network and nonlinear demand) do not differ from the single product, linear demand models of Section 3.1 and 3.2. To avoid redundancy, we will not dedicate much time to these examples.

2.3.1 Using true distribution sampling: data-driven approach

In order to isolate the differences between the models, we will reduce the problem to a single product case ($n = m = 1$) and linear demand functions with additive uncertainty. The reader should note that this is used here for simplicity of notation only and the use of other demand and uncertainty models can be easily incorporated.

Define the **MaxMin** model as in the optimization problem formulated in (G.1), in Appendix G. For the regret-based models, we can solve the hindsight revenue $\Pi^*(\delta^{(i)})$ from (G.2) for each data point $\delta^{(i)}$ of the uncertainty sample by solving a simple convex optimization problem. The **MinMax Regret** is defined in (G.3) by subtracting the term $\Pi^*(\delta^{(i)})$ from the right hand side in the objective constraint (the first constraint). Similarly, the **MaxMin Ratio** is defined by dividing the right hand side by the term $\Pi^*(\delta^{(i)})$, as in the optimization problem (G.4).

Robust optimization is often criticized for generating very conservative policies. In other words, it provides good performance in the worst-case revenue (or regret/competitive ratio), but bad overall average performance. For this reason, we will compare the solution of our robust models with a more traditional stochastic optimization framework, where we want to maximize the expected value of the revenue. If we define the pricing decision as a function of the previous realized uncertainty, we have an adjustable stochastic optimization problem, which is the same as the adjustable robust models above, except that the objective approximates the expected value of the revenue by taking the average of each revenue point: $\frac{1}{N}\sum_{i=1}^{N}\left[\sum_{t=1}^{T}\sum_{j=1}^{n} p_{j,t}(a_t - b_t p_{j,t} + \delta_t^{(i)}) - \sum_{l=1}^{m} o_l {w_l}^{+} + g_l {w_l}^{-}\right]$. This solution concept is what we refer to as Sample Average Approximation (**SAA**) and we detailed the optimization model (G.5).

In Figure 2-1, we observe the out-of-sample revenue performance of our four proposed models with different sample sizes varying from 5 to 200. The revenue 'boxplots' displayed throughout this section can be explained as follows: The middle line indicates the median; The boxes contain 50% of the observed points (25% quantile to the 75% quantile); The lines extending from the boxes

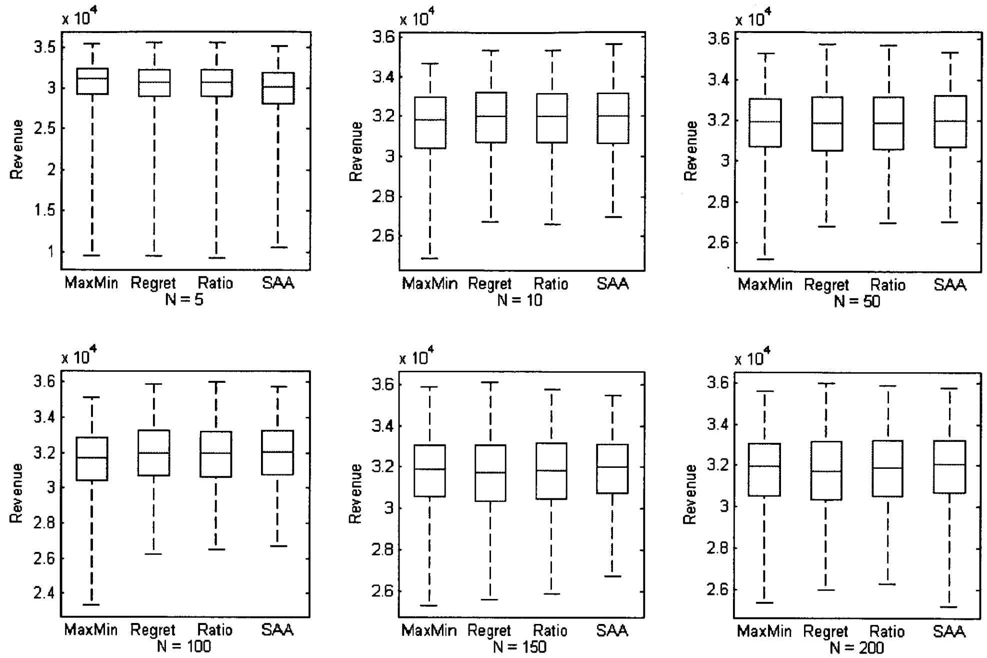

Figure 2-1: Example of out-of-sample revenue distribution for single product case. Sample sizes vary from 5 to 200 samples.

display the range of the remaining upper 25% and lower 25% of the revenue distribution. Not much can be concluded from this particular instances' revenue snapshot, but it illustrates how the revenue outcomes can be quite different between models when sample sizes are too small. The intrinsic volatility of the revenues creates too much noise for us to make precise qualitative judgements between the different models. This issue can be fixed by performing paired tests between different models, which compares the difference in revenues between models taken at each data sample. In these paired tests, the intrinsic volatility of revenues is reduced by the difference operator, which allows us to compare the models' performances.

Using the paired comparisons, we will measure the average difference in revenues and construct an empirical confidence interval in order to draw conclusions about which model performs better on the average revenue.

Another useful statistic to measure robust models is known as Conditional Value-at-Risk (CVaR), which is also known in finance as mean shortfall. Taken at a given quantile α, the CVaR_α is the expected value of a random variable conditional that the variable is below it's α quantile. This statistic is as a way to measure the tail of a distribution and determine how bad things can go when they go wrong. Since we are not looking at one particular revenue distribution, but instead the difference between two revenues, it is only natural to look at the positive tail of the distribution, when the second comparing model is at risk. In a slight abuse of notation, define for the positive side of the distribution, $\alpha > 50\%$, the CVaR_α to be the expected value of a variable conditional on it being above the α quantile.

$$\text{For } \alpha \leq 50\%, \text{ CVaR}_\alpha(X) = E\left[X | X \leq \alpha\right]$$
$$\text{For } \alpha > 50\%, \text{ CVaR}_\alpha(X) = E\left[X | X \geq \alpha\right]$$

Another important issue to be observed is how the different models behave when provided small sample sizes versus large sample size. For this reason, we will perform the experiments in the following order, starting with $N = 5$ and increasing until $N = 200$:

1. Take a sample of size N uncertainty data points;

2. Optimize for different models to obtain pricing policies;

3. Measure out-of-sample revenues for 2000 new samples, keeping track of pairwise differences in revenues;

4. Record statistics of average difference and CVaR of the difference in revenues for each pair of models;

5. Repeat from steps 1-4 for 1000 iterations to build confidence intervals for the average difference and CVaR of the difference;

6. Increase the sample size N and repeat 1-5.

In the first test we compare the revenues of the SAA and the MaxMin Ratio model. In Figure 2-2, we display a histogram of the difference of out-of-sample revenues for a particular instance of the single product problem, using $N = 200$ samples. More specifically, for each of the 2000 out-of-sample revenue outcomes generated, we calculate the difference between the revenue under SAA policy minus the revenue under MaxMin Ratio model. By calculating the statistics of this histogram, we obtain a mean revenue difference of 133 and standard deviation of 766. In fact, we observe there is an indication that the SAA obtains a small revenue advantage on a lot of cases, but can perform much worse in a few of the bad cases.

We also obtained $\text{CVaR}_{5\%} = -2233$ and the $\text{CVaR}_{95\%} = 874$. One can clearly see that the downside shortfall of using the SAA policy versus the MaxMin Ratio can be damaging in the bad cases, while in the better cases the positive upside is quite limited.

Up to this point, we have only demonstrated the revenue performances for a single illustrative

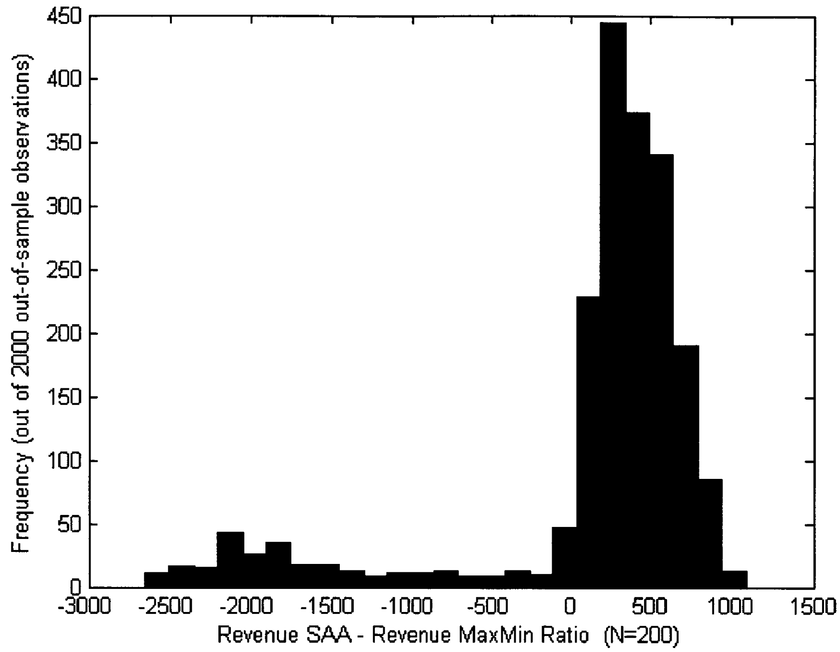

Figure 2-2: Paired comparison between SAA and MaxMin Ratio

instance of the problem. The next step is to back-up these observations by repeating the experiment multiple times and building confidence intervals for these estimates. We are also interested in finding out if these differences appear when the number of samples changes. More particularly, we want to know if the average revenue difference between the SAA and the MaxMin Ratio is significantly different from zero and if the $\mathrm{CVaR}_{5\%}$ of this difference is significantly larger than $\mathrm{CVaR}_{95\%}$ (in absolute terms).

In Figure 2-3, we plot the distribution of the average difference in the SAA revenue minus the MaxMin Ratio revenue. More specifically, for each sample size N, we sampled 1000 times the set of in-sample data of size N. Each one of these data sets is what we call an iteration of the experiment. For each of these iterations, we solved the pricing policies and evaluated the revenues for another 2000 out-of-sample data points and evaluated the difference in revenues between the policies for each of these points. This average difference obtained over those 2000 points correspond to one particular iteration. The histograms in Figure 2-3 show the distribution of this average difference obtained over all 1000 iterations of this experiment.

One can initially observe that for $N = 200$ the average difference lies clearly on the positive side of the histogram, showing that the SAA performs better on average than the MaxMin Ratio.

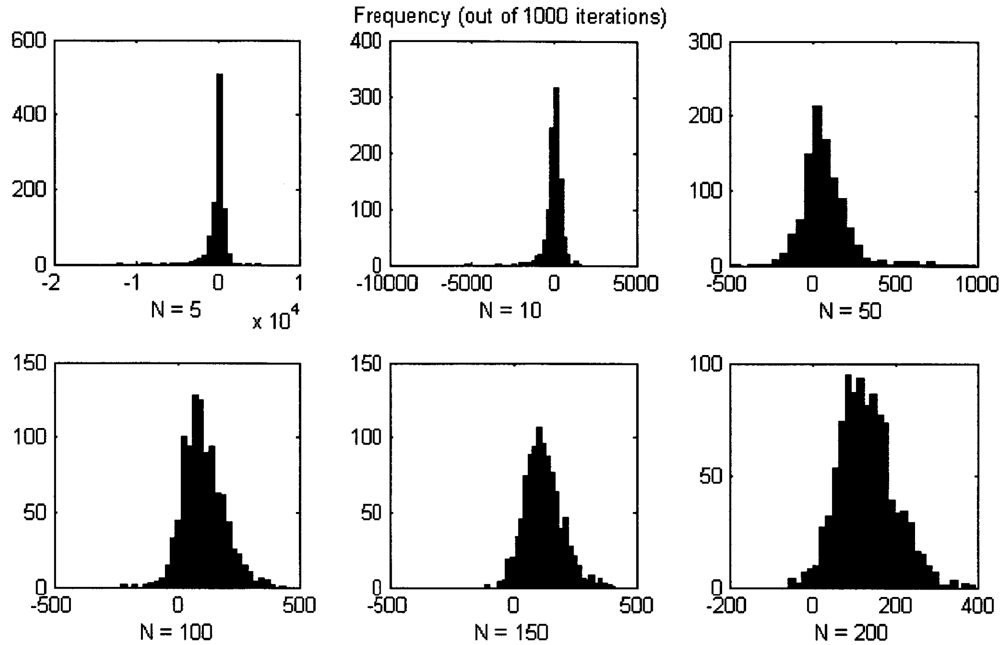

Figure 2-3: Average of the revenue difference: SAA - MaxMin Ratio

When N is reduced, the histograms move closer to zero. When $N = 5$, the average difference seems to lie somewhere around zero, with a bias towards the negative side, implying that when the sample size is small the MaxMin Ratio will often perform better on average. These observations are confirmed by Table 2.1, where we display the mean average difference in the second column, with a 90% confidence interval (third column) calculated from the percentile of the empirical distribution formed over 1000 iterations. When $N = 5$ and $N = 10$, the average difference in revenues between the SAA and the Maxmin Ratio is in fact negative on the average of all iterations of the experiment, although not statistically significant, since the confidence interval contains zero. In other words, the confidence intervals suggest that there is not enough evidence to say the average difference in revenues is either positive or negative. On the other hand, when $N = 150$ or $N = 200$, we can see that with a 90% confidence level, the average difference is positive, meaning the SAA policy will provide better average revenues than the MaxMin Ratio.

In Table 2.1, we also present the lower 5% and upper 95% tails of the revenue difference distribution. We can observe in this case that when the sample size is small, the two sides of the distribution behave rather similarly. As N increases, we can observe the lower tail of the revenue difference distribution is not very much affected and the length of the confidence interval seems to

	Average		5% CVaR		95% CVaR	
N	Mean	90% CI	Mean	90% CI	Mean	90% CI
5	-79	[-1763,909]	-2686	[-11849,-198]	2469	[676,6437]
10	-29	[-837,549]	-2345	[-7195,-348]	2113	[609,4115]
50	78	[-115,324]	-2260	[-4798,-382]	962	[300,1953]
100	103	[-17,252]	-2355	[-4124,-735]	773	[310,1372]
150	122	[8,253]	-2378	[-3758,-992]	721	[344,1174]
200	130	[31,249]	-2418	[-3606,-1257]	691	[363,1068]

Table 2.1: Average and CVaR of the difference: Revenue SAA - Revenue MaxMin Ratio

decrease around similar mean values of $CVaR_{5\%}$. The upper tail of the distribution, on the other hand, seems to shrink as we increase N. When $N = 200$, the upside benefits of using SAA rather than MaxMin Ratio will be limited by an average of 691 on the 5% better cases, while the 5% worst cases will fall short by an average of -2418. Moreover, with a 90% confidence level, we observe that the $CVaR_{5\%}$ will be larger in absolute value than the $CVaR_{95\%}$, since the confidence intervals do not intersect. In other words, we have a statistically significant evidence that the SAA revenues will do much worse than the MaxMin Ratio in the bad cases, while not so much better in the better cases.

To summarize the paired comparison between the SAA and the MaxMin Ratio, we observed that in small sample sizes, the MaxMin Ratio seems to perform both better on average and in the worst cases of the revenue distribution, but the statistical testing for these results can't either confirm or deny this conclusion. For large sample sizes, the results are more clear: with 90% confidence level, the SAA will obtain a better average revenue than the MaxMin Ratio, but will pay a heavy penalty in the worst revenue cases. This last conclusion can be viewed as intuitive, since the SAA model tries to maximize average revenues, while the robust models like the MaxMin Ratio try to protect against bad revenue scenarios. On the other hand, the better performance of the robust model over the SAA using small sample sizes is not intuitive and understanding this behavior leads to interesting directions for future research.

Before proceeding to the next experiments among the robust models, we should make a note that similar paired tests of the the SAA with the other robust models (MaxMin and MinMax Regret) were also performed and produced similar results as the ones displayed above using the MaxMin Ratio. In the interest of conciseness, we decided not to display them here as they would not add any additional insight.

In the following experiments we will perform paired comparisons between the different robust models. In Figure 2-4, we display the histogram of the paired comparison between the revenues of the MaxMin Ratio and the MaxMin models. This histogram displays the distribution over 2000 out-of-sample data points in the difference in revenues for one particular iteration of the experiment. Although the distribution seems skewed towards the negative side, the average is actually positive, due to the long tail in the right-hand side. We would like to know for multiple iterations of the experiment if this average remains positive and if the tail of the distribution is much larger on the positive side, as displayed in this histogram. Moreover, we would like to compare these results for multiple initial sample sizes.

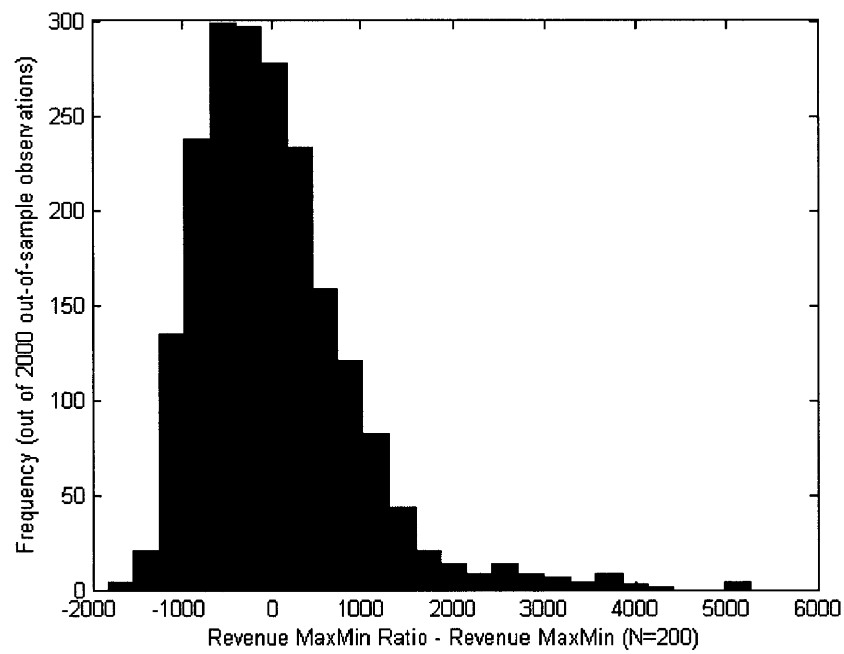

Figure 2-4: Paired comparison between the MaxMin Ratio and the MaxMin

	Average		5% CVaR		95% CVaR	
N	Mean	90% CI	Mean	90% CI	Mean	90% CI
5	577	[-322,2454]	-1369	[-4072,-245]	3912	[997,10884]
10	341	[-202,1126]	-1178	[-2847,-400]	3382	[1556,7367]
50	80	[-266,472]	-1149	[-1924,-614]	3267	[2014,5132]
100	50	[-200,381]	-1125	[-1619,-655]	3025	[1914,4116]
150	21	[-206,310]	-1109	[-1522,-687]	3031	[1981,4079]
200	23	[-204,289]	-1107	[-1481,-713]	2934	[1848,3925]

Table 2.2: Average and CVaR of the difference: Revenue MaxMin Ratio - Revenue MaxMin

The results of Table 2.2 show that the MaxMin Ratio does have a consistently better average performance than the MaxMin model for all sample sizes tested, but it is a small margin and the confidence intervals cannot support or reject this claim. Also, we can see that the MaxMin Ratio is more robust than the MaxMin, as the 5% best cases for the MaxMin Ratio (95% CVaR column) are significantly better than the 5% best cases of the MaxMin (5% CVaR column). Note that for $N \geq 50$ the confidence intervals for the 95% and 5% CVaR do not intersect, suggesting that with a 90% confidence level, the MaxMin Ratio is in fact more robust than the MaxMin under this 5% CVaR measure.

In summary, the MaxMin Ratio model outperforms the MaxMin in all areas. It displayed better average revenues and shortfall revenues (5% CVaR) than the MaxMin for all sample sizes. It is rather counter intuitive that the MaxMin Ratio model can be at the same time more robust and have better average performance. This could be a result of the sampling nature of these models. More specifically, the worst case revenues for the pricing problem are usually located in the borders of the uncertainty set, which is a hard place to estimate from sampling. When using regret based models, like the MaxMin Ratio, we adjust the revenues by the hindsight revenues. This shift in the position of the worst-case deviations of demand can make it easier for a sampling procedure to work.

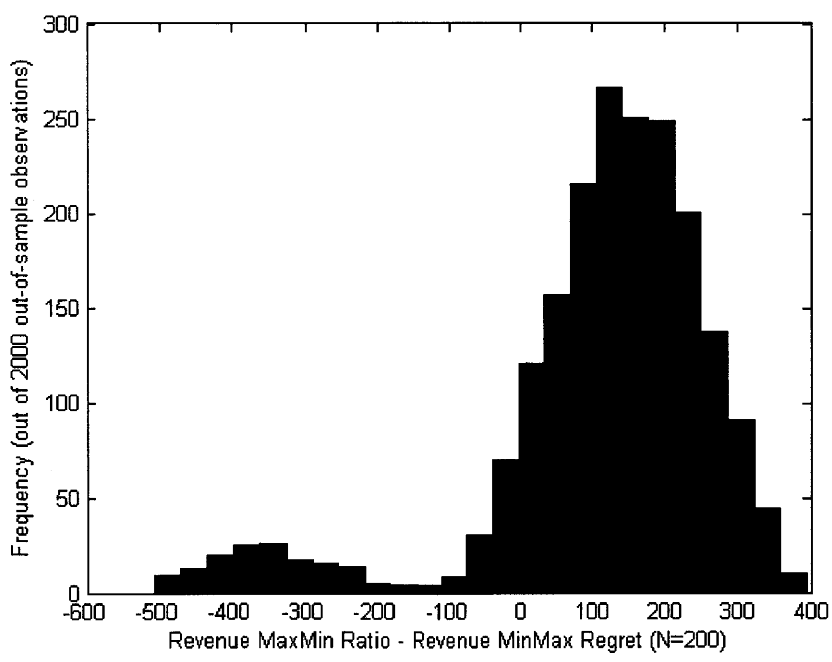

Figure 2-5: Paired comparison between the MaxMin Ratio and the MinMax Regret

In the next experiment, we compare the two regret-based robust models: MaxMin Ratio and

MinMax Regret. Figure 2-5 displays the difference distribution in revenues from the MaxMin Ratio and MinMax Regret for one particular iteration of the experiment. We can see that there is a strong bias towards the positive side, showing that the MaxMin Ratio performs better on average than the MinMax Regret, but there is a small tail on the left with much higher values. When observing the difference in revenues of these two models for 1000 iterations, we obtain the following Table 2.3 of statistics.

	Average		5% CVaR		95% CVaR	
N	Mean	90% CI	Mean	90% CI	Mean	90% CI
5	-16	[-97,32]	-232	[-796,-25]	106	[10,321]
10	-13	[-90,41]	-341	[-2847,-91]	138	[29,350]
50	15	[-48,100]	-483	[-1924,-231]	296	[76,415]
100	54	[-6,140]	-458	[-1619,-267]	240	[94,386]
150	68	[9,159]	-456	[-1522,-292]	260	[102,405]
200	83	[19,163]	-462	[-1481,-309]	285	[115,409]

Table 2.3: Average and CVaR of the difference: Revenue MaxMin Ratio - Revenue MinMax Regret

We should first note that the scale of the numbers in Table 2.3 is much smaller than in the previous paired tests, as we would expect the two regret models to share more common traits than the other models we compared with the MaxMin Ratio.

The average difference between the MaxMin Ratio and the MinMax Regret is initially negative, for small sample sizes. It later becomes positive for large samples and we notice the 90% confidence interval for this average difference is strictly positive at $N = 200$. In other words, for large samples the MaxMin Ratio performs better on average revenues than the MinMax Regret. On the other hand, the 5% CVaR is approximately twice as large (in absolute value) as the 95% CVaR, although without the statistical validation from the confidence intervals. This suggests the MinMax Regret could be the more robust model between the two regret-based models.

2.3.2 Using random scenario sampling

In this experiment, we observe the behavior of the scenario based models when the sampling distribution is not the same as the true underlying distribution. Consider the following parameters for the model: $T = 2$, $n = m = 1$, $a = [200, 200]$, $b = [0.5, 0.5]$, $C = 120$, $o = 1000$, $s = 10$. For simplicity, we will also consider a box-type uncertainty set, where $U = \{\delta : |\delta_t| \leq 15, \forall t = 1, 2\}$. In fact these are the same parameters as in the previous section. In contrast with the previous section, in this experiment we will use a truncated normal distribution with a biased average as the true underlying distribution. More specifically, as the uncertainty set is defined with a box-shaped set $\delta_t \in [-15, 15]$, the underlying true distribution is sampled from a normal with mean 5 and standard deviation 7.5 and truncated over the range $[-15, 15]$. On the other hand, we will assume this is not known to the firm and we will be sampling scenarios from a uniform distribution over the uncertainty set. These uniformly sampled points will then be used to run our sampling based pricing models.

	Average		5% CVaR		95% CVaR	
N	Mean	90% CI	Mean	90% CI	Mean	90% CI
5	-165	[-3753,1649]	-3620	[-15124,-318]	2997	[903,6895]
10	-6	[-1726,925]	-2400	[-8952,-401]	2720	[760,5236]
50	-88	[-472,112]	-1462	[-3748,-262]	1195	[280,2588]
100	-31	[-192,59]	-941	[-2848,-200]	961	[285,2117]
150	-21	[-118,36]	-692	[-2008,-203]	827	[177,1849]
200	-29	[-158,19]	-872	[-2370,-195]	670	[200,1474]

Table 2.4: Using artificial sampling distribution: Average and CVaR of the difference Revenue SAA - Revenue MaxMin Ratio

Table 2.4 displays the paired comparison of the SAA revenues minus the MaxMin Ratio revenues, in a similar fashion as the comparisons performed in Section 3.2. Note that the average of the difference in revenues favors the negative side, i.e. the side where MaxMin Ratio solution outperforms the SAA, suggesting that this robust solution will perform better than the SAA even on the average. This result, although counter intuitive, appears consistent for both small and large sample sizes, with the confidence intervals moving away from the positive side as the sample size increases. The CVaR comparisons, on the other hand, are inconclusive.

	Average		5% CVaR		95% CVaR	
N	Mean	90% CI	Mean	90% CI	Mean	90% CI
5	875	[-417,2704]	-1326	[-4438,-227]	4064	[1156,10306]
10	701	[22,1601]	-1385	[-3824,-447]	3603	[2020,6349]
50	580	[247,1077]	-1191	[-2603,-703]	3052	[1722,4614]
100	644	[355,1089]	-1001	[-1553,-737]	2989	[1727,4718]
150	672	[361,1098]	-962	[-1194,-724]	3015	[1764,4753]
200	659	[389,1082]	-987	[-1338,-798]	2995	[1702,4783]

Table 2.5: Using artificial sampling distribution: Average and CVaR of the difference Revenue MaxMin Ratio - Revenue MaxMin

Table 2.5 displays the results of the paired comparison of the revenues in the MaxMin Ratio model minus the MaxMin model. We can clearly see that in this case, where we used an artificial sampling distribution, the MaxMin Ratio performs both better on average and in the worst-cases.

The confidence interval of the average difference in revenues is positive even for the small sample size of $N = 10$, distancing away from zero as N increases. As for the CVaR comparision, note that for $N \geq 100$ the 5% better cases of the MaxMin Ratio are on average better by 2989, while the 5% better cases for the MaxMin are only better by an average of 1001. Note that the confidence intervals around these values do not overlap, strongly suggesting that the MaxMin Ratio is more robust than the MaxMin. Since the MaxMin is clearly dominated by the MaxMin Ratio, we will not expose here the comparison of the MaxMin with the other models.

	Average		5% CVaR		95% CVaR	
N	Mean	90% CI	Mean	90% CI	Mean	90% CI
5	-33	[-157,57]	-312	[-926,-50]	147	[15,479]
10	-35	[-133,51]	-434	[-1205,-155]	227	[47,582]
50	2	[-86,82]	-532	[-931,-376]	296	[124,419]
100	22	[-48,98]	-526	[-801,-415]	327	[232,417]
150	31	[-26,92]	-518	[-708,-446]	349	[266,409]
200	31	[-17,93]	-533	[-682,-434]	390	[308,410]

Table 2.6: Using artificial sampling distribution: Average and CVaR of the difference Revenue MaxMin Ratio - Revenue MinMax Regret

In the experiment of Table 2.6, we display the results for the revenues of the MaxMin Ratio model minus the MinMax Regret model. Note that the MinMax Regret model appears to obtain a better average difference than the MaxMin Ratio for sample sizes up to 10. As the sample size increases to 50 and above, the MaxMin Ratio performs better on average. Besides this initial evidence, either one of these results cannot be fully supported by the confidence intervals, which are containing zero in all instances of sample size. On the other hand, we notice that theses confidence intervals are moving away from the negative side as the sample size increases, which agrees with our conclusion.

On the other hand, the CVaR comparison indicates that the MinMax Regret is in general more robust than the MaxMin Ratio, as the 5% better cases of the MinMax Regret (negative side) are on average larger than the 5% better cases of the MaxMin Ratio. This result is in fact supported by the confidence intervals, which do not overlap (in absolute value) for sample sizes of $N = 150$ and $N = 200$.

We should mention that comparing the MinMax Regret model with the SAA led to the same results as the comparison between MaxMin Ratio and the SAA and therefore these results are not displayed here as they do not add any additional insight.

To summarize the experiments performed using the artificial sampling distribution, where the samples provided come from a uniform distribution, while the true distribution is a biased truncated normal, we observed that the SAA usually performs worse than the MaxMin Ratio on the average revenue. The MaxMin Ratio model clearly outperforms the MaxMin model, both in average and worst-case differences in revenues. When comparing the two regret type models, we have evidence that the MaxMin Ratio should perform better on average than MinMax Regret for large enough sample sizes, but the MinMax Regret is clearly more robust than the MaxMin Ratio.

2.4 Testing the model with real airline data

In this section we perform a study using real airline data, collected from a flight booking website. Tracking a specific flight with capacity of 128 seats over a period of two months for 7 different departures dates, we acquired the posted selling price and the inventory remaining for that given flight. We assumed there is no overbooking allowed and salvage value is zero, since we don't have any information about these parameters. Note that for each departure date, we have the same product being sold over and over again. To make the data more consistent, we aggregated this data into 3 periods of time: 6 or more weeks prior to departure; 5-3 weeks; and 1-2 weeks. For each of these 3 periods, we averaged the observed price and derived from the inventory data the demand observed by the company during the respective period. To summarize, we obtained 7 observations of price and demand for a 3 period horizon. Figure 2-6 displays the data points observed and a possible model estimation using a linear demand model. By multiplying price times demand for each period in each flight we obtain an estimate for the flight's revenue of $\$29,812 \pm 3,791$ (notation for Average $\pm$ Standard Deviation). Note here that since we have few data points, it does not make sense to talk about the CVaR measure that was previously used to measure robustness in Sections 3.1 and 3.2. Instead, the concept of sample standard deviation will be used to express volatility of the revenue outcomes and therefore some form of robustness.

Using these 7 sample points, we perform a leave-one-out procedure to test the performance of our pricing models, where for each iteration we select one of the samples to be left out for testing and the remaining 6 points are used for training. During the training phase, a linear demand-price relation is estimated for the 6 training points and the residuals will be the deviation samples $\delta^{(i)}$, used to find the respective pricing strategies in our sample-based models. Using the test point that was left out, we obtain the difference between the point and the line estimated with the 6 training points, defining the out-of-sample deviation. It is important to emphasize that the estimation of the

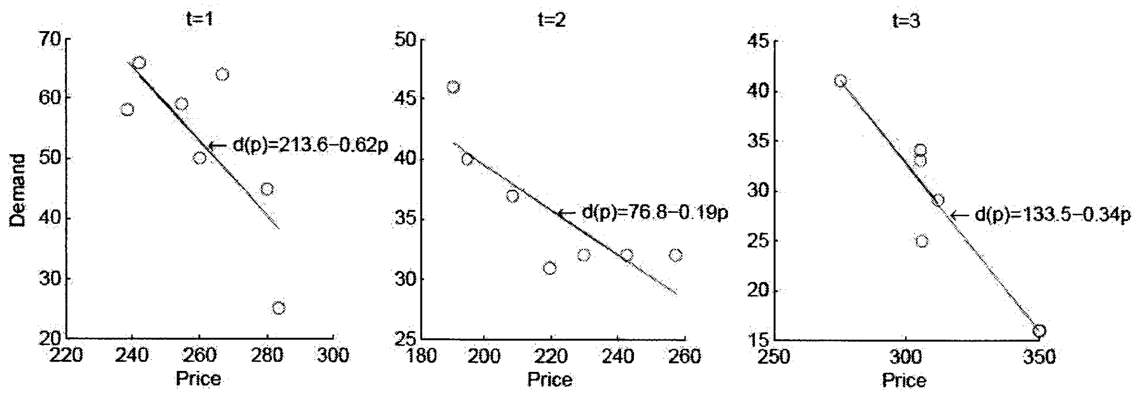

Figure 2-6: Price x Demand Estimation

Model	Average	Standard Deviation
Actual Revenue	29,812	3,791
MaxMin	31,774	1,740
MinMax Regret	31,891	1,721
MaxMin Ratio	31,906	1,713
Sample Average Approximation	31,848	1,984

Table 2.7: Airline case: Revenue performance for each pricing model

demand curve needs to be redone every time we leave a point out of the sample in order to obtain the independence between the testing sample and the estimated demand function. After solving the pricing models for each demand function, we can measure the out-of-sample revenue for each pricing strategy using the left-out demand point. By repeating this procedure 7 times, we obtain 7 out-of-sample revenue results, which are summarized in Table 2.7.

The MaxMin Ratio outperformed all other models, demonstrating the highest out-of-sample average revenue and lowest standard deviation. Compared to the actual revenue of the airline, assuming there are no external effects, the MaxMin Ratio obtained a 7% increase in the average revenue and a 55% decrease in standard deviation.

This study demonstrates, in a small scale problem, the procedure for applying our models on a real-world situation. On the other hand, there are some aspects of the airline industry that were not captured in this data set and are not considered by our models. When airlines make their pricing decisions, they take into consideration, for example, network effects between different flights and competitors' prices, which should be incorporated in future extensions of this model. A few other facts that we overlook is the aggregation done over the periods of time, due to the lack of data, or the fact that not every seat is sold at the same price at any given time, given that some people choose to buy their tickets with or without restrictions and connections. Therefore, these price-demand points are not an exact depiction of customers' response to price, but it is the best approximation we can achieve given the information that was collected. For these reasons, we cannot claim a 7% actual increase for the revenue of this flight, but instead this study suggests that there may be room for improvement in the current pricing techniques used by the airline. For instance, the methodology proposed here can be an easy improvement over the deterministic models that are used to approximate revenue-to-go functions in revenue management heuristics.

2.5 Conclusions

In this chapter, we developed a framework to solve the network dynamic pricing problem using a sampling based approach to find approximate-closed-loop robust pricing policies.

Throughout the chapter, we have shown how the adjustable robust pricing problem can easily turn into a non-convex problem, which is intractable to solve using traditional robust optimization techniques. Instead we proposed the sampling based approach where we can easily solve the sampled problem as a convex optimization model. One question that arises from this approach is how many samples do we need to have a performance guarantee on our scenario-based solution. To answer this question, we introduced the notion of ϵ-robustness and found the sample size needed to achieve an ϵ-robust solution with some confidence level. Moreover, we were able to extend this result from the data-driven optimization framework to a random scenario sampling approach using a relatively mild assumption on the underlying distribution.

Notice how the random scenario sampling bound is a powerful concept, given that we are now able to use a data-driven methodology to solve a robust problem. Data-driven and robust optimization have been generally considered two separate fields, since they use very different initial assumptions of information about the problem. In this research, we are building a bridge between data-driven optimization and robust optimization, which has not been explored very much. We started with an intractable robust model and illustrated how to solve a close approximation to the robust solution without any previous data, only with a likelihood bound on the underlying distribution.

Besides the theoretical guarantees mentioned above, we explored numerically the revenue performance of the different models. Some of the key managerial insights of our numerical experiments are:

- The SAA has a risk-neutral perspective, therefore maximizing average revenues, regardless of the final revenue distribution. It gets the best average revenue performance if the number of samples is large enough and if they are reliably coming from the true uncertainty distribution. On the other hand, the revenue outcome is subject to the largest amount of variation, with a large shortfall in the bad revenue cases when compared to the robust models.

- The MinMax Regret model presented the more robust revenue performance when compared with other models. The MaxMin Ratio model strikes a balance between the conservativeness of the robust MinMax Regret and the aggressiveness of the SAA.

- All the Robust models, namely the MaxMin, MinMax Regret and MaxMin Ratio, will perform just as well as the SAA in average revenues for small sample sizes. For large sample sizes, the robust models will have smaller average revenues, but also with smaller shortfalls for the extreme revenue cases, in other words good for risk-averse firms. This advantage of having more stable revenue distributions are present even for small sample sizes.

- The MaxMin model tends to be too conservative. It is usually dominated by the MaxMin Ratio model both in worst-case revenue performance and in average revenues.

- The MaxMin Ratio and MinMax Regret have the most consistent performance when the data set provided does not come from the true underlying distribution of the demand uncertainty. As in the previous case, where the sampling distribution was the true distribution, the MinMax Regret appears to be more conservative than the MaxMin Ratio.

Note that the methodological framework presented in this chapter is quite general and is not limited to dynamic pricing. For instance, (Perakis 2009) developed a dynamic traffic assignment model using a very similar framework. Because of the practicality of these models, we believe these sampling based optimization models can easily be applied to many more application areas.

It would also be interesting to further research the Sample Average Approximation approach, by possibly finding a theoretical performance guarantee based on the number of samples used and comparing this with the bounds for the robust framework. We find that a specially promising direction of research will be to better understand the differences in behavior between the robust and the SAA models when using small sample sizes. Also, more research can be done on showing under which conditions can the affine policies achieve the optimal value of the fully closed-loop solution.

2.6 Appendix

2.6.1 Notation.

We summarize here the notation that is used in throughout the chapter:

C : Vector of initial inventories for each resource.

T : Length of the time horizon.

δ : Vector of the realization of the demand uncertainty, containing one component for each product and time.

U : Uncertainty set, $\delta \in U$.

s : Vector of pricing decisions.

S : Strategy space, independent of the demand uncertainty.

$p_{j,t}(s, \delta)$: Price of product j at time t as a function of the pricing strategy and the uncertainty realized up to time t.

$\tilde{d}_{j,t}(s, \delta)$: Realized demand of product j at time t, given the pricing strategy and uncertainty.

o : Vector of overbooking penalties per unit of resource.

g : Vector of salvage values per unit of resource.

$\tilde{w}(s, \delta)$: Vector function that maps the pricing strategy and the realized uncertainty into overbooked/left-over inventory.

w : Vector of slack variables to account for overbooked/left-over inventory. At optimality, $w_l = \tilde{w}_l(y, \delta)$, while w^+ is the amount of overbooked units and w^- is the left-over inventory.

$\Pi(s, \delta)$: Actual revenue given a pricing policy and realized uncertainty.

$\Pi^*(\delta)$: Perfect hindsight revenue, given some uncertainty vector.

$h^{obj}(s, \delta)$: The objective function for any given pricing strategy s, uncertainty δ.
obj can be one of three objectives: MaxMin, Regret or Ratio
$h^{MaxMin}(s, \delta) = \Pi(s, \delta)$
$h^{Regret}(s, \delta) = \Pi(s, \delta) - \Pi^*(\delta)$
$h^{Ratio}(s, \delta) = \Pi(s, \delta)/\Pi^*(\delta)$

z : Robust objective variable.

$\delta^{(i)}$: i^{th} sampled realization of the demand uncertainty

N : Number of scenario samples, as in $\delta^{(1)}, ..., \delta^{(N)}$

x : Set of all decision variables, which is the pricing strategy s and the robust objective z.

n_x : Dimension of the strategy space.

$f(x, \delta)$: Short notation for all constraints in the pricing model.

$\mathcal{Q}$: True underlying distribution of the uncertainty with support in U.

$\mathcal{P}$: Sampling distribution, from which we obtain scenarios of the uncertainty δ.

k : Bounded likelihood ratio factor between the sampling distribution $\mathcal{P}$ and the true distribution $\mathcal{Q}$.

$V_{\mathcal{Q}}(x)$: Probability of violation, i.e., probability under the true distribution $\mathcal{Q}$, that the realized uncertainty will violate some constraints given a policy x.

ϵ : Robust level.

β : Risk of failure, or alternatively $(1-\beta)$ is the confidence level.

2.6.2 Illustration of the non-convexity issue.

Consider MaxMin robust problem with 2 time periods, single resource, single product, without salvage value or overbooking, a linear demand function and a linear pricing policy, the function $h(u, v, \delta_1, \delta_2)$ for any fixed u and v is:

$$h(\delta_1, \delta_2) = u_1(a_1 - b_1u_1 + \delta_1) + (u_2 + v_{2,1}\delta_1)(a_2 - b_2(u_2 + v_{2,1}\delta_1) + \delta_2)$$

The eigenvalues of the Hessian matrix of this function are $(-b_2v_{2,1} + \sqrt{{b_2}^2{v_{2,1}}^2 + ^2})v_{2,1}$, $(-b_2v_{2,1} - \sqrt{{b_2}^2{v_{2,1}}^2 + ^2})v_{2,1}$. These will always have one positive and one negative eigenvalue. Therefore the function will neither be convex nor concave.

2.6.3 Existence of an optimal solution.

In this section we prove there exists an optimal solution to (5). Note that the function $f(s, \delta)$ is continuous with respect to s and δ from the initial assumption on the demand function, a property that is not modified by introducing the deviation component to the function. We first need to show that $\bar{f}(s) = \min_{\delta \in U} f(s, \delta)$ is also a continuous function of s. For some $\Delta > 0$:

$$\bar{f}(s + \Delta) - \bar{f}(s) = \min_{\delta \in U} f(s + \Delta, \delta) - \min_{\delta \in U} f(s, \delta)$$

Define the deviations δ_1, δ_2 the deviations that attain the minimum at the two optimization problems above, $\delta_1 = argmin_{\delta \in U} f(s+\Delta, \delta)$, $\delta_2 = argmin_{\delta \in U} f(s, \delta)$. These give us upper and lower bounds:

$$f(s+\Delta, \delta_1) - f(s, \delta_1) \le \bar{f}(s+\Delta) - \bar{f}(s) \le f(s+\Delta, \delta_2) - f(s, \delta_2)$$

By letting $\Delta \to 0$, both sides will go to 0. It follows that:

$$\lim_{\Delta \to 0} |\bar{f}(s+\Delta) - \bar{f}(s)| = 0$$

Which implies that $\bar{f}(s)$ is a continuous function. Since S is a compact set, we use Weierstrass Theorem to guarantee existence of solution to (5).

2.6.4 Proof of Theorem 1.

Before proving Theorem 1, we first introduce Lemma D.1 below that bounds the density of the mode of a log-concave distribution.

Lemma 2.6.1 *Let* f *be the density of a one-dimensional random variable* X *with standard deviation* σ *and mode* M.
(a) If f *is log-concave, then* $\mathrm{f}(M) \le \frac{e^{\sqrt{6}}}{\sigma\sqrt{2}}$.
(b) If f *is log-concave and symmetric, then* $\mathrm{f}(M) \le \frac{1}{\sigma\sqrt{2}}$.

Proof. Let F be the cumulative density of the random variable X and m be the median of X. Define the isoperimetric constant as $Is = \inf_{\mathrm{x}} \frac{\mathrm{f(x)}}{\min\{\mathrm{F(x)},(1-\mathrm{F(x)})\}}$, see Bobkov (1999) for more details about the properties of the isoperimetric constant.

For any $\mathrm{x} \le m$, we have that $\mathrm{F(x)} \le 1-\mathrm{F(x)}$. Therefore, $Is \le \frac{\mathrm{f(x)}}{\mathrm{F(x)}}$ and from the log-concavity of f, we know that the reversed hazard rate $\frac{\mathrm{f(x)}}{\mathrm{F(x)}}$ is decreasing, then the minimum will be achieved at the median. Similarly for the other side, for any $\mathrm{x} \ge m$, we have that $\mathrm{F(x)} \ge 1-\mathrm{F(x)}$. Therefore, $Is \le \frac{\mathrm{f(x)}}{1-\mathrm{F(x)}}$ and from the log-concavity of f, we have an increasing failure rate $\frac{\mathrm{f(x)}}{1-\mathrm{F(x)}}$. Again, the minimum will be achieved at the median. Since at the median $\mathrm{F}(m) = 1-\mathrm{F}(m) = 1/2$, then $Is = \frac{\mathrm{f}(m)}{\mathrm{F}(m)} = 2\mathrm{f}(m)$.

From Bobkov (1999), Proposition 4.1, we have that $Is^2 \le \frac{2}{\sigma^2}$. For symmetric distributions the median is equal to the mode, which directly implies the result for the symmetric case of Lemma

2.6.1(b):

$$(2\mathrm{f}(M))^2 = (2\mathrm{f}(m))^2 \leq \frac{2}{\sigma^2} \Rightarrow \mathrm{f}(M) \leq \frac{1}{\sigma\sqrt{2}}$$

For non-symmetric distributions, Lemma 2.6.1(a), we must find a relation between $\mathrm{f}(M)$ and $\mathrm{f}(m)$. Note that since $\log \mathrm{f}(.)$ is a concave function, the first order approximation will always upper-bound the original function (as seen in Lemma 3.3 of Levi et al. (2010)). Therefore:

$$\log \mathrm{f}(M) \leq \log \mathrm{f}(m) + \frac{\mathrm{f}'(m)}{\mathrm{f}(m)}(M - m)$$

For non-differentiable distributions, note this relation is true for all sub-gradients $\mathrm{f}'(m)$ at m. From the decreasing reversed hazard rate property of log-concave distributions, we have that: $\frac{d}{dx}\frac{\mathrm{f(x)}}{\mathrm{F(x)}} \leq 0$, which implies that $\frac{\mathrm{f'(x)}}{\mathrm{f(x)}} \leq \frac{\mathrm{f(x)}}{\mathrm{F(x)}}$. When $M \geq m$, we get:

$$\log \mathrm{f}(M) \leq \log \mathrm{f}(m) + \frac{\mathrm{f}(m)}{\mathrm{F}(m)}(M - m)$$

From the increasing failure rate property of log-concave distributions, we have that $\frac{d}{dx}\frac{\mathrm{f(x)}}{1-\mathrm{F(x)}} \geq 0$, which implies $\frac{\mathrm{f'(x)}}{\mathrm{f(x)}} \geq -\frac{\mathrm{f(x)}}{1-\mathrm{F(x)}}$. Then for the case $M < m$, we have:

$$-\log \mathrm{f}(M) \geq -\log \mathrm{f}(m) + \frac{\mathrm{f}'(m)}{\mathrm{f}(m)}(m - M) \geq -\log \mathrm{f}(m) - \frac{\mathrm{f}(m)}{1-\mathrm{F}(m)}(m - M)$$

Combining the relations above and the fact that $\mathrm{F}(m) = 1 - \mathrm{F}(m) = 1/2$, we obtain for either $M \geq m$ or $M < m$:

$$\log \mathrm{f}(M) \leq \log \mathrm{f}(m) + 2\mathrm{f}(m)|M - m|$$

Since any log-concave distribution is unimodal, we can use the relations in Basu and DasGupta (1992), Corollary 4, to bound the distance between the mode and the median: $|M - m| \leq \sigma\sqrt{3}$. Therefore, $\log \mathrm{f}(M) \leq \log \mathrm{f}(m) + 2\mathrm{f}(m)\sigma\sqrt{3}$. Since $(2\mathrm{f}(m))^2 \leq \frac{2}{\sigma^2}$, we have that $\mathrm{f}(m) \leq \frac{1}{\sigma\sqrt{2}}$, which implies $\log \mathrm{f}(M) \leq \log\frac{1}{\sigma\sqrt{2}} + 2\frac{1}{\sigma\sqrt{2}}\sigma\sqrt{3} = \log\frac{1}{\sigma\sqrt{2}} + \sqrt{6}$. Taking the exponential on both sides we get $\mathrm{f}(M) \leq \frac{e^{\sqrt{6}}}{\sigma\sqrt{2}}$, which concludes the proof of Lemma 2.6.1(a). □

Using Lemma 2.6.1, we can directly obtain the likelihood ratio bound k on Theorem 1 by using a bound on the density of the mode $f(M)$ for log-concave distributions. Note that the uniform sampling distribution has a constant density value of $1/(ub - lb)$ for each product $j = 1, ..., n$ and time period $t = 1, ..., T$, which implies a ratio bound of $\left[(ub - lb)\frac{e^{\sqrt{6}}}{\sigma\sqrt{2}}\right]$ for each product and time.

Using the independence property, we can just multiply this bound nT times to get the desired result. The symmetric case follows from the same argument.

2.6.5 Proof of Theorem 2.

Define x^N as the solution of the scenario based problem for a sample of size N. The probability of violation under the true uncertainty measure $\mathcal{Q}$ is defined as $V_{\mathcal{Q}}(x^N) \doteq \mathbf{P}_{\mathcal{Q}}\{\delta \in \Delta : f(x^N, \epsilon) > 0\}$.

Define B as the set of "bad" sampling outcomes that will cause the probability of violation of the sampled solution to be greater than the ϵ robust level: $B \doteq \{(\delta^{(1)}, ..., \delta^{(N)}) : V_{\mathcal{Q}}(x^N) > \epsilon\}$. The goal of this proof is to bound the size of B under the sampling measure $\mathcal{P}$, *i.e.*, $\mathbf{P}_{\mathcal{P}}(B)$.

Let n_x be the dimension of the strategy space and let I be a subset of indexes of size n_x from the full set of scenario indexes $\{1, ..., N\}$. Let $\mathcal{I}$ be the set of all possible choices for I, which contains $\binom{N}{n_x}$ possible subsets. Let x^I be the optimal solution to the model using only those I scenarios. Define $B_I \doteq \{(\delta^{(1)}, ..., \delta^{(N)}) : V_{\mathcal{Q}}(x^I) > \epsilon\}$. Also based on this subset I, define the set of all possible sampling outcomes that will result in the same optimal solution as the solution using only the I scenarios: $\Delta_I^N \doteq \{(\delta^{(1)}, ..., \delta^{(N)}) : x^I = x^N\}$.

From Calafiore and Campi (2006), we have that $B = \bigcup_{I \in \mathcal{I}} \left(B_I \bigcap \Delta_I^N\right)$. Therefore we need to bound the sets $B_I \bigcap \Delta_I^N$ under the measure $\mathcal{P}$.

Note that using the Bounded Likelihood Ratio assumption, we have that for each point in B, the condition $V_{\mathcal{Q}}(x^N) > \epsilon$ implies $V_{\mathcal{P}}(x^N) > \epsilon/k$. Similarly for each point in B_I, we have that $V_{\mathcal{P}}(x^I) > \epsilon/k$.

Fix any set I, for example, the first indexes $I = 1, ..., n_x$, then note that the set B_I is a cylinder with base on the first n_x constraints, since the constraint $V_{\mathcal{Q}}(x^I) > \epsilon$ only depends on the first n_x samples. Now take any point $(\delta^{(1)}, ..., \delta^{(n_x)}, \delta^{(n_x+1)}, ..., \delta^{(N)})$ in the base of B_I, since we know it will satisfy $V_{\mathcal{P}}(x^I) > \epsilon/k$. For this point to be in $B_I \bigcap \Delta_I^N$, we only need to make sure each of the samples $\delta^{(n_x+1)}, ..., \delta^{(N)}$ do not violate $f(x^I, \delta) \leq 0$, otherwise $x^I = x^N$ will not hold as needed in Δ_I^N. Since we have that each $N - n_x$ samples are taken independently according the sampling distribution $\mathcal{P}$ and $V_{\mathcal{P}}(x^I) > \epsilon/k$, we have that:

$$\begin{aligned}\mathbf{P}_{\mathcal{P}}\left(B_I \bigcap \Delta_I^N\right) &< (1 - \epsilon/k)^{N-n_x} \mathbf{P}_{\mathcal{P}}(\delta^{(1)}, ..., \delta^{(n_x)} \in \text{ base of } B_I) \\ &\leq (1 - \epsilon/k)^{N-n_x}\end{aligned}$$

A bound on the whole set B can then be found by summing that bound over all possible subsets in $\mathcal{I}$.

$$\mathbf{P}_{\mathcal{P}}(B) < \binom{N}{n_x}(1-\epsilon/k)^{N-n_x} \doteq \beta(N,\epsilon)$$

2.6.6 Sample size bound for sampling with true distribution

The following theorem states the confidence level bound for the data-driven case where the sampling distribution is the true distribution.

Theorem 2.6.1 *From Campi and Garatti (2008), given the scenarios of uncertainty* $\delta^{(1)}, ..., \delta^{(N)}$ *drawn from the true uncertainty distribution* $\mathcal{Q}$. *Let* n_x *be the dimension of the strategy space* X *and* ϵ *be the robust level parameter. Then define a "risk of failure" parameter* $\beta(N,\epsilon)$ *as:*

$$\beta(N,\epsilon) \doteq \sum_{i=0}^{n_x-1} \binom{N}{i} \epsilon^i (1-\epsilon)^{N-i} \tag{2.7}$$

Then with probability greater than $(1-\beta(N,\epsilon))$, *the solution found using this method is* ϵ*-level robustly feasible, i.e.,*

$$\mathbf{P}_{\mathcal{Q}}((\delta^{(1)}, ..., \delta^{(N)}) : V_{\mathcal{Q}}(x^N) \leq \epsilon) \geq (1-\beta(N,\epsilon))$$

If we are given a desired confidence level β and a robust level ϵ, it is natural to ask what is the minimum number of samples required. This design question can be answered by numerically solving for N the equation in (2.7). Calafiore (2009) has shown that by using the Chernoff bound on the lower binomial tail, we can obtain an explicit formula for N that will guarantee that (2.7) is satisfied.

Corollary 2.6.1 *From Calafiore (2009), if the sample size* N *follows:*

$$N \geq \frac{2}{\epsilon}\left(\ln\frac{1}{\beta} + n_x\right) + 1$$

Then with probability greater than $1-\beta$ *the solution to the sampled problem will be* ϵ*-level robustly feasible.*

Numerically testing the bound:

The goal of this experiment is to validate the results of Theorem 2.6.1, which is provably a tight confidence/robustness bound for the data-driven robust problem. The bound from Theorem 2 is less tight, i.e., requires more samples to achieve the same level of confidence. For this reason, it is sufficient for us to test the tighter bound of Theorem 2.6.1 in order to verify the effectiveness of both these bounds. Note that this experiment is done under the assumption that the data samples provided are from the true distribution of uncertainty, the data-driven case.

To test the bound we consider a small instance of the MaxMin robust model, with a single product ($n = m = 1$), two time periods ($T = 2$), a linear demand model and affinely adjustable pricing policies. Using an adjustable pricing policy, the set of pricing decision variables has dimension $n_x = 4$, with one static price for each time period u_1, u_2, an adjustable component v_2 and the robust level z. Note that we will use dummy variables for the net inventory $w^{(i)}$ that are completely determined by the other variables, therefore they do not count for the dimension of the strategy space in our confidence/robust bound calculation. More specifically, the demand at each period is given as a linear function of the price, where a is a constant demand term, b is the price elasticity of demand and δ is an additive uncertain component: $\tilde{d}_t(u, v, \delta) = a_t - b_t(u_t + v_t \sum_{k=1}^{t-1} \delta_k^{(i)}) + \delta_t$. For simplicity of notation, we will not separate the case for the first time period, $t = 1$, where there is no adjustability to be considered. Therefore, wherever you consider the case of $t = 1$, define the summation from $k = 1$ to $t - 1$ as an empty sum, with zero value. Moreover, we will assume δ lies within a box-type uncertainty set, which is widely used in the robust optimization literature, defined as: $U = \{(\delta_1, \delta_2, ..., \delta_T) : |\delta_t| \leq \Gamma_t, \forall t = 1, ..., T\}$. The experiments in this section can be performed with any type of convex uncertainty set.

The parameters used for this test were $T = 2$, $a = [200, 200]$, $b = [0.5, 0.5]$, $C = 100$ and $\Gamma = [15, 30]$. For simplicity, we will consider a model virtually without overbooking or salvage value: $o = 10000$, $s = 10$. We assume for this test that the uncertainty is uniformly distributed over the uncertainty set. Moreover, we are sampling our data points from the same distribution. In other words, we are testing the data-driven case, where the scenario samples $\delta^{(1)}, ..., \delta^{(N)}$ come from the true distribution, possibly from historical data.

At each experiment, we draw $N = 200$ uncertainty scenarios to use in the optimization model to determine the pricing policy and another 2000 scenarios to evaluate the out-of-sample performance of this sampling-based pricing policy. Define the estimated probability of violation $\hat{V}$ as the proportion of the out-of-sample scenarios for which the constraints are violated under the scenario based policy.

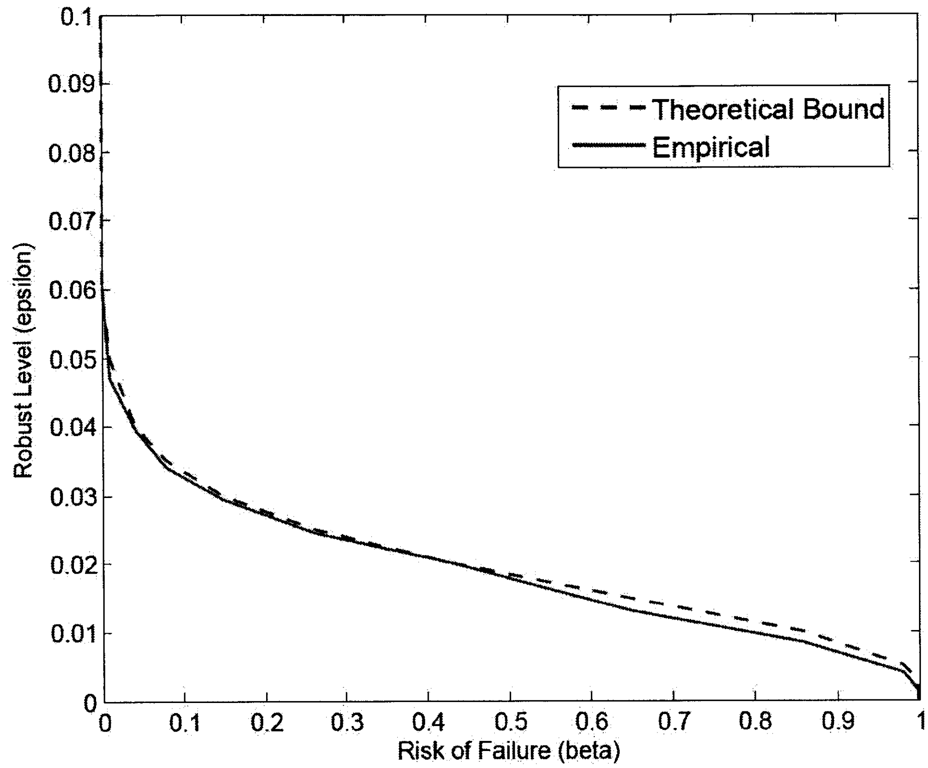

Figure 2-7: Bound Test for $N = 200$: $\epsilon \times \beta$

We repeated the experiment 300 times in order to obtain multiple instances of the problem, each with a different probability of violation. In other words, we drew other samples of $N = 200$ scenarios and repeated the procedure described before. We counted the probability of violation $\hat{V}$ for each experiment and ordered the experiments in a decreasing vector. For any given failure level $\hat{\beta}$, define the $\hat{\beta}$-quantile of this vector as the estimated robust level $\hat{\epsilon}$ associated with $\hat{\beta}$. This means that a proportion $\hat{\beta}$ of the experiments we ran, have demonstrated a probability of violation greater than $\hat{\epsilon}$.

$$\mathbf{P}(\hat{V} > \hat{\epsilon}) \doteq \hat{\beta}(\hat{\epsilon})$$

From Theorem 2.6.1, we also have a theoretical guarantee on this "risk of failure" β:

$$\mathbf{P}(V_Q > \epsilon) \leq \beta(N, \epsilon) \doteq \sum_{i=0}^{n_x - 1} \binom{N}{i} \epsilon^i (1-\epsilon)^{N-i}$$

Using $N = 200$ samples, Figure 2-7 shows, with a dashed line, the theoretical bound on the ϵ-robust level (y-axis) for each possible failure level β (x-axis). The solid line shows the estimated

$\hat{\epsilon}$ for the same failure levels β. For example, one should read this graph as follows: for $\beta = 0.0395$, i.e., confidence level of 96.05%, the solution for the MinMax model is 4%-robustly feasible, according to Theorem 2.6.1. In other words, assuming we have a good sample, there is at most a 4% chance that the actual competitive ratio will be lower than the solution provided by scenario based solution. When observing the number of violations that actually happened, the estimated probability of violation, 3.96% of the simulations had a ratio below the sample based ratio. At confidence level 99.97%, the theoretical robust level is 7% and our estimated probability of violation was 6.05%. Note how the estimated probabilities of violation are fairly close to the theoretical bound. In Figure 2-7, we present the outcome of the test with 300 experiments, but it is interesting to note that the empirical line converges to the theoretical bound roughly within 100 experiments. Overall, we have a strong empirical evidence that the bound in Theorem 2.6.1 is effective, if not in fact tight, for the pricing problem.

2.6.7 Summary of optimization models in the numerical experiments.

Throughout the modeling section of the chapter, Section 2, we described the general framework of the pricing models using simplified concise notation. In this section of the appendix, we will summarize the full model instances used to run the simulations in Section 3, trying to be as clear as possible in case the reader wishes to implement any of these models in a convex optimization solver. Define the following model as the **MaxMin**:

$$\begin{aligned}
&\max_{z,u,v,p,w,w^+} \quad z \\
&s.t. \left\{ \begin{array}{l}
z \leq \sum_{t=1}^{T} p_t^{(i)} \left(a_t - b_t p_t^{(i)} + \delta_t^{(i)} \right) - \sum_{l=1}^{m} \left((o_l - g_l) w_l^{+(i)} + g_l w_l^{(i)} \right) \\
\sum_{t=1}^{T} M \left(a_t - b_t p_t^{(i)} + \delta_t^{(i)} \right) \leq C + w^{(i)} \\
p_t^{(i)} = u_{j,t} + \sum_{l=1}^{m} v_{j,l,t} \sum_{j'=1}^{n} \sum_{t'=1}^{t-1} M_{l,j'} \delta_{j',t'}^{(i)} \geq 0, \qquad \forall t = 1, ..., T \\
w_l^{(i)} \leq w_l^{+(i)}, \qquad \forall l = 1, ..., m \\
w_l^{+(i)} \geq 0, \qquad \forall l = 1, ..., m
\end{array} \right\} \\
&\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad \forall i = 1, ..., N
\end{aligned} \tag{2.8}$$

For the regret based models, define the perfect hindsight revenue as the solution of the following model:

$$
\begin{aligned}
\Pi^*(\delta) = \quad & \max_{u,v,p,w,w^+} && \sum_{t=1}^{T} p_t\left(a_t - b_t p_t + \delta_t\right) - \sum_{l=1}^{m}\left((o_l - g_l)w_l^+ + g_l w_l\right) \\
& s.t. && \sum_{t=1}^{T} M\left(a_t - b_t p_t + \delta_t\right) \leq C + w \\
& && p_t \geq 0, \qquad \forall t = 1, ..., T \\
& && w_l \leq w_l^+, \qquad \forall l = 1, ..., m \\
& && w_l^+ \geq 0, \qquad \forall l = 1, ..., m
\end{aligned}
\tag{2.9}
$$

Using the solution of the hindsight revenue $\Pi^*(\delta^{(i)})$ for every sample of $\delta^{(i)}$ provided, we can define the **MixMax Regret** model as:

$$
\begin{aligned}
& \max_{z,u,v,p,w,w^+} \quad z \\
& s.t. \left\{
\begin{aligned}
& z \leq \sum_{t=1}^{T} p_t^{(i)}\left(a_t - b_t p_t^{(i)} + \delta_t^{(i)}\right) - \sum_{l=1}^{m}\left((o_l - g_l)w_l^{+(i)} + g_l w_l^{(i)}\right) - \Pi^*(\delta^{(i)}) \\
& \sum_{t=1}^{T} M\left(a_t - b_t p_t^{(i)} + \delta_t^{(i)}\right) \leq C + w^{(i)} \\
& p_t^{(i)} = u_{j,t} + \sum_{l=1}^{m} v_{j,l,t} \sum_{j'=1}^{n}\sum_{t'=1}^{t-1} M_{l,j'}\delta_{j',t'}^{(i)} \geq 0, \qquad \forall t = 1, ..., T \\
& w_l^{(i)} \leq w_l^{+(i)}, \qquad \forall l = 1, ..., m \\
& w_l^{+(i)} \geq 0, \qquad \forall l = 1, ..., m
\end{aligned}
\right\} \\
& \qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad \forall i = 1, ..., N
\end{aligned}
\tag{2.10}
$$

Similarly, the **MaxMin Ratio** model is defined as:

$$
\begin{aligned}
& \max_{z,u,v,p,w,w^+} \quad z \\
& s.t. \left\{
\begin{aligned}
& z \leq \left(\sum_{t=1}^{T} p_t^{(i)}\left(a_t - b_t p_t^{(i)} + \delta_t^{(i)}\right) - \sum_{l=1}^{m}\left((o_l - g_l)w_l^{+(i)} + g_l w_l^{(i)}\right)\right) / \Pi^*(\delta^{(i)}) \\
& \sum_{t=1}^{T} M\left(a_t - b_t p_t^{(i)} + \delta_t^{(i)}\right) \leq C + w^{(i)} \\
& p_t^{(i)} = u_{j,t} + \sum_{l=1}^{m} v_{j,l,t} \sum_{j'=1}^{n}\sum_{t'=1}^{t-1} M_{l,j'}\delta_{j',t'}^{(i)} \geq 0, \qquad \forall t = 1, ..., T \\
& w_l^{(i)} \leq w_l^{+(i)}, \qquad \forall l = 1, ..., m \\
& w_l^{+(i)} \geq 0, \qquad \forall l = 1, ..., m
\end{aligned}
\right\} \\
& \qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad \forall i = 1, ..., N
\end{aligned}
\tag{2.11}
$$

Alternatively, we will also use the **Sample Average Approximation (SAA)** model as a benchmark, which can be defined as:

$$\max_{z,u,v,p,w,w^+} \frac{1}{N}\sum_{i=1}^{N}\left[\sum_{t=1}^{T} p_t^{(i)}\left(a_t - b_t p_t^{(i)} + \delta_t^{(i)}\right) - \sum_{l=1}^{m}\left((o_l - g_l)w_l^{+(i)} + g_l w_l^{(i)}\right)\right]$$

$$s.t. \left\{ \begin{array}{l} \sum_{t=1}^{T} M\left(a_t - b_t p_t^{(i)} + \delta_t^{(i)}\right) \leq C + w^{(i)} \\ p_t^{(i)} = u_{j,t} + \sum_{l=1}^{m} v_{j,l,t} \sum_{j'=1}^{n}\sum_{t'=1}^{t-1} M_{l,j'}\delta_{j',t'}^{(i)} \geq 0, \qquad \forall t = 1, ..., T \\ w_l^{(i)} \leq w_l^{+(i)}, \qquad \forall l = 1, ..., m \\ w_l^{+(i)} \geq 0, \qquad \forall l = 1, ..., m \end{array} \right\} \quad \forall i = 1, ..., N \tag{2.12}$$

2.6.8 Static vs. Adjustable policies.

In Figure 2-8, we display the out-of-sample revenues between the static and the adjustable policies for the MinMax robust model. The static counterpart is defined the same way as in (11), but with the price function consisting only of the static component $u_{j,t}$. We can clearly see that the static policy is dominated by the adjustable policy, which performs a better worst case revenue as well as in overall revenue distribution. This is observation is exactly as we expected, since the adjustable policy gives the model an extra degree of freedom, which should only improve the performance of the model.

2.6.9 Network example: multiple products and resources.

For this experiment, consider a simple network problem, with $m = 2$ resources and $n = 3$ products. The incidence matrix used is:

$$M = \begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & 1 \end{pmatrix}$$

This can be viewed as the airline problem with two legs who is selling either the single leg itineraries or a combination of both legs. We will again use linear demand functions with additive uncertainty for simplicity. The reader should note that this is used here for simplicity of notation only and the use of other demand and uncertainty models can be easily incorporated. We further assume the pricing policies for each product/time-period to be affine functions of the accumulated uncertainty for the corresponding resources consumed over the previous time-periods. For a given uncertainty real-

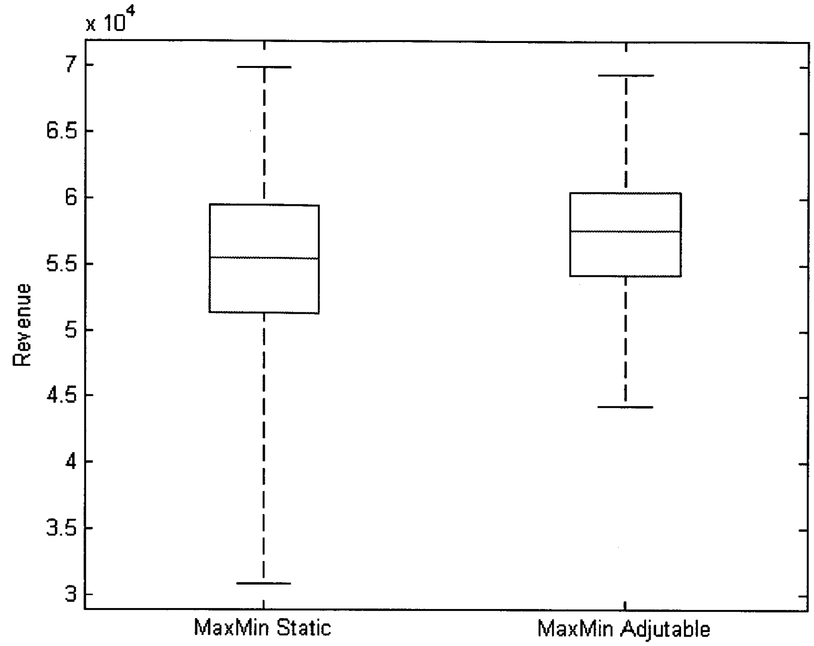

Figure 2-8: Example of out-of-sample revenue distribution: Static vs. Adjustable policies

ization $\delta^{(i)}$ and a set of pricing decisions u and v, define the prices at each period according to (3). Instead of using more complex policies, which could include cross-product and time dependencies, we use here a simpler policy as suggested in (3)(i.e. a policy for each product that depends only on its own remaining capacity) because we simply want to illustrate how to apply this framework in a network problem.

Note that the $\sum_{j'=1}^{n} \sum_{t'=1}^{t-1} M_{l,j'} \delta_{j',t'}^{(i)}$ is the amount of unpredicted resources of type l consumed by time t due to the realized uncertainty $\delta^{(i)}$ up to time $t-1$. The adjustable component of the pricing strategy, $v_{j,l,t}$ is the price impact on product j at time t from the consumption of resource l.

We compare the trade-offs between the worst-case performance and the average performance for each of these four models. The particular instance of the problem used in this experiment is: $T = 2$, $a_{j,t} = 200$, $b_{j,t} = 0.5$ for all products j and time t, $C_l = 120$, $o_l = 1000$, $g_l = 10$ for each resource l. These numbers were chosen to be somewhat similar to the case study displayed in Section 4. We defined a box type uncertainty set, $U = \{(\delta_1, \delta_2) : |\delta_t| \leq 15, \forall t = 1, 2\}$, and sampled data points using a normal distribution with zero mean and standard deviation of 15 and then truncated over the uncertainty set U. For the optimization part, we used 200 samples to obtain pricing policies for each method and simulated the out-of-sample performance of these policies using 2000 samples. Figure

2-9 displays the revenue distribution for an instance of this problem with our four models. The difference between the robust models does not seem significant, while the SAA presents a better average and lower worst case revenues. In general, the results for the network cases are consistent with the single product cases displayed in Section 3.

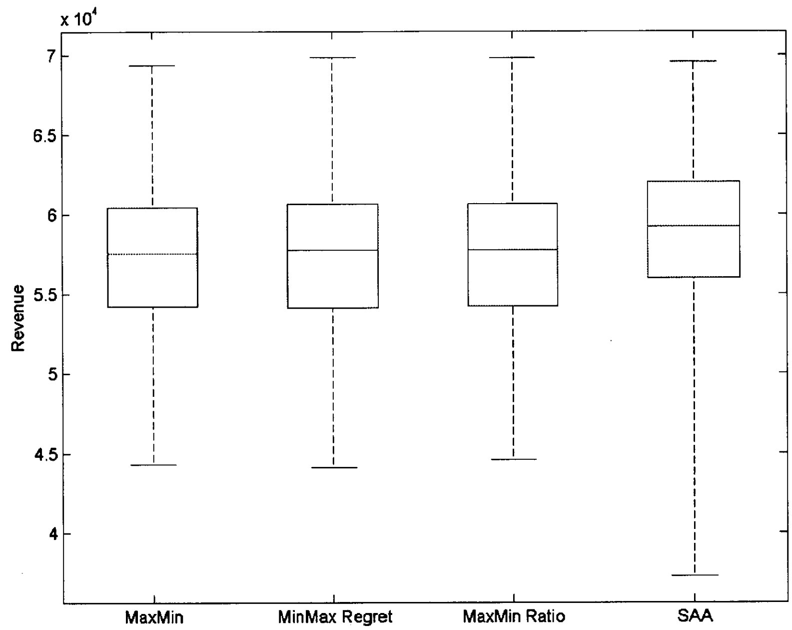

Figure 2-9: Example of out-of-sample revenue distribution for the network problem.

2.6.10 Nonlinear demand models.

Note that the numerical experiments of Section 3 were done using only linear demand functions, although our framework can also deal with nonlinear demands. In this section of the appendix, we simply demonstrate an application of the sampling based framework with a logarithmic demand function defined as $\tilde{d}(p) = a - b\log(p) + \delta$, where the demand is a linear function of the logarithm of the price plus some deviation δ. Using this demand function in a single product problem, we can apply the same underlying models defined in the previous section. For example, define the **MinMax**

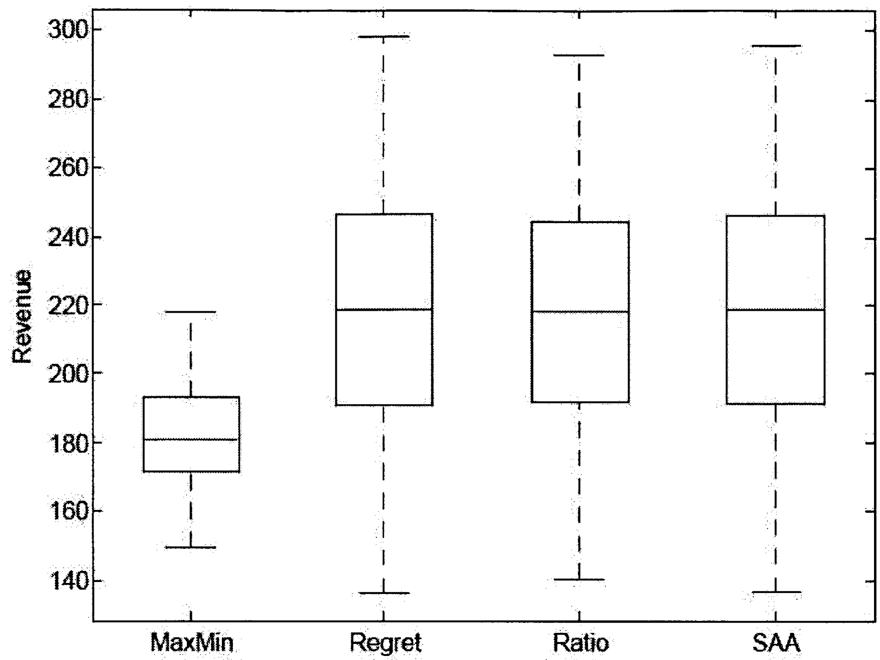

Figure 2-10: Logarithmic Demand Model: $T = 2, C = 10, a_t = 10, b_t = 2$

Regret as:

$$
\begin{aligned}
\max_{p,u,v,z} \quad & z \\
s.t. \quad & z \leq \sum_{t=1}^{T} p_t^{(i)}(a_t - b_t \log(p_t^{(i)}) + \delta_t^{(i)}) - o(w^{(i)})^+ - g(w^{(i)})^- - \Pi^*(\delta^{(i)}), \\
& \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \forall i = 1, ..., N \\
& \sum_{t=1}^{T}(a_t - b_t \log(p_t^{(i)}) + \delta_t^{(i)}) \leq C + w^{(i)}, \qquad \forall i = 1, ..., N \\
& p_t^{(i)} = u_t + v_t \sum_{k=1}^{t-1} \delta_k^{(i)} \geq 0, \qquad \forall t = 1, ..., T, \forall i = 1, ..., N
\end{aligned}
$$

Figure 2-10 displays the overall performance of the four different models under logarithmic demand. As we have shown in this section, there is no particular difference in the methodology needed to implement different demand models, as long as they satisfy Assumption 1. If we are provided with historical data of price and demand realizations, the demand function should be chosen according to the best possible fit to the data set and the uncertainty realizations will be the estimation errors.

Chapter 3

Consumer Choice Model For Forecasting Demand And Designing Incentives For Solar Technology

3.1 Introduction

Solar photovoltaic (PV) technology has greatly improved over the last two decades. With today's technology, cheap and efficient rooftop solar panels are available to residential consumers with the help of public subsidy programs. As the market for these solar panels develops worldwide, many questions remain unanswered about how this technology is adopted by customers and how to design incentives for its adoption. In this chapter we will develop a framework to model and control the adoption process of solar technology. Furthermore, we will test our framework by developing an empirical study based on the history of the German solar market.

Forecasting demand can be particularly challenging for new technologies that are not fully mature yet. More specifically, the cost of a new solar panel installation decreases as more solar panels are sold, mainly due to improvements in the installation network and manufacturing of the PV modules. These cost improvements are mainly induced by demand-driven economies of scale and competition effects that stimulate cost cuts and research. Additionally, consumer awareness about the technology will improve with the number of installed systems, which creates a second positive feedback in the adoption process. In other words, the cost improvements and the information spread through the consumer market will reinforce the demand process over time. These are commonly referred to as network externalities. These effects are particularly dominant at the early stages of technological development. Also because of these network effects, governments interested in accelerating this adoption process may often want to subsidize early adopters.

In particular, one of the reasons why governments subsidize the installation of solar panels is to promote the development of the solar technology so that it will become cost competitive with traditional sources of generation, therefore economically self-sustainable. The point where electricity generated from a solar installation reaches the electricity grid price is usually called grid-parity, either at the wholesale market price or at the higher end-consumer retail price. Estimates for reaching this grid-parity point can be as early as 2013-2014 at the retail level or 2023-2024 at the wholesale level (see (Bhandari and Stadler 2009)). Therefore, the main question that needs to be addressed by policy makers today is not whether solar technology will eventually takeoff, but rather when it will takeoff and whether something should be done to accelerate this process.

Most of the incentives devised today for solar technology come from governmental subsidies in the form of installation rebates, feed-in-tariffs or subsidized loans. Recent reports by (IEA 2004), (EEG 2007), (BMU 2008) and (EPIA 2009) provide great insights about the current status of solar PV technology as well as the history of the subsidy programs behind it.

In this chapter, we study the problem from the perspective of the policy maker, where the goal is to find the optimal subsidy value to offer customers willing to adopt these rooftop solar panels. More specifically, we assume the government has a particular adoption target, within a given time frame, and is able to offer a rebate to the customers adopting the technology. These targets are very common among policy makers, as in the "1,000 rooftops" and "100,000 rooftops" programs in Germany. (Jager-Waldau 2007) summarizes some target levels for renewable energy production and photovoltaic adoption through Europe, as decided by the European Commission. One particular example of such policies comes from (EUC 1997), a white paper from the European Commission, which proposed a target adoption level of 500,000 PV solar installations across the European Union by 2010. This number was surpassed by Germany alone, as the proposed target was very conservative. In our modeling framework we do not focus on the origin of the adoption targets, but rather assume they are given.

Once these adoption targets are determined, understanding the purchasing behavior of potential solar panel customers and how this will affect the overall spread of the technology is crucial for minimizing the subsidy costs paid by the tax-payers. We provide a modeling framework to tackle the policy design problem and also develop an empirical study of the German solar market to validate our assumptions and demonstrate how to apply this framework in a realistic practical setting. More specifically, we calibrate our model using data from the solar market in Germany, which has a history of strong subsidy policies. We further demonstrate how to forecast the future adoption levels in Germany and how they can use this model to find optimal subsidy levels as a function of their future adoption targets, as well as quantify the trade-off between adoption levels and subsidy costs. Finally, we investigate the efficiency of the current subsidy policy in Germany.

The outline of this chapter can be described as: In the remainder of Section 3.1, we summarize the main contributions of this chapter and discuss some of the relevant literature for this research area. In Section 4.1, we define the demand model and the policy optimization problem. In Section 3.3, we conduct the empirical study on the German market. In Section 3.4, we analytically explore the structure of the policy design problem to develop insights about the optimal rebate policy. In Section 3.5, we summarize the results of this chapter and outline some possible directions of research that extend this work.

3.1.1 Contributions

In summary, the goal of this chapter is to propose a framework for designing subsidy policies for the adoption of new technologies, while considering the impact of consumer behavior and cost evolution dynamics in the overall adoption process. Our first contribution is the development of a new practical policy optimization tool that policy-makers can apply to create comprehensive policy recommendations. Furthermore, we test the applicability of this model with an empirical study of the German solar market, where we demonstrate how to estimate the model parameters using real data, solve the policy design problem and discuss insights about the state of the current policy in Germany.

In the empirical study of the German solar market, we estimate our demand model using market data from 1991-2007. We then use the fitted model to produce adoption level forecasts until the year 2030 and solve hypothetical policy design questions in order to obtain insights about the structure of the system cost, the optimal solution and the efficiency of the current policy. We show that the system cost of subsidizing is a convex function of the adoption target level. In other words, the cost of subsidizing becomes increasingly more expensive as the adoption target increases. This is partially due to the fact that the network externality benefits of early subsidies become saturated. We observe this effect empirically in the German market study. We also prove it for the general setting during our analytical study of the optimal solution.

Finally, we demonstrate in our empirical analysis that the current subsidy policies in Germany are not economically efficient. We show this by optimizing the policy while trying to replicate the current forecast of the adoption path. For every possible adoption level in this baseline adoption forecast, there exists a way to reach that given level and still achieve a lower cost for the system. In particular, this is done by raising earlier subsidies and lowering future subsidies. This means that if the German government is trying to achieve a certain target adoption level at some point in the future, the current subsidies are suboptimal. We further prove that because of the decreasing nature of the optimal rebate policy, if the German government is actually trying to optimize a social welfare objective, the current policies are still suboptimal. We would like to bring special attention to the novelty of this adoption target optimization approach for studying welfare efficiency without any knowledge about the actual welfare function or solving the optimal welfare problem. In summary, these results definitely raise a warning about the efficiency of the current policies in Germany. The current Feed-in-Tariff program in place today in Germany is already in a phasing-out stage (see Figure 3-1 or the report (EEG 2007) for further details). Our assessment indicates that these subsidies should be higher now and lower in the future, while the magnitude of this change depends on the actual objective of the German government. In other words, to achieve higher adoption levels at lower cost, they can increase current subsidies and also the rate of decay for future subsidies (faster phase-out).

3.1.2 Literature Review

Historically, the economics literature has focused on models for policy design with a social welfare objective, not adoption target levels (see (Ireland and Stoneman 1986), (Joskow and Rose 1989) and (Stoneman and Diederen 1994) for further references). On the other hand, the marketing literature has focused primarily on diffusion models for new technologies without a policy design focus. For further reference in diffusion models, see the seminal work of (Rogers 1962) and (Bass 1969) or more recent review papers by (Mahajan, Muller, and Bass 1990) and (Geroski 2000). In this work, we combine ideas from these fields to develop a policy optimization framework for solar technology using a choice model of demand that incorporates network externalities. To the best of our knowledge, this is the first research to approach the policy-making problem from the target adoption level perspective and apply it to real market data in an empirical setting.

The models for innovation with network externalities are quite familiar to economists. (Farrell and Saloner 1985), (Farrell and Saloner 1986), (Katz and Shapiro 1986) and (Katz and Shapiro

1992) began exploring the issue of technology standardization and compatibility, which provide positive network externalities at first, but may later inhibit innovation. This effect can be particularly important in the adoption of computer software and telecommunications, but not so much in photovoltaic technology. (Chou and Shy 1990) argue that returns to scale on the production level can also produce the same behavior as observed in cases of consumers with preferences that are affected by direct network externalities. The cost reductions that follow the adoption process is commonly known as the learning-by-doing effect. This effect has been widely studied since the seminal paper by (Arrow 1962), and also more specifically for the case of photovoltaics, see (Harmon 2000), (IEA 2000), (McDonald and Schrattenholze 2001), (Nemet 2006) and (Bhandari and Stadler 2009).

The way that information spreads through the network of customers is another important effect that has been given a lot of attention. How consumers become aware of a new product or gather information about it's quality may determine the successful take-off of a new technology. (Ellison and Fudenberg 1995) propose a model of word-of-mouth information spread across consumers, where the structure of the social network may lead to an inefficient herding behavior. (Vives 1997) develops a model for social learning of public information where he shows that the rate of information gathering is slower than socially optimal, with examples both in a learning-by-doing case and consumers learning about a new product. In particular, (Vives 1997) develops a theoretical model for social learning. In this model, the precision of public knowledge increases at a rate $t^{1/3}$, where t is the number of time periods. He admits that this particular functional form is a direct result of his modeling choices, but the general idea that the amount of information gathered is concave over time should remain valid regardless of the model. This is similar to one of our modeling assumptions. (Ulu and Smith 2009) have recently developed a model for how information spread affects the adoption behavior of consumers, where they concluded that better information sources increases the consumers' value function for adopting the technology but perhaps induces them to wait longer for information gathering. (Aral, Muchnik, and Sundararajan 2009) argue that traditional methods for estimating social contagion may be overestimating this network effect, while homophily between consumers can explain a large portion of the technology spread.

The theoretical models mentioned above can provide intuition for the information spread effect, but are generally not applicable for designing a control policy for the adoption process. When studying this effect in a practical empirical setting, it is common to assume a functional form for how consumer utility is affected by some proxy measure of the information spread. For example, (Shurmer and Swann 1995) advocate for the use of either a linear or log relation between the network size and consumer utility in a simulation based study of the spreadsheet software market, whereas they note that a basic model without this effect makes very bad market forecasts. (Berndt, Pindyck, and Azoulay 2003) also considers both a linear and a log effect of depreciated adoption levels when estimating the diffusion of a pharmaceutical product. (Doganoglu and Grzybowski 2007) used a linear function to model the network externality effect on consumers' utility for an empirical study of the German mobile telephony market. (Swann 2002) also studies the functional form of network externalities in consumers' utility of adoption, proposing conditions for linear or S-shaped functions, focusing on cases where there is a concrete benefit of network size after the purchase, like in telecommunication networks. Using a more detailed model of the network, (Tucker 2008) uses individual measures of the agents' position on the social network to analyze their impact on the overall adoption of video-conference technology within a bank. Her empirical study demonstrates how agents that are "central" and/or "boundary-spanners" (between disconnected groups) become more influential in the technological adoption process by creating a larger network externality. (Goolsbee and Klenow 2002) develop an empirical study of the diffusion of home computers in the US, with emphasis on local network externality effects using geographic information of adoption. In particular, they show that people were more likely to buy a computer if more people in their local area had already adopted the technology. (Jager 2006) reaches a similar conclusion through a behavioral study among adopters of solar PV technology, using a survey of residents of a city in the Netherlands. In our model, we will use the log effect, as suggested in (Shurmer and Swann 1995) and (Berndt, Pindyck, and Azoulay 2003), because it satisfies the concave behavior that we want to model. We additionally tried other similar functional forms during our empirical study, but the log effect presented the best fit.

Within the broader marketing literature, (Hauser, Tellis, and Griffin 2006) enumerates the multiple directions of future research that should to be explored by the marketing community. One such direction is to improve our understanding of consumer response to innovation. We address this issue by using a diffusion model based on the logit demand model. Our particular diffusion model can be placed within the broader class of proportional hazard rate models. Developed by (Cox 1972) for modeling the life time of an agent in a larger population, the hazard rate model has been widely used in biostatistics and its application in marketing has been well documented in (Helsen and Schmittlein 1993). In our case, the agent's life-time duration is the moment he/she makes the purchase decision and adopts the technology. We diverge from the original Cox model in the particular functional form of the adoption probability, where we use the logit demand derivation for the probability of purchase. All these models will result in the familiar S-shaped diffusion pattern. The particular functional form chosen for our model is derived from the characteristics of consumer purchasing behavior. Additionally, it was chosen because it provides good estimation results with the German market data and analytical tractability.

Recent work by (Benthem, Gillingham, and Sweeney 2008) estimated a demand model with a similar learning-by-doing effect on a study of the California Solar Initiative and (Wand and Leuthold 2010) used the same model on a study of the German market. These two papers assume there is a known environmental externality cost that is avoided by PV installations and assume the government tries to maximize the net social welfare of the system when deciding the subsidy policies. On the other hand, these environmental costs are mostly deduced from the global impact of climate change and it is not clear why governments would have an incentive to pay these costs. Furthermore, according to their model, the net social welfare cost of solar incentives is mainly determined by the difference in the interest rate used by the government and consumers. If these rates were the same, there would be no cost for increasing the subsidies, only positive externalities. This social welfare maximization approach assumes a perfectly efficient welfare redistribution system. In our research, we don't need this assumption because we do not attempt to quantify the social welfare benefits of solar technology. Instead, we directly use adoption targets as our policy optimization objective, which we believe is a more realistic setting. To the best of our knowledge, our research is the first work to apply a consumer choice modeling approach to understand the adoption of solar technology with fixed adoption targets. We are also able to use this framework of adoption targets to evaluate the welfare efficiency of the current German policy without making assumptions on the environmental benefits of solar installations.

3.2 Model

We consider a modeling framework that can be divided in two parts: the demand model and the policy-maker's problem. In Section 3.2.1 we develop a solar panel demand model based on the customers' purchasing behavior and the network externalities that increase the incentives for future purchases the more consumers adopt the technology. In Section 3.2.2, we propose an optimization model for solving the subsidy policy design question. We have included a notation summary in Appendix 3.6.1, which may be a useful reference while reading this modeling section.

3.2.1 Demand Model

The first step to understand the adoption process of a certain technology is to understand the customer behavior. At each time step (for example each year), we consider every household as a potential customer who is given a choice between purchasing a solar panel or not. Let M_t be the market size (number of households) and x_t be the number of customers at a given time t that have already adopted the technology, in this case rooftop photovoltaic solar panels. Define r_t as the rebate level offered by the government, which is the policy maker's decision variable in order to control the adoption rate of the technology. Let $q_t(x_t, r_t)$ be the demand for solar panels at time t. The technology adoption in the population is given by the discrete time diffusion process: $x_{t+1} = x_t + q_t(x_t, r_t)$.

The demand model we propose in this chapter is centered around the average consumer's utility profile, namely $V_t(x_t, r_t)$. In order to maintain tractability and also due to the lack of additional data for the empirical study, we will make the following assumptions:

Assumption 3.2.1

a) At each time period t, a customer will either buy an average sized solar installation (denoted $AvgSize$) or not buy anything;

b) After purchase, this customer is out of the market: no depreciation, resell or additional purchase options;

c) The solar yield (electricity generation) and installation costs are homogenous across the entire country;

d) Demand q_t for solar panels at time t follows a logit demand model, which is a function of the utility that consumers have for purchasing a solar panel at time t.

In particular, Assumption 3.2.1.d defines the demand q_t as a logit demand function, which is equal to the number of remaining potential customers times the probability that each of these customers will make the purchase decision at time t. This customer purchase probability is what we also call adoption or diffusion rate. For a review on diffusion models, see (Mahajan, Muller, and Bass 1990).

The motivation behind the logit demand model comes from customers being rational utility maximizing agents. With this in mind, define $V_t(x_t, r_t)$ as the nominal utility that the average consumer has for purchasing a solar panel at time t. It is a function of the current state of the system x_t and the rebate levels r_t. Additionally, define $\epsilon_{t,i}$ as a random utility factor for a given customer i at t that captures the heterogeneity of consumers' utility. It represents the utility impact of all characteristics that different consumers have, for example geographic location, household sizes, discount rate differences, or environmental conscience. Let $U_{t,i}$ be customer i's perceived utility for purchasing a solar panel at time t. This is given by:

$$U_{t,i} = V_t(x_t, r_t) + \epsilon_{t,i} \tag{3.1}$$

The logit demand model is one of the most common demand models used in the discrete choice literature (see for example (Ben-Akiva and Lerman 1993; Train 2003) for further references on discrete choice models). This is mainly due to its analytical tractability. This is also the reason we use this model in our framework. The logit model assumes consumers are utility maximizing agents and the heterogenous component $\epsilon_{t,i}$ comes from an extreme value distribution. Therefore, at each point in time, customers are given a random draw $\epsilon_{t,i}$ and will decide to purchase the solar panel if the utility of purchase $U_{t,i}$ is greater than zero (where zero is the utility of no purchase). Therefore, the probability of adoption for any given consumer can be obtained by integrating the distribution of $\epsilon_{t,i}$ over the region $U_{t,i} > 0$. This gives us the well-known logit demand model:

$$q_t(x_t, r_t) = (M_t - x_t)\frac{e^{V_t(x_t,r_t)}}{1 + e^{V_t(x_t,r_t)}} \tag{3.2}$$

The first term $(M_t - x_t)$ represents the number of left-over consumers who have not purchased a solar panel yet at time t and the remaining term is the probability of adoption for any of these customers. Additionally, we need to assume the following:

Assumption 3.2.2

a) Consumers do not make a strategic timing decision when purchasing the panel. If their private signal $\epsilon_{t,i}$ is strong enough so that $U_{t,i} > 0$, then they will make the purchase at that time t.

b) The heterogeneity random components $\epsilon_{t,i}$ are uncorrelated across time periods.

One may argue that Assumption 3.2.2.a is a strong assumption, as consumers might be tempted to wait for panels to become cheaper. Note that (Kapur 1995; Goldenberg, Libai, and Muller 2010) show that the presence of network externalities might encourage agents to wait and delay overall adoption process. On the other hand, (Choi and Thum 1998) suggest that in the presence of network externalities consumers do not wait enough to adopt a new technology, settling for what is available instead of waiting for an improvement. Nevertheless, as we observed in the German market data of Section 3.3 (see Figure 3-2), the Feed-in-Tariffs offered by the government are decreasing faster than the installation costs after 2005. This gives the incentive for consumers to not be strategic about their timing decisions. Before 2005, this was not the case, which suggests that strategic timing behavior of consumers might have influenced the demand for panels and a more complex model of consumer behavior might be necessary. On the other hand, using more complex models may lead to estimation and tractability problems. Assumption 3.2.2.a can be interpreted as consumers being short-sighted agents that can only maximize their utility within a given time period.

Assumption 3.2.2.b considers that each customer is given a draw of it's private utility shock $\epsilon_{t,i}$ for that year, which is independent from the private shocks in previous years ($\epsilon_{\tau,i}$ for $\tau < t$). This can be a stringent assumption, as people who have not adopted solar panels because they have lower income or live in a region that has low solar irradiation will tend to not adopt in future periods as well. These problems can be reduced by introducing demographic data into the demand model, which would make $\epsilon_{t,i}$ capture less of these fixed components of heterogeneity. For example, one way to use demographic data into the demand model would be to introduce random coefficients as in the BLP model (introduced in (Berry 1994) and (Berry, Levinsohn, and Pakes 1995)). This approach has been quite popular in the industrial organization literature recently. The main idea behind the random coefficients model is that consumer's sensitivity to the monetary value of the solar panel, the NPV of the project, should be heterogenous due to those demographic differences between consumers. By using the distribution of each of these demographic components within the entire population, we can use computational methods to find better estimates of the probability of adoption. Although this would possibly lead to a more precise demand model, we choose to avoid the BLP approach for the following reasons: we do not have sufficient data points to introduce more parameters to estimate, given that we are working with yearly data over a span of 17 years; we wish to maintain the analytical closed-form tractability of the demand function, as it will allow us to explore insights about the structure of the optimal solution of the policy design problem (see Section 3.4).

To fully specify the demand function described in (3.2), we need to define the nominal utility of the average consumer, denoted by $V_t(x_t, r_t)$. Consumers' perceived utility for adopting a solar panel should be a function of many parameters, including the monetary value of the investment and the awareness that customers have of the given technology. The first component of $V_t(x_t, r_t)$ is the monetary value of an average solar installation purchased at time t, namely NPV_t. This component is equal to the government installation rebate to consumers r_t plus the future discounted cash flows d_t minus the installation cost k_t, all this is multiplied by the size of an average household installation, denoted by $AvgSize$. That is:

$$NPV_t(x_t, r_t) = (-k_t(x_t) + r_t + d_t)AvgSize \tag{3.3}$$

In particular, we model the installation cost $k(x_t)$ as a decreasing function of the number of solar panels sold x_t. This is consistent with the learning-by-doing effect, that is, the more people adopt a given technology, the cheaper this technology becomes in the future. In other words, the installation cost can be expressed as a decreasing function of the installed capacity. For further references of learning-by-doing in photovoltaics, see (Harmon 2000; IEA 2000; McDonald and Schrattenholze 2001; Nemet 2006; Sderholm and Sundqvist 2007; Bhandari and Stadler 2009; Yu, van Sark, and Alsema 2009).

In our model, we represent the solar installation costs with a single term $k_t(x_t)$, expressed in €/Wp of installed capacity (nominal solar installation sizes are measured in Watt-peak, i.e. Wp, which represents the electricity peak generation capacity in Watts under standard laboratory conditions). In practice, there are many different parts in an installation of a solar panel. These include the solar module, additional electronic components (also known as Balance-Of-System) and labor costs. Ideally, we would want to model the evolution of each cost separately, since the module costs evolve according to the global demand for solar panels, while the other costs decrease with the number of local installations. Nevertheless, given that we only have information on the total installation costs for our empirical study, we simplify the cost evolution dynamics by defining a single cost function for the solar installation. This cost function decreases with the number of installations in the country. In particular, the log-log learning curve is the standard model in the learning-by-doing literature. Let a_I and b_I be the installation cost parameters and ν_t represent a random technological improvement factor for time t. Then the cost dynamics can be described as:

$$\log(k_t(x_t)) = a_I + b_I \log(x_t) + \nu_t \tag{3.4}$$

Finally, the discounted cash flow d_t denotes the present value of the cash payments the customer will receive after purchasing the panel at time t. Note that in countries like Germany, where a Feed-in-Tariff system is implemented, the customer will lock the price of the tariff on the year he purchases the panel and will keep selling electricity at that price for the duration of the contract. For example, in Germany this contract lasts 20 years (see the report (EEG 2007) for further reference). Most estimates for the lifetime of solar panels suggest that they would last 30 years or more, but given that this value is discounted to the present and the Feed-in-Tariff expires in 20 years, possibly bringing the selling price to retail levels, the residual cash flow after 20 years will be very small compared to earlier ones. For simplicity, we choose to consider the discounted cash flow of the panel only for the duration of the Feed-in-Tariff contract. Let FIT_t (in €/kWp) be the revenue earned at each year for a panel bought at time t. This is the value of the Feed-in-Tariff contract times the average annual electricity output of a 1 kWp nominal capacity solar panel. Besides the electricity selling revenue, the consumer needs to pay yearly operation and maintenance (OM_t) costs (this is about 2% of k_t every year). We further assume a discount rate of δ_c (which is about 3 to 5%). Then the discounted cash flow is given by:

$$d_t = \sum_{\tau=1}^{T_{mod}} \frac{1}{(1+\delta_c)^\tau}[FIT_t - OM_t] \tag{3.5}$$

With the discounted cash flow described in (3.5), the installation costs given in (3.4) and the government installation rebate level r_t, we obtain the net present value of the installation NPV_t (see definition in (3.3)). For the remainder of this chapter, we consider d_t as a given constant, as defined in (3.5). In other words, we take the Feed-in-Tariff subsidies as data and the government can further subsidize only by introducing upfront rebates r_t. Negative rebates can also be used in our model. This would imply a tax increase on the sale of the panels. Either way, for a fixed discount rate δ_c, both forms of subsidies are equivalent from an average consumer's perspective. Behaviorally, different forms of subsidies evoke diverse responses among customers and the net present value might not be the best way to capture how consumers perceive this investment. Another possible measure of investment quality would be the payback period (time until investment cost is recovered). We will not focus on the discussions about the cognitive limitations of consumers or possible trade-offs between different forms of subsidy. We use the more common economic definition of utility that consumers are directly affected by the net present value of their investment.

As mentioned before, the second component that affects the consumer's perceived utility toward the solar panel purchase is the awareness level of the customer about the technology. In particular, there are two network externalities that we want to emphasize in our model: learning-by-doing and information spread. Because of these network effects, it might be cheaper for the government to subsidize the early adopters in order to promote a faster overall adoption of the technology. The first network externality, learning-by-doing on the installation costs, is modeled in (3.4).

The second externality is what we call information spread effect, or imitating customer behavior, and it can be usually observed for most new technology adoption processes. It has been well documented in the marketing literature (see for example (Mahajan, Muller, and Bass 1990)) and in the behavioral economics literature (see for example (Jager 2006)). We provide a more in depth discussion of the literature on information spread in Section 3.1.2. In summary, this effect happens because consumers become increasingly more aware about a new technology as more people buy the product. In our case, the more rooftop panels are adopted in a neighborhood, other consumers in the same neighborhood will be more likely to adopt the technology as well (see (Goolsbee and Klenow 2002; Jager 2006)). On the other hand, the marginal impact of this information spread on the remaining customers should naturally decrease with the number of adoptions (see (Vives 1997)). Therefore, the effect of this externality on consumer purchases should be a concave function of the number of customers that have already adopted this technology.

As mentioned in Section 3.1.2, theoretical agent-based models for information spread through a network provide useful insights about the overall effect, but are generally not practical for empirical applications. In order to conduct empirical work on this subject, it is often common to assume a particular functional form for how such network externalities affects consumer utility and aggregate purchase behavior, as in (Shurmer and Swann 1995; Swann 2002; Berndt, Pindyck, and Azoulay 2003; Doganoglu and Grzybowski 2007). In this part of the model, the effect we want to capture is the development of consumer awareness and how it affects consumers' perceived utility of purchase. In particular, we model this effect as a penalty function on the proportion of adopted customers x_t/M_t, which lies between 0 and 1. We propose the following limiting conditions for this penalty function: If nobody has adopted the new technology, consumers are generally unaware of the product and their perceived utility of purchase should go to $-\infty$; If everyone has adopted the technology, $x_t/M_t = 1$, then this penalty should go to zero. Together with the concavity condition mentioned before, we propose the use of a logarithmic relation between average consumer's perceived utility and the adopted share of the population: $V_t \sim \log(x_t/M_t)$. This functional form is consistent with previous empirical work on technology adoption, in particular (Shurmer and Swann 1995; Berndt, Pindyck, and Azoulay 2003), where they also test a logarithmic relation between network externalities and consumer utility of purchase.

Assumption 3.2.3 *The average customer's perceived utility for purchasing a solar panel is proportional to the log of the ratio of adopted customers in the population, due to information spread and consumer awareness of the technology:* $V_t \sim \log(x_t/M_t)$

This functional relation in Assumption 3.2.3 is by no means the only choice for modeling the information spread effect while satisfying the concavity and limiting conditions. It was chosen mainly due to the good fit demonstrated in our empirical study of the German solar market, compared to other functional forms we tested (for example $V_t \sim 1 - (x_t/M_t)^{-1}$). Its simplicity and tractability are additional advantages of this modeling choice, which are important both for estimation purposes in Section 3.3.1 and for the analytical results of Section 3.4.

Gathering all the utility components described so far, define the average customer's perceived utility for purchasing a solar panel at time t given by:

$$V_t(x_t, r_t) = a_D NPV_t(x_t, r_t) + b_D \log(x_t/M_t) + c_D + \xi_t \tag{3.6}$$

The first part $a_D NPV_t(x_t, r_t)$ denotes the monetary component of the utility, $b_D \log(x_t/M_t)$ denotes the impact of the information spread in the consumer's utility, c_D denotes the baseline utility for making a solar panel purchase, and finally ξ_t is a random demand shock for year t. Note that a_D, b_D and c_D are demand parameters that need to be estimated, while ξ_t is a random utility component. In particular, ξ_t represents all unobserved demand shocks for a given year that affect all consumers and cannot be captured in our data. This could for example represent a demand shock due to a strong advertising campaign in that year. The definition of consumer i's utility function is then given by adding the nominal average consumer utility with the heterogeneity component:

$$U_{t,i} = a_D NPV_t(x_t, r_t) + b_D \log(x_t/M_t) + c_D + \xi_t + \epsilon_{t,i} \tag{3.7}$$

As defined before in (3.2), we can now explicitly write the demand model as:

$$q_t(x_t, r_t) = (M_t - x_t) \frac{e^{a_D NPV_t(x_t,r_t) + b_D \log(x_t/M_t) + c_D + \xi_t}}{1 + e^{a_D NPV_t(x_t,r_t) + b_D \log(x_t/M_t) + c_D + \xi_t}} \tag{3.8}$$

In order to use the demand model defined in (3.8) and obtain statistically significant estimation results, we need to make some assumptions.

Assumption 3.2.4

a) ξ_t *is not correlated with* $NPV_t(x_t, r_t)$ *or* $\log(x_t/M_t)$*;*

b) ξ_t *is not autocorrelated with* ξ_τ*, for all* $\tau < t$*;*

The correlation described in Assumption 3.2.4.a can be a problem for the estimation procedure, but is usually treated with the use of instrumental variables. Autocorrelation, as described in Assumption 3.2.4.b, is a very common problem when estimating time series data. This problem can usually be solved by fitting an auto-regressive model for these demand shocks together with the demand model (for example, $\xi_t = \alpha\xi_{t-1} + \eta_t$). In order to maintain simplicity of the model and minimize the number of parameters to be estimated, we have assumed correlations are zero. Furthermore, we have tested this assumption in the empirical study of the German market data and the estimation output demonstrated no significant correlation.

In summary, the full demand model can be described by the following set of equations, for all $t = 1, ..., T-1$:

$$\begin{aligned}
&\text{Diffusion Process:} && x_{t+1} = x_t + q_t(x_t, r_t) \\
&\text{Logit Demand:} && q(x_t, r_t) = (M_t - x_t)\frac{e^{a_D NPV_t(x_t,r_t)+b_D \log(x_t/M_t)+c_D+\xi_t}}{1 + e^{a_D NPV_t(x_t,r_t)+b_D \log(x_t/M_t)+c_D+\xi_t}} \\
&\text{Net Present Value:} && NPV_t(x_t, r_t) = (-k_t(x_t) + r_t + d_t)AvgSize \\
&\text{Learning-by-Doing:} && k(x_t) = e^{a_I + b_I \log(x_t) + \nu_t}
\end{aligned} \tag{3.9}$$

3.2.2 Policy-Maker's Problem

The discrete choice model framework described in the previous section can be used to solve a variety of management and policy problems. In particular, policy makers are traditionally faced with the problem of setting subsidy levels to stimulate the adoption of a technology up to a target level within a certain time frame. (Jager-Waldau 2007) provides some examples of renewable energy/photovoltaic target levels for the European Union, although without clear implementation guidelines. As an example of these target policies within a country, there was Germany's pioneer "1000 Solar Rooftops" program in the early 90's. The next program, known as "100,000 Solar Rooftops", started in 1999 with subsidized loans and expected to install around 300MWp of solar panels within 6 years. The program ended before schedule in 2003 when the target was reached, suggesting the subsidy might have been higher than necessary. This seems to suggest that without further understanding of customer behavior and the dynamics of the adoption process, the policies can become short-sighted, possibly under/over-subsidizing.

Using our demand model defined in (3.9), the policy maker should find the optimal rebate levels r_t for $t = 1, ..., T-1$ in order to minimize the total present value of the rebate costs, while still achieving the target adoption level x_T at the end of the planning horizon. In this chapter, we only consider a deterministic model and therefore the random components ν_t and ξ_t (from equations (3.4) and (3.8)) will be set to zero. This work, to the best of our knowledge, is the first one to deal with target policy optimization with a choice model approach and network externalities. Introducing uncertainty into the policy optimization framework adds an extra level of complexity to the model that would overshadow some of the insights that we are trying to obtain. Nevertheless, we believe this is actually a very promising direction to extend this work. In Chapter 4, we consider a subsidy design problem with uncertain demand. For this chapter we will focus only on the deterministic counterpart of the policy problem defined in the following optimization model:

$$
\begin{aligned}
&Cost_1(x_1, x_T) = \\
&\min_{r_1,...,r_{T-1}} \sum_{t=1}^{T-1} \delta_g^{t-1} r_t q_t(x_t, r_t) \\
&s.t. \quad x_{t+1} = x_t + q_t(x_t, r_t), && \forall t = 1, ..., T-1 \\
&\qquad q(x_t, r_t) = (M_t - x_t)\frac{e^{a_D(r_t - k(x_t) + d_t) + b_D \log(x_t/M_t) + c_D}}{1 + e^{a_D(r_t - k(x_t) + d_t) + b_D \log(x_t/M_t) + c_D}}, && \forall t = 1, ..., T-1 \\
&\qquad k(x_t) = e^{a_I + b_I \log(x_t)}, && \forall t = 1, ..., T-1
\end{aligned}
\tag{3.10}
$$

Note that parameters M_t, d_t, δ_g and x_1 are given data, denoting respectively the market size at time t, the discounted future cash flow of solar installations purchased at time t, the government's discount rate, and the initial number of household solar installations sold before time t. The set of parameters (a_I, b_I) and (a_D, b_D, c_D) allow us to define the cost evolution dynamics and the demand function, respectively. These parameters need to be estimated using a historical data set, as we demonstrate in Section 3.3.1. Note that we replace the $NPV_t(x_t, r_t)$ by $(r_t - k(x_t) + d_t)$, in order to make the notation more concise. The average installation size originally included in the NPV definition of (3.3) can be suppressed because we are estimating a_D and this only causes a proportional shift in the estimate.

The problem described in (3.10) can be solved numerically using a dynamic programming reformulation, where the state of the system is x_t (the number of solar panels sold up to time t). The adoption target condition can be enforced with a terminal system cost of zero if the target adoption level has been achieved and set to infinity otherwise.

$$Cost_T(y, x_T) = \left\{ \begin{array}{l} 0, \; if y \geq x_T \\ \infty, \; o.w. \end{array} \right\} \tag{3.11}$$

At each step, the policy maker decides the rebate level r_t. The immediate rebate cost observed at each period is the rebate times the amount of people who adopted at that given time step: $r_t q_t(x_t, r_t)$. The objective of the government at each time step is to minimize the immediate rebate cost plus discounted future rebate costs. Define for $t = 1, ..., T-1$ the following cost-to-go functions:

$$Cost_t(x_t, x_T) = \min_{r_t} r_t q_t(x_t, r_t) + \delta_g Cost_{t+1}(x_t + q_t(x_t, r_t), x_T) \tag{3.12}$$

It is easy to see that the solution of the dynamic program in (3.12) leads to a solution of the original problem in (3.10). This is due to the fact that the state variable x_t decouples the problem across multiple time periods. The second term in the cost-to-go function, x_T, is used here as a fixed parameter, i.e., some policy target that is known beforehand. We use this notation because we will later explore the implications of changing the target adoption levels in the overall system cost (see Section 3.4).

Note that the DP formulation in (3.12) can be numerically solved by discretizing the state-space. Given that we only carry one state variable, we can perform a line search over r_t at each period t and efficiently compute the cost-to-go functions by backwards induction.

We can easily add further complexity levels to the model in (3.12), such as constraints on the rebate levels and quantity caps on the number of subsidized panels (these are actually commonplace in many countries). For example, in order to avoid strategic timing behavior of the customers, we have argued that subsidy levels decrease at a faster rate than the costs improve. In order to maintain that argument, we might need to introduce a decreasing rebate constraint $r_t \leq r_{t-1}$. For that, we need to add another dimension to the state-space of the dynamic program to keep track of previous rebate levels. Nevertheless, this is not much harder to solve, as 2-state DP is still numerically tractable. In the remainder of this chapter, we focus only on the base model defined in (3.10), without these extensions of the problem. We have implemented and tested some of these constraints. Nevertheless, they do not add additional insight into the policy design problem that we are dealing with.

3.3 Empirical Analysis of the German Solar Market

In this section we perform an empirical study of the German solar market by estimating the demand model described in Section 3.2.1 and using this model to produce forecasts for future adoption levels of the solar technology and validate the model. Furthermore, we use these forecasts and the DP formulation of the policy-maker's problem described in Section 3.2.2 to produce policy recommendations.

We have gathered the following information on the German PV solar market data:

a) Number of households in Germany from 1991 to 2007

b) Forecasted number of households in Germany from 2008 to 2030

c) Feed-in-Tariff rates (€/kWh) from 1991 to 2007

d) Feed-in-Tariff forecasted rates (€/kWh) from 2008 to 2030

e) Average solar installation cost (€/kWp)from 1991 to 2007

f) Nameplate peak capacity (MWp) of solar panels installed in Germany from 1991 to 2007

g) Distribution of PV solar installation sizes made in 2009

h) Discount rate used by customers and government

i) Average annual PV solar electricity yield (annual kWh/kWp)

Sources for the data collected for this study include (IEA 2004; Schaeffer, Alsema, Seebregts, Beurskens, de Moor, van Sark, Durstewitz, Perrin, Boulanger, Laukamp, and Zuccaro 2004; Wissing 2006; EEG 2007; PVPS 2008; Frondel, Ritter, and Schmidt 2008; EPIA 2009; Bhandari and Stadler 2009), as well as the databases of the Eurostat (European Commission Statistical Office) and the Federal Statistics Office of Germany.

The data for the Feed-in-Tariff rates both past and forecasted can be seen in Figure 3-1. The average solar installation cost k_t is displayed in Figure 3-2, together with the discounted cash flow d_t for a solar installation. This discounted cash flow data is displayed in Figure 3-2 and can be obtained using the Feed-in-Tariff rates, discount rates and the annual solar yield data. The discount rate used by customers and government is assumed to be $\delta_c = \delta_g = 95\%$ (approximately equivalent to the 5% interbank interest rate). Finally, we assume that average annual solar yield is 750 kWh

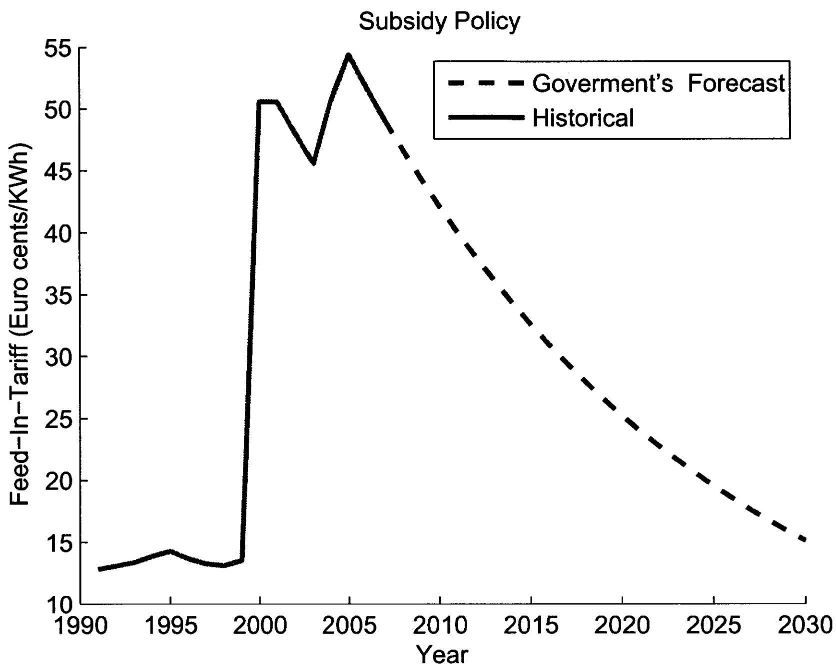

Figure 3-1: Current subsidy policy in Germany

per kWp of installed peak capacity, which is derived by the average total amount of PV electricity generated divided by the installed capacity in each year.

The nameplate peak capacity (MWp) of solar panels installed from 1991-2007, together with the distribution of installations in 2009, will be used to estimate the number of residential installations done between 1991-2007. The resulting estimated number of solar household installations together with the number of households in Germany is displayed in Figure 3-3. The details of these calculations will be discussed next.

The data obtained for the amount of historical solar installations was in the total cumulative installed capacity of PV solar panels in Germany, including both rooftop and open-space installations. In general, these two types of installations are very different both in terms of size and incentive tariffs. Modeling the adoption of both types of installations with the same demand model can be inaccurate, but we could not obtain differentiated data about the size of installations from 1991-2007. According to (Reichmuth, Erfurt, Schiffler, Kelm, and Schmidt 2010), such information was not even collected for rooftop systems for this data range. Starting in 2009, Germany's Federal Network Agency (Bundesnetzagentur) requires all new PV installations to register their installed capacity. Using this new database, we obtain the size of all new solar installations performed

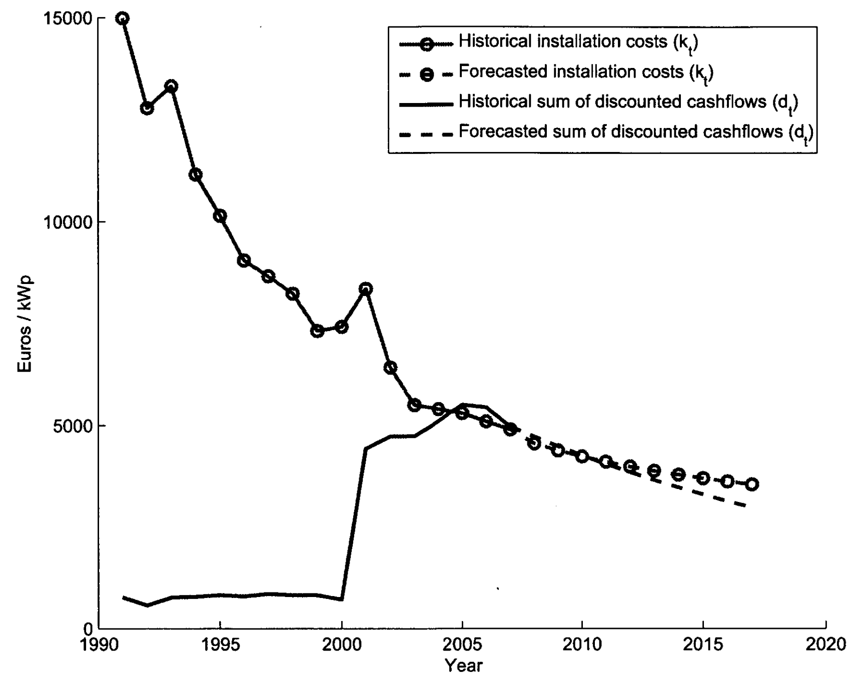

Figure 3-2: Installation costs (k_t) vs. Discounted cash flows (d_t)

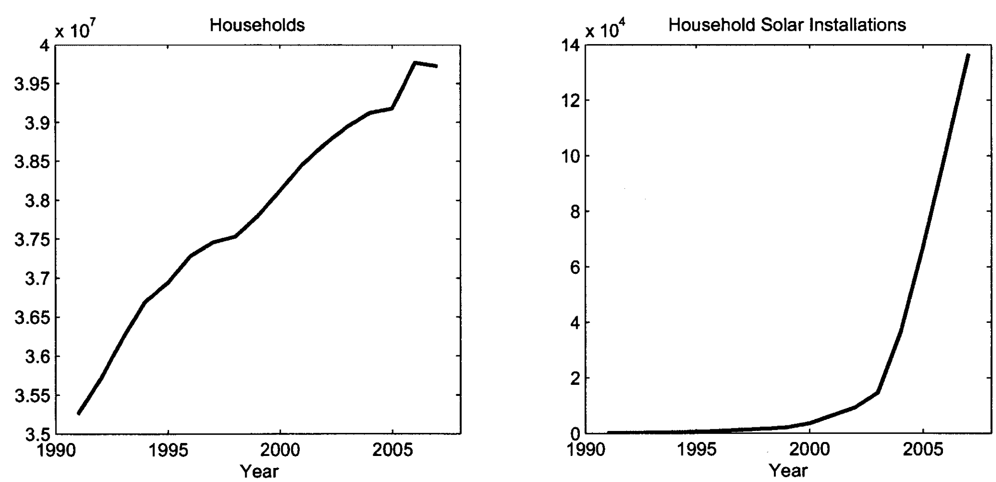

Figure 3-3: Number of households (M_t) and number household of solar installations (x_t) in Germany

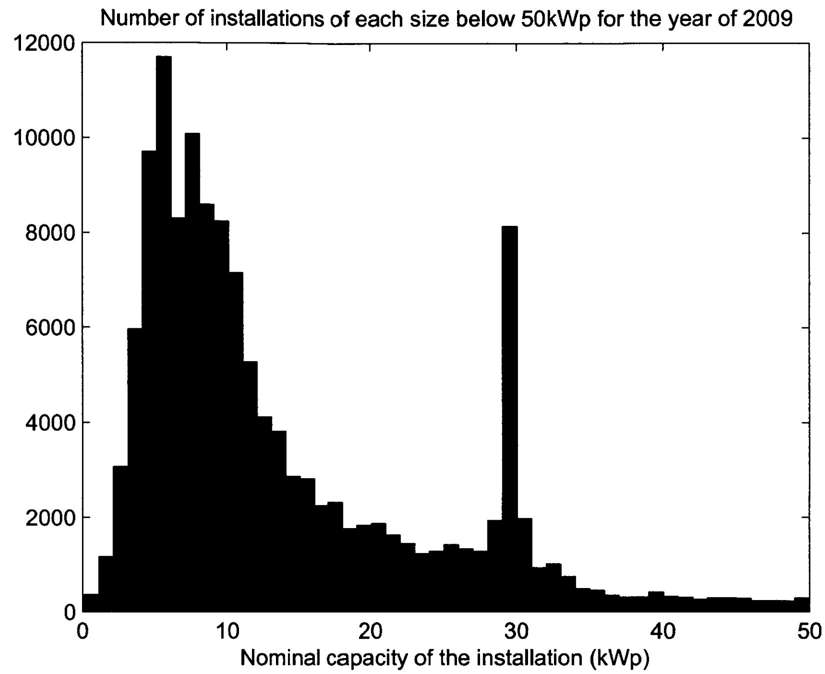

Figure 3-4: Histogram of PV installations in 2009

in 2009. In Figure 3-4, we display a histogram with the installations in 2009 for each size ranging from $1kW_p$ to $50kW_p$. In fact the full data for 2009 includes very large installations, including some solar farms of approximately $50MW_p$ and $20MW_p$. Residential rooftop installations are usually considered to be under $30kW_p$. This is also the criterion used in the Renewable Energy Sources Act (EEG) to define the feed-in-tariffs for small scale installations. In fact, we observe in Figure 3-4 a sharp increase in installation numbers exactly at $30kW_p$, as customers have a strong incentive not to go over this limit in order to obtain the higher feed-in rates.

In order to differentiate residential installations from open-space installations, we will use the proportion of installations sizes in 2009 to infer the number of residential installations from the total aggregate installed capacity from 1991 to 2007. We understand this is a strong assumption, but it the best we can do with the information that is available. In fact, the sum of all PV systems installed in 2009 is $3,429kW_p$, while 42.84% of these were from installations under $30kW_p$. The total number of installations under $30kW_p$ was $122,863$ (out of a total of $141,278$ new PV systems) and the average size of these residential systems was $11.95kW_p$. To put things in perspective, these new rooftop installations in 2009 broke yet another record for the number of installations in the country and yet covered approximately only 0.32% of the households in Germany.

Using the historical series of total installed nameplate capacity of solar panels in Germany (both residential and not) together with the ratio of residential installations of 42.84% and the average system size of $11.95kW_p$, we extrapolate the historical x_t adoption level, i.e. the number of residential customers that had purchased a solar panel before each year between 1991-2007. The result is displayed in Figure 3-3.

3.3.1 Fitting the Installation Cost and Demand Model

There are basically two estimations to be made from the data that was gathered: the cost function and the consumer utility model. In particular, we need to estimate those five coefficients (a_I, b_I) and (a_D, b_D, c_D). Note that the cost function appears inside the consumer utility model through the NPV of the solar installation. If we try to estimate both relations together, the estimation will have problems with the endogeneity of the system cost evolution in the adoption process. Therefore, we can first estimate the dynamics for the cost of solar installations and then estimate the utility model afterwards. The cost improvement function was estimated with a simple regression on the log-log relationship between k_t and x_t, as defined in (3.4). Table 3.1 displays the estimation results. The

	Estimate	Std. Error
a_I	3.05***	(0.0635)
b_I	−0.127***	(0.0065)
R^2	0.907	
$\bar{R}^2$	0.893	

Table 3.1: Estimation Results for the Installation Cost (Learning-by-Doing effect)

results of the cost dynamics fitted above in Table 3.1 can be translated into a perhaps more common terminology of Learning Rate (LR) and Progress Ratio (PR). In particular, $PR = 2^{b_I} = 92\%$ and $LR = 1 - PR = 8\%$. The significance level for the estimates in Table 3.1 are satisfactory, as indicated by a p-value less than 0.1% (marked *** on the table).

The demand model defined in (3.8) can be expressed as a linear function of the utility parameters that we want to estimate. Let $\lambda_t = \frac{x_{t+1}-x_t}{M_t-x_t} = \frac{e^{a_D NPV_t + b_D \log(x_t/M_t) + c_D + \xi_t}}{1+e^{a_D NPV_t + b_D \log(x_t/M_t) + c_D + \xi_t}}$. Then $1 - \lambda_t = \frac{M_t - x_{t+1}}{M_t - x_t} = \frac{1}{1+e^{a_D NPV_t + b_D \log(x_t/M_t) + c_D + \xi_t}}$. Therefore:

$$\log\left(\frac{x_{t+1}-x_t}{M_t - x_{t+1}}\right) = \log\left(\frac{\lambda_t}{1-\lambda_t}\right) = a_D NPV_t + b_D \log(x_t/M_t) + c_D + \xi_t \tag{3.13}$$

We consider the unobserved demand shock ξ_t as the error measure and use a generalized method of moments approach to estimate the relation in (3.13). Note that a necessary condition for (3.13) to hold is $M_t > x_{t+1} > x_t$. This condition is generally true for the adoption of any new technology, since demand is always positive and the market size is still far from the number of adopted customers (see Figure 3-3). Table 3.2 displays the estimation results. The significance level for the estimates in Table 3.2 are indicated by a p-value less than 1% for b_D (marked **) and less than 10% for a_D and c_D (marked *).

	Estimate	Std. Error
a_D	0.164*	(0.112)
b_D	0.657**	(0.240)
c_D	−2.891*	(1.592)
R^2	0.957	
$\bar{R}^2$	0.950	

Table 3.2: Estimation Results for the Demand Model

The estimation results from Tables 3.1 and 3.2 seem to present a good fit to the historical data for the cost and demand curves. We can now use our calibrated model to forecast future adoption levels and solve the policy-making problem to obtain insights about the situation of the German market.

3.3.2 Forecasting and Policy Optimization

Using the model estimated in Section 3.3.1, in Figure 3-5 we forecast the future baseline adoption levels, x_t^B, using forecasts of the number of households and future feed-in rates. Define this baseline adoption path as the natural adoption process if we do not intervene on the subsidy levels ($r_t = 0$).

Using the installation distribution in 2009, we can infer the total (both residential and non-residential) PV installed capacity for the following few years, 2008-2013. In Figure 3-6, we compare our results with a well recognized forecast benchmark from the European Photovoltaic Industry Association (EPIA). Our baseline predictions for the total installed solar generation capacity in Germany by 2013 are 11.5% above the EPIA conservative (status-quo) forecast and 10.4% below the EPIA aggressive (stronger policy) forecast. This comparison serves as a sanity check for us to trust the forecasting ability of our model.

Using the estimated model, we also demonstrate how to use the policy design tool developed in Section 3.2.2 with a hypothetical adoption target. Starting from 2008, consider the target adoption level for 2030 to be at our baseline adoption forecast at $x_T^B = 12.3\%$. In this case, we observe that by readjusting the current subsidy policy, we can obtain net present value savings of 32.5 billion

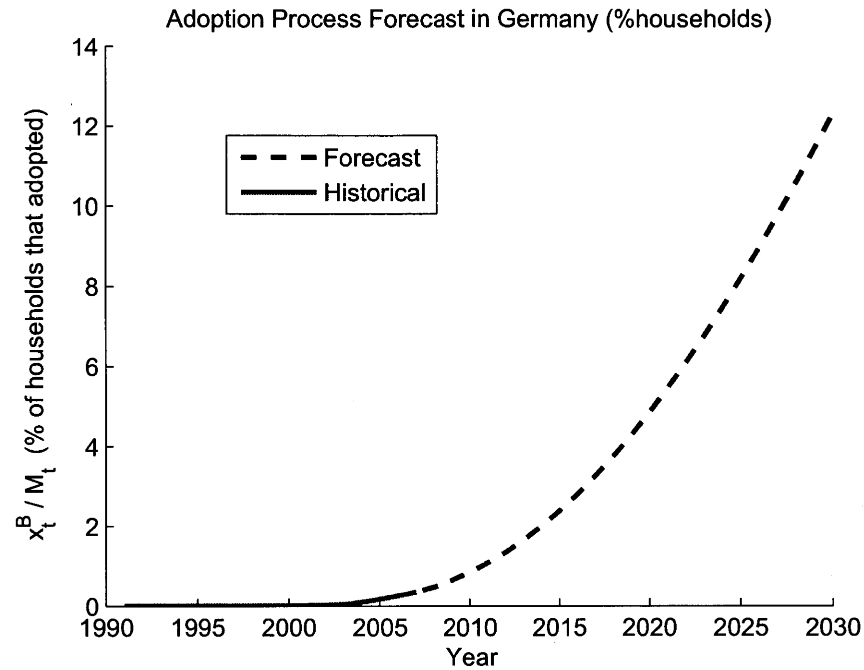

Figure 3-5: Baseline forecast of adoption ($r_t = 0$) in Germany as a proportion of the households

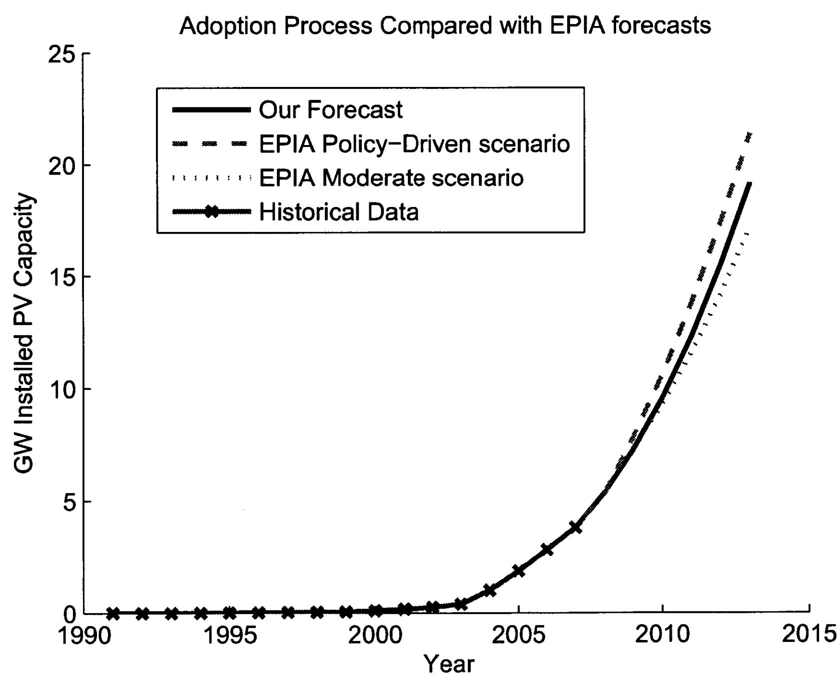

Figure 3-6: Comparing our baseline forecast with EPIA benchmark

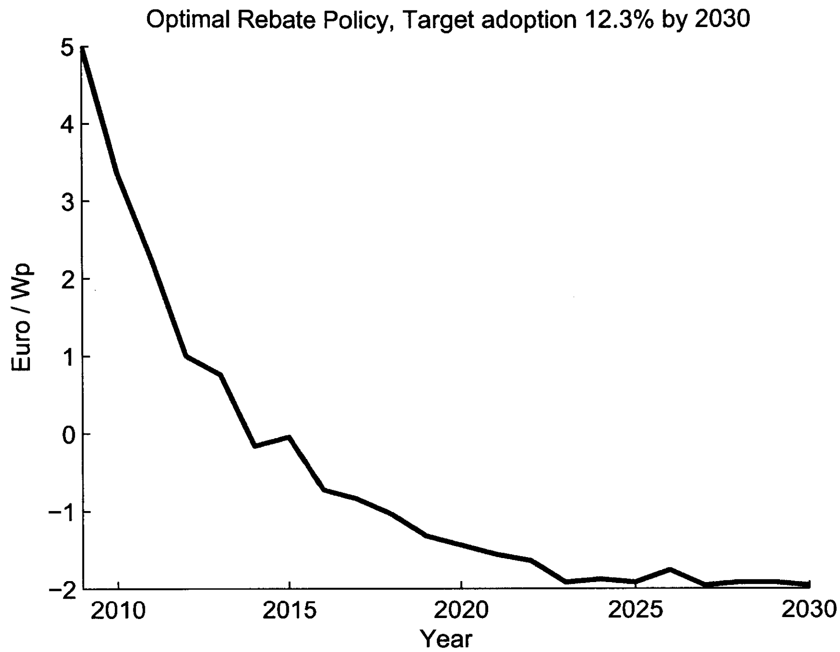

Figure 3-7: Optimal Rebate Policy: Target 12.3% by 2030

Euros, over the next 22 years. In Figure 3-7, we display the optimal rebate strategy. This strategy is computed by numerically solving the dynamic program in (3.12) and it displays positive rebates in the early stages and negative rebates after 2015. In other words, this rebate structure could translate into an increase in subsidies in the first few years and the removal of some of the current subsidies later on (possibly smaller Feed-in-Tariffs or higher sales taxes). The jerkiness of this plot is due to the rough discretization used to solve the dynamic program. Note also that the rebate structure is decreasing over time. This is consistent with the assumption that consumers should have no incentives to be strategic about their purchase timing decision.

By looking at the structure of the optimal rebate path in Figure 3-7, we can see that there are three forces defining this optimal policy: The first one increases the subsidies at the beginning of the planning horizon, in order to kick-start the effects of the positive network externalities. The second contradicting force comes from the discounted nature of the problem, favoring later subsidies. The third force is a free-riding effect, where subsidizing is cheaper at later periods because network externalities have already taken effect. The combination of these three effects will make the optimal rebate path distribute rebates in a non-trivial manner. In other words, it is not optimal to waste all our subsidizing efforts at the first stage, but instead there is an efficient way to distribute the rebates along the time horizon with minimal cost to the system.

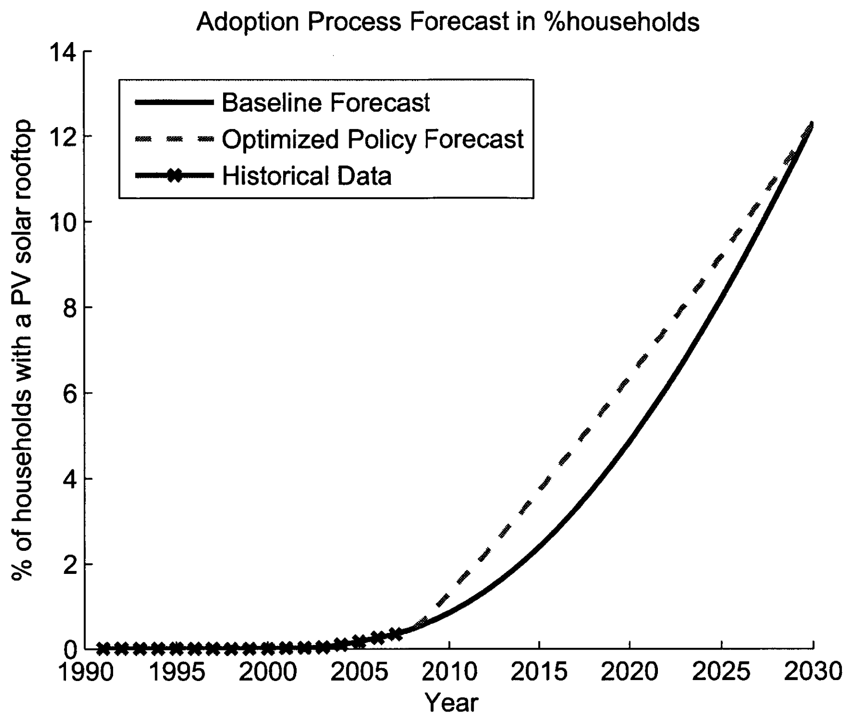

Figure 3-8: Adoption Forecast: Target 12.3% by 2030

Figure 3-8 displays the adoption forecast using this optimal rebate policy and compare it to the baseline adoption path. Additionally, Figure 3-9 displays the forecasted evolution of installation costs under the baseline path and with optimized policies.

In order to understand the trade-off between the adoption target level established for the year 2030 and the cost it will incur for the government, we ran the policy optimization for multiple target levels, ranging from 1% to 25%. In Figure 3-10, we observe that below a 16.3% adoption level, the government can actually save money by optimally managing the subsidy policy. This is consistent with Figures 3-7-3-8, where we display the optimal rebate and adoption path for a particular adoption target of 12.3%.

In Section 3.4, we explore further the analytical structure of the optimal policy and target cost function. More specifically, in Theorem 3.4.1 we prove that the government's cost function is convex as a function of the adoption target, as we observed in Figure 3-10. This convexity result requires some mild assumptions on the installation cost and demand model parameters which are clearly satisfied for our empirical study. We also further analyze in Section 3.4.1 the behavior of the optimal solution as we change the target adoption level for a two-stage problem, where we conclude that one of the reasons why increasing the adoption targets becomes increasingly more expensive is

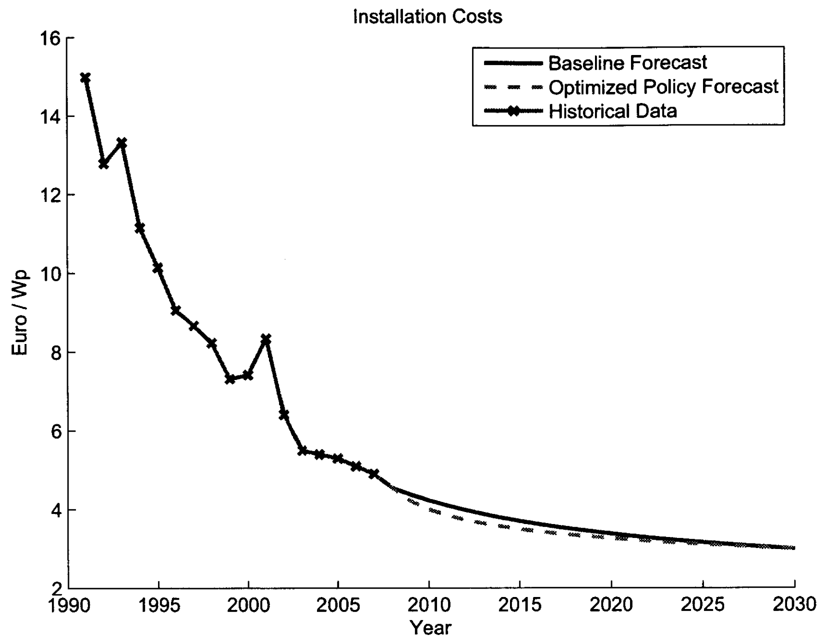

Figure 3-9: Installation Costs: Target 12.3% by 2030

because of the saturation of the network externality benefits.

For our last experiment, we varied the target adoption deadline T from 2009 to 2030 and assumed that the target adoption rate was our baseline estimate for the adoption level at that given year, x_T^B, as seen in Figure 3-5. We then optimized the subsidy policy for that given target x_T^B and observed the government cost for achieving that same target level by time T. The motivation of this experiment is to reverse engineer what could potentially be the government's current subsidy policy motivation or determine if the current policy is suboptimal. If the government was in fact optimally designing the current subsidies to reach an adoption target at any of these years, in theory, our baseline should forecast the optimal adoption path to that target. In other words, the optimal rebate from the optimization model should be $r_t = 0$, for all $t \leq T$, and the potential cost improvement of changing the subsidy policy should also be zero.

In fact, we observe that for any target deadline T between 2009 and 2030, there is a cheaper way to achieve the same adoption level as the baseline forecast predicts, x_T^B. Figure 3-11 displays the result of this experiment. We note that the potential cost savings is always positive for any target adoption level in the baseline forecasted adoption path. This indicates that the current design of the subsidies are not optimally designed for any potential adoption target.

Another hypothesis is that the government is actually maximizing some measure of social wel-

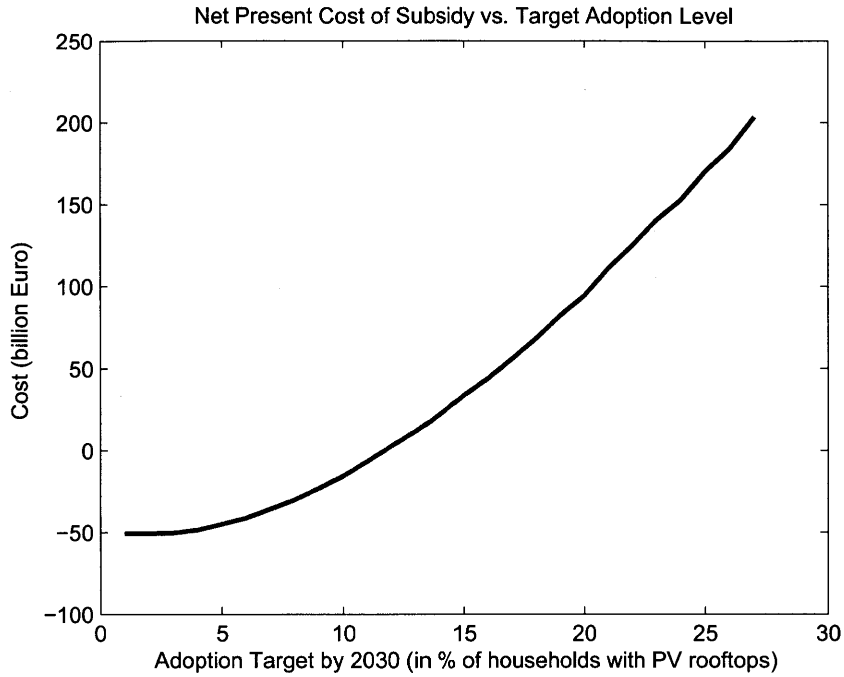

Figure 3-10: Trade-off between adoption target by 2030 and net system cost (or savings)

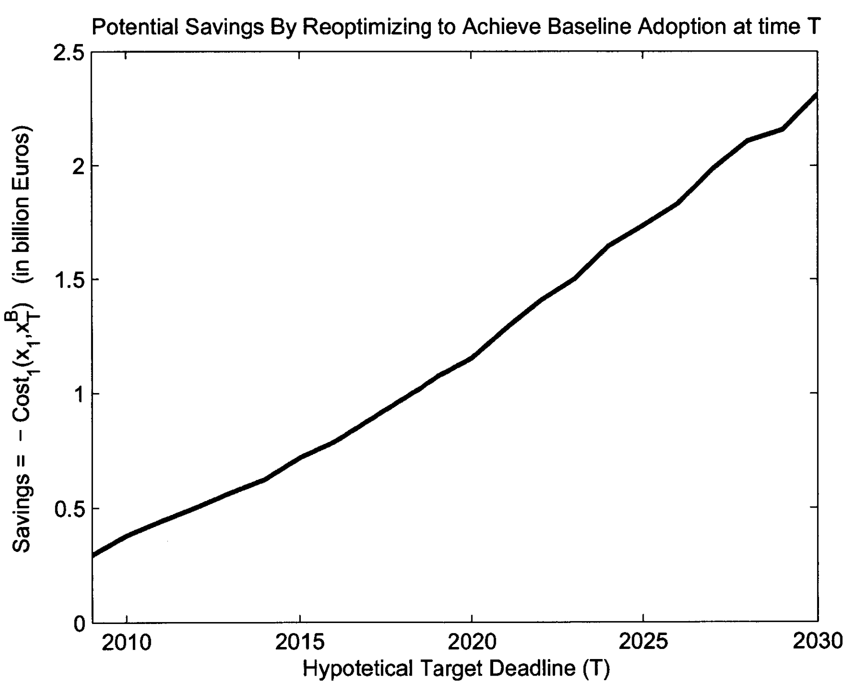

Figure 3-11: Cost saving by optimizing subsidies with target adoption x_T set at the baseline x_T^B

fare, as opposed to trying to achieve some given target, and could potentially be optimal under that objective. We discuss the welfare problem in more detail in Section 3.4.2. From Theorem 3.4.2, we show that because the optimal rebate structure found by our dynamic program is decreasing in time, i.e., $r_t \geq r_{t+1}$ (see for example, Figure 3-7), then our optimized adoption path is always above the baseline adoption path $x_t^* \geq x_t^B$ (see Figure 3-8). Therefore the cumulative welfare benefits from solar panel adoption will always be higher in our optimized solution. Our new solution will not be necessarily the optimal welfare solution, but it will certainly provide higher social welfare than the current baseline path. This shows that the government is still acting suboptimally, even from a social welfare perspective.

The result that the current subsidy policy is suboptimal, both from a target adoption and from a welfare perspective, needs to be evaluated carefully. Throughout our modeling process we have made many assumptions about consumer behavior and the demand structure that need further exploration. Also, a few simplifying assumptions were made simply because of the lack of detailed data on the solar market. That being said, we believe this empirical study developed a first step in analyzing this issue and raises a clear warning sign about the economic efficiency of the current policy. We believe these experiments can be further improved by updating the data set and possibly using a more detailed demand model.

3.4 Analysis of the Policy-Maker's Problem

In this section, we explore some of the theoretical insights that can be obtained by analyzing the structure of the optimization model we developed in Section 3.2.2 for the policy design problem. Consider the problem faced by the policy maker in (3.10), where x_1 is the initial number of solar panels sold and x_T is the given adoption target. As before, we control the adoption levels by adjusting the rebate rates r_t. For a full notation summary, see Appendix 3.6.1. We will make a few technical assumptions about the parameters of the model that are necessary for the analysis:

Assumption 3.4.1

a) $M_t > x_{t+1} > x_t$, *for all* $t = 1, ..., T-1$.

b) $b_I < 0$.

c) $a_D > 0$, $b_D > 0$, *and* $a_D + b_D \leq 1$.

Assumption 3.4.1.a means that we are only concerned with the problems where the market is still under development and the potential market size is greater than the number of panels sold at any point of our decision horizon. Assumption 3.4.1.b is true by the nature of the learning-by-doing effect, which decreases installation cost with the number of installed panels. Assumption 3.4.1.c is a technical assumption that we use to obtain convexity of the government's cost function. Assumptions $a_D > 0$ and $b_D > 0$ hold due to the nature of the demand model, but the last part $a_D + b_D \leq 1$ is not as obvious. The top-level intuition behind it is that the benefits of network externalities such as learning-by-doing and information spread may have a concave impact on the system cost. Without these network effects, the nature of the logit model alone could guarantee convexity for the cost function. This assumption, $a_D + b_D \leq 1$, guarantees that the concave network effects do not overshadow the convexity of the demand model. In fact, this condition is easily satisfied in the empirical study of Section 3.3.1.

We will now show that the total present system cost $Cost_1(x_1, x_T)$ is convex in the future target adoption level. As the demand function $q(x_t, r_t)$ is monotonically increasing in r_t, we can easily invert the relation and express the rebate as a function of the desired demand, q_t.

$$r_t(x_t, q_t) = k(x_t) - d_t - \frac{b_D \log(x_t/M_t) + c_D + \xi_t - \log(-q_t/(q_t - M_t + x_t))}{a_D}$$

Note also, that demand is determined for a given a adoption path $q_t = x_{t+1} - x_t$

$$r_t(x_t, x_{t+1}) = k(x_t) - d_t - \frac{b_D \log(x_t/M_t) + c_D + \xi_t - \log\left(\dfrac{x_{t+1} - x_t}{M_t - x_{t+1}}\right)}{a_D} \tag{3.14}$$

For simplicity, consider the following 3-period model (T=3), where x_1 is the initial state and x_3 is the final target state. The only decision to be made is where the middle state x_2 should be placed. Once x_2 is decided, the rebates for both periods will be determined by $r_1(x_1, x_2)$ and $r_2(x_2, x_3)$ according to equation (3.14). By controlling directly the adoption path and not the rebates, we can deal with a single variable unconstrained problem, instead of a two-variable problem with balance constraints. Define the inside cost function:

$$J_1(x_2, x_3) = r_1(x_1, x_2)(x_2 - x_1) + \delta r_2(x_2, x_3)(x_3 - x_2)$$

Then the policy maker's problem can be reformulated as:

$$Cost_1(x_1, x_3) = \min_{x_2} J_1(x_2, x_3) \tag{3.15}$$

The following lemma will be used to show the convexity of the total system cost function.

Lemma 3.4.1 *$J_1(x_2, x_3)$ is jointly convex in x_2 and x_3.*

The proof of Lemma 3.4.1 is very heavy in algebraic manipulations and therefore placed in Appendix 3.6.2 to streamline the reading. Given the convexity of $J_1(x_2, x_3)$ in x_2, we know that the optimal solution $x_2^*(x_3)$ comes from the solution of the first order optimality condition, $\frac{dJ}{dx_2}(x_2^*(x_3), x_3) = 0$. From the joint convexity of J_1, we can also obtain the following results.

Corollary 3.4.1 *Let $x_2^*(x_3)$ be the optimal adoption path for a given target x_3. Then $Cost_1(x_1, x_3)$ $= J_1(x_2^*(x_3), x_3)$ is a convex function of x_3.*

The proof of this Corollary 3.4.1 is a well known result from convex analysis and it comes directly from the joint convexity of the inner function $J_1(x_2, x_3)$ (for further reference see (Boyd and Vandenbergh 2004)). We will use Corollary 3.4.1 in the proof of convexity for the T-period case in Theorem 3.4.1.

Another interesting outcome of Lemma 3.4.1 is stated below in Corollary 3.4.2. This corollary is derived from the Implicit Function Theorem (see (Bertsekas 1995))

Corollary 3.4.2 *Let $x_2^*(x_3)$ be the optimal adoption path for given target x_3. Then the first order optimality condition on x_3 implies: $\frac{dx_2^*}{dx_3}(x_3) = -\frac{d^2J}{dx_2 dx_3}(x_2^*(x_3), x_3)\left(\frac{d^2J}{dx_2^2}(x_2^*(x_3), x_3)\right)^{-1}$.*

The above result, Corollary 3.4.2, will be later used to develop insights about the structure of the optimal solution. For now, we will focus on the convexity result. In the original T-period problem, we obtain the following result.

Theorem 3.4.1 *$Cost_1(x_1, x_T)$ is convex in x_T.*

The intuition behind the proof of Theorem 3.4.1 is to show convexity for an additional period $T = 4$ and $Cost_1(x_1, x_4) = \min_{x_3} Cost_1(x_1, x_3) + \delta^2 r_3(x_3, x_4)(x_4 - x_3)$. By induction, we can show that the cost for any time horizon T is a convex function of the target. Once again, the derivation of this proof is relegated to Appendix 3.6.3.

This result may be useful in order to extend this model into many future research directions, including solving the problem with uncertainty (with randomness in demand and/or technological progress) or introducing multiple products (for example, different installation sizes). In these cases, one possible approach would be to use approximate dynamic programming, which may require some convexity structure of the value function.

3.4.1 Insights on the optimal solution

By examining at the 3-period problem defined in (3.15), we can obtain some insights on the structure of the optimal solution and optimal system cost. Consider the first order condition: $\frac{dJ}{dx_2}(x_2^*(x_3), x_3) = 0$. By rearranging the terms of this equation, we obtain:

$$\frac{M_1 - x_1}{a_D(M_1 - x_2^*(x_3))} + r_1(x_1, x_2^*(x_3)) + \delta\left[k'(x_2^*(x_3)) - \frac{b_D}{a_D x_2^*(x_3)} - \frac{1}{a_D} - r_2(x_2^*(x_3), x_3)\right] = 0$$

This can be used to find the optimal mid-point adoption level $x_2^*(x_3)$, located where the marginal cost of increasing the level in the first period is the same as the marginal benefit from the second period. In the first period, for each marginal unit of x_2 that we increase over the optimal, we incur a marginal cost of the rebate price $r_1(x_1, x_2^*(x_3))$, plus a rebate adjustment of $\frac{M_1 - x_1}{a_D(M_1 - x_2^*(x_3))}$ needed to meet the higher demand in this first period. In the second period, the marginal unit increase in x_2 will lower the overall system cost (with discounting δ) by the rebate $r_2(x_2^*(x_3), x_3)$ adjusted for the network externalities gain $k'(x_2^*(x_3)) - \frac{b_D}{a_D x_2^*(x_3)}$ and also for the fact that we need to serve a lower demand, which also affects the rebate level by $\frac{1}{a_D}$. Strict convexity of $J(x_2, x_3)$ in x_2 guarantees that this equation has a monotonically increasing left hand side, which means that there is a unique optimal solution and it can be easily computed numerically.

It is only natural to ask how the optimal mid-point adoption level changes with the adoption target level. From Corollary 3.4.2, we have:

$$\frac{dx_2^*}{dx_3}(x_3) = -\frac{d^2 J}{dx_2 dx_3}(x_2^*(x_3), x_3)\left(\frac{d^2 J}{dx_2^2}(x_2^*(x_3), x_3)\right)^{-1}$$

We can then show the following properties of the optimal solution:

Proposition 3.4.1 *Let $x_2^*(x_3)$ be the optimal mid-point level. Then $\frac{dx_2^*}{dx_3}(x_3) > 0$, which implies that the optimal x_2 is strictly increasing in the target x_3. If we also have that $x_2 - x_1 \leq M_2 - x_3$, then we can also show $\frac{dx_2^*}{dx_3}(x_3) < 1$.*

See Appendix 3.6.4 for a proof of Proposition 3.4.1. We use this result to get intuition about how the system cost changes as a function of the adoption target.

Consider the variation in the optimal rebate levels. We can express them as:

$$\begin{aligned}
\frac{dr_1}{dx_3}(x_1, x_2^*(x_3)) = & \left[\frac{1}{a_D(x_2^*(x_3)-x_1)} + \frac{1}{a_D(M_1-x_2^*(x_3))}\right]\frac{dx_2^*}{dx_3}(x_3) \\
\frac{dr_2}{dx_3}(x_2^*(x_3), x_3) = & \frac{1}{a_D(x_3-x_2^*(x_3))} + \frac{1}{a_D(M_2-x_3)} \\
& + \left[k'(x_2^*(x_3)) - \frac{b_D}{a_D x_2^*(x_3)} - \frac{1}{a_D(x_3-x_2^*(x_3))}\right]\frac{dx_2^*}{dx_3}(x_3)
\end{aligned}$$

The derivative of the system cost can also be expressed as:

$$\begin{aligned}
\frac{dCost_1}{dx_3}(x_1, x_3) = & \frac{dr_1}{dx_3}(x_1, x_2^*(x_3))(x_2^*(x_3) - x_1) + r_1(x_1, x_2^*(x_3))\frac{dx_2^*}{dx_3}(x_3) \\
& + \delta\frac{dr_2}{dx_3}(x_2^*(x_3), x_3)(x_3 - x_2^*(x_3)) + \delta r_2(x_2^*(x_3), x_3)\left(1 - \frac{dx_2^*}{dx_3}(x_3)\right)
\end{aligned}$$

In order to develop intuition about the optimal cost variation, assume we are working under the regime where $M_t - x_{t+1} >> x_{t+1} - x_t$, for both $t = 1, 2$. This is the case in any solar market today and for the foreseeable future. For example, in the German case studied in the previous section the amount of solar capacity installed is not even 1% of the potential market size. For this reason we reformulate the optimal rebate derivatives with an approximation, where the terms $\frac{1}{a_D(M_1-x_2^*(x_3))}$ and $\frac{1}{a_D(M_2-x_3)}$ go to zero. The cost derivative can be approximated by:

$$\begin{aligned}
\frac{dCost_1}{dx_3}(x_1, x_3) \cong & \left[\frac{1}{a_D} + r_1(x_1, x_2^*(x_3))\right]\frac{dx_2^*}{dx_3}(x_3) + \\
& \delta\left[\frac{1}{a_D} + r_2(x_2^*(x_3), x_3)\right]\left(1 - \frac{dx_2^*}{dx_3}(x_3)\right) + \\
& \delta(x_3 - x_2^*(x_3))\left[k'(x_2^*(x_3)) - \frac{b_D}{a_D x_2^*(x_3)}\right]\frac{dx_2^*}{dx_3}(x_3)
\end{aligned} \quad (3.16)$$

Each new marginal unit of target adoption level x_3 will need to be distributed into the first and second period of sales determined by $\frac{dx_2^*}{dx_3}(x_3)$ and $(1 - \frac{dx_2^*}{dx_3}(x_3))$ respectively. Note from the first term in the above equation that a marginal increase in the target level will increase the mid-point level by $\frac{dx_2^*}{dx_3}(x_3)$ and each additional unit of mid-point adoption level will cost an additional $\frac{1}{a_D} + r_1(x_1, x_2^*(x_3))$ to the system, where $\frac{1}{a_D}$ is the rebate adjustment due to increased demand in the first period. The cost of the additional target level units allocated to the second period will be $\left[\frac{1}{a_D} + r_2(x_2^*(x_3), x_3)\right]\left(1 - \frac{dx_2^*}{dx_3}(x_3)\right)$, where $\frac{1}{a_D}$ is the adjustment in the second rebate price due to higher demand. On the other hand, each unit of mid-point level increase will save the system some money on the second period because of network externalities, which is represented by the last term on the equation. The externality benefits affect all sales made in the second period $(x_3 - x_2^*(x_3))$, not just the new additional units required for the marginal target increase.

With the relation in (3.16), the policy maker can obtain the trade-offs of raising the target adoption level, without having to resolve the entire system cost. If x_3 is already very high, it is likely that the cost benefits due to the network externalities are saturated, as $k'(x_2^*(x_3)) - \frac{b_D}{a_D x_2^*(x_3)}$ will increase and approach zero as we increase x_3 (note that $k'(x) < 0$ and $k''(x) > 0$). Then raising the target levels become increasingly more expensive, which is one reason why the cost function $Cost_1(x_1, x_3)$ is convex, as we concluded in Theorem 3.4.1.

3.4.2 Welfare Maximization

Perhaps more common in the economics literature, the objective of a policy optimization problem can be expressed as a social welfare maximization problem. In this particular case, it is debatable how one should quantify the benefits of developing the solar industry for a particular government. There is obviously a global benefit for clean electricity generation, but the local benefits from avoiding carbon emissions cannot be rewarded to a single state or country unless we develop an efficient global carbon market. Other pollutants have more local impact, but in general pollution avoidance cannot solely justify solar technology, as there are other technologies that are much more cost efficient (from wind generation to building retro-fitting). On the other hand, there are less tangible benefits of stimulating the solar technology by a particular government. These include generation portfolio diversification, peak-load reduction, development of a local solar manufacturing and installation industry.

For the reasons above, we have so far waived the welfare discussion and assumed that policy-makers have a given strategic adoption target. We further demonstrated how the cost behaves for different target levels, which could potentially aid policy-makers when setting such targets. If we could quantify all the benefits of solar adoption, then we could use welfare maximization to find these targets with a slight modification of our optimization model.

Suppose we are given a benefit function $Benefit_t(x_t)$ which depends on the realized adoption at each stage. This would be the case if there was a given price for CO_{2_t} at time t and every solar panel installed saves the country that amount of money for avoiding carbon emissions. See for example (Benthem, Gillingham, and Sweeney 2008) for an example of policy optimization with carbon externality costs. Naturally, this function $Benefit_t(x_t)$ should be increasing on the adoption level at x_t. Define the social welfare objective as:

$$\mathit{Welfare}(x_1, ..., x_T) = \sum_{t=1}^{T-1} \left[\mathit{Benefit}_t(x_t) - \delta^{t-1} r_t(x_t, x_{t+1})(x_{t+1} - x_t)\right]$$

The optimal welfare problem can be solved by maximizing the expression above, which can be done numerically by solving a dynamic program. Next, we define a useful result connecting the optimization for adoption targets with social welfare efficiency.

Theorem 3.4.2 *Let x_t^B be the baseline adoption path, where $r_t = 0$, for all $t = 1, ..., T-1$. Consider x_T^B as the adoption target in the optimization model (3.10) and let x_t^* be the optimal adoption path for this model. If the overall system cost is negative, $Cost_1(x_1, x_T^B) < 0$, and the optimal rebate path is non-increasing, $r_t(x_t^*, x_{t+1}^*) \geq r_{t+1}(x_{t+1}^*, x_{t+2}^*)$, then the welfare of the new optimized path is greater than the welfare under the baseline path, i.e.,*

$$\mathit{Welfare}(x_1, x_2^*, ..., x_{T-1}^*, x_T^B) > \mathit{Welfare}(x_1, x_2^B ..., x_{T-1}^B, x_T^B)$$

Note that Theorem 3.4.2 connects with our findings in the empirical study at the end of Section 3.3.2. In that part of the study, we argue that the current subsidy policy for the German solar market is suboptimal, even from a welfare maximization perspective. The proof of Theorem 3.4.2 is displayed in Appendix 3.6.5.

3.5 Conclusions

In summary, we model the adoption of solar photovoltaic technology as a diffusion process where customers are assumed to be rational agents following a discrete choice model. We show how this framework can be used by a policy maker to design optimal incentives in order to achieve a desired adoption target with minimum cost for the system. In particular, this policy design model takes into consideration network externalities such as information spread and cost improvements through learning-by-doing. To demonstrate the applicability of this framework, we develop an empirical study of the German photovoltaic market and show how this model can be fitted to actual market data and how it can be used for forecasting and subsidy policy design. Finally, we analyze the structure of the optimal solution of the subsidy design problem to obtain insights about the government's subsidizing cost and to understand how this adoption target optimization problem can be related to the welfare maximization problem.

We show in our numerical experiments that in the early stages of the adoption process, it is optimal for the government to provide strong subsidies, which take advantage of network externalities to reach the target adoption level at a lower cost. As the adoption level increases, these network externalities become saturated and the price paid for raising the adoption target becomes increasingly more expensive. In particular, we are able to prove analytically that the system cost is a convex function of the adoption target. This convex trade-off between adoption targets and subsidy cost was also evident in our empirical study. We believe that this framework for quantifying the cost of adoption targets could be a very useful tool for the policy-makers that design these targets.

We also observe in this empirical study that the current subsidy policy in Germany is not being efficiently managed. There is a large potential real-world impact in this policy recommendation, which can be further improved with a more detailed demand model and data collection. With the current data and modeling set-up, we can argue the suboptimality of the German policy because for every possible adoption target chosen along the baseline adoption path, there is a better way to reach the same target at a lower cost for the system. More specifically, we can achieve that by raising early subsidies and lowering future subsidies. Finally, we proved that even if the government is maximizing social welfare instead of minimizing cost of achieving a target, the current subsidy policy is still suboptimal.

3.6 Appendix

3.6.1 Notation

The following notation summary is useful for reference throughout the chapter.

M_t	: Market size at year t, equal to number of households
x_t	: Number of household solar panels installed up to year t
x_T	: Target adoption level
T	: Length of the policy time horizon, also known as target adoption deadline
q_t	: Demand for household solar panels installed at year t (kWp)
r_t	: Government installation rebate (€/kWp)
k_t	: Solar installation cost at year t, including labor, module and hardware (€/Wp)
δ_c, δ_g	: Discount rate of the customers and government, respectively
d_t	: Discounted future cash flows of a solar installation (€/kWp)
FIT_t	: Feed-in-Tariff at year t times the average electricity output (€/kWp)
OM_t	: Operation and Maintenance cost (€/kWp)
$AvgSize$	: Average household installation size (kWp)
T_{mod}	: Lifetime of the module
NPV_t	: Net Present Value of an average sized solar installation purchased at t (€)
a_I, b_I	: Installation cost parameters, from learning-by-doing effect
a_D, b_D, c_D	: Demand model parameters or consumer utility function parameters
ξ_t	: Unobserved demand shocks at time t
$\epsilon_{t,i}$	: Random utility component for customer i at time t
$U_{t,i}$	: Utility of purchasing a solar panel for customer i at time t
V_t	: Nominal utility that the average consumer has for purchasing a solar panel
$Cost_t(x_t, x_T)$	: Subsidy cost at time t, starting from adoption level x_t and reaching target x_T

3.6.2 Proof of Lemma 3.4.1

The definition of $J_1(x_2, x_3)$ is given by:

$$J_1(x_2, x_3) = r_1(x_1, x_2)(x_2 - x_1) + \delta r_2(x_2, x_3)(x_3 - x_2) = f_1(x_2) + \delta f_2(x_2, x_3)$$

We can prove the joint convexity by parts. The first part: $f_1(x_2) = r_1(x_1, x_2)(x_2 - x_1)$ needs only to be proved convex in x_2. By taking the second derivative of this term in x_2 we obtain:

$$\frac{d^2 f_1}{dx_2^2} = (M_1 - x_1)^2/(a_D(M_1 - x_2)^2(x_2 - x_1)) > 0$$

To prove joint convexity of $J_1(x_2, x_3)$ it remains to show that the principal components of the Hessian of $f_2(x_2, x_3)$ are also positive. In particular we will split up the function f_2 in two parts. Choose $\alpha_1, \alpha_2 > 0$ such that $\alpha_1 + \alpha_2 = 1$.

$$g_1(x_2, x_3) = \left(k(x_2) + \frac{\alpha_1}{a_D} \log\left(\frac{x_3 - x_2}{M_2 - x_3}\right)\right)(x_3 - x_2)$$

$$g_2(x_2, x_3) = \left(-\frac{b_D \log(x_2)}{a_D} + \frac{\alpha_2}{a_D} \log\left(\frac{x_3 - x_2}{M_2 - x_3}\right)\right)(x_3 - x_2)$$

We can rewrite $f_2(x_2, x_3)$ as:

$$f_2(x_2, x_3) = g_1(x_2, x_3) + g_2(x_2, x_3) - \left(d_2 + \frac{c_D}{a_D} - \frac{b_D \log(M_2)}{a_D}\right)(x_3 - x_2)$$

At this point, we just need to prove joint convexity of g_1 and g_2. In particular, we will show that the principal components of the Hessian of g_1 and g_2 are positive. For the first term, the first second derivative is given by:

$$\frac{d^2 g_1}{dx_2^2} = k''(x_2) - 2k'(x_t) + \frac{\alpha_1}{a_D(x_3 - x_2)} > 0$$

This is positive simply by verifying each term. Note here that $k''(x_2) > 0$ and $k'(x_2) < 0$ because of the decreasing and convex nature of the learning function given by $b_I < 0$.

It remains to show that the determinant of the Hessian of g_1 is positive:

$$\frac{d^2 g_1}{dx_2^2}\frac{d^2 g_1}{dx_3^2} - \left(\frac{d^2 g_1}{dx_2 dx_3}\right)^2 \geq 0$$

With some algebraic manipulations, we can show that the expression above reduces to:

$$k''(x_2)\frac{\alpha_1}{a_D}\left(\frac{M_2 - x_2}{M_2 - x_3}\right)^2 - 2k'(x_2)\frac{\alpha_1}{a_D}\frac{M_2 - x_2}{(M_2 - x_3)^2} - (k'(x_2))^2 \geq 0$$

In particular, $k''(x_2) = \frac{b_I(b_I - 1)}{x_2^2}k(x_2)$ and $k'(x_2) = \frac{b_I}{x_2}k(x_2) = \frac{x_2}{b_I - 1}k''(x_2)$. Also $k'(x_2)^2 = \frac{b_I}{b_I - 1}k''(x_2)$ Then the expression above becomes:

$$k''(x_2)\frac{\alpha_1}{a_D}\left(\frac{M_2-x_2}{M_2-x_3}\right)^2 - 2\frac{x_2}{b_I-1}k''(x_2)\frac{\alpha_1}{a_D}\frac{M_2-x_2}{(M_2-x_3)^2} - \frac{b_I}{b_I-1}k''(x_2) \geq 0$$

This can be reduced to:

$$\frac{\alpha_1}{a_D}\frac{M_2-x_2}{(M_2-x_3)^2}\left(M_2-x_2-\frac{2x_2}{b_I-1}\right) \geq \frac{b_I}{b_I-1}$$

In particular $-\frac{2x_2}{b_I-1} > 0$, then $\frac{M_2-x_2}{(M_2-x_3)^2}\left(M_2-x_2-\frac{2x_2}{b_I-1}\right) > \left(\frac{M_2-x_2}{M_2-x_3}\right)^2 > 1$ where the last inequality comes from $M_2 > x_3 > x_2$ of Assumption 3.4.1.a. Also from Assumption 3.4.1.b, we have that $1 > \frac{b_I}{b_I-1}$. Therefore the equation we want to prove can be implied if $\frac{\alpha_1}{a_D} \geq 1$, since:

$$\frac{\alpha_1}{a_D}\frac{M_2-x_2}{(M_2-x_3)^2}\left(M_2-x_2-\frac{2x_2}{b_I-1}\right) > \frac{\alpha_1}{a_D} \geq 1 > \frac{b_I}{b_I-1}$$

Let $\alpha_1 = a_D$ and we have proven the joint convexity of g_1. For the joint convexity of g_2 we need to do a similar algebraic manipulation. Note that the first term in g_2 is $-b_D\log(x_2)/a_D$ is the equivalent of $k(x_2)$ in the proof of g_1. The first component of the Hessian matrix will be:

$$\frac{d^2g_2}{dx_2^2} = b_D/(a_Dx_2^2) + 2b_D/(a_Dx_2) + \frac{\alpha_2}{a_D(x_3-x_2)} > 0$$

This is also positive. The determinant of the Hessian of $g_2(x_2,x_3)$ can be expressed as:

$$(a_D/b_D)(-b_D/(a_Dx_2))^2\frac{\alpha_2}{a_D}\left(\frac{M_2-x_2}{M_2-x_3}\right)^2$$
$$-2(-b_D/(a_Dx_2))\frac{\alpha_2}{a_D}\frac{M_2-x_2}{(M_2-x_3)^2} - (-b_D/(a_Dx_2))^2 \geq 0$$

This is equivalent to:

$$(a_D/b_D)(-b_D/(a_Dx_2))\frac{\alpha_2}{a_D}\left(\frac{M_2-x_2}{M_2-x_3}\right)^2 - 2\frac{\alpha_2}{a_D}\frac{M_2-x_2}{(M_2-x_3)^2} - (-b_D/(a_Dx_2)) \leq 0$$

$$(-b_D/(a_Dx_2))\left(\frac{(a_D/b_D)\alpha_2}{a_D}\left(\frac{M_2-x_2}{M_2-x_3}\right)^2 - 1\right) - 2\frac{\alpha_2}{a_D}\frac{M_2-x_2}{(M_2-x_3)^2} \leq 0$$

The expressions above will be true if $\left(\frac{(a_D/b_D)\alpha_2}{a_D}\left(\frac{M_2-x_2}{M_2-x_3}\right)^2 - 1\right) \geq 0$, which can be reformulated into:

$$\frac{\alpha_2}{b_D} \geq \left(\frac{M_2-x_3}{M_2-x_2}\right)^2$$

Since $M_2 > x_3 > x_2$, we need to show only that $\alpha_2 \geq b_D$, then:

$$\frac{\alpha_2}{b_D} \geq 1 > \left(\frac{M_2 - x_3}{M_2 - x_2}\right)^2$$

In particular we had chosen $\alpha_2 = 1 - \alpha_1 = 1 - a_D$. From Assumption 3.4.1.c, we have that $a_D + b_D \leq 1$, which implies that $\alpha_2 \geq b_D$, which concludes our proof that $g_2(x_2, x_3)$ is jointly convex in (x_2, x_3). Together with the earlier proofs that g_1 and f_1 are jointly convex as well, we have that $J_1(x_2, x_3)$ is jointly convex.

3.6.3 Proof of Theorem 3.4.1

Consider the 4-period problem ($T = 4$):

$$\begin{aligned} Cost_1(x_1, x_4) = \quad & \min_{x_2, x_3} r_1(x_1, x_2)(x_2 - x_1) + \delta r_2(x_2, x_3)(x_3 - x_2) + \delta^2 r_3(x_3, x_4)(x_4 - x_3) \\ & \min_{x_3} Cost_1(x_1, x_3) + \delta^2 r_3(x_3, x_4)(x_4 - x_3) \end{aligned}$$

We need to show that $Cost_1(x_1, x_3) + \delta^2 r_3(x_3, x_4)(x_4 - x_3)$ is jointly convex in (x_3, x_4). We already know that $Cost_1(x_1, x_3)$ is convex in x_3 from Corollary 3.4.1. Joint convexity of $r_3(x_3, x_4)(x_4 - x_3)$ can be proven in the say way as we did for $f_2(x_2, x_3)$ (see Appendix 3.6.2). The combination of these results implies that $Cost_1(x_1, x_4)$ is convex in x_4. Define $Cost_1(x_1, x_T)$:

$$Cost_1(x_1, x_T) = \min_{x_{T-1}} Cost_1(x_1, x_{T-1}) + \delta^{T-2} r_{T-1}(x_{T-1}, x_T)(x_T - x_{T-1})$$

By induction we can easily show that $Cost_1(x_1, x_T)$ is convex in x_T.

3.6.4 Proof of Proposition 3.4.1

From Corollary 3.4.2, we have that $\frac{dx_2^*}{dx_3}(x_3) = -\frac{d^2 J}{dx_2 dx_3}(x_2^*(x_3), x_3)\left(\frac{d^2 J}{dx_2^2}(x_2^*(x_3), x_3)\right)^{-1}$. In particular:

$$\frac{d^2 J}{dx_2 dx_3}(x_2^*(x_3), x_3) = \delta\left(k'(x_2) - \frac{b_D}{a_D x_2} - \frac{1}{a_D(M_2 - x_3)} - \frac{1}{a_D(x_3 - x_2)}\right) < 0$$

Furthermore,

$$\frac{d^2 J}{dx_2^2}(x_2^*(x_3), x_3) =$$

$$\delta\left(k''(x_2) - 2k'(x_2) + \frac{b_D}{a_D x_2^2} + 2\frac{b_D}{a_D x_2} - \frac{1}{a_D(x_3 - x_2)}\right) + \frac{(M_1 - x_1)^2}{a_D(x_2 - x_1)(M_1 - x_2)^2} > 0$$

These equations will directly imply $\frac{dx_2^*}{dx_3}(x_3) > 0$. In order to obtain $\frac{dx_2^*}{dx_3}(x_3) < 1$, we need to show $\frac{d^2 J}{dx_2^2}(x_2^*(x_3), x_3) > -\frac{d^2 J}{dx_2 dx_3}(x_2^*(x_3), x_3)$, which can be seen by analyzing each term separately:

$$k''(x_2) - 2k'(x_2) > -k'(x_2)$$

$$\frac{b_D}{a_D x_2^2} + 2\frac{b_D}{a_D x_2} > \frac{b_D}{a_D x_2}$$

$$\frac{(M_1 - x_1)^2}{a_D(x_2 - x_1)(M_1 - x_2)^2} \geq \frac{\delta}{a_D(M_2 - x_3)}$$

The first two expressions are easy to verify. The third one is generally true from the dimensions of the problem we are dealing with, but we introduce one more assumption to be rigorous: $x_2 - x_1 \leq M_2 - x_3$. We know that $\frac{(M_1 - x_1)^2}{(M_1 - x_2)^2} > 1 > \delta$, therefore with the additional assumption $x_2 - x_1 \leq M_2 - x_3$, we conclude that $\frac{dx_2^*}{dx_3}(x_3) < 1$.

3.6.5 Proof of Theorem 3.4.2

Let x_t^B be the baseline adoption path, where $r_t = 0$ for all $t \geq 1$. Pick some arbitrary time $T \geq 3$ and set the adoption target for the optimization model (3.10) at x_T^B and let x_t^* be the new optimal adoption path, with negative system cost $Cost_1(x_1, x_T^B) < 0$, and the optimal rebate path is non-increasing $r_t(x_t^*, x_{t+1}^*) \geq r_{t+1}(x_{t+1}^*, x_{t+2}^*)$.

If the system cost is negative, there must be a negative rebate along the rebate path. If all rebates are negative, then the new adoption path x_t^* is strictly below the baseline path x_t^B, which doesn't reach the target adoption at x_T^B, therefore is a contradiction. Then there must be a rebate that is positive as well.

From monotonicity of rebates, we can infer that the first rebate must be positive and the last rebate negative and that there is a cross point in time $1 \leq \bar{t} \leq T$ such that all rebates $r_t > 0$ for all $t \leq \bar{t}$ and all rebates are non-positive $r_t \leq 0$ for $t > \bar{t}$. We can further infer that if the adoption level at the cross point is lower than the baseline level, $x_{\bar{t}}^* < x_{\bar{t}}^B$, then from the non-positive rebates after the cross point we cannot reach the target adoption x_T^B, which is a contradiction. For the same reason, we can infer that $x_t^* \geq x_t^B$ for $t \geq \bar{t}$. Since all the rebates are positive before the cross point, we can further infer that $x_t^* \geq x_t^B$ for $t \leq \bar{t}$ as well. This implies that the new the adoption path dominates the baseline path at every time step: $x_t^* \geq x_t^B$.

From the definition of the welfare function, we have that:

$$\textit{Welfare}(x_1, ..., x_T) = \sum_{t=1}^{T-1} \left[\textit{Benefit}_t(x_t) - \delta^{t-1} r_t(x_t, x_{t+1})(x_{t+1} - x_t)\right]$$

We know that:

$$Cost_1(x_1, x_T^B) = \sum_{t=1}^{T-1} \left[\delta^{t-1} r_t(x_t^*, x_{t+1}^*)(x_{t+1}^* - x_t^*)\right] < 0 = \sum_{t=1}^{T-1} \left[\delta^{t-1} r_t(x_t^B, x_{t+1}^B)(x_{t+1}^B - x_t^B)\right]$$

Furthermore, we know that $x_t^* \geq x_t^B$ and the welfare benefit functions are increasing, therefore:

$$\textit{Benefit}_t(x_t^*) \geq \textit{Benefit}_t(x_t^B)$$

By adding all the welfare benefit terms and the cost, we obtain the original expression and conclude the proof:

$$\textit{Welfare}(x_1, x_2^*, ..., x_{T-1}^*, x_T^B) > \textit{Welfare}(x_1, x_2^B ..., x_{T-1}^B, x_T^B)$$

Chapter 4

Designing Subsidies With Industry Response Dynamics: Commitment vs. Flexibility

4.1 Introduction

Subsidy programs for emerging technologies have come under special scrutiny over the last few years. As the amount of money invested in these programs become more substantial in times of financial hardship, researchers are trying to understand what are the key cost drivers for these policies. In this chapter, we argue that uncertainty around future subsidy policy decisions is one of these cost drivers. In fact, policy flexibility may increase this uncertainty and cause unintended adverse effects on the supply-chain of the new technology we want to promote. A few examples of this problem can be observed in the solar subsidy programs implemented in Europe over the past two decades.

In 1991, Germany created a pioneer incentive program that effectively created a market for solar photovoltaic panels when the technology was not yet affordable to household consumers. By 2001, the German government introduced a new feed-in-tariff program that would pay solar panel owners 50.62 Euro cents per kWh of electricity produced, more than 3 times the average retail electricity price ((EEG 2007)). This new feed-in-tariff system kick-started a new solar industry and by the end of 2011 there were more than 24.7 GW of installed photovoltaic capacity in Germany, which represents roughly 37% of the total installed capacity worldwide ((EPIA 2011)). The subsidy programs that supported this technological development have been readjusted every few years in order to adapt to ever changing political interests, technological improvements and market conditions. The effect that these policy adjustments have caused on the solar industry is not yet clear and this research aims to shed a new light into this problem.

The Spanish solar program was perhaps less fortunate in that aspect. With a strong incentive system put in place since the mid-90's, the entire country's industry of solar installations was brought to a near collapse when the program was abruptly reduced amidst the 2008-2009 financial crisis.

The goal of this research is to study the effects of policy changes in the production decisions from the industry. We will not focus on a large scale policy disruption, such as the Spanish case. Instead, we will focus on a more controlled type of policy updating, like the German case.

Consider a government that has an adoption target for the technology to be achieved within a given time horizon. The government introduces a different consumer subsidy level for each year of the horizon. This subsidy is awarded to any consumer who purchases the new good during that given year.

We develop in this chapter a two-period game model, where the government chooses a consumer subsidy level for each period and the industry supplier chooses a production level for each period. There is an exogenous demand for the product in each period that is uncertain, but with a known probability distribution. The government can boost the demand level by increasing the subsidies. The sequence of decisions made by the government and the supplier is different under what we call committed and flexible settings. Under a committed setting, the government makes both subsidy decisions at the beginning of the horizon, followed by the supplier. Under a flexible setting, each period is played sequentially for both government and supplier. The key difference for the flexible case is that the government can adjust his subsidy levels after observing the first period demand.

The objective of the government is to minimize the total expected government spending, while achieving an expected target adoption level at the end of the horizon. The supplier wants to simply maximize profits. We model this supplier as multi-period profit maximizing newsvendor, with inventory carry-over. Note that both players have a somewhat aligned incentive to promote the technology.

4.2 Contributions

Under some conditions on the price and cost evolution for the technology, we are able to obtain the sub-game perfect equilibrium production and subsidy levels under both flexible and committed settings. We observe that the optimal production levels for the first period are smaller under the flexible setting. Furthermore, the subsidy levels are on average smaller for the committed setting in both periods.

In order to analyze the total cost of the subsidy program for the government, we compute the total spending as the combination of both sales and subsidies. The first counter-intuitive result we obtained is that under a flexible setting, the government spending is more exposed to the variance of the demand uncertainty. Instead of a hedging effect, the option of policy adjustment (flexibility) decreases initial production levels, which makes the first period of the game less significant. By putting more weight in the second period, the spending level becomes less diversified across the two periods.

As a consequence, when looking at total average spending, we observed that under a flexible setting the government has to pay a higher cost for achieving the same target level. This difference becomes even larger as demand volatility increases.

The intuition behind this result can be explained as following: If too many customers acquire the product early on, the government may be tempted to reduce the subsidy program in order to avoid future costs. At first glance, it may seem that this policy flexibility is valuable for the government to reduce the expected cost of the subsidy program. On the other hand, by preempting such a policy change in the future, the industry might not want to produce too much in the early periods. With our model, we show that by committing to a future subsidy level, independent of the early sales, the government sends a signal to the industry to increase production levels in the early stages of the game. When aggregating over the time horizon, the expected government spending for the committed policy is smaller than the flexible policy.

The main contributions of this model can be summarized as:

- We expand the traditional multi-period inventory management (newsvendor) model to incorporate government incentives.
- Under policy commitment, the early supply-side service level is higher.
- Average spending levels of a committed subsidy policy are lower than a flexible policy.

- Flexible policies are also more vulnerable (in average cost) to early demand volatility.

4.3 Model

In this section we develop the model described before. The government is choosing a rebate level to offer consumers, followed by the supplier that chooses the production quantities. At the end of each period, the uncertain demand is realized and the remaining inventory (if any) is left over for the next period. The two settings mentioned before, committed and flexible, differ only on the timing of the government's decision which will be explained in detail later. Under a committed setting, the government chooses rebate levels for all periods before the horizon begins. In the flexible setting, the rebate levels are decided at the beginning of each period, possibly changing as a function of previously realized production and demand levels.

For ease of exposition, we will only consider a two period horizon model, $t \in \{1, 2\}$. It is possible that the intuition built for the two period model can be reproduced for longer horizons, as the different periods decouple given the state of the system, namely the left-over inventory and the realized sales level. Because of that, it is not clear whether a longer time horizon would be significantly more complicated to solve, but the comparison of policy commitment versus flexibility should be evident even within a single time-period transition.

At the beginning of the horizon, the government has an average target level q of consumers that should purchase the product by the end of the time horizon. This policy target is public information, known to consumers and the industry. In order to achieve this target, the government will set consumer rebates r_t, for each time period t. Any consumer who purchases the product at that time period, will be awarded that rebate. At each period $t \in \{1, 2\}$, the supplier chooses production quantities u_t as a function of the current level of inventory x_t and the rebate levels r_t announced by the government. The number of available units to be sold at each period is given by: $\mathit{Supply}_t = x_t + u_t$.

Demand for the product at time t is realized as a function of the rebate levels r_t and the nominal uncertain demand ϵ_t. We call the random variable ϵ_t the nominal demand because it is the intrinsic demand for the product if no rebate was offered ($r_t = 0$). Assume that for each additional rebate dollar in r_t, we obtain an additional b_t units of demand. This value b_t is the demand sensitivity to the rebate, which we assume to be constant and known by both government and supplier. This demand can be formally defined as: $\mathit{Demand}_t = b_t r_t + \epsilon_t$.

Assumption 4.3.1 *We assume that the cumulative distribution of the uncertainty ϵ_t is given by the continuous function F_t, which is known by both the government and the supplier. This random variable is bounded, $\epsilon_t \in [A_t, B_t]$, always positive $A_t > 0$, and uncorrelated across time periods. We further assume that the government target is larger than the average nominal demand at the second period $q > E[\epsilon_2]$.*

Note that this distribution does not need to be the same across time periods. Also, the lack of correlation between time periods is not necessary for solving the optimal policies, but only for calculating the expected spending levels in Proposition 4.3.4. The technical assumption that the target is larger than the second period nominal demand, $q > E[\epsilon_2]$, is required for the proof of Theorem 4.3.1 and should be true for any practical application. If the target was too small, $q \leq E[\epsilon_2]$, the rebate program might not be necessary to boost demand levels. Instead, we will focus our model in applications where the rebate program is needed to drive demand levels up to the desired target.

The sales level s_t for a given period t will be determined by the rebate level, the production decisions of the supplier, and the realization of uncertainty. Given a supply level and a demand realization at period t, the number of units sold s_t is the minimum of supply and demand: $s_t = \min(\mathit{Supply}_t, \mathit{Demand}_t) = \min(x_t + u_t, b_t r_t + \epsilon_t)$. The inventory left for the next period will then be $x_{t+1} = x_t + u_t - s_t$.

The objective of the government when implementing this policy is to minimize total expected expenditures while still satisfying the given average sales target, q. More formally, for our two period model, the objective is to minimize $E[r_1 s_1 + r_2 s_2]$ subject to an average sales target constraint: $E[s_1 + s_2] \geq q$.

The objective of the supplier is to maximize total expected profits by choosing production levels u_t. There is a fixed linear production cost per unit c_t for each unit produced, u_t. The unit selling price p_t is assumed to be exogenous and fixed before the beginning of the time horizon. Units not sold by the end of the horizon ($t = 2$) get sold for a salvage value denoted by p_3. More formally, the supplier's objective can be written as: $E[p_1 s_1 - c_1 u_1 + p_2 s_2 - c_2 u_2 + p_3 x_3]$.

In applications of designing subsidies for emerging technologies it is common to observe decreasing costs and prices over time. We formalize this assumption as following.

Assumption 4.3.2 *The prices and costs are decreasing over time. Also, prices are higher than costs* *and the salvage value p_3 is smaller than any production cost.*

$$p_1 > p_2$$
$$c_1 > c_2$$
$$p_1 > c_1$$
$$p_2 > c_2$$
$$c_2 > p_3$$

This is the base of the model definitions. We next formalize the difference between the two proposed settings, committed and flexible, according to the timing of the government's decision.

4.3.1 Sequence of events: Committed vs. Flexible

In the committed setting, the government decides both rebate levels r_1 and r_2 before the beginning of the time horizon. Following the rebate level decision, the manufacturer decides on the production quantity for the first period u_1. After the first production decision is made, the first period demand is realized ϵ_1. The amount of inventory not sold is carried over to the second period. The manufacturer then decides the second period production quantity u_2. Then the second period demand ϵ_2 is realized.

In the flexible setting, the government chooses only the first rebate level r_1. The supplier then follows by choosing a production quantity u_1 and the first period demand is realized ϵ_1. After the first period is played, the government chooses the rebate level for period two, r_2, followed by the supplier's decision u_2 and the demand realization ϵ_2.

The sequence os events for these two settings is displayed in Figure 4-1. For the remainder of the chapter, we use the superscript c and f respectively to represent these two settings. We then formalize the optimization problems faced by each player: government and supplier.

4.3.2 Committed setting

In the committed setting described before, the government moves first by choosing rebates and the industry follows by the choosing production quantities. The optimal decisions by each party can be viewed as a bi-level optimization problem. In the first stage, the government chooses a rebate policy r_1 and r_2 subject to the optimal production decision of the supplier. The optimal supplier policy is

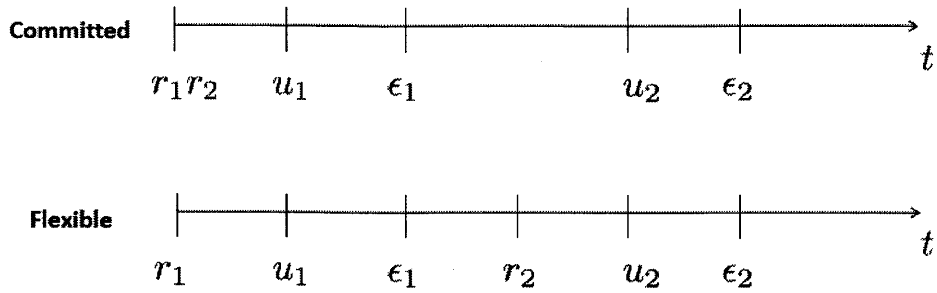

Figure 4-1: Sequence of events under committed and flexible setting

itself the solution to a profit maximization problem.

$$
\begin{aligned}
\min_{r_1,r_2\geq 0} \quad & r_1E[s_1(x_1,u_1^{c*}(x_1,r_1),r_1,\epsilon_1)]+r_2E[s_2(x_2,u_2^{c*}(x_2,r_2),r_2,\epsilon_2)] \\
s.t. \quad & E[s_1(x_1,u_1^{c*}(x_1,r_1),r_1,\epsilon_1)]+E[s_2(x_2,u_2^{c*}(x_2,r_2),r_2,\epsilon_2)]\geq q \\
& u_1^{*}(x_1,r_1)^c, u_2^{c*}(x_2,r_2)=\text{argmax}_{u_1,u_2\geq 0} \;\; p_1E[s_1(x_1,u_1,r_1,\epsilon_t)]-c_1u_1 \\
& \qquad\qquad +p_2E[s_2(x_2,u_2,r_2,\epsilon_2)]-c_2u_2 \qquad (4.1)\\
& \qquad\qquad +p_3E[x_3] \\
\text{Where:} \quad & s_t(x_t,u_t,r_t,\epsilon_t)=\min(x_t+u_t,b_tr_t+\epsilon_t) \\
& x_{t+1}=x_t+u_t-s_t(x_t,u_t,r_t,\epsilon_t)
\end{aligned}
$$

The formulation above misrepresents the fact that the production quantity $u_2^{c*}(x_2,r_2)$ is a function of first period demand uncertainty. To be more formal, the problem above should be formulated as a 3-stage optimization problem, where the first stage is the rebate decisions, the second stage the first period production decision, and the third stage is the second period production decision. The second and third stage decisions are made by the same player, the manufacturer, but the first period demand realization occurs between these two stages. In order to solve this problem, we will start with the last period optimization problem of the manufacturer.

Define $h^c(x_2)$ as the second period profit-to-go of the manufacturer under the committed setting, given the current inventory level x_2. At the beginning of the second period, the manufacturer solves the following problem to maximize the current profit:

$$h_2^c(x_2)=\max_{u_2\geq 0} p_2E[\min(x_2+u_2,b_2r_2+\epsilon_2)]-c_2u_2+p_3E[\max(x_2+u_2-b_2r_2-\epsilon_2,0)] \quad (4.2)$$

At the first period, the manufacturer will solve the following problem in order to maximize the expected first period profit plus the profit-to-go of the second period.

$$\max_{u_1 \geq 0} p_1 E[\min(x_1 + u_1, b_1 r_1 + \epsilon_1)] - c_1 u_1 + E[h_2^c(x_2(u_1, r_1, \epsilon_1)))] \tag{4.3}$$

Define the sales at period t as $s_t(x_t, u_t, r_t, \epsilon_t) = \min(x_t + u_t, b_t r_t + \epsilon_t)$. The left-over inventory for the second period will be $x_1 + u_1 - s_1(x_1, u_1, r_1, \epsilon_1)$

Also define $u_1^{c*}(r_1)$ and $u_2^{c*}(x_2, r_2)$ as the optimal ordering strategies under a committed setting.

Given the optimal ordering strategy $u_1^{c*}(r_1)$, the inventory at the second period should be $x_2(r_1, \epsilon_1) = x_1 + u_1^{c*}(r_1) - s_1(x_1, u_1^{c*}(r_1), r_1, \epsilon_1)$.

The government's optimization problem can be defined as the following (4.4):

$$\begin{aligned} \min_{r_1, r_2 \geq 0} \quad & r_1 E[s_1(x_1, u_1^{c*}(r_1), r_1, \epsilon_1)] + r_2 E[s_2(x_2(r_1, \epsilon_1), u_2^{c*}(x_2(r_1, \epsilon_1), r_2), r_2, \epsilon_2)] \\ s.t. \quad & E[s_1(x_1, u_1^{c*}(r_1), r_1, \epsilon_1) + s_2(x_2(r_1, \epsilon_1), u_2^{c*}(x_2(r_1, \epsilon_1), r_2), r_2, \epsilon_2)] \geq q \end{aligned} \tag{4.4}$$

In order to solve the above optimization problems (4.5.2), (4.19) and (4.4), we define the following notation:

$$\begin{aligned} k_1^c &= F_{\epsilon_1}^{-1}\left(1 - \tfrac{c_1 - c_2}{p_1 - c_2}\right) \\ k_2^c &= F_{\epsilon_2}^{-1}\left(1 - \tfrac{c_2 - p_3}{p_2 - p_3}\right) \\ v_1^c &= E[\min(k_1^c, \epsilon_1)] \\ v_2^c &= E[\min(k_2^c, \epsilon_2)] \end{aligned} \tag{4.5}$$

Note that k_1^c and k_2^c are quantiles of the uncertain demand distribution. They are defined in a similar way to the more common newsvendor quantile, denoting the optimal supply levels desired by the manufacturer. The following lemma formalizes how they are used in the optimal ordering quantities.

We will further assume that the manufacturer does not stay idle for the second time period, $u_2^{c*} > 0$. This would not be the case when demand realized in the first period is too small, leaving inventory too high for the second period. We will not consider these cases and this condition is imposed mainly for technical reasons. If we did consider the case where the left-over inventory x_2 is allowed to be larger than the desired supply level, the optimal first period production level u_1^{c*} would not have a closed-form solution. On the other hand, with a few simple conditions on the demand uncertainty, price and costs, we can ensure that the left-over inventory is not larger then the next period's desired supply level. These conditions are stated in the following Assumption 4.3.3.

Assumption 4.3.3 *In order to assure the manufacturer does not idle, we assume:*

$$k_2^c \geq k_1^c - A_1$$

$$k_1^c \geq x_1$$

This Assumption 4.3.3 is not necessary, but it is sufficient to guarantee our no idling condition. Since we had further assumed that the nominal demand ϵ_1 is always positive, $A_1 > 0$. One can satisfy Assumption 4.3.3 if $k_2^c \geq k_1^c$. That means the target "newsvendor" service level of the second period is larger than in the first period. In other words, in the absence of a rebate policy, the manufacturer would try to serve a larger number of customers in the second period simply from demand, cost and price conditions. This is clearly true in applications of developing technologies, where markets are growing over time.

The second assumption, $k_1^c \geq x_1$, guarantees the first period production level is positive. If there is too much inventory to begin the time horizon, the problem becomes too trivial.

With the assumptions stated above, we can derive the optimal ordering quantities.

Lemma 4.3.1 *The optimal ordering quantities from problems (4.5.2) and (4.19) can be described as follows:*

$$\begin{aligned} u_1^{c*}(x_1, r_1) &= b_1 r_1 + k_1^c - x_1 \\ u_2^{c*}(x_2, r_2) &= b_2 r_2 + k_2^c - x_2 \end{aligned} \tag{4.6}$$

With the definition of the optimal ordering quantity, the sales level can be reduced to: $s_t = \min(u_t + x_t, b_t r_t + \epsilon_t) = b_t r_t + \min(k_t^c, \epsilon_t)$. The last term is a random variable, defined by the minimum of the demand distribution quantile and the actual realization of demand. When taking the expectation of sales, we refer to the notation v_1^c and v_2^c as the expected sales quantile.

We will further assume that the target adoption level q is large enough, otherwise we may have to deal with the less interesting trivial cases where $r_t = 0$. More formally, we want the target adoption level to be at least twice as big as v_t^c for each period.

Assumption 4.3.4 *The target adoption level is sufficiently large:*
$q \geq 2v_1^c$ *and* $q \geq 2v_2^c$

With the above assumption, we can define the optimal rebates:

Lemma 4.3.2 *The optimal committed rebate policy from problem (4.4), can be described as follows:*

$$\begin{aligned} r_1^{c*} &= \frac{q}{b_1+b_2} - \frac{v_1^c(2b_1+b_2)}{2b_1(b_1+b_2)} - \frac{v_2^c}{2(b_1+b_2)} \\ r_2^{c*} &= \frac{q}{b_1+b_2} - \frac{v_2^c(b_1+2b_2)}{2b_2(b_1+b_2)} - \frac{v_1^c}{2(b_1+b_2)} \end{aligned} \tag{4.7}$$

All proofs have been postponed to the appendix for ease of exposition. With the solution to the committed problem, we will now describe the flexible problem, followed by a discussion of the difference between them.

4.3.3 Flexible setting

In the flexible setting, the government moves first by choosing only the first period rebate and the industry follows by the choosing production quantities. The optimal decisions by each party can be viewed as a multi-tiered optimization problem. In the first stage, the government chooses a rebate policy r_1 subject to the optimal production decision of the supplier u_1. That production quantity is subject to the government's policy for the second period rebate r_2. That second rebate is subject to the second period production decision. Because of this sequential dependency, we will define the optimization problems of the flexible setting beginning with the end of the horizon.

Given a certain rebate level r_2 and inventory x_2 at the the second period, define $h(x_2, r_2)$ as the profit-to-go of the manufacturer. This value can be determined by solving the following problem.

$$h^f(x_2, r_2) = \max_{u_2 \geq 0} p_2 E[\min(x_2+u_2, b_2 r_2+\epsilon_2)] - c_2 u_2 + p_3 E[\max(x_2+u_2-b_2 r_2-\epsilon_2, 0)] \tag{4.8}$$

The optimal production level $u_2^*(x_2, r_2)$ can be determined for any set (x_2, r_2). Knowing the strategy of the manufacturer, the government will set the rebate level r_2 that minimizes the cost of subsidizing. Given a level of sales in the first period s_1 and remaining inventory for next period x_2, define $g(s_1, x_2)$ as the the second period cost-to-go of the government. The optimal second period rebate level $r_2^*(s_1, x_2)$ and the cost-to-go $g(s_1, x_2)$ are determined by the following problem.

$$\begin{aligned} g(s_1, x_2) = \quad & \min_{r_2} \quad r_2 E[s_2(x_2, u_2^*(x_2, r_2), r_2, \epsilon_2)] \\ & s.t. \quad\;\; s_1 + E[s_2(x_2, u_2^{c*}(x_2, r_2), r_2, \epsilon_2)] \geq q \end{aligned} \tag{4.9}$$

At the first period, the manufacturer knows only the first period rebate decided by the government r_1. The optimal first period order quantity $u_1^*(x_1, r_1)$ should maximize both the immediate expected profit plus the expected second period profit-to-go.

$$\begin{aligned}\max_{u_1 \geq 0} \quad & p_1 E[s_1(x_1, u_1, r_1, \epsilon_1)] - c_1 u_1 \\ & + E[h^f(x_2(x_1, u_1, r_1, \epsilon_1), r_2^*(s_1(x_1, u_1, r_1, \epsilon_1), x_2(x_1, u_1, r_1, \epsilon_1)))]\end{aligned}$$

Sales in the first period are given by $s_1(x_1, u_1, r_1, \epsilon_1) = \min(x_1 + u_1, b_1 r_1 + \epsilon_1)$ and inventory in the second period is defined as $x_2(x_1, u_1, r_1, \epsilon_1) = x_1 + u_1 - s_1(x_1, u_1, r_1, \epsilon_1)$.

Knowing the contingent order strategy of the manufacturer, $u_1^*(r_1)$, the government must then find the optimal first period rebate r_1^* that minimizes both the immediate cost and the second period cost-to-go.

$$\max_{r_1} r_1 E[s_1(u_1^*(r_1), r_1, \epsilon_1)] + E[g(s_1(u_1^*(r_1), r_1, \epsilon_1), x_2(u_1^*(r_1), r_1, \epsilon_1))] \tag{4.10}$$

In order to solve the above equations, we define the following notation in a similar way as we did for the committed setting. We now use the superscript f to denote the flexible setting. Note that the only difference between the committed and flexible definitions are in the denominator part of k_1^c and k_1^f. We will further explore this difference in the following section.

$$\begin{aligned} k_1^f &= F_{\epsilon_1}^{-1}\left(1 - \frac{c_1 - c_2}{p_1 - p_2}\right) \\ k_2^f &= F_{\epsilon_2}^{-1}\left(1 - \frac{c_2 - p_3}{p_2 - p_3}\right) \\ v_1^f &= E[\min(k_1^f, \epsilon_1)] \\ v_2^f &= E[\min(k_2^f, \epsilon_2)] \end{aligned} \tag{4.11}$$

We will again use a similar assumption to guarantee the manufacturer does not idle production. The reasons are similar as we developed for Assumption 4.3.3.

Assumption 4.3.5 *In order to assure the manufacturer does not idle, we assume:*

$$k_2^f \geq k_1^f - A_1$$

$$k_1^f \geq x_1$$

Using the notation above we obtain the solution for the optimization problems described before.

Lemma 4.3.3 *The optimal ordering quantities and rebate levels from from problems (4.23)-(4.26) can be described as follows:*

$$\begin{aligned}
u_1^{f*}(x_1, r_1) &= \max(b_1 r_1 + k_1^f - x_1, 0) \\
u_2^{f*}(x_2, r_2) &= \max(b_2 r_2 + k_2^f - x_2, 0) \\
r_1^{f*} &= \frac{q}{b_1 + b_2} - \frac{v_1^f(2b_1 + b_2)}{2b_1(b_1 + b_2)} - \frac{v_2^f}{2(b_1 + b_2)} \\
r_2^{f*}(s_1, x_2) &= \frac{q - s_1 - v_2^f}{b_2}
\end{aligned} \tag{4.12}$$

In particular, using the optimal rebate and ordering for the first period defined above, we calculate the expectation of the sales level $E[s_1] = b_1 r_1^{f*} + v_1^f$. Using this, the expected rebate level for the second period will be:

$$E[r_2^{f*}(s_1, x_2)] = \frac{q}{b_1 + b_2} - \frac{v_2^f(b_1 + 2b_2)}{2b_2(b_1 + b_2)} - \frac{v_1^f}{2(b_1 + b_2)} \tag{4.13}$$

4.3.4 Insights

In this section, we develop the key insights extracted from the model. All the proofs have been postponed to the appendix.

When comparing the flexible and committed settings we introduced above, the first thing to notice is the difference in the optimal stocking levels for any fixed rebate level. Since $p_2 > c_2$, we can state the following result:

Proposition 4.3.1 *The relation between the two settings' optimal stocking levels and expected sales will be governed by the following equations:*

$$\begin{aligned}
k_1^c &> k_1^f \\
k_2^c &= k_2^f \\
v_1^c &> v_1^f \\
v_2^c &= v_2^f
\end{aligned} \tag{4.14}$$

Proposition 4.3.1 is the building block of all our insights. This difference in ordering levels will drive differences in sales and rebate level. Note that the key difference between the k_1^c and k_1^f is that the denominator inside the inverse distribution function is given by $p_1 - p_2$ in the flexible case and $p_1 - c_2$ in the committed case. This denominator represents the marginal profit from selling an extra unit at that given period. In the flexible case, this profit is given by the price in the first period p_1, offset by the opportunity cost of selling in the second period, p_2. In the committed case, the first period profit is offset by the second period cost c_2.

With this key difference in the optimal stocking levels, we can compare the expected sales outcomes.

Proposition 4.3.2 *The total expected sales meets the government target:* $E[s_1^c+s_2^c] = E[s_1^f+s_2^f] = q$. *Additionally, the expected difference in sales from flexible to committed setting will be:*

$$\begin{aligned} E[s_1^c - s_1^f] &= \tfrac{b_2}{2(b_1+b_2)}(v_1^c - v_1^f) > 0 \\ E[s_2^c - s_2^f] &= -\tfrac{b_2}{2(b_1+b_2)}(v_1^c - v_1^f) < 0 \end{aligned} \tag{4.15}$$

The fact that average total sales equals the target is not surprising, as the government uses this condition within it's optimization problem. Proposition 4.3.2 above quantifies the average amount of sales that is postponed to the second period when the game dynamics is changed from a committed setting to a flexible setting. In order to understand the effect of this postponement on the total government spending, we need to further analyze the optimal rebate levels.

Proposition 4.3.3 *Under a committed setting, the rebates are on average lower than under a flexible setting.*

$$\begin{aligned} r_1^{c*} &< r_1^{f*} \\ r_2^{c*} &< E[r_2^{f*}(s_1)] \end{aligned} \tag{4.16}$$

Using the results from Proposition 4.3.2 and Proposition 4.3.3, we can compare the expected level of spending from the government. Under a committed setting, the spending will be given by $E[Spending^c] = E[s_1^c]r_1^c + E[s_2^c]r_2^c$, as rebate levels are set deterministically before hand. Under a flexible setting, the government spending is defined as $E[Spending^f] = E[s_1^f]r_1^f + E[s_2^f r_2^f]$. Note that the rebate for the second period under flexible setting is now a random variable and therefore cannot be taken outside the expectation trivially. The following Proposition 4.3.4 derives the average spending levels for the government under the two settings, after a few algebraic manipulations.

Proposition 4.3.4 *The average government spending levels can be defined as follows:*

$$\begin{aligned} E[Spending^c] &= (b_1 r_1^c + v_1^c)r_1^c + (b_2 r_2^c + v_2^c)r_2^c \\ E[Spending^f] &= (b_1 r_1^f + v_1^f)r_1^f + (b_2 E[r_2^f] + v_2^f)E[r_2^f] + \tfrac{VAR(\min\{k_1^f, \epsilon_1\})}{b_2} \end{aligned} \tag{4.17}$$

Note that as demand volatility increases, namely on the variance of the first period demand, the spending level of the government increases only under a flexible setting. This is easy to see when defined the spending levels for the second period in the flexible case: $E[s_2^f r_2^f]$. Since they are both a function of the first period demand ϵ_1, the expectation of the product will depend on the second moment of ϵ_1. This result from Proposition 4.3.4 presents the first key counter intuitive insights. Flexibility is not a good hedging mechanism, as it is making the government more vulnerable to demand volatility. We formalize next the difference in total average spending from the government under a committed setting and a flexible setting.

Theorem 4.3.1 *The governments spends on average more money under a flexible setting than under a committed one.*

$$E[Spending^f] > E[Spending^c]$$

More specifically, the difference in spending levels will be given by:

$$\begin{aligned} E[Spending^f] - \quad E[Spending^c] &= \frac{VAR(\min\{k_1^f, \epsilon_1\})}{b_2} + \\ &+ \frac{1}{4b_1(b_1+b_2)}\left[2b_1(v_1^c - v_1^f)(2q - v_2^c) + b_2(v_1^c)^2 - b_2(v_1^f)^2\right] \end{aligned}$$

The proof is described in the Appendix section. Note that the difference in average spending levels will increase with the variance in the first period demand.

4.4 Conclusions

Flexibility is often seen as an asset in most operations management applications. In the problem we are studying in this chapter, where the government is designing consumer subsidies, policy flexibility is a liability for the government. This counter intuitive result comes from the fact that industry is strategically responding to the policy design and will under invest in the early stages of the horizon if uncertain about the policy level at later stages.

This result carries a very significant qualitative insight for policy makers. The constant readjustment of the subsidy policies used in the past can cause serious adverse effects in the production incentives. Perhaps Germany should not be revising their policies as often as they did.

Among the many possible improvements of the model, it would be good to investigate if these results extend for more than two time periods and with multiple competing suppliers. It would also be important to study the case where prices and costs are endogenous parameters of the system, possibly evolving together with the sales levels. We conjecture that extending the time horizon should not affect our main insight about positive value of policy commitment. It is not clear how competition and endogeneity will affect the outcome of the game.

This model still has a lot of improvement opportunities, but it provides already a very important lesson about the value of commitment. Using the current framework, it would be also useful to understand the effect of the rebate policy on consumers and industry surplus. This chapter focuses only on the government spending level. The combination of these effects will lead to a better understanding of the economic efficiency of the subsidy program.

4.5 Appendix

4.5.1 Notation Summary

The following notation summary is useful for reference throughout the chapter.

$$
\begin{aligned}
x_t &: \text{Initial inventory at the begining of period } t \\
u_t &: \text{Production quantity that supplier chooses for period } t \\
r_t &: \text{Rebate level that governement chooses for period } t \\
b_t &: \text{Demand sensitivity to rebates at period } t \\
\epsilon_t &: \text{Nominal demand for the product in the absense of rebates at period } t \\
F_t &: \text{Demand distribution, } \epsilon_t \sim F_t \\
A_t, B_t &: \text{Lower and upper bound of the demand distribution, } \epsilon_t \in [A_t, B_t] \\
p_t &: \text{Selling price of product at time } t \in \{1,2\} \\
p_3 &: \text{Salvage value for remaining inventory at the end of horizon} \\
c_t &: \text{Production cost per unit at time } t \\
Supply_t &: \text{On-hand supply level at time } t\text{, } Supply_t = x_t + u_t \\
Demand_t &: \text{Demand level at time } t\text{, } Demand_t = b_t r_t + \epsilon_t \\
s_t &: \text{Sales level at time } t\text{, } s_t = \min\{Supply_t, Demand_t\} \\
k_t^j &: \text{Ordering quantile at period } t \text{ under setting } j \in \{c, f\}\text{, committed or flexible} \\
v_t^j &: \text{Expected sales quantile at period } t \text{ under setting } j \in \{c, f\}, v_t^j = E[\min(k_t^j, \epsilon_t)] \\
Spending^j &: \text{Expected governement spending under setting } j \in \{c, f\}
\end{aligned}
$$

4.5.2 Proof of Lemma 4.3.1

Define $h_2^c(x_2)$ as the second period supplier's profit-to-go function and $H_2^c(x_2, u_2)$ as the second period supplier's objective function. Consider the supplier's problem second period problem defined below:

$$
\begin{aligned}
h_2^c(x_2) &= \max_{u_2 \geq 0} \quad H_2^c(x_2, u_2) \\
&= \max_{u_2 \geq 0} \quad p_2 E[\min(x_2 + u_2, b_2 r_2 + \epsilon_2)] - c_2 u_2 \\
&\qquad\qquad +p_3 E[\max(x_2 + u_2 - b_2 r_2 - \epsilon_2, 0)]
\end{aligned} \tag{4.18}
$$

Let $\hat{u}_2$ be the solution of the first order condition of the problem above, which follows:

$$p_2 P(x_2 + \hat{u}_2 \leq b_2 r_2 + \epsilon_2) - c_2 + p_3 P(x_2 + \hat{u}_2 - b_2 r_2 - \epsilon_2 \geq 0) = 0$$

Which is equivalent to:

$$p_2(1 - F_{\epsilon_2}(x_2 + \hat{u}_2 - b_2 r_2)) - c_2 + p_3 F_{\epsilon_2}(x_2 + \hat{u}_2 - b_2 r_2) = 0$$

The unique solution to the first order condition can be isolated from the equation above:

$$\hat{u}_2 = b_2 r_2 - x_2 + F_{\epsilon_2}^{-1}\left(\frac{p_2 - c_2}{p_2 - p_3}\right) = b_2 r_2 - x_2 + k_2^c$$

Additionally, the second derivative of the objective function will be given by:

$$\frac{d^2 H_2^c(x_2, u_2)}{du_2^2} = -p_2 f_{\epsilon_2}(x_2 + u_2 - b_2 r_2) + p_3 f_{\epsilon_2}(x_2 + u_2 - b_2 r_2) \leq 0$$

Note here that f_{ϵ_2} is the pdf of ϵ_2, which is always positive. Since $p_2 > p_3$, the second order condition is satisfied and $\hat{u}_2$ is the maximizer of the unconstrained problem.

From Assumption 4.3.2, we have that $c_2 > p_3$. If this was not the case, the supplier could produce an infinite number of units during period $t = 2$ at a cost below the salvage value, making infinite profit. Since $c_2 > p_3$,n the limit $u_2 \to \infty$, the objective function goes to: $H_2^c(x_2, u_2) \to -\infty$. From continuity of the objective function $H_2^c(x_2, u_2)$, there is a solution of the maximization problem above, which must be either at the boundary $u_2 = 0$ or at the first order condition $u_2 = \hat{u}_2$.

Furthermore, since the objective value is finite at $u_2 = 0$ and $-\infty$ at $u_2 \to \infty$, the objective function $H_2^c(x_2, u_2)$ must be non-increasing in u_2 for $u_2 \geq \hat{u}_2$. Therefore, the optimal second period ordering level in the committed setting is given by:

$$u_2^*(x_2, r_2) = \max(b_2 r_2 - x_2 + k_2^c, 0)$$

At the first period, the manufacturer will solve the following problem in order to maximize the expected first period profit plus the profit-to-go of the second period.

$$\max_{u_1 \geq 0} p_1 E[\min(x_1 + u_1, b_1 r_1 + \epsilon_1)] - c_1 u_1 + E[h_2^c(x_2(x_1, u_1, r_1, \epsilon_1)))] \tag{4.19}$$

Define the following first period production quantity:

$$\hat{u}_1 = b_1 r_1 - x_1 + k_1^c$$

We will next show that this quantity satisfies the first order condition of the problem in (4.19). First, note that under this policy and Assumption 4.3.3, we obtain the no idling condition. If there is any left over inventory, it will be given by:

$$x_2 = x_1 + \hat{u}_1 - b_1 r_1 - \epsilon_1 = k_1^c - \epsilon_1 \le k_2^c \le b_2 r_2 + k_2^c$$

The first inequality above comes from Assumption 4.3.3 and the second from the non-negativity of the rebates. Therefore, the optimal second period ordering policy simplifies to $u_2^*(x_2, r_2) = b_2 r_2 - x_2 + k_2^c$.

Under the optimal ordering policy, the profit-to-go is given by:

$$h_2^c(x_2) = p_2(b_2 r_2 + E[\min(k_2^c, \epsilon_2)] - c_2(b_2 r_2 - x_2 + k_2^c) + p_3 E[\max(k_2^c - \epsilon_2, 0)]$$

Define the derivative of the expected profit-to-go function:

$$\frac{dE[h_2^c]}{du_1} = E\left[\left(\frac{dh_2^c}{dx_2}\right)\left(\frac{dx_2}{du_1}\right)\right] = (c_2)\left(F_{\epsilon_1}(x_1 + u_1 - b_1 r_1)\right)$$

Therefore the unconstrained first order condition of (4.19) can be expressed as:

$$p_1(1 - F_{\epsilon_1}(x_1 + u_1 - b_1 r_1)) - c_1 + c_2 F_{\epsilon_1}(x_1 + u_1 - b_1 r_1) = 0$$

Note that $\hat{u}_1$ is the unique solution to the expression above. Additionally, the second order condition is satisfied as:

$$-p_1 f_{\epsilon_1}(x_1 + u_1 - b_1 r_1)) + c_2 f_{\epsilon_1}(x_1 + u_1 - b_1 r_1) < 0$$

The second order condition above is satisfied since $p_1 > p_2 > c_2$ and the pdf f_{ϵ_1} is always positive. Therefore, the optimal solution is either $u_1 = \hat{u}_1$ or in the boundary $u_1 = 0$. From

Assumption 4.3.3 we obtain that $\hat{u}_1 > 0$, therefore:

$$u_1^*(x_1, r_1) = b_1 r_1 - x_1 + k_1^c$$

4.5.3 Proof of Lemma 4.3.2

Consider the committed government problem:

$$\begin{aligned}
\min_{r_1, r_2 \geq 0} \quad & r_1 E[s_1(x_1, u_1^{c*}(x_1, r_1), r_1, \epsilon_1)] + r_2 E[s_2(x_2, u_2^{c*}(x_2, r_2), r_2, \epsilon_2)] \\
s.t. \quad & E[s_1(x_1, u_1^{c*}(x_1, r_1), r_1, \epsilon_1)] + E[s_2(x_2, u_2^{c*}(x_2, r_2), r_2, \epsilon_2)] \geq q \\
\text{Where:} \quad & s_t(x_t, u_t, r_t, \epsilon_t) = \min(x_t + u_t, b_t r_t + \epsilon_t) \\
& x_{t+1} = x_t + u_t - s_t(x_t, u_t, r_t, \epsilon_t)
\end{aligned} \tag{4.20}$$

Using the optimal production quantities, $u_1^{c*}(x_1, r_1)$ and $u_2^{c*}(x_2, r_2)$, defined in Lemma 4.3.1, we obtain the expected sales level: $E[s_t(x_t, u_t, r_t, \epsilon_t)] = b_t r_t + E[\min(k_t^c, \epsilon_t)] = b_t r_t + v_t^c$. The optimization problem then reduces to:

$$\begin{aligned}
\min_{r_1, r_2 \geq 0} \quad & r_1(b_1 r_1 + v_1^c) + r_2(b_2 r_2 + v_2^c) \\
s.t. \quad & (b_1 r_1 + v_1^c) + (b_2 r_2 + v_2^c) \geq q
\end{aligned} \tag{4.21}$$

The objective function is increasing in both r_1 and r_2, and the expected sales is a continuous function. Therefore, the optimal solution must occur when the constraint is tight. We can solve the problem by first isolating $r_1 = (q - v_1^c - b_2 r_2 - v_2^c)/b_1$. The remaining problem is:

$$\min_{r_2 \geq 0} \quad \tfrac{q - v_1^c - b_2 r_2 - v_2^c}{b_1}(q - b_2 r_2 - v_2^c) + r_2(b_2 r_2 + v_2^c) \tag{4.22}$$

Note that the objective function is clearly convex in r_2. By taking the first order condition, we obtain:

$$r_2^{c*} = \frac{q}{b_1 + b_2} - \frac{v_2^c(b_1 + 2b_2)}{2b_2(b_1 + b_2)} - \frac{v_1^c}{2(b_1 + b_2)}$$

The first period rebate value follows from the target constraint:

$$r_1^{c*} = \frac{q}{b_1 + b_2} - \frac{v_1^c(2b_1 + b_2)}{2b_1(b_1 + b_2)} - \frac{v_2^c}{2(b_1 + b_2)}$$

Note that these values will be positive from Assumption 4.3.4. Therefore, r_1^{c*} and r_2^{c*} are the optimal committed rebate policies.

4.5.4 Proof of Lemma 4.3.3

Similar to the optimal solution of the committed case, we will solve the flexible problem by iterating backwards. The derivation of the last stage, namely the second period production level u_2, will be the same as in the proof of Lemma 4.3.1. In particular, the problem is described by:

$$h_2^f(x_2, r_2) = \max_{u_2 \geq 0} p_2 E[\min(x_2+u_2, b_2r_2+\epsilon_2)] - c_2u_2 + p_3E[\max(x_2+u_2-b_2r_2-\epsilon_2, 0)] \quad (4.23)$$

The optimal ordering quantity can be expressed:

$$u_2^*(x_2, r_2) = \max(b_2r_2 - x_2 + k_2^f, 0)$$

The government second period decision will be given by:

$$\begin{aligned} g(s_1, x_2) = \quad & \min_{r_2} \quad r_2E[s_2(x_2, u_2^*(x_2, r_2), r_2, \epsilon_2)] \\ & s.t. \quad\quad s_1 + E[s_2(x_2, u_2^{c*}(x_2, r_2), r_2, \epsilon_2)] \geq q \end{aligned} \quad (4.24)$$

Note that using the optimal ordering quantity, we obtain:

$$E[s_2(x_2, u_2, r_2, \epsilon_2)] = b_2r_2 + E[\min(k_2^f, \epsilon_2)] = b_2r_2 + v_2^f$$

It is easy to see that the objective function is increasing in r_2 and also the left hand side of the constraint. Therefore the optimal solution will occur when the constraint is tight:

$$r_2^{f*}(s_1, x_2) = \frac{q - s_1 - v_2^f}{b_2}$$

The manufacturer will solve the first period production quantity with knowledge of the government rebate strategy for the second period. The problem can be described as:

$$\begin{aligned} \max_{u_1 \geq 0} \quad & p_1E[s_1(x_1, u_1, r_1, \epsilon_1)] - c_1u_1 \\ & + E[h_2^f(x_2(x_1, u_1, r_1, \epsilon_1), r_2^*(s_1(x_1, u_1, r_1, \epsilon_1), x_2(x_1, u_1, r_1, \epsilon_1)))] \end{aligned} \quad (4.25)$$

As in the committed setting, we will also assume that the manufacturer does not idle in the second period and later verify that this is indeed the case. Note under the optimal second period ordering policy $h_2^f = p_2(b_2r_2+v_2^f) - c_2(b_2r_2 - x_2 + k_2^f) + p_3E[\max(k_2^f - \epsilon_2, 0)]$. When combined with the second period rebate policy, we obtain: $h_2^f = p_2(q-s_1) - c_2(q-s_1-x_2) + p_3E[\max(k_2^f - \epsilon_2, 0)]$.

Note that $s_1 = \min(x_1+u_1, b_1r_1+\epsilon_1)$. Then $dE[h_2^f]/du_1 = -p_2(1-F_{\epsilon_1}(x_1+u_1-b_1r_1))+c_2$. By taking the first derivative of the objective function in (4.5.4), we obtain:

$$p_1(1 - F_{\epsilon_1}(x_1 + u_1 - b_1r_1)) - c_1 - p_2(1 - F_{\epsilon_1}(x_1 + u_1 - b_1r_1)) + c_2 = 0$$

Therefore, the first order condition reduces to: $(p_1-p_2)F_{\epsilon_1}(x_1+u_1-b_1r_1)) = p_1-c_1-p_2+c_2$. Note also that the second derivative is negative, since $p_1 > p2$, which implies that the objective function is concave.

It is easy to see that the following first period production quantity uniquely satisfies the first order condition propose above:

$$\hat{u}_1^f = b_1r_1 - x_1 + k_1^f$$

Since from Assumption 4.3.5 we have that $x_1 \leq k_1^f$, we know the optimal solution is positive: $u_1^{f*} = \hat{u}_1^f$.

Under this policy, $s_1(x_1, \hat{u}_1, r_1, \epsilon_1) = b_1r_1 + \min(k_1^f, \epsilon_1)$. Additionally, $x_2 = x_1 + u_1 - s_1 = k_1^f - \min(k_1^f, \epsilon_1)$. The second period rebate policy will be determined as: $r_2^{f*}(s_1, x_2) = \dfrac{q - b_1r_1 + \min(k_1^f, \epsilon_1) - v_2^f}{b_2}$

The first period government problem is defined as:

$$\max_{r_1} r_1E[s_1(x_1, u_1^*(r_1), r_1, \epsilon_1)] + E[g(s_1(x_1, u_1^*(r_1), r_1, \epsilon_1))] \tag{4.26}$$

Note that from the solution of the second period rebate problem, we have:

$$g(s_1) = \frac{q - s_1 - v_2^f}{b_2}(q - s_1) = \frac{q - (b_1r_1 + \min(k_1^f, \epsilon_1)) - v_2^f}{b_2}(q - (b_1r_1 + \min(k_1^f, \epsilon_1)))$$

The optimal rebate quantity for the first period can be obtained by solving the first order condition of the problem (4.26). We further note that the second derivative is always positive, indicating the function is convex. The solution to the first order condition is stated below:

$$r_1^{f*} = \frac{q}{b_1 + b_2} - \frac{v_1^f(2b_1 + b_2)}{2b_1(b_1 + b_2)} - \frac{v_2^f}{2(b_1 + b_2)}$$

4.5.5 Proof of Proposition 4.3.1

From the definitions of the ordering quantiles in (4.5) and (4.11) we have: $k_1^c = F_{\epsilon_1}^{-1}\left(1 - \frac{c_1 - c_2}{p_1 - c_2}\right)$ and $k_1^f = F_{\epsilon_1}^{-1}\left(1 - \frac{c_1 - c_2}{p_1 - p_2}\right)$.

We assumed the distribution of ϵ to be a continuous distribution with full support on $[A_t, B_t]$. Since the function $F_{\epsilon_1}^{-1}$ is increasing, we need only to show. $1 - \frac{c_1 - c_2}{p_1 - c_2} > 1 - \frac{c_1 - c_2}{p_1 - p_2}$. This can be implied from $p_1 - c_2 > p_1 - p_2$, which is true from Assumption 4.3.2: $p_2 > c_2$. Therefore, $k_1^c > k_1^f$. For the second time period, the relationship is trivially true from the definition $k_2^c = k_2^f$.

Additionally, from the definition of the expected sales quantiles in (4.5) and (4.11), we have: $v_1^c = E[\min(k_1^c, \epsilon_1)]$ and $v_1^f = E[\min(k_1^f, \epsilon_1)]$. Note that $\min(k_1^c, \epsilon_1) \geq \min(k_1^f, \epsilon_1)$ for any value of ϵ_1. Since the distribution is fully supported in $[A_1, B_1]$ and we know that $A_1 < k_1^f < k_1^c < B_1$, so there will be some measurable part of the distribution where the inequality is strict: $\min(k_1^c, \epsilon_1) > \min(k_1^f, \epsilon_1)$. Therefore, we obtain $v_1^c > v_1^f$. The relationship $v_2^c = v_2^f$ is trivially true.

4.5.6 Proof of Proposition 4.3.2

Using the optimal ordering policies and rebate levels from Lemmas 4.3.1, 4.3.2 and 4.3.3, we obtain the expected sales levels as;

$$E[s_t] = E[\min\{x_t + u_t, b_t r_t + \epsilon_t\}] = b_t E[r_t^{j*}] + E[\min\{k_t^j, \epsilon_t\}] = b_t E[r_t^{j*}] + v_t^j$$

Where j can be either c or f for committed or flexible setting respectively. The first and second period expected sales level will be given by:

$$E[s_1^j] = \frac{b_1 q}{b_1 + b_2} + \frac{b_2 v_1^j}{2(b_1 + b_2)} - \frac{b_1 v_2^j}{2(b_1 + b_2)}$$

$$E[s_2^j] = \frac{b_2 q}{b_1 + b_2} + \frac{b_1 v_2^j}{2(b_1 + b_2)} - \frac{b_2 v_1^j}{2(b_1 + b_2)}$$

Note that the average sales maintain the same structure between the two different settings $j \in \{c, f\}$. The only difference is on the first expected sales quantile v_1^c and v_1^f. We can now calculate the difference of expected sales, where the terms without v_1^c and v_1^f will cancel each other therefore obtaining the desired result:

$$\begin{aligned} E[s_1^c - s_1^f] &= \tfrac{b_2}{2(b_1+b_2)}(v_1^c - v_1^f) > 0 \\ E[s_2^c - s_2^f] &= -\tfrac{b_2}{2(b_1+b_2)}(v_1^c - v_1^f) < 0 \end{aligned} \tag{4.27}$$

It is also easy to see that $E[s_1^j] + E[s_2^j] = q$, for both $j \in \{c, f\}$, which is a good sanity check for the solution. This is to be expected since the government uses this condition as a tight constraint in the optimization problem.

4.5.7 Proof of Proposition 4.3.3

From the definitions of the optimal rebate levels in Lemmas 4.3.2 and 4.3.3, we obtain the difference between the first period rebate in the flexible and committed settings is given by:

$$r_1^{c*} - r_1^{f*} = -\frac{2b_1 + b_2}{2b_1(b_1 + b_2)}[v_1^c - v_1^f] < 0$$

From Proposition 4.3.1, we know the second term is positive, $v_1^c > v_1^f$. Therefore, the rebate level in the committed setting is smaller: $r_1^{c*} - r_1^{f*} < 0$. For the second period rebates, we calculate the expected difference in rebates:

$$r_2^{c*} - E[r_2^{f*}(s_1)] = -\frac{v_1^c}{2(b_1 + b_2)} + \frac{v_1^f}{2(b_1 + b_2)} < 0$$

The difference is driven by the $v_1^c > v_1^f$ from Proposition 4.3.1.

4.5.8 Proof of Proposition 4.3.4

Under a committed setting, the expected spending levels for each period is easily obtained since the rebate levels are deterministic. $Spending^c = E[s_1]r_1^{c*} + E[s_2]r_2^{c*}$. The expected sales will be given by $E[s_t] = \min\{x_t + u_t, b_t r_t + \epsilon_t\}$. Under the optimal ordering policy of Lemma 4.3.1, we obtain $E[s_t] = b_t r_t^{c*} + E[\min\{k_t^c, \epsilon_t\}] = b_t r_t^{c*} + v_t^c$. The first relationship is proven. Under a flexible setting, we obtain a $E[s_1]r_1^{f*} = b_1 r_1^{1*} + v_1^f$ in a similar way for the first time period. For the second period, note that both rebate and sales are random variables. Therefore, using $s_2 =$

$b_f r_2^{f^*} + \min\{k_2^f, \epsilon_2\}$, we obtain the expectation of the product will be given by:

$$\begin{aligned}
E[s_2 r_2^{f^*}(s_1)] &= E\left[\left(b_2(q - s_1 - v_2^f)/b_2 + \min\{k_2^f, \epsilon_2\}\right)(q - s_1 - v_2^f)/b_2\right] \\
&= E\left[(q - s_1 - v_2^f)^2\right]/b_2 + E\left[\min\{k_2^f, \epsilon_2\}(q - s_1 - v_2^f)\right]/b_2 \\
&= E\left[q - s_1 - v_2^f\right]^2/b_2 + VAR[q - s_1 - v_2^f]/b_2 \\
&\quad + E\left[\min\{k_2^f, \epsilon_2\}(q - v_2^f)\right]/b_2 - E\left[\min\{k_2^f, \epsilon_2\}(s_1)\right]/b_2
\end{aligned}$$

Using $s_1 = b_1 r_1^{f^*} + \min\{k_1^f, \epsilon_1\}$, we obtain:

$$\begin{aligned}
E[s_2 r_2^{f^*}(s_1)] &= \left(q - E[s_1] - v_2^f\right)^2/b_2 + VAR[s_1]/b_2 + v_2^f(q - v_2^f)/b_2 \\
&\quad - E\left[\min\{k_2^f, \epsilon_2\}(b_1 r_1^{f^*} + \min\{k_1^f, \epsilon_1\})\right]/b_2 \\
&= \left(q - E[s_1] - v_2^f\right)^2/b_2 + VAR[s_1]/b_2 + v_2^f(q - v_2^f)/b_2 - v_2^f(b_1 r_1^{f^*})/b_2 \\
&\quad - E\left[\min\{k_2^f, \epsilon_2\}\min\{k_1^f, \epsilon_1\}\right]/b_2
\end{aligned}$$

From independence of ϵ_1 and ϵ_2, we get:

$$E[s_2 r_2^{f^*}(s_1)] = \left(q - E[s_1] - v_2^f\right)^2/b_2 + VAR[s_1]/b_2 + v_2^f(q - v_2^f)/b_2 - v_2^f(b_1 r_1^{f^*})/b_2 - v_2^f v_1^f/b_2$$

With some algebraic manipulation, we obtain:

$$E[s_2 r_2^{f^*}(s_1)] = \left(q - E[s_1] - v_2^f\right)(q - E[s_1])/b_2 + VAR[\min\{k_1^f, \epsilon_1\}]/b_2$$

By using the definition of $r_2^{f^*} = \left(q - s_1 - v_2^f\right)/b_2$ we get:

$$E[s_2 r_2^{f^*}(s_1)] = (b_2 E[r_2^f(s_1)] + v_2^f)E[r_2^f(s_1)] + VAR[\min\{k_1^f, \epsilon_1\}]/b_2$$

. This concludes the proof of the proposition.

4.5.9 Proof of Theorem 4.3.1

Since $v_2^f = v_2^c$, we will replace both of them by simply v_2. With some algebraic manipulations, we obtain:

$$E[Spending^f] = \frac{VAR(\min\{k_1^f, \epsilon_1\})}{b_2} + \frac{-v_2^2 b_1^2 - 4qb_2v_2b_1 - 4qv_1^f b_2b_1 + 4b_1b_2q^2 + 2v_1^f b_2v_2b_1 - (v_1^f)^2 b_2^2}{4b_1b_2(b_1+b_2)}$$

Similarly, for the committed setting we obtain:

$$E[Spending^c] = \frac{-v_2^2 b_1^2 - 4qb_2v_2b_1 - 4qv_1^c b_2b_1 + 4b_1b_2q^2 + 2v_1^c b_2v_2b_1 - (v_1^c)^2 b_2^2}{4b_1b_2(b_1+b_2)}$$

When calculating the difference in spending, many terms cancel each other, leaving the following relation:

$$E[Spending^f] - E[Spending^c] = \frac{VAR(\min\{k_1^f, \epsilon_1\})}{b_2} + \frac{1}{4b_1(b_1+b_2)}\left[2b_1(v_1^c - v_1^f)(2q - v_2^c) + b_2(v_1^c)^2 - b_2(v_1^f)^2\right]$$

Note that $VAR(\min\{k_1^f, \epsilon_1\}) > 0$ is trivially true. We also know that from Proposition 4.3.1: $v_1^c > v_1^f > 0$. Therefore, we get both $2b_1(v_1^c - v_1^f) > 0$ and $b_2(v_1^c)^2 - b_2(v_1^f)^2 > 0$. The remaining middle term $(2q - v_2^c)$ is also positive from the assumption that the target is large enough that the rebate solution is non-trivial: $q > E[\epsilon_2]$, which is itself larger than the expected sales quantile: $q > E[\epsilon_2] > E[\min(k_2^c, \epsilon_2)] = v_2^c$. Therefore, we obtain the concluding result.

$$E[Spending^f] > E[Spending^c]$$

Chapter 5

Conclusions

The initial question that was posed in the introduction of this thesis is one that will always intrigue researchers, as it carries very significant implications for the real world. *"How much should a company charge for its product?"* That same question can be asked by a variety of decision makers, with different objectives in mind. In this thesis, we first we examined at how firms adjust their prices to maximize profits. When promoting new technologies, it is often the government that has a vested interest in boosting sales by offering price incentives. In all, this thesis tries to shed some light on new modeling frameworks that can potentially benefit different types of decision makers.

The literature on dynamic pricing and revenue management is quite vast. The research exposed in Chapter 2 builds on top of two different optimization methodologies that are often used separately for this pricing problem: robust and data-driven optimization. By building a bridge between these two fields, we construct a tractable and practical sampling-based solution framework that can be used effectively to price complex multi-resource products within a fixed time horizon. The contributions are both theoretical and practical. We showed how robust the sampling-based solution is depending on the number of samples used. When samples are not available, we proposed a new scenario generating technique that enables the manager to get provably robust solutions using the same data-driven optimization model. From the practical side, we showed how the robust solution performs well for risk averse managers, especially when using a small sample size. This can be of great value to managers, who can now confidently use this optimization tool to better protect themselves from the downside risk of their pricing policies.

On Chapter 3, we explored how policy-makers should design price incentives for solar panels. We used a discrete choice model to forecast adoption levels while considering the important network effects like learning-by-doing and information spread. Stimulating these network effects is one of the main reasons why governments invest in these subsidy programs. We proposed a new framework for policy optimization that minimizes the cost for the government while still achieving desired adoption levels. This framework was validated through an empirical study of the German solar market. Using historical data to estimate the demand model parameters, we analyzed the forecasted adoption levels and re-optimized the future subsidy levels. The main insight of the empirical study was that the current policy is not efficient. We further proved that there is no possible social welfare measure that could potentially justify the current program. Because of the network effects, our model suggests that Germany could save money by increasing near-term subsidy levels and phasing out the program faster. The adoption levels would still be higher than current forecasts and the overall program will be a cheaper for the government. This empirical result should be taken as a warning flag rather than a policy recommendation, as it obviously depends on many modeling assumptions and the limited amount of data we obtained. Nevertheless, it demonstrates the practicality of our policy-making tool and how policy efficiency can be evaluated and improved.

As a drawback to our modeling framework in Chapter 3, we only consider the learning-by-doing effect as a supply-side implications of the policy decisions. When firms are strategically responding to policy-maker's decision, the policy optimization problem can be much more complicated. In Chapter 4, we address this issue from a theoretical perspective. We propose a two-period game setting, where the government designs price incentives while a supplier (e.g. solar panel manufacturer) reacts to the subsidy policy by adjusting the production levels to maximize profits. In particular, we focus on the trade-offs between commitment and flexibility for the policy-maker. When faced with uncertain demand, flexibility to postpone decisions is usually seen as an operational advantage. We show this is not the case for the government in this setting. By committing to a long-term policy before hand, the government encourages the supplier to produce more in the early stages. The total average cost of the subsidy program will be lower under a committed setting than a flexible setting while achieving the same adoption targets. This research has only scratched the surface on the way policy-makers and firms interact, but it already provides some potentially useful policy recommendation. This counter-intuitive message that flexibility is a liability is a powerful insight that policy-makers should be aware of when designing future policy programs.

Bibliography

Adida, E. and G. Perakis (2005). A robust optimization approach to dynamic pricing and inventory control with no backorders. *Mathematical Programming Special Issue on Robust Optimization.*

Akan, M. and B. Ata (2009). Bid-price controls for network revenue management: Martingale characterization of optimal bid prices. *Mathematics of Operations Research 34*(4), 912–936.

Aral, S., L. Muchnik, and A. Sundararajan (2009). Distinguishing influence-based contagion from homophily-driven diffusion in dynamic networks. *Proceedings of the National Academy of Sciences 106*(51), 21544–21549.

Araman, V. and R. Caldentey (2009). Dynamic pricing for non-perishable products with demand learning. *Operations Research.* (to appear).

Arrow, K. J. (1962). The economic implications of learning by doing. *The Review of Economic Studies 29*(3), pp. 155–173.

Bagnoli, M. and T. Bergstrom (1989). Log-concave probability and its applications.

Ball, M. O. and M. Queyranne (2006). Toward robust revenue management: Competitive analysis of online booking. Robert H. Smith School Research Paper No. RHS 06-021. Available at SSRN: http://ssrn.com/abstract=896547.

Bass, F. M. (1969). A new product growth for model consumer durables. *Management Science 15*(5), 215–227.

Bayraksan, G. and D. P. Morton (2006, September). Assessing solution quality in stochastic programs. *Mathematical Programming 108*(2-3), 495–514.

Belobaba, P. P. (1987). *Air travel demand and airline seat inventory management.* Ph. D. thesis, MIT.

Belobaba, P. P. (1992). Optimal versus heuristic methods for nested seat allocation. *AGIFORS Reservations and Yield Management Study Group Meeting.*.

Ben-Akiva, M. and S. Lerman (1993). *Discrete choice analysis: theory and application to travel demand.* MIT, Cambridge, MA.

Ben-Tal, A., B. Golany, A. Nemirovski, and J. Vial (2005). Retailer-supplier flexible commitments contracts: A robust optimization approach. *Manufacturing and Service Operations Management 7*(3), 248–271.

Ben-Tal, A., A. Goryashko, E. Guslitzer, and A. Nemirovski (2004). Adjustable robust solutions of uncertain linear programs. *Mathematical Programming 99*(2), 351–376.

Ben-Tal, A. and A. Nemirovski (1998). Robust convex optimization. *Mathematics of Operations Research 23*, 769–805.

Ben-Tal, A. and A. Nemirovski (1999). Robust solutions to uncertain linear programs. *Operations Research Letters 25*, 1–13.

Ben-Tal, A. and A. Nemirovski (2000). Robust solutions of linear programming problems contaminated with uncertain data. *Mathematical Programming 88*, 411–424.

Benthem, A., K. Gillingham, and J. Sweeney (2008). Learning-by-doing and the optimal solar policy in california. *The Energy Journal 29*(3), 131–151.

Berndt, E. R., R. S. Pindyck, and P. Azoulay (2003). Consumption externalities and diffusion in pharmaceutical markets: Antiulcer drugs. *The Journal of Industrial Economics 51*(2), 243–270.

Berry, S., J. Levinsohn, and A. Pakes (1995). Automobile prices in market equilibrium. *Econometrica 63*(4), pp. 841–890.

Berry, S. T. (1994). Estimating discrete-choice models of product differentiation. *The RAND Journal of Economics 25*(2), 242–262.

Bertsekas, D. (1995). *Nonlinear Optimization.* Athena Scientific.

Bertsimas, D. and S. de Boer (2005). Simulation-based booking limits for airline revenue management. *Operations Research 53*(1), 90–106.

Bertsimas, D., D. Iancu, and P. Parrilo (2009). Optimality of affine policies in multi-stage robust optimization. working paper, MIT.

Bertsimas, D. and I. Popescu (2003, August). Revenue management in a dynamic network environment. *Transportation Science 37*(3).

Bertsimas, D. and M. Sim (2004). The price of robustness. *Operations Research 52*(1), 35–53.

Bertsimas, D. and A. Thiele (2006). A robust optimization approach to supply chain management. *Operations Research 54*(1).

Besbes, O. and A. Zeevi (2008). Dynamic pricing without knowing the demand function: Risk bounds and near-optimal algorithms. *Operations Research.* (to appear).

Bhandari, R. and I. Stadler (2009). Grid parity analysis of solar photovoltaic systems in germany using experience curves. *Solar Energy 83*(9), 1634 – 1644.

Birbil, S. I., J. B. G. Frenk, J. A. S. Gromicho, and S. Zhang (2006). An integrated approach to single-leg airline revenue management: The role of robust optimization. *Econometric Institute Report.* E.I. 2006-20.

Bitran, G. and R. Caldentey (2003). An overview of pricing models and revenue management. *Manufacturing and Service Operations Management 5*(3), 203–229.

BMU (2008). Lead study 2008 - further development of the strategy to increase the use of renewable energies within the context of the current climate protection goals of germany and europe. Study commissioned by the German Federal Ministry for the Environment, Nature Conservation and Nuclear Safety (BMU).

Bobkov, S. G. (1999). Isoperimetric and analytic inequalities for log-concave probability measures. *The Annals of Probability* (4).

Boyd, S. and L. Vandenbergh (2004). *Convex Optimization.* Cambridge University Press.

Calafiore, G. (2009, November). On the expected probability of constraint violation in sampled convex programs. *Journal Journal of Optimization Theory and Applications 143*(2), 405–412.

Calafiore, G. and M. Campi (2005). Uncertain convex programs: randomized solutions and confidence levels. *Mathematical Programming 102*, 25–46.

Calafiore, G. and M. Campi (2006, May). The scenario approach to robust control design. *IEEE Transactions On Automatic Control 51*(5), 742– 753.

Campi, M. and S. Garatti (2008). The exact feasibility of randomized solutions of uncertain convex programs. *SIAM Journal on Optimization 19*(3), 1211–1230.

Caramanis, C. (2006, May). *Adaptable Optimization: Theory and Algorithms.* PhD dissertation, Department of Electrical Engineering and Computer Science, MIT.

Chiralaksanakul, A. and D. P. Morton (2004). Assessing policy quality in multi-stage stochastic programming. *Stochastic Programming E-Print Series (www.speps.org).*

Choi, J. P. and M. Thum (1998). Market structure and the timing of technology adoption with network externalities. *European Economic Review 42*(2), 225 – 244.

Chou, C. and O. Shy (1990). Network effects without network externalities. *International Journal of Industrial Organization 8*(2), 259 – 270.

Cox, D. R. (1972). Regression models and life-tables. *Journal of the Royal Statistical Society. Series B (Methodological) 34*(2), 187–220.

Doganoglu, T. and L. Grzybowski (2007). Estimating network effects in mobile telephony in germany. *Information Economics and Policy 19*(1), 65 – 79.

EEG (2007). Renewable energy sources act (EEG) - progress report 2007.

Ellison, G. and D. Fudenberg (1995). Word-of-mouth communication and social learning. *The Quarterly Journal of Economics 110*(1), 93–125.

Elmaghraby, W. and P. Keskinocak (2003). Dynamic pricing: Research overview, current practices and future directions. *Management Science 49*, 1287–1309.

EPIA (2009). Global market outlook for photovoltaics until 2013.

EPIA (2011). European photovoltaic industry association - market report.

Eren, S. S. and C. Maglaras (2009). Monopoly pricing with limited demand information. Working paper, Columbia University.

EUC (1997). Energy for the future: renewable sources of energy white paper for a community strategy and action plan.

Farias, V. and B. Van Roy (2009). Dynamic pricing with a prior on market response. *Operations Research.* (to appear).

Farrell, J. and G. Saloner (1985). Standardization, compatibility, and innovation. *The RAND Journal of Economics 16*(1), 70–83.

Farrell, J. and G. Saloner (1986). Installed base and compatibility: Innovation, product preannouncements, and predation. *The American Economic Review 76*(5), 940–955.

Frondel, M., N. Ritter, and C. Schmidt (2008). Germany's solar cell promotion: Dark clouds on the horizon. *Ruhr Economic Papers.*

Gallego, G. and I. Moon (1993). The distribution free newsboy problem: Review and extensions. *The Journal of the Operational Research Society 44*(8).

Geroski, P. A. (2000). Models of technology diffusion. *Research Policy 29*(4-5), 603 – 625.

Goldenberg, J., B. Libai, and E. Muller (2010). The chilling effects of network externalities. *International Journal of Research in Marketing 27*(1), 4 – 15.

Goldfarb, D. and G. Iyengar (2003). Robust convex quadratically constrained programs. *Mathematical Programming 97*(3), 495–515.

Goolsbee, A. and P. J. Klenow (2002). Evidence on learning and network externalities in the diffusion of home computers. *Journal of Law and Economics 45*(2), 317–343.

Harmon, C. (2000). Experience curves of photovoltaic technology.

Hauser, J., G. J. Tellis, and A. Griffin (2006). Research on innovation: A review and agenda for marketing science. *Marketing Science 25*(6), 687 – 717.

Helsen, K. and D. C. Schmittlein (1993). Analyzing duration times in marketing: Evidence for the effectiveness of hazard rate models. *Marketing Science 12*(4), 395–414.

IEA (2000). Experience curves for energy technology policy.

IEA (2004). Renewable energy - market & policy trends in IEA countries.

Ireland, N. and P. Stoneman (1986). Technological diffusion, expectations and welfare. *Oxford Economic Papers 38*(2), pp. 283–304.

Jager, W. (2006). Stimulating the diffusion of photovoltaic systems: A behavioural perspective. *Energy Policy 34*(14), 1935 – 1943.

Jager-Waldau, A. (2007). Photovoltaics and renewable energies in europe. *Renewable and Sustainable Energy Reviews 11*(7), 1414 – 1437.

Joskow, P. L. and N. L. Rose (1989). The effects of economic regulation. *Handbook of Industrial Organization 2*, 1449 – 1506.

Kapur, S. (1995). Technological diffusion with social learning. *The Journal of Industrial Economics 43*(2), 173–195.

Katz, M. L. and C. Shapiro (1986). Technology adoption in the presence of network externalities. *The Journal of Political Economy 94*(4), 822–841.

Katz, M. L. and C. Shapiro (1992). Product introduction with network externalities. *The Journal of Industrial Economics 40*(1), 55–83.

Kleywegt, A., A. Shapiro, and T. Homem-de Mello (2001). The sample average approximation method for stochastic discrete optimization. *SIAM Journal on Optimization, 12*(2), 479–502.

Lan, Y., H. Gao, M. Ball, and I. Karaesmen (2006). Revenue management with limited demand information. *Management Science.* (to appear).

Levi, R., G. Perakis, and J. Uichanco (2010). New bounds and insights on the data-driven newsvendor model. working paper, MIT.

Levi, R., R. O. Roundy, and D. B. Shmoys (2007). Provably near-optimal sampling-based policies for stochastic inventory control models. *Mathematics of Opreations Research 32*(4), 821–839.

Lim, A. and J. Shanthikumar (2007). Relative entropy, exponential utility, and robust dynamic pricing. *Operations Research 55*(2), 198–214.

Lim, A., J. Shanthikumar, and T. Watewai (2008). Robust multi-product pricing. Working paper, Department of Industrial Engineering and Operations Research, University of California, Berkeley, CA.

Littlewood, K. (1972). Forecasting and control of passenger bookings. *Proceedings of the Twelfth Annual AGIFORS Symposium.*

Mahajan, V., E. Muller, and F. M. Bass (1990). New product diffusion models in marketing: A review and directions for research. *The Journal of Marketing 54*(1), 1–26.

McDonald, A. and L. Schrattenholze (2001). Learning rates for energy technologies. *Energy Policy.*

Nemet, G. F. (2006). Beyond the learning curve: factors influencing cost reductions in photovoltaics. *Energy Policy 34*(17), 3218 – 3232.

Nguyen, D. T. and G. Perakis (2005). Robust competitive pricing for multiple perishable products. Working paper, MIT.

Pagnoncelli, B., S. Ahmed, and A. Shapiro (2009). Computational study of a chance constrained portfolio selection problem. *Journal of Optimization Theory and Applications.* (to appear).

Perakis, G. (2009). Alleviating travel delay uncertainties in traffic assignment and traffic equilibrium. *Proceedings for NSF Engineering Research and Innovation Conference.*

Perakis, G. and G. Roels (2008). Regret in the newsvendor model with partial information. *Operations Research 56*(1), 188–203.

Perakis, G. and A. Sood (2006). Competitive multi-period pricing for perishable products. *Mathematical Programming.*

PVPS (2008). IEA PVPS - annual report.

Reichmuth, M., I. Erfurt, A. Schiffler, T. Kelm, and M. Schmidt (2010). Analyse zur mglichen anpassung der eeg-vergtung fr photovoltaik-anlagen.

Roels, G. and G. Perakis (2007). Robust controls for network revenue management. *Manufacturing & Service Operations Management.* (to appear).

Rogers, E. M. (1962). *Diffusion of Innovations.* Free Press, New York.

Rusmevichientong, P., B. Van Roy, and P. W. Glynn (2006). A non-parametric approach to multiproduct pricing. *Operations Research 54*(1), 89–98.

Savage, L. J. (1951). The theory of statistical decisions. *American Statistical Association Journal 46*, 55–67.

Scarf, H. E. (1958). A min-max solution to an inventory problem. *Studies in Mathematical Theory of Inventory and Production*, 201–209.

Schaeffer, G., E. Alsema, A. Seebregts, L. Beurskens, H. de Moor, W. van Sark, M. Durstewitz, M. Perrin, P. Boulanger, H. Laukamp, and C. Zuccaro (2004). Learning from the sun - analysis of the use of experience curves for energy policy purposes: The case of photovoltaic power. final report of the photex project.

Sderholm, P. and T. Sundqvist (2007). Empirical challenges in the use of learning curves for assessing the economic prospects of renewable energy technologies. *Renewable Energy 32*(15), 2559 – 2578.

Shurmer, M. and P. Swann (1995). An analysis of the process generating de facto standards in the pc spreadsheet software market. *Journal of Evolutionary Economics 5*(2), 119.

Soyster, A. (1973). Convex programming with set-inclusive constraints and applications to inexact linear programming. *Operations Research 21*, 1154–1157.

Stoneman, P. and P. Diederen (1994). Technology diffusion and public policy. *The Economic Journal 104*(425), pp. 918–930.

Swann, G. M. P. (2002). The functional form of network effects. *Information Economics and Policy 14*(3), 417 – 429.

Talluri, K. and G. van Ryzin (1998). An analysis of bid-price controls for network revenue management. *Management Science 44*(11), 1577–1593.

Talluri, K. and G. van Ryzin (2004). *The Theory and Practice of Revenue Management.* Kluwer Academic Publishers.

Thiele, A. (2006, January). Single-product pricing via robust optimization. Working paper, Lehigh University.

Train, K. (2003). *Discrete Choice Methods with Simulation.* Cambridge University Press.

Tucker, C. (2008). Identifying formal and informal influence in technology adoption with network externalities. *Management Science 54*(12), 2024 – 2038.

Ulu, C. and J. E. Smith (2009). Uncertainty, Information Acquisition, and Technology Adoption. *Operations Research 57*(3), 740–752.

van Ryzin, G. and J. McGill (2000). Revenue management without forecasting or optimization: An adaptive algorithm for determining airline seat protection levels. *Management Science 46*(6), 760–775.

van Ryzin, G. and G. Vulcano (2008). Simulation-based optimization of virtual nesting controls of network revenue management. *Operations Research 56*(4), 865–880.

Vives, X. (1997). Learning from others: A welfare analysis. *Games and Economic Behavior 20*(2), 177 – 200.

Wand, R. and F. Leuthold (2010). Feed-in tariffs for photovoltaics: Learning by doing in germany? Working Paper.

Wissing, L. (2006). Exchange and dissemination of information on pv power systems. *National Survey Report of PV Power Applications in Germany.* Prepared on behalf of BMU German Federal Ministry for the Environment, Nature Conservation and Nuclear Safety.

Yu, C., W. van Sark, and E. Alsema (2009). Incorporation of input prices and scale effects in a multi-factor learning curve for photovoltaic technology. *Proceedings of the 10th IAEE European Conference, Vienna.*.

Yue, J., B. Chen, and W. M.-C. (2006). Expected value of distribution information for the newsvendor problem. *Operations Research 54*(6), 1128–1136.

Zhan, R. and Z. Shen (2005). Newsvendor problem with pricing: properties, algorithms, and simulation. In *WSC '05: Proceedings of the 37th conference on Winter simulation*, pp. 1743–1748. Winter Simulation Conference.

Zhang, L. (2006, September). Multi-period pricing for perishable products: Uncertainty and competition. MS dissertation, Computation Design and Optimization, MIT.