Inverse Optimization, Incentive Design and Healthcare Policy

by

Auyon Siddiq

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Engineering – Industrial Engineering and Operations Research

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Zuo-Jun (Max) Shen, Co-Chair
Professor Anil Aswani, Co-Chair
Professor Lee Fleming
Professor Terry Taylor

Spring 2018

Abstract

Inverse Optimization, Incentive Design, and Healthcare Policy

by

Auyon Siddiq

Doctor of Philosophy in Industrial Engineering and Operations Research

University of California, Berkeley

Professor Zuo-Jun (Max) Shen, Co-Chair
Professor Anil Aswani, Co-Chair

This dissertation presents mathematical models and algorithms that draw from optimization and statistics and are motivated by practical problems in operations management. We discuss theoretical properties of the proposed models as well as their relevance to practice. In particular, we focus on the role these models can play in addressing challenges in healthcare operations and policy.

In Chapter 2, we address the problem of building models of agent behavior from observational data regarding the agent's decisions. Concretely, we consider the *inverse optimization* problem, which refers to the estimation of unknown model parameters of a convex optimization problem from observations of its optimal solutions. First, we provide a new formulation for inverse optimization, which takes the form of a bi-level program where the optimality conditions of the lower level program are expressed using strong duality. In contrast to existing methods, we show that the parameter estimates produced by our formulation are statistically consistent under appropriate conditions. Second, we propose two solution algorithms based on our duality-based formulation: an enumeration algorithm that is applicable to settings where the dimensionality of the parameter space is modest, and a semiparametric approach that combines nonparametric statistics with a modified version of our formulation. These numerical algorithms are shown to maintain the statistical consistency of the underlying formulation. Lastly, using both synthetic and real data, we demonstrate that our approach performs competitively when compared with existing heuristics.

In Chapter 3, we employ an inverse optimization approach to redesign a class of Medicare contracts. We formulate the existing contract between Medicare and a provider as a principal-agent model. We then propose an alternate contract, which we show to dominate the status quo contract under reasonable conditions by producing a strictly higher expected payoff for both Medicare and the provider. We then propose an estimator based on inverse optimization for estimating a model of provider behavior, using a dataset containing the financial performance of a group of Medicare providers that account for 7 million beneficiaries and over $70 billion in Medicare spending. We estimate a performance improvement – in terms of savings accrued by Medicare – of 40% under the alternate contract, which suggests significant room for improvement in the status quo.

In Chapter 4, we propose a data-driven modeling approach to facility location in a setting where the location of demand points is subject to uncertainty. The model is motivated by

the problem of placing automated external defibrillators in public locations in anticipation of sudden cardiac arrest. We propose a distributionally robust optimization approach where the demand distribution is continuous in the plane and uncertain. We propose a solution technique based on row-and-column generation that exploits the structure of the uncertainty set and allows us to solve practical-sized instances of the defibrillator deployment problem. Using real cardiac arrest data, we conduct an extensive numerical study and find that hedging against cardiac arrest location uncertainty can produce defibrillator deployments that outperform a intuitive sample average approximation by 9 to 15%. Our findings suggest that accounting for cardiac arrest location uncertainty can lead to improved accessibility of defibrillators during cardiac arrest emergencies and the potential for improved survival outcomes.

Contents

List of Figures

List of Tables

Acknowledgements

I am grateful to my advisors, Professors Anil Aswani and Zuo-Jun Max Shen, for their unabating support, patience and wisdom during the development of this thesis. Through countless meeting and discussions over the last five years, they have pushed me beyond my comfort zone to become a more critical thinker and effective researcher, for which I am thankful. I am also indebted to Professors Terry Taylor and Lee Fleming for their invaluable mentorship, and Professor Timothy C.Y. Chan for many helpful discussions regarding the third chapter. Lastly, I extend endless gratitude to my family for their continuing love and support.

Chapter 1

Introduction

This dissertation presents a collection of mathematical models at the intersection of optimization, statistics, and operations management. Particular emphasis is given on the application of these models to contemporary challenges in healthcare operations and policy. We focus on analyzing the theoretical properties of the proposed models as well as their practical relevance to related policy areas. While each of the three main chapters of this dissertation can be digested as an independent study, they are unified in their emphasis on the application of data analytics and optimization for the purposes of improving healthcare policy and practice. In this section, we provide a brief summary of the main ideas contained within each chapter, including their main contributions to the operations management literature.

In Chapter 2, we propose a new method for inverse optimization, which refers to the estimation of unknown parameters of a convex optimization problem from (potentially noisy) observations of its optimal solutions. This work is motivated by the recognition that a significant share of real world data represents *decisions* made by individuals or firms. This observed "decision data" can often be characterized as approximate solutions to suitably defined optimization problems. Accordingly, estimating the parameters of these latent optimization problems can shed insight into the agent's behavior, and also enable the design of contracts or mechanisms that induce desirable behavior from the agent.

Our key contribution is a new estimator for inverse optimization in the setting where observations of the optimal solution are noisy. We show that our approach produces parameter estimates that attain two statistical properties, which we call *risk consistency* and *estimation consistency*. In short, risk consistency refers to asymptotically optimal out-of-sample predictive performance, and estimation consistency refers to accurate inference of the unknown model parameters. While the formulation achieves these valuable statistical properties, it also takes the form of a bi-level program, which are in general non-convex and therefore difficult to solve to optimality. Therefore, our secondary contribution in Chapter 2 is to propose two algorithms for solving the bilevel formulation, both of which leverage a novel duality-based reformulation. We also prove that existing convex formulations for inverse optimization from the literature are statistically inconsistent, which implies that these approaches will generate parameter estimates that converge to incorrect values. This is perhaps not unexpected, because we also prove that the problem of inverse optimization with noisy data is NP-hard.

Additionally, we propose a novel reformulation for bilevel programs that contains a convex subproblem. This convex structure is absent if standard approaches (i.e., Karush-Kuhn-

Tucker conditions) are used to describe the lower-level problem. In particular, we propose to upper bound the objective function by its dual, which enforces a zero duality gap and therefore describes an optimal solution. The benefit of our optimality condition is it has convexity properties that support the design of our solution algorithms.

In Chapter 3, we consider the data-driven redesign of a class of contracts currently in use by the Centers for Medicare and Medicaid Services, known as the Medicare Shared Savings Program (MSSP). The MSSP is a voluntary program that aims to control costs by incentivizing Medicare providers to reduce healthcare spending through increasing the efficiency of healthcare delivery. Providers are assigned financial benchmarks based on their historical spending and are given bonus payments for reducing spending to below their financial benchmark. A key step in our analysis is to build a model of provider behavior from financial data from the MSSP. We leverage the key idea from Chapter 2 by formulation an estimator for provider behavior as an inverse optimization problem.

Our specific contributions in Chapter 3 are as follows. First, we prove that, under reasonable conditions, introducing a performance-based subsidy to the MSSP can increase total Medicare savings, despite the additional payments that must be made to the provider. Since the provider always benefits from subsidy payments, this result implies the existence of a subsidized contract that *dominates* all possible contracts within the current program. In other words, we show that introducing a subsidy can boost Medicare savings in addition to increasing provider payments, making the MSSP more attractive to providers. We also show that providers that are ineffective at generating savings may drop out of the MSSP under the current contract, which provides further support in favor of introducing a subsidy to the MSSP.

Second, based on our structural estimation, we find that optimizing the shared savings rate of the existing MSSP contract results in minimal improvements to Medicare savings. Moreover, optimizing the shared savings rate has the undesirable effect of decreasing provider payments, which can further jeopardize provider participation. By contrast, we estimate that under the proposed contract, Medicare savings and provider payments may increase by 43% and 8%, respectively, which supports the dominance result discussed above, and provides evidence in favor of incorporating provider investment subsidies into a revised MSSP contract. We also find that the proposed subsidized contract performs favorably in comparison to a fully non-parametric contract, which represents the theoretical best performance that can be expected under the estimated model.

Third, our estimation also reveals that providers with high benchmark expenditures are more likely to be effective at generating savings than providers with low benchmarks. On average, we find that providers with benchmark expenditures of less than $10,000 per beneficiary were unable to generate savings on a per beneficiary basis, whereas providers with benchmarks greater than $14,000 reduced spending by $260 per beneficiary. This disparity in provider performance may be explained by the fact that the financial benchmarks are calculated based on the provider's historical spending. As a consequence, an provider with historically high spending may have relatively more "room for improvement" than an provider that has historically been cost efficient. Since provider payments are proportional to savings, our finding suggests that providers with low benchmarks are at an increased risk of abandoning the MSSP, which highlights a need for adjusting the benchmarking methodology.

In Chapter 4, we consider a more localized healthcare policy problem: improving sur-

vival outcomes from public location cardiac arrest by equipping bystanders with automated external defibrillators (AED) and empowering them to act as "lay responders". Our contributions in Chapter 4 are as follows. First, we present a robust optimization model for AED deployment that accounts for uncertainty in the spatial distribution of cardiac arrests. Our model permits cardiac arrests to arrive anywhere within a continuous service area, which we approximate through a fine discretization. We use historical cardiac arrest data to construct a distributional uncertainty set that consists of a set of generally-shaped uncertainty regions, where only the probability of the next cardiac arrest event occurring within each uncertainty region is known. We also present an auxilliary result that shows that our choice of uncertainty set and parameters unifies the canonical p-median and p-center problems, as well as their robust analogues.

Second, we present a bound on the approximation error induced by the discretization of the service area, which depends on the granularity of the discretization and the configuration of the uncertainty regions. The bound can be used to quantify the suboptimality of the solutions produced by our discretized model with respect to the underlying continuous problem.

Third, we propose a new row-and-column generation algorithm for solving the robust AED deployment problem. Our proposed algorithm operates by exploiting the structure of the uncertainty set, which allows it to substantially outperform a standard Benders-based algorithm and a single-stage MIP reformulation. Notably, our algorithm decouples the size of the master problem from the number of scenarios used to discretize the service area, which allows us to efficiently solve problems where the discretization is extremely fine. Combined with the bounds on the discretization error, the row-and-column generation algorithm allows us to efficiently obtain provably near-optimal solutions to the underlying continuous problem.

Lastly, we present results from an extensive numerical study on the placement of public AEDs using real cardiac arrest data from Toronto, Canada. Our results demonstrate that accounting for uncertainty in cardiac arrest locations can decrease the distance between cardiac arrest victims and the nearest AED by 9-15%, under both typical and worst-case realizations of the uncertainty.

Chapter 2

Inverse Optimization with Noisy Data

2.1 Background and motivation

An appreciable share of real-world data represents *decisions* by individuals or firms, which can often be characterized as the solutions of correspondingly-defined optimization problems. Estimating the parameters of these latent optimization problems has the potential to enable the prediction of future decisions and support the design of data-driven mechanisms. Examples of domains where this is important include health systems engineering (Aswani et al., 2016), energy systems engineering (Ratliff et al., 2014a), and marketing (Green and Srinivasan, 1990), where such estimation may lead to new approaches that enable the individualization of products and incentives. For example, consider a single homeowner who *each day* observes an electricity price and weather forecast and then adjusts the temperature set-point for their home's air-conditioner. By modeling this homeowner's decision as being generated from an optimization problem, we can directly estimate the price elasticity of comfort – as measured by a standardized function of the temperature set-point and the outside temperature (ASHRAE, 2013) – for this particular homeowner. This information is valuable for designing personalized incentive bonus schemes that encourage participation in demand-response programs (Aalami et al., 2010) or promote energy-efficiency (Aswani and Tomlin, 2012).

This chapter considers the problem of estimating unknown model parameters of an optimization problem based on noisy measurements of its optimal solutions, which is often referred to as *inverse* optimization. In particular, we provide the first *statistical inference* perspective on the inverse optimization problem. This is important because real-world decision data is noisy, either because (i) the data collection process introduces measurement noise, (ii) the decision-maker deviates from optimal decisions – phenomena often referred to as *bounded rationality* (Tversky and Kahneman, 1981), or (iii) there is mismatch between the parametric form of the model and the true underlying decision-making process.

Noisy data make inverse optimization challenging because noise in the solution data can preclude the existence of a single set of model parameters that renders all observed solutions exactly optimal. In this setting, the goal of inverse optimization is to find a set of model parameters that achieves a good "fit" with respect to the solution data. More specifically, we

are interested in two statistical questions. First, how can we generate estimates of unknown model parameters that asymptotically provide the best possible predictions from the chosen parametric form of the model? In statistics, this property is known as *risk consistency* (Bartlett and Mendelson, 2002, Greenshtein and Ritov, 2004, Chatterjee, 2014). Second, when the chosen model matches the true model that is is generating the solution data, how can we generate estimates that asymptotically converge to the true value of the unknown parameters? In statistics, this property is known as *consistency* (Wald, 1949, Jennrich, 1969, Bickel and Doksum, 2006). We will use the term *estimation consistency* to distinguish this concept from risk consistency. Note that estimation consistency generally implies risk consistency.

Restated, a risk consistent estimate asymptotically achieves the lowest possible prediction error (out of all possible predictions permitted by the class of models considered). Hence, risk consistency and estimation consistency allow us to be confident that prediction and estimation accuracy, respectively, will generally improve with additional data. By contrast, an estimator that fails to be risk consistent (so-called *inconsistent* estimators) may yield poor predictions, even if a large amount of data is available. Proving consistency of an estimator is an important topic in the theory of statistical inference (cf. Wald (1949), Jennrich (1969), Bartlett and Mendelson (2002), Greenshtein and Ritov (2004), Bickel and Doksum (2006), Chatterjee (2014), Aswani (2015)), and consistency is considered to be a minimal requirement for an estimator (Bickel and Doksum, 2006).

We provide an overview of the main sections of this chapter. In Section 2.3, we describe the statistical and computational challenges of inverse optimization with noisy data. The section begins by formally defining a (convex) forward optimization problem and its corresponding inverse optimization problem. We specifically formulate the inverse optimization problem such that its solution has the desired statistical consistency properties. Our approach is conceptually similar to least squares regression in the sense that we also employ a sum-of-squares loss function to fit a parametric model to noisy data. The substantive difference is that inverse optimization involves estimating the (possibly multi-valued) solution set of a general convex optimization problem, whereas regression typically involves estimating a (single-valued) function which has a closed form expression. Due to these differences, much of the classical statistical theory on least-squares regression (Jennrich, 1969) is invalid in the inverse optimization setting, and thus new analysis is required. We also note that our approach is not restricted to the use of an ℓ_2 norm: Results similar to those in this chapter can be proved for other loss functions, such as absolute deviation or a likelihood function, but we do not consider those extensions in this work.

In Section 2.4, we prove that our inverse optimization formulation produces statistically consistent estimates of the unknown model parameters. The key technical difficulty in proving these results is dealing with continuity issues. In particular, the risk measures are not continuous in the general case, but are rather lower semicontinuous. As alluded to above, this precludes the use of the typical statistical machinery used to prove consistency results (namely the uniform law of large numbers (Jennrich, 1969) and related uniform bounds (Bartlett and Mendelson, 2002, Greenshtein and Ritov, 2004)). To circumvent this difficulty, we define a regularized version of the inverse optimization problem that smooths out any discontinuities. We show that the regularization parameter can be adjusted to approximate the original inverse optimization problem to arbitrary accuracy. The regularized version of

our formulation enables us to prove the desired statistical consistency results.

Section 2.5 provides two numerical algorithms for solving our formulation of the inverse optimization problem. The first numerical algorithm is an enumeration algorithm that is applicable to settings where the dimensionality of the parameter space is modest (i.e., at most four or five parameters). The second numerical algorithm is a semiparametric approach that combines nonparametric statistics with a modified version of our formulation of the inverse optimization problem. The statistical consistency of these two numerical algorithms are shown using the results from Section 2.4. Lastly, in Section 2.6, we use synthetic and real data to demonstrate the competitiveness of our approach when compared to existing heuristics (Keshavarz et al., 2011, Bertsimas et al., 2015).

2.2 Related literature

Existing inverse optimization models differ based on their specification of the *loss function*, and the different models can be broadly categorized into either (i) deterministic settings, or (ii) noisy settings. The work in the deterministic setting has primarily focused on single observation situations, wherein a single optimal solution is observed and then used to estimate parameters of the optimization problem. However, in the noisy setting past work has considered situations with either a single observation and multiple observations.

We begin by describing some of the work in the deterministic setting: Ahuja and Orlin (2001a) consider the estimation of objective function coefficients of general linear programs given a single optimal solution. The feasible region of the inverse problem is formulated using the constraints of the dual program and complemetary slackness conditions. Since the observed solution is assumed to be optimal, feasibility of the inverse problem is guaranteed. Iyengar and Kang (2005) and Zhang and Xu (2010) extend inverse optimization to certain conic forward problems using conic duality theory. Inverse optimization models have also been studied in the context of integer programs (Schaefer, 2009, Wang, 2009) and network problems (Burton and Toint, 1992, Hochbaum, 2003, Zhang and Liu, 1996). With respect to applications, inverse optimization models has been employed in many different domains, including healthcare (Erkin et al., 2010, Chan et al., 2014), energy (Ratliff et al., 2014b, Saez-Gallego et al., 2016), finance (Bertsimas et al., 2012), production planning (Troutt et al., 2006), demand management (Carr and Lovejoy, 2000, Bajari et al., 2007), auction design (Beil and Wein, 2003), telecommunication (Faragó et al., 2003) and geoscience (Burton and Toint, 1992). We refer the reader to Heuberger (2004) for a survey of inverse optimization methods.

The noisy setting has been less studied. Chan et al. (2014) propose a generalized approach to inverse optimization for linear programs where the (single) observed solution may be suboptimal or infeasible. Instead of complementary slackness, the authors use dual feasibility and strong duality to formulate the inverse problem. To accommodate noise, the strong duality constraint is relaxed to guarantee feasibility of the inverse problem. Saez-Gallego et al. (2016) also consider inverse optimization for linear programs, and formulate the inverse problem using KKT conditions. Keshavarz et al. (2011) formulates the inverse problem using the KKT conditions of the optimization problem. To accommodate noise, the KKT conditions are relaxed by introducing slack variables to allow the data to "approximately" satisfy the KKT conditions. Similarly, Bertsimas et al. (2015) consider inverse problems where the

observed data are assumed to be in an equilibrium. The authors enforce optimality conditions using a variational inequality, and similarly relax the optimality conditions by introducing slack variables to allow the data to "approximately" satisfy the variational inequality.

Our work in this chapter is most closely related to the noisy setting with multiple observations that has been previously considered by Keshavarz et al. (2011) and Bertsimas et al. (2015). The key distinction between our work and these two previous approaches is in the choice of the loss function. In (Keshavarz et al., 2011) and (Bertsimas et al., 2015), the loss function is measured by the amount of slack required to make the measured data satisfy an approximate optimality condition (either the KKT conditions (Keshavarz et al., 2011) or a variational inequality describing optimality (Bertsimas et al., 2015)). In contrast, our approach is to jointly estimate (i) the parameters of the optimization problem, and (ii) the denoised versions of the measured data (i.e. the true underlying optimal solutions). By performing this joint estimation, we are able to define our loss function to be the average discrepancy between the measured data and the (estimated) denoised data. As we will show, this difference in loss function leads to significantly improved statistical performance. A secondary distinction is that we propose the use of a novel optimality condition: specifically, we upper bound the objective function of a convex optimization problem by its dual – thereby enforcing a zero duality gap and guaranteeing optimality. An important benefit of using this alternate optimality condition is that it has favorable convexity and continuity properties (which are not available when using KKT conditions or variational inequalities to represent optimality) that enable design of numerical algorithms for solving the inverse optimization problem.

Lastly, we note here that the work in this chapter coincides with the paper by Aswani et al. (2018b).

2.3 Model

This section begins by formalizing the notation for the forward problem, before defining the noisy inverse optimization problem. For the case where we have access to data (rather than the underlying distributions), we formulate a related sample average approximation of the inverse optimization problem. We show that both these inverse problems are NP-hard. We conclude by showing that existing heuristic approaches for solving the inverse optimization problem are statistically inconsistent, meaning that in the limit of infinite data, these heuristic approaches converge to incorrect estimates of the unknown model parameters and produce sub-optimal estimates with respect to the expected prediction error.

We briefly summarize some of the less usual aspects of our notation. We use $\|\cdot\|$ to denote the ℓ_2-norm. The indicator function $\mathbf{1}(p)$ is defined to be

$$\mathbf{1}(p) = \begin{cases} 1, & \text{if condition } p \text{ is satisfied} \\ 0, & \text{otherwise} \end{cases} \tag{2.1}$$

The notation $[r] = \{1, \ldots, r\}$ refers to sequential set. The Kuratowski limit superior of a sequence of sets $\mathcal{C}_\nu \subseteq \mathbb{R}^d$ is defined as

$$\limsup_\nu(\mathcal{C}_\nu) = \{x \in \mathbb{R}^d : \liminf_\nu \operatorname{dist}(x, \mathcal{C}_\nu) = 0\}, \tag{2.2}$$

where $\text{dist}(x, \mathcal{C}) = \inf\{\|x - c\| \mid c \in \mathcal{C}\}$. We similarly define $\text{dist}(\mathcal{B}, \mathcal{C}) = \inf\{\text{dist}(x, \mathcal{C}) \mid x \in \mathcal{B}\}$.

2.3.1 Forward optimization problem

Let $x \in \mathbb{R}^d$ be the decision variable, $u \in \mathbb{R}^m$ be the external input variable, and $\theta \in \mathbb{R}^p$ be the parameter vector. Then the forward optimization problem is given by

$$\min_x \{f(x, u, \theta) \mid g(x, u, \theta) \leq 0\}, \qquad \textsf{FOP}$$

where $f : \mathbb{R}^d \times \mathbb{R}^m \times \mathbb{R}^p \to \mathbb{R}$ is a function and $g : \mathbb{R}^d \times \mathbb{R}^m \times \mathbb{R}^p \to \mathbb{R}^q$ is a vector-valued function. The solution set of FOP is the set-valued function given by $\mathcal{S}(u, \theta) = \arg\min_x\{f(x, u, \theta) \mid g(x, u, \theta) \leq 0\}$. The value function of FOP is given by

$$V(u, \theta) = \min_x\{f(x, u, \theta) \mid g(x, u, \theta) \leq 0\},$$

and the feasible set is defined as $\Phi(u, \theta) = \{x \in \mathbb{R}^d : g(x, u, \theta) \leq 0\}$.

2.3.2 Inverse optimization problem

Suppose $(u, y) \in \mathbb{R}^m \times \mathbb{R}^d$ is a vector-valued random variable that is distributed according to some unknown but fixed joint distribution $\mathbb{P}_{(u,y)}$. Let $\mathcal{U} \times \mathcal{Y} \subseteq \mathbb{R}^m$ be the support of this distribution, meaning the smallest set that satisfies the property $\mathbb{P}_{(u,y)}(\mathcal{U}, \mathcal{Y}) = 1$. If we define the function

$$Q(\theta) = \mathbb{E}\Big(\min_{x \in \mathcal{S}(u,\theta)} \|y - x\|^2\Big), \qquad \textsf{RISK}$$

then the inverse optimization problem is given by

$$\min\{Q(\theta) \mid \theta \in \Theta\}, \qquad \textsf{IOP}$$

where $\Theta \subseteq \mathbb{R}^p$ is a known set. We make the following assumptions:

A1. The functions $f(x, u, \theta)$ and $g(x, u, \theta)$ are continuous in x, u, θ and convex in x for fixed u, θ.

A2. The set Θ is convex.

These assumptions are fairly mild. **A1** is equivalent to stating that FOP is a convex optimization problem. Although **A2** is necessary for the semiparametric algorithm presented in Section 4 (because it ensures polynomial time computability of the algorithm), it is not necessary for our main results regarding statistical consistency because these results only require that Θ is well-posed. Hence, **A2** is one way to ensure Θ is well-posed. An alternative assumption under which our statistical consistency results can be shown to hold is if Θ is discrete-valued and finite.

When the joint distribution $\mathbb{P}_{(u,y)}$ is unknown, we cannot solve IOP without additional information. Fortunately, we can leverage the iid measurements (u_i, y_i) for $i \in [n]$. In

principle, we can solve IOP using the following sample average approximation:

$$\min \{Q_n(\theta) \mid \theta \in \Theta\}, \qquad \text{IOP–SAA}$$

where

$$Q_n(\theta) = \min_{x_i} \frac{1}{n} \sum_{i=1}^{n} \|y_i - x_i\|^2 \qquad \text{RISK–SAA}$$
$$\text{s.t. } x_i \in \mathcal{S}(u_i, \theta), \qquad \forall i \in [n]$$

In the context of a decision-making agent, u_i may be interpreted as an external signal the agent responds to and y_i as a noisy observation of the corresponding decision of the agent. Note that in the expression RISK, the variable x is constrained to be an optimal solution of the forward problem. Similarly, we may interpret x_i as representing an underlying optimal solution (unperturbed by noise) of FOP in the i^{th} instance. Note also that while the u_i and θ are both parameters of FOP, they are different in that the u_i are known and may vary across the n observations, whereas θ is unknown and is fixed across all instances.

For a concrete example, consider the numerical experiments presented in Section 5.4, where we estimate an individual's utility function capturing the tradeoff between maintaing a comfortable indoor temperature versus the amount of energy consumption (and implicitly the air conditioning energy costs) required to cool the room. In that example, the u represents the outside air temperature, θ_1 captures the decision-maker's (unknown) tradeoff between comfort and energy consumption, θ_2 parameterizes their (unknown) preferred temperature (i.e., the preferred temperature is $\theta_2 + u$), x represents the true optimal temperature setpoint (for the given u and θ), and y represents the temperature set-point that we observe.

2.3.3 NP-hardness

Although all the functions and sets involved in FOP and IOP are convex, solving IOP is NP-hard.

Theorem 2.1. *If* **A1** *and* **A2** *hold, then* IOP *is NP-hard.*

Inapproximability results for IOP can be shown under the setting where Θ is allowed to be a discrete set (i.e, **A1** holds, but **A2** does not hold). In particular, there is a straightforward reduction from the shortest vector problem. This implies that IOP is NP-hard to approximate to within any factor up to $2^{(\log d)^{1-\epsilon}}$, for any $\epsilon \geq 0$ (Haviv and Regev, 2012).

Polynomial-time solvability of IOP is possible in very specific settings. For instance, if FOP is a QP with the solution set $\mathcal{S}(u, \theta) = \arg\min_x\{x^2 - 2(\theta + u) \cdot x\} = \theta + u$ or an LP with the solution set $\mathcal{S}(u, \theta) = \arg\min_x\{x : x = \theta + u\} = \theta + u$, then IOP is a QP: $\min_{\theta \in \Theta} \{\mathbb{E}((y - \theta - u)^2)\}$, and its minimizer is $\theta^* = \mathbb{E}(y - u)$. If FOP is a general linear program, then the inverse problem IOP takes the form of a quadratic bilevel program, which are generally NP-hard (Audet et al., 1997). Branch-and-bound algorithms have been proposed for solving such bilevel programs (Bard and Moore, 1990). In general, since $\mathcal{S}(u_i, \theta)$ is the optimal solution sets to FOP under input u_i, the problem IOP–SAA is a bilevel program, which are usually difficult to solve (Dempe et al., 2015). In fact, IOP–SAA is also NP-hard to solve.

Corollary 2.1. *If* **A1** *and* **A2** *hold, then* IOP–SAA *is NP-hard.*

Polynomial-time solvability of IOP–SAA is possible in very specific settings. For instance, the constructions in Remark **??** lead to instances of IOP–SAA that are quadratic programs.

2.4 Consistent estimation in inverse optimization

We begin with two statistical definitions of consistency: risk consistency and estimation consistency. These definitions are stated in order of increasing stringency, meaning that risk consistency is necessary (in situations with sufficient continuity) for estimation consistency. The first definition relates to the best predictions possible using the given forward optimization problem.

Definition 2.1 (Risk consistency)**.** An estimate $\hat{\theta}_n \in \Theta$ is risk consistent if

$$Q(\hat{\theta}_n) \xrightarrow{p} \min\{Q(\theta) \mid \theta \in \Theta\}. \tag{2.3}$$

We should interpret the function $Q(\theta)$ as the expected prediction error when the parameter values are θ, where the prediction is the solution set $\mathcal{S}(u, \theta)$. And so the above definition is stating that an estimator θ_n is risk consistent if the expected prediction error of the estimate θ_n converges in probability to the minimum prediction error possible when we use the forward optimization model described by FOP and constrain θ to belong to Θ. In other words, an estimator is risk consistent if it asymptotically provides the best predictions possible.

The second statistical definition relates to the situation where the forward optimization model described by FOP is correct and there is a *true* parameter. In particular, it applies to situations where the below identifiability condition is satisfied. Briefly summarized, the identifiability condition is satisfied when FOP is such that two different parameter values θ_1 and θ_2 lead to two different distributions for measurements of the decision data y_i. More details and clarifying examples are found in Appendix 2.4.4.

IC. There exists a unique $\theta_0 \in \Theta$ such that the following three sub-conditions hold: (i) $y = \xi + w$, where $\xi \in \mathcal{S}(u, \theta_0)$, $\mathbb{E}(w) = 0$, $\mathbb{E}(w^2) < +\infty$, and u, ξ are independent of w, (ii) for all $\theta \in \Theta \setminus \theta_0$ there exists $\mathcal{U}(\theta) \subseteq \mathcal{U}$ such that $\mathbb{P}(u \in \mathcal{U}(\theta)) > 0$ and $\text{dist}(\mathcal{S}(u, \theta), \mathcal{S}(u, \theta_0)) > 0$ for each $u \in \mathcal{U}(\theta)$, and (iii) for each fixed $\theta \in \Theta$ we have $\mathbb{P}(\{u : \mathcal{S}(u, \theta) \text{ is multivalued}\}) = 0$.

The first sub-condition of the identifiability condition is stating that the solution data y_i is a noisy measurement (with noise random variable w) of a point that belongs to the solution set $\mathcal{S}(u_i, \theta_0)$, and the second sub-condition is stating that when θ is different from θ_0 then this leads to different solution sets. This second sub-condition is necessary, because otherwise we could not distinguish the predictions of FOP when the parameters θ differ from θ_0. The third sub-condition eliminates pathological cases that occur when the solution set at a fixed θ is so large that it approximately encompasses all possible solutions. Note that this third sub-condition is mild, and examples where it is satisfied include when (i) FOP is strictly convex, or when (ii) FOP is a linear program with random coefficients drawn from a continuous distribution; it holds for other examples as well. The second statistical definition is related to this identifiability condition.

Definition 2.2 (Estimation consistency)**.** Suppose **IC** holds. An estimate $\hat{\theta}_n \in \Theta$ is estimation consistent if

$$\hat{\theta}_n \xrightarrow{p} \theta_0. \tag{2.4}$$

Stated in words, an estimate $\hat{\theta}_n$ is estimation consistent if it converges in probability to the true parameter values θ_0. This is the classical notion of consistency of a statistical estimator (Bickel and Doksum, 2006).

Though these statistical notions of consistency are quite natural, it is the case that existing heuristic approaches for solving the inverse optimization problem are statistically inconsistent. We will use VIA to refer to the variational inequality method of Bertsimas et al. (2015), and we refer to the KKT conditions approach of Keshavarz et al. (2011) as KKA.

Proposition 2.1. *Suppose* **A1**,**A2** *and* **IC** *hold. Then* VIA *(Bertsimas et al., 2015) and* KKA *(Keshavarz et al., 2011) are not estimation consistent.*

Corollary 2.2. *Suppose* **A1**,**A2** *hold. Then* VIA *(Bertsimas et al., 2015) and* KKA *(Keshavarz et al., 2011) are not risk consistent.*

The intuition for why VIA and KKA are statistically inconsistent is that they are minimizing an incorrect measure of error: These approaches generate an estimated set of parameters that minimizes the level of suboptimality of the measured solution data. However, this leads to biased estimates because suboptimality is measured by (i) deviations in the value of the objective function of FOP and (ii) the amount of constraint violation of FOP, whereas noise directly perturbs the solution data. This is in contrast to our approach (as exemplified by IOP-SAA) which generate an estimated set of parameters that minimizes the deviation between predicted and measured solution data. This distinction between suboptimality and deviations in the solution data becomes most apparent (and critical) in problems with constraints.

Given the statistical inconsistency of existing heuristics, we propose to solve the noisy inverse optimization problem by instead solving SAA-IOP. First, we will need to impose a regularity condition to ensure that FOP and IOP–SAA are numerically well-posed:

R1. For each $u \in \mathcal{U}$ and $\theta \in \Theta$, the feasible set $\Phi(u, \theta)$ is closed, bounded, and has a nonempty interior (i.e., $\text{int}(\Phi(u, \theta)) \neq \emptyset$). The feasible set $\Phi(u, \theta)$ is also absolutely bounded, meaning there exists $M > 0$ such that $\|x\| \leq M$, for all $x \in \Phi(u, \theta)$, $u \in \mathcal{U}$, and $\theta \in \Theta$.

Condition **R1** is equivalent to requiring FOP to have a strictly feasible point (i.e., Slater's condition holds), and that the feasible set of FOP is closed and bounded. The first subcondition requiring the feasible set be closed and bounded is needed to ensure the existence of well-posed primal and dual solutions, and it could be replaced by more general conditions. For instance, we could have instead assumed FOP satisfies the uniform level-boundedness condition (Rockafellar and Wets, 1998). We use the above for simplicity of stating the results. The condition that $\Phi(u, \theta)$ has a nonempty interior[1] is needed to ensure continuity of $\mathcal{S}(u, \theta)$ through application of the Berge Maximum Theorem (Berge, 1963).

[1]

[1] **R1** can be relaxed to requiring a nonempty *relative* interior if the affine constraints of FOP are of the form $Mx + \zeta(u, \theta) = 0$, where M is a matrix and ζ is a continuous function. The reason is that our proofs

The simplest case of statistical consistency of SAA-IOP occurs when the function $f(x, u, \theta)$ is strictly convex, because of the following result:

Proposition 2.2. *Suppose* **A1**,**A2** *and* **R1** *hold. If $f(x, u, \theta)$ is strictly convex in x for fixed $u \in \mathcal{U}$ and $\theta \in \Theta$, then $Q_n(\theta)$ is continuous.*

In this case, we can prove risk and estimation consistency using standard arguments (Jennrich, 1969, van der Vaart, 2000, Bickel and Doksum, 2006) from statistics that use the uniform law of large numbers (Jennrich, 1969). However, this approach cannot be applied to the more general case where $f(x, u, \theta)$ is not strictly convex. In particular, when $f(x, u, \theta)$ is not strictly convex, the function $Q_n(\theta)$ will not generally be continuous. And so a different argument is required because the uniform law of large numbers does not apply to discontinuous functions.

Our approach will be to use a statistical consistency result originally due to Wald (1949) that uses a one-sided bounding argument. The advantage of this approach is that it only requires lower semicontinuity, which we show always holds for $Q_n(\theta)$. However, this result only implies the estimates $\hat{\theta}_n$ converge in probability to the set of minimizers of $Q(\theta)$. This cannot imply risk consistency in the general case because $Q_n(\theta)$ is lower semicontinuous, which means that $Q(\hat{\theta}_n)$ can remain bounded from the minimum $Q(\theta)$. And so for the general case, we will show that a weak risk consistency result holds.

To develop the statistical consistency results for the most general case, we will develop a regularized version of RISK–SAA that is guaranteed to be continuous. The first step of this construction involves proposing a new reformulation for bilevel programs that we call a duality-based reformulation. Next, we use this reformulation to construct a regularized version of RISK–SAA and prove its continuity. We use this regularized version to prove statistical consistency results about IOP–SAA and a regularized version of IOP–SAA.

2.4.1 Duality-based reformulation

One approach to solving bilevel problems (such as IOP–SAA) is to reformulate the problem as a normal (i.e., single level) optimization problem by replacing the constraints $x_i \in \mathcal{S}(u_i, \theta)$ with an optimality condition (Dempe et al., 2015). One possibility is to replace $x_i \in \mathcal{S}(u_i, \theta)$ by the KKT conditions of FOP, and another possibility is to upper bound the objective function using the value function $f(x_i, u_i, \theta) \leq V(u_i, \theta)$. Unfortunately, these approaches often encounter numerical difficulties. The KKT approach leads to a nonlinear program with combinatorial complexity, because of the complimentary slackness in KKT. The value function approach is difficult to implement because closed-form expressions for the value function are not available except for very special cases.

Here, we present a new optimality condition. Given the numerical difficulties of existing approaches, we propose to solve bilevel programs (such as IOP–SAA) by using the Lagrangian dual function to upper bound the objective function. The following proposition shows that our idea of using the dual as an upper bound represents a novel optimality condition.

make use of a result (Example 5.10 of Rockafellar and Wets (1998)) on the continuity of parametrized convex constraints with a nonempty interior, and this result can be generalized for the above case through minor modifications (using corresponding results on relative interiors from Section 2.H of Rockafellar and Wets (1998)) to ensure continuity of the feasible set of FOP with a nonempty relative interior. Generalizing Example 5.10 of Rockafellar and Wets (1998) or our results to cases with more complex affine constraints will require further study.

Proposition 2.3. *Suppose* **A1** *and* **R1** *hold. Then $x \in \mathcal{S}(u,\theta)$ if and only if there exists a corresponding $\lambda \in \mathbb{R}^q$ for which x, λ satisfy the inequalities*

$$\begin{aligned} & f(x,u,\theta) - h(\lambda,u,\theta) \leq 0 \\ & g(x,u,\theta) \leq 0 \\ & \lambda \geq 0 \end{aligned} \tag{2.5}$$

where $h(\lambda,u,\theta)$ is the Lagrangian dual function of FOP*.*

Proposition 2.3 is a consequence of strong duality for convex optimization problems. We can now exactly reformulate RISK–SAA as the following optimization problem:

$$\begin{aligned} Q_n(\theta) = \min_{x_i,\lambda_i} \ & \frac{1}{n}\sum_{i=1}^{n} \|y_i - x_i\|^2 & & \\ \text{s.t.} \ & f(x_i,u_i,\theta) - h(\lambda_i,u_i,\theta) \leq 0, & \forall i \in [n] & \qquad \text{DB–RISK–SAA} \\ & g(x_i,u_i,\theta) \leq 0, & \forall i \in [n] & \\ & \lambda_i \geq 0, & \forall i \in [n] & \end{aligned}$$

It should be noted that the formulation DB–RISK–SAA requires the Lagrangian dual function $h(\lambda,u,\theta)$ to be computable in closed form, which is the case for a large class of convex (e.g., linear, quadratic, conic) optimization problems that arise in practice (Boyd and Vandenberghe, 2009). In cases where the dual function does not have an analytical representation, we may still solve DB–RISK–SAA by developing an algorithm that computes $h(\lambda,u,\theta)$ numerically, although designing such an algorithm is beyond the scope of this work.

One important feature of this reformulation is that it is a convex optimization problem for fixed values of θ.

Proposition 2.4. *Suppose* **A1** *and* **R1** *hold. Then* DB–RISK–SAA *is a convex optimization problem for fixed θ.*

Proposition 2.4 follows directly from **A1** and the concavity of the dual function in λ.

2.4.2 Regularized formulation

Recall that $Q_n(\cdot)$ is generally not continuous even when **A1**,**A2**, and **R1** hold. Consequently, we develop a regularized version of the duality-based problem that is guaranteed to be continuous. We define the ϵ-regularized version of the duality-based problem to be

$$\begin{aligned} Q_n(\theta;\epsilon) = \min_{x_i,\lambda_i} \ & \frac{1}{n}\sum_{i=1}^{n} \|y_i - x_i\|^2 & & \\ \text{s.t.} \ & f(x_i,u_i,\theta) - h(\lambda_i,u_i,\theta) \leq \epsilon, & \forall i \in [n] & \qquad \text{R–DB–RISK–SAA} \\ & g(x_i,u_i,\theta) \leq \epsilon, & \forall i \in [n] & \\ & \lambda_i \geq 0, & \forall i \in [n] & \end{aligned}$$

We associate this to a regularized version of the sample average approximation of the inverse optimization problem:

$$\min\{Q_n(\theta;\epsilon) \mid \theta \in \Theta\}. \qquad \text{R–IOP–SAA}$$

The idea of this regularization is that we relax the optimality conditions to allow points x_i to be an ϵ-optimal solution. Recall that a point

$$x^\epsilon \in \epsilon\text{-}\arg\min\{f(x) \mid g(x) \leq 0\}, \tag{2.6}$$

if $f(x^\epsilon) - f^* \leq \epsilon$ and $g(x^\epsilon) \leq \epsilon$, where $f^* = \min\{f(x) \mid g(x) \leq 0\}$.

Proposition 2.5. *Suppose* **A1** *and* **R1** *hold. Then a point x is an ϵ-optimal solution if and only if there exists a corresponding $\lambda \in \mathbb{R}^q$ for which x, λ satisfy the inequalities*

$$\begin{aligned} &f(x,u,\theta) - h(\lambda,u,\theta) \leq \epsilon \\ &g(x,u,\theta) \leq \epsilon \\ &\lambda \geq 0 \end{aligned} \tag{2.7}$$

where $h(\lambda,u,\theta)$ is the Lagrangian dual function of FOP.

Proposition 2.5 follows from continuity of $f(x,u,\theta)$ and $g(x,u,\theta)$. One benefit of this regularization is that it ensures convexity of R–DB–RISK–SAA when θ is fixed.

Proposition 2.6. *Suppose* **A1**,**A2** *and* **R1** *hold. Then* R–DB–RISK–SAA *is a convex optimization problem for fixed θ.*

Proposition 2.6 follows immediately from the convexity of FOP and concavity of the dual function $h(\lambda,u,\theta)$ in λ. Though the above propositions show that the regularization is equivalent to replacing optimality conditions with ϵ-optimality conditions while maintaining convexity for fixed values of θ, the main benefit of the regularization is that it ensures the function $Q_n(\theta;\epsilon)$ defined in R–DB–RISK–SAA is continuous in θ, ϵ for any $\epsilon > 0$.

Proposition 2.7. *Suppose* **A1**,**A2** *and* **R1** *hold. Then the function $Q_n(\theta;\epsilon)$ is jointly continuous in θ, ϵ for any $\epsilon > 0$.*

A point of note is that within the above proof, we show that the set of ϵ-optimal solutions of a parametric convex optimization problem $\mathcal{S}(u_i,\theta;\epsilon)$ is continuous with respect to the parametrization θ; this is in contrast to the solution set of a parametric convex optimization problem $\mathcal{S}(u_i,\theta)$, which is in general only upper hemicontinuous with respect to the parametrization θ. The case of a parametric strictly convex optimization problem is the exception, which as shown in the proof of Proposition 2.2 has a continuous (with respect to the parametrization θ) solution set.

The function $Q_n(\theta;\epsilon)$ will not be jointly continuous in θ, ϵ at $\epsilon = 0$. However, it satisfies another property that is useful for solving IOP–SAA:

Proposition 2.8. *Suppose* **A1**,**A2** *and* **R1** *hold, and let $\epsilon_\nu > 0$ be a monotone decreasing sequence with $\epsilon_\nu \to 0$. Then we have $\min\{Q_n(\theta;\epsilon_\nu) \mid \theta \in \Theta\} \to \min\{Q_n(\theta) \mid \theta \in \Theta\}$ and*

$$\limsup_\nu(\arg\min\{Q_n(\theta;\epsilon_\nu) \mid \theta \in \Theta\}) \subseteq \arg\min\{Q_n(\theta) \mid \theta \in \Theta\}. \tag{2.8}$$

If $z_\nu > 0$ is a monotone decreasing sequence with $z_\nu \to 0$, then we also have

$$\limsup_\nu (z_\nu\text{-}\arg\min\{Q_n(\theta;\epsilon_\nu) \mid \theta \in \Theta\}) \subseteq \arg\min\{Q_n(\theta) \mid \theta \in \Theta\}. \tag{2.9}$$

Corollary 2.3. *Suppose* **A1**,**A2** *and* **R1** *hold. Given any $d > 0$, there exists $E, Z > 0$ such that if $\hat{\theta}_n \in z\text{-}\arg\min\{Q_n(\theta;\epsilon) \mid \theta \in \Theta\}$ for any $0 \le z \le Z$ and $0 \le \epsilon \le E$, then $\text{dist}(\hat{\theta}_n, \arg\min\{Q_n(\theta) \mid \theta \in \Theta\}) < d$.*

Note that Corollary 2.3 is simply a restatement of Proposition 2.8. These results say that approximately solving R–IOP–SAA is equivalent to approximately solving IOP–SAA.

2.4.3 Statistical consistency

In order to prove statistical consistency, we will need to impose an additional regularity condition that ensures expectations of corresponding random variables exist.

R2. The set Θ is closed and bounded, and $\mathbb{E}(y^2) < +\infty$.

This regularity assumption ensures that the law of large numbers (Wald, 1949, Jennrich, 1969, van der Vaart, 2000) holds in our setting. The above expectation condition holds in many situations, including when $\mathcal{Y}$ is bounded or when y has a sub-exponential distribution (Vershynin, 2012). This allows for settings where **IC** holds with measurement noise that is Gaussian, Bernoulli, bounded support, Laplacian, Exponential, and many other distributions.

Our first statistical consistency result is that solving R–IOP–SAA is risk consistent. To state the result, we must formally define the regularized version of the inverse optimization problem. The regularized risk is

$$Q(\theta;\epsilon) = \mathbb{E}\Big(\min_{x \in \mathcal{S}(u,\theta;\epsilon)} \|y - x\|^2\Big), \tag{R–RISK}$$

where $\mathcal{S}(u,\theta;\epsilon) = \{x \in \mathbb{R}^d : f(x,u,\theta) \le V(u,\theta) + \epsilon,\ g(x,u,\theta) \le \epsilon\}$ is the set of ϵ-optimal solutions to FOP. For given $\epsilon > 0$, we define the regularized inverse optimization problem to be

$$\min\{Q(\theta;\epsilon) \mid \theta \in \Theta\}. \tag{R–IOP}$$

The first statistical consistency result specifically concerns nearly-optimal solutions of R–IOP–SAA. We say that a sequence of solutions $\hat{\theta}_n$ is nearly-optimal for R–IOP–SAA with fixed $\epsilon > 0$ in probability if for any $\delta > 0$ we have

$$\lim_{n\to\infty} \mathbb{P}\Big(\text{dist}(\hat{\theta}_n, \arg\min\{Q_n(\theta;\epsilon) \mid \theta \in \Theta\}) > \delta\Big) = 0. \tag{2.10}$$

Theorem 2.2. *Suppose* **A1**,**A2** *and* **R1**,**R2** *hold. Given any fixed $\epsilon > 0$, if $\hat{\theta}_n$ is nearly-optimal for* R–IOP–SAA *in probability, then we have $Q(\hat{\theta}_n;\epsilon) \xrightarrow{p} \min\{Q(\theta;\epsilon) \mid \theta \in \Theta\}$.*

This result says that if we choose any $\epsilon > 0$ and solve R–IOP–SAA to generate an estimate $\hat{\theta}_n$, then the predictions given by the ϵ-optimal solutions to FOP (i.e., $\mathcal{S}(u,\hat{\theta}_n;\epsilon)$) are asymptotically the best possible set of predictions when the error of predictions is measured using

R-RISK. A stronger risk consistency result is not possible in the general setting because $Q(\theta)$ is typically discontinuous, and so the above result can be interpreted as a weak consistency result.

A stronger risk consistency result is possible in the case where $f(x, u, \theta)$ is strictly convex. We say that a sequence of solutions $\hat{\theta}_n$ is nearly-optimal for IOP–SAA in probability if for any $\delta > 0$ we have

$$\lim_{n\to\infty} \mathbb{P}\Big(\text{dist}(\hat{\theta}_n, \arg\min\{Q_n(\theta) \mid \theta \in \Theta\}) > \delta\Big) = 0. \tag{2.11}$$

Theorem 2.3. *Suppose* **A1**,**A2** *and* **R1**,**R2** *hold. If $f(x, u, \theta)$ is strictly convex in x (for fixed $u \in \mathcal{U}$ and $\theta \in \Theta$) and $\hat{\theta}_n$ is nearly-optimal for* IOP–SAA *in probability, then we have $Q(\hat{\theta}_n) \xrightarrow{p} \min\{Q(\theta) \mid \theta \in \Theta\}$.*

This result says that when FOP is a strictly convex optimization problem and we solve IOP–SAA to generate an estimate $\hat{\theta}_n$, then the predictions given by the solutions to FOP (i.e., $\mathcal{S}(u, \hat{\theta}_n)$) are asymptotically the best possible set of predictions when the error of predictions is measured using RISK. The reason it is possible to show risk consistency in this case is that $Q(\theta)$ will be continuous in this setting. Our final statistical consistency result is that solving IOP–SAA is estimation consistent when **IC** holds.

Theorem 2.4. *Suppose* **A1**,**A2** *and* **R1**,**R2** *and* **IC** *hold. If $\hat{\theta}_n$ is nearly-optimal for* IOP–SAA *in probability, then we have $\hat{\theta}_n \xrightarrow{p} \theta_0$.*

2.4.4 Identifiability in inverse optimization

Estimation consistency in any statistical setting (including inverse optimization with noisy data) requires that an identifiability condition holds, and such identifiability conditions can be stated under a variety of different mathematical formulations (Wald, 1949, Jennrich, 1969, Bartlett and Mendelson, 2002, Greenshtein and Ritov, 2004, Bickel and Doksum, 2006, Chatterjee, 2014, Aswani, 2015). The intuition for these different formulations is the same: Essentially, an identifiability condition states that the output of the model is different for two distinct sets of model parameters. It is important to note that identifiability is a statistical property of the model and the error metric used. Consequently, it is possible for an estimator to be statistically inconsistent, even when an identifiability condition holds (see for instance Proposition proposition:estincon). In the context of inverse optimization with noisy data, we define an identifiability condition **IC**.

Showing that **IC** holds is complicated by the presence of constraints in FOP. To illustrate this, consider two related instances of FOP with $x \in \mathbb{R}$ and $\theta \in [0, 2]$. The first $\min(x - \theta)^2$ is FOP-I, and the second $\min\{(x - \theta)^2 \mid x \leq 1\}$ is FOP-II. Since these two problems are strictly convex, their minimizers are unique. Next, suppose we would like to estimate θ given a (noiseless) measurement y_i of the minimizer. Observe that FOP-I is identifiable because we must have $\theta = y_i$. However, FOP-II is not identifiable because if $y_i = 1$, then we may have any $\theta \in [1, 2]$. Thus, the constraint $x \leq 1$ renders FOP-II unidentifiable, and precludes the possibility of **IC** holding for FOP-II.

[2]Note that this notion of near-optimality is defined with respect to IOP–SAA, whereas the definition of near-optimality given in (15) is with respect to the regularized formulation R–IOP–SAA.

Though FOP-II is not identifiable, a related problem is identifiable because of external inputs. In particular, consider an FOP-III with $x \in \mathbb{R}$ and $\theta \in [0, 2]$ that is given by $\min\{(x - \theta - u)^2 \mid x \leq 1\}$. This problem is strictly convex, and so its minimizer is unique for each fixed value of u. In fact, the minimizer is given by $y_i = \min\{(\theta + u_i), 1\}$. And so a sufficient condition for identifiability of FOP-III is if $\mathbb{P}(u_i \leq -1) > 0$. For instance, if $u_i = -1$ then $y_i = \theta - 1$ and so θ is uniquely determined by y_i. The presence of the input parameter u ensures identifiability of FOP-III.

2.5 Solution methods

Solving IOP–SAA with $Q_n(\theta)$ as formulated in DB–RISK–SAA is still difficult because it is a nonconvex problem even under **A1**,**A2**,**R1**. We will propose two approaches to solving this problem. The first is an enumeration algorithm that is applicable to situations where p is modest (i.e., the $\theta \in \mathbb{R}^p$ parameter has between 1 to 5 dimensions). The second approach we describe is a semiparametric algorithm, and it can be used in cases where $\theta \in \mathbb{R}^p$ is higher-dimensional and the noise term w has a specific distribution. For both algorithms, we will prove that the estimates computed by these methods satisfy the conditions required for statistical consistency.

The difference in the two algorithms is how they trade-off computational and statistical performance. The enumeration algorithm requires computation exponential in p, while the semiparametric algorithm needs computation polynomial p computation. But the statistical performance of the methods will be the opposite. The estimates and risk of the enumeration algorithm are anticipated to converge at faster rate (with respect to the number of data points) than those of the semiparametric algorithm. The reason is that the semiparametric algorithm makes use of a nonparametric step (via the L2NW estimator), which is well-known to generally converge at a slower rate than a fully parametric approach. Precisely characterizing the statistical convergence rates of the two algorithms is left open for future work.

Though the enumeration algorithm needs exponential in p computation, it is still practical for many real-world problems. Many principal-agent problems (e.g. Zhang and Zenios (2008), Crama et al. (2008)) use models where the parameter set is modest in dimensionality (i.e., utility functions with 2 or 3 *type* parameters). We demonstrate the practicality of the enumeration algorithm in Section 5 through an energy-related example using real data.

2.5.1 Enumeration algorithm

The main idea of this algorithm is that computing $Q_n(\theta)$ and $Q_n(\theta; \epsilon)$ for fixed values of θ can be done in polynomial time since DB–RISK–SAA and R–DB–RISK–SAA are convex optimization problems by Propositions 2.4 and 2.6, respectively. This approach enumerates over different fixed values of θ and solves a series of polynomial time problems. However, Θ is a continuous set since because it is convex by **A2**. To enable enumeration, we discretize Θ using a δ-net of Θ, which we will call $\mathcal{T}(\delta)$. (Here, we define this to mean that $\mathcal{T}(\delta)$ is a finite set such that $\max_{\theta \in \Theta} \min_{t \in \mathcal{T}(\delta)} \|t - \theta\| \leq \delta$.) We then compute $Q_n(\theta; \epsilon)$ for all $\theta \in \mathcal{T}(\delta)$. And our approximate solution is finally given by $\hat{\theta}_n = \arg\min\{Q_n(\theta; \epsilon) \mid \theta \in \mathcal{T}(\delta)\}$.

Algorithm 1: Enumeration Algorithm

Data: fixed $\delta > 0$ and $\epsilon \geq 0$

Result: estimate $\hat{\theta}_n$

set $\mathcal{T}(\delta)$ to be δ-net of Θ;

foreach $\theta \in \mathcal{T}(\delta)$ **do**

  compute $Q_n(\theta; \epsilon)$ by solving R–DB–RISK–SAA;

set $\hat{\theta}_n \in \arg\min\{Q_n(\theta; \epsilon) \mid \theta \in \mathcal{T}(\delta)\}$;

This approach requires continuity of $Q_n(\theta; \epsilon)$ because otherwise performing an enumeration via the δ-net $\mathcal{T}(\delta)$ may not get sufficiently close to the optimal value. However, $Q_n(\theta; \epsilon)$ is only guaranteed to be continuous at $\epsilon = 0$ when $f(x, u, \theta)$ is strictly convex for fixed u, θ by Proposition 2.2 and since $Q_n(\theta; 0) = Q_n(\theta)$ by definition. Hence, we require $\epsilon > 0$ for cases where $f(x, u, \theta)$ is *not* strictly convex to ensure continuity of $Q_n(\theta; \epsilon)$ by Proposition 2.7. Of course, when $f(x, u, \theta)$ is strictly convex we can set $\epsilon = 0$ and maintain continuity of $Q_n(\theta; \epsilon)$.

This approach is formally presented in Algorithm 1. Importantly, it can be shown that this enumeration algorithm generates nearly-optimal solutions of IOP–SAA and R–IOP–SAA. This means the solutions computed by this algorithm satisfy the conditions in Theorems 2, 3, and 4 that are needed for statistical consistency. In practice, ϵ is chosen to be $\epsilon = 0$ when FOP is strictly convex, and otherwise ϵ is chosen to be a small positive value that controls the desired precision of the resulting estimate. An appropriate approach to choose ϵ and δ is to use cross-validation, which is a standard data-driven approach from statistics for choosing such parameters (Hastie et al., 2009).

Theorem 2.5. *Suppose* **A1**,**A2** *and* **R1** *hold. Given any $d > 0$, there exists $E, \Delta > 0$ such that if $\hat{\theta}_n$ is computed using the enumeration algorithm (i.e., Algorithm 1) for any $0 < \epsilon \leq E$ and $0 < \delta \leq \Delta$, then* $\text{dist}(\hat{\theta}_n, \arg\min\{Q_n(\theta) \mid \theta \in \Theta\}) < d$.

Theorem 2.5 states that the estimate obtained using the enumeration algorithm will be at most a distance of d from the set of optimal solutions to IOP-SAA. It immediately follows that for small d, the solution of the enumeration algorithm will retain the desirable statistical properties of the solutions to IOP-SAA. As mentioned above, in the special case where FOP is a strictly convex optimization problem we can simplify the algorithm by setting $\epsilon = 0$. We have a corresponding result about the correctness of the algorithm in this case.

Theorem 2.6. *Suppose* **A1**,**A2** *and* **R1** *hold. If $f(x, u, \theta)$ is strictly convex in x (for fixed $u \in \mathcal{U}$ and $\theta \in \Theta$), then given any $d > 0$ there exists $\Delta > 0$ such that if $\hat{\theta}_n$ is computed using the enumeration algorithm for $\epsilon = 0$ and any $0 < \delta \leq \Delta$, then* $\text{dist}(\hat{\theta}_n, \arg\min\{Q_n(\theta) \mid \theta \in \Theta\}) < d$.

2.5.2 Semiparametric algorithm

Our second approach to solving IOP–SAA is a semiparametric approach. We will need to make an additional assumption about the structure of the problem, as well as impose two more regularity conditions, in order to be able use this approach. We begin with the addi-

tional assumption.

A3. The constraint function $g(x, u, \theta)$ is independent of θ, meaning it can be written as $g(x, u, \theta) = g_0(x, u)$. The objective function $f(x, u, \theta)$ is affine in θ, meaning it can be written as

$$f(x, u, \theta) = f_0(x, u) + \sum_{j=1}^{p} \theta_j f_j(x, u). \tag{2.12}$$

Independence of the constraint g from θ is required because the semiparametric approach relies on fully knowing the feasible region of the forward problem. We note that this is not a particularly strong assumption, since in utility estimation settings one would expect the unknown parameters to appear in the objective function of the forward problem. Keshavarz et al. (2011) and Bertsimas et al. (2015) also assume that the feasible region of the forward problem is independent of the unknown parameters. The second part of **A3** ensures that the Lagrangian dual function $h(\lambda, u, \theta)$ is concave in θ. This will enable efficient computation in our semiparametric approach. Next, we describe the two additional regularity conditions. The first is

R3. The objective function $f(x, u, \theta)$ is strictly convex in x (for fixed $u \in \mathcal{U}$ and $\theta \in \Theta$) and twice continuously differentiable in x, u, θ, and the constraints $g(x, u, \theta)$ are continuously differentiable in x, u, θ.

Condition **R3** ensures smoothness in the objective function and constraints. The reason we also include a strict convexity assumption is that it acts as a regularity condition: Strictly speaking, we require uniqueness of solutions to FOP (which is needed for the de-noising step in our semiparametric algorithm) and a second-order growth condition

$$f(x, u, \theta) \geq V(u, \theta) + c \cdot [\text{dist}(x, \mathcal{S}(u, \theta))]^2, \tag{2.13}$$

for some $c > 0$ and all $x \in \Phi(u, \theta)$ (which ensures Hölder continuity of the solution set $\mathcal{S}(u, \theta)$ with degree $1/2$ (Bonnans and Shapiro, 2000)). Unfortunately, this growth condition can be difficult to directly check even though it has been completely characterized for convex optimization problems (Bonnans and Ioffe, 1995b). Fortunately, strict convexity with Slater's constraint qualification (which holds under **R1**) implies both uniqueness of solutions to FOP and this second-order growth condition (Bonnans and Ioffe, 1995a). Hence **R3** is sufficient for proving statistical convergence using our algorithm. We also note that our results could be extended to the case where the problem satisfies the first-order growth condition

$$f(x, u, \theta) \geq V(u, \theta) + c \cdot \text{dist}(x, \mathcal{S}(u, \theta)), \tag{2.14}$$

for some $c > 0$ and all $x \in \Phi(u, \theta)$. Under this alternate growth condition, the solution set is Hölder continuous with degree 1 (instead of $1/2$). This affects the bound expression in Proposition 2.9 slightly, but otherwise does not qualitatively change our results.

R4. The noise random variable w has a sub-exponential distribution, meaning there exists

$c > 0$ such that $\mathbb{P}(|w| > t) \leq \exp(1 - t/c)$. Also, the probability density function $\mu(u)$ of u is continuously differentiable and is bounded from zero (i.e., $\min_{u \in \mathcal{U}} \mu(u) > 0$).

This regularity condition ensures the distribution of the random variables w, u are not extreme. Most commonly used heavy-tailed noise distributions are sub-exponential distributions, and so **R4** is satisfied by Gaussian, Bernoulli, bounded support, Laplacian, Exponential, and many other distributions (Vershynin, 2012). Also, the regularity condition on $\mu(u)$ implies $\mathcal{U}$ is bounded.

The idea behind the semiparametric approach is the observation that R–DB–RISK–SAA is convex in θ for fixed x when **A3** holds. However, because the y_i are measured with noise, we cannot simply make the substitution $x_i = y_i$. To overcome this difficulty, we first de-noise the y_i using a nonparametric estimator. Specifically, we define the ℓ_2-regularized Nadaraya-Watson (L2NW) estimator (Aswani et al., 2013) as

$$\overline{x}_i = \frac{\gamma^{-m} \cdot \frac{1}{n} \sum_{j=1}^{n} y_j \cdot K(\frac{u_j - u_i}{\gamma})}{\sigma + \gamma^{-m} \cdot \frac{1}{n} \sum_{j=1}^{n} K(\frac{u_j - u_i}{\gamma})}, \tag{2.15}$$

where $\gamma > 0$ is the *bandwidth* parameter, $\sigma > 0$ is the *ℓ_2-regularization* parameter, and $K : \mathbb{R}^m \to \mathbb{R}$ is a *kernel function* that satisfies the following properties (i) $K(u) \geq 0$, (ii) $K(u) = 0$ for $\|u\| > 1$, (iii) $K(u) = K(-u)$, and (iv) $\int K(u)du = 1$. A common example of a kernel function is the Epanechnikov kernel, which is defined as the function

$$K(u) = \begin{cases} \frac{3}{4} \cdot (1 - \|u\|^2), & \text{if } \|u\| \leq 1 \\ 0, & \text{otherwise} \end{cases} \tag{2.16}$$

The L2NW estimator (2.15) is computed in polynomial time, and it serves to de-noise the x_i in the manner described by the following proposition.

Proposition 2.9. *Suppose* **A1** *and* **R1**–**R4** *hold. If $\gamma = O(n^{-2/(8m+1)})$ and $\sigma = O(\gamma)$, then $\mathcal{S}(u, \theta)$ consists of a single point, and for sufficiently large n we have we have*

$$\mathbb{P}\Big(\max_{i \in [n]} \|\overline{x}_i - \mathcal{S}(u_i, \theta_0)\| > n^{-1/(18m)} \Big) \leq k_1 \exp\Big(- k_2 n^{1/4} \Big), \tag{2.17}$$

where $k_1, k_2 > 0$ are constants. In particular, this implies $\max_{i \in [n]} \|\overline{x}_i - \mathcal{S}(u_i, \theta_0)\| \xrightarrow{p} 0$.

Before we present our algorithm, we need one more result that provides additional understanding for the semiparametric approach. Consider the following optimization problem

$$\min_{\theta} \max_{\epsilon \geq 0} \{ Q_n(\theta; \epsilon) \mid \theta \in \Theta \}, \tag{ROBUST–IOP–SAA}$$

Proposition 2.10. *Suppose* **A1**,**A2** *and* **R1** *hold. Then the solution sets in θ of* ROBUST–IOP–SAA *and* IOP–SAA *are equivalent, and the optimal value of* ROBUST–IOP–SAA *occurs at $\epsilon = 0$.*

Proof. Let $\mathcal{C}_n(\theta, \epsilon)$ be the feasible set of R–DB–RISK–SAA. As shown in the proof for

Algorithm 2: Semiparametric Algorithm

Data: fixed $\gamma > 0$ and $\sigma > 0$
Result: estimate $\hat{\theta}_n$
foreach $i \in [n]$ **do**
  compute $\overline{x}_i$ using using (2.15);
compute $\hat{\theta}_n$ using SP–IOP–RISK–SAA;

Proposition 2.8, the feasible set satisfies

$$\mathcal{C}_n(\theta, 0) \subseteq \mathcal{C}_n(\theta, \epsilon), \tag{2.18}$$

for all $\epsilon \geq 0$. As a result, we must have that $Q_n(\theta; 0) \geq Q_n(\theta; \epsilon)$ for all $\epsilon \geq 0$. This means that $\max_{\epsilon \geq 0} Q_n(\theta; \epsilon) = Q_n(\theta; 0)$. The result holds because $Q_n(\theta; 0) = Q_n(\theta)$ by definition. □

Given the above relationship that the optimal value of ROBUST–IOP–SAA occurs at $\epsilon = 0$, we propose to solve the inverse optimization problem using the following formulation:

$$\begin{aligned} \hat{\theta}_n \in \arg\min \; & \frac{1}{n} \sum_{i=1}^{n} \epsilon_i & & \\ \text{s.t.} \; & f(\overline{x}_i, u_i, \theta) - h(\lambda_i, u_i, \theta) \leq \epsilon_i, & \forall i \in [n] & \\ & \lambda_i \geq 0, & \forall i \in [n] & \end{aligned} \qquad \text{SP–IOP–RISK–SAA}$$

where the $\overline{x}_i$ are as defined in (2.15). This is a convex optimization problem.

Proposition 2.11. *Suppose* **A1**–**A3** *and* **R1** *hold. Then* SP–IOP–RISK–SAA *is a convex optimization problem.*

We now have the elements to construct our semiparametric algorithm, which is a two-step approach. In the first step, we de-noise the y_i data using the L2NW estimator given in (2.15). This de-noising step produces an estimate of the true underlying optimal solution, which we represent by $\bar{x}_i$. While the estimates $\bar{x}_i$ are asymptotically (in n) optimal (cf. Proposition 2.9), they may be suboptimal at finite n. Therefore, in the second step, we solve SP–IOP–RISK–SAA, which produces a parameter estimate $\hat{\theta}_n$ that minimizes the suboptimality of $\bar{x}_i$. This approach maintains statistical consistency because the $\bar{x}_i$ are denoised, and it is formally presented in Algorithm 2. Importantly, it can be shown that this semiparametric algorithm generates nearly-optimal solutions of IOP–SAA. This means the solutions computed by this algorithm satisfy the conditions in Theorems 2, 3, and 4 that are needed for statistical consistency. In practice, the values of σ and γ can be can be chosen using cross-validation, which is a standard data-driven approach from statistics for choosing such parameters (Hastie et al., 2009).

Theorem 2.7. *Suppose* **A1**–**A3** *and* **R1**–**R4** *and* **IC** *hold. If* $\sigma = O(n^{-2/(8m+1)})$, $\lambda = O(\sigma)$, *and* $\hat{\theta}_n$ *is computed using the semiparametric algorithm (i.e., Algorithm 2) ; then* $\hat{\theta}_n$ *is nearly-optimal for* IOP–SAA *in probability.*

Theorem 2.7 states that the semiparametric algorithm produces estimates that are statistically consistent under the appropriate conditions. In the next section, we present several

numerical experiments which validate our theoretical results as well as the performance of the enumeration and semiparametric algorithms.

2.6 Numerical results

We present numerical results that demonstrate the statistical consistency of our algorithms for inverse optimization with noisy data, and the results show our algorithms perform competitively against KKA (Keshavarz et al., 2011) and VIA (Bertsimas et al., 2015). We begin by conducting two types of tests using synthetic data. The first type is where the model is kept fixed and the number of data points increases, and the purpose is to demonstrate either estimation consistency or risk consistency of our algorithms. The second type is where the number of data points is kept fixed and the number of the parameters in the model increases, and the purpose is to demonstrate the feasibility of using our algorithms on large-scale problems. We then apply our framework to a real data set, where we estimate a utility function that describes the tradeoff made between occupant comfort and energy consumption when setting a thermostat temperature setpoint for air-conditioning.

2.6.1 Enumeration algorithm with synthetic data

In the first experiments, we generate data using a given FOP and then use the same set of equations in SAA-IOP. In other words, the first set of experiments are situations where the model whose parameters are being identified exactly match the model that generates the data. As a result, this setting consists of situations where **IC** is satisfied. The first example is where: (i) FOP-A is $\min\{(\theta + u) \cdot x \mid x \in [-1, 1]\}$, (ii) u has a uniform distribution with support $[-1, 1]$, (iii) the measurement noise w has a normal distribution with zero mean and unit variance, (iv) the data is generated with $\theta_0 = 1$, and (v) the enumeration algorithm (i.e., Algorithm 1) was applied with $\epsilon = 0.001$, $\delta = 0.01$, and $\Theta = [-1, 1]$. The second example is where: (i) FOP-B is $\min\{x^2 - (\theta + u) \cdot x \mid x \in [0, 1]\}$, (ii) u has a uniform distribution with support $[0, 2]$, (iii) the measurement noise w has a normal distribution with zero mean and unit variance, (iv) the data is generated with $\theta_0 = \frac{1}{2}$, and (v) the enumeration algorithm (i.e., Algorithm 1) was applied with $\epsilon = 0$, $\delta = 0.01$, and $\Theta = [0, 2]$.

The results averaged over 100 repetitions of sampling $n \in \{10, 30, 50, 100, 300, 500, 1000\}$ data points and then estimating the parameter θ are summarized in Table 2.1. We label the enumeration algorithm (i.e., Algorithm 1) as ENA in the table. These results display estimation consistency of the enumeration algorithm since estimation error is decreasing to zero. To further illustrate estimation consistency, we conducted an experiment with the two examples above where the data was generated with a θ_0 that was randomly chosen from a uniform distribution with support $[-1, 1]$ and $[0, 2]$ for the first and second examples, respectively. A plot comparing the estimates $\hat{\theta}_n$ to the true parameter θ_0 for the first situation when $n = 1,000$ is shown in Figure 2.1, and a plot comparing the estimates $\hat{\theta}_n$ to the true parameter θ_0 for the second situation when $n = 10,000$ is shown in Figure 2.2. Consistent estimates should line up along the diagonal, and hence these plots demonstrate the estimation consistency (inconsistency) of the enumeration algorithm (KKA and VIA). Recall from the discussion in Section 2 that KKA and VIA are inconsistent because they minimize an incorrect measure of error, and this discrepancy is most significant for points where the optimal solution

Table 2.1: Estimation error $|\hat{\theta}_n - \theta_0|$ of enumeration algorithm (ENA) and benchmark algorithms (KKA and VIA) on two synthetic instances (n increasing, $p = 1$).

	n	10	30	50	100	300	500	1000
Data: FOP-A Model: FOP-A	ENA	0.2616	0.0926	0.0380	0.0211	0.0055	0.0030	0.0009
	KKA	0.8686	0.8293	0.8182	0.8257	0.8130	0.8231	0.8170
	VIA	0.5552	0.4976	0.4829	0.4887	0.4807	0.4846	0.4780
Data: FOP-B Model: FOP-B	ENA	0.4577	0.2481	0.1510	0.0501	0.0222	0.0123	0.0063
	KKA	0.5065	0.2281	0.1595	0.0751	0.0398	0.0342	0.0238
	VIA	0.9488	0.7051	0.6344	0.4284	0.3145	0.3810	0.2962

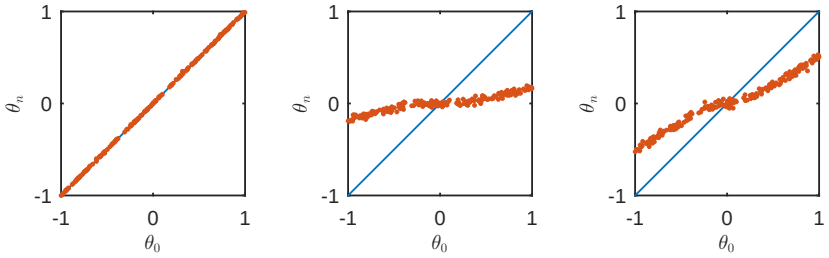

(a) ENA (b) KKA (c) VIA

Figure 2.1: Scatter plot comparing estimated parameter $\hat{\theta}_n$ versus true parameter θ_0 as computed by ENA, KKA and VIA algorithms at $n = 1,000$, when the data and model are both FOP-A.

of FOP lies on the boundary of the feasible set. KKA and VIA perform more poorly for FOP-A than for FOP-B because FOP-A is a linear program, which has almost all of its optimal solutions on the boundary of the feasible set, whereas FOP-B is a quadratic program, which has more optimal solutions within the strict interior of the feasible set.

In the second set of experiments, we generate data using a given model that is different than the FOP used to formulate SAA-IOP. In other words, this set of experiments are situations where the model whose parameters are being identified does not match the model that generates the data. As a result, this setting consists of situations where **IC** is *not* satisfied. The first example is where: (i) the data is generated by FOP-C which is $\min\{\frac{3}{2} \cdot x^2 - (1+u) \cdot x \mid x \in [0,1]\}$, (ii) the model estimated by IOP-SAA is FOP-B, (iii) u has a uniform distribution with support $[0,5]$, (iv) the measurement noise w has a normal distribution with zero mean and unit variance, and (v) the enumeration algorithm (i.e., Algorithm 1) was applied with $\epsilon = 0$, $\delta = 0.01$, and $\Theta = [0,2]$. The second example is where: (i) the data is generated by the statistical model SQR-1 given by $y_i = \min\{\max\{\sqrt{u_i}, 0\}, 1\} + w_i$, (ii) the model estimated by IOP-SAA is FOP-B, (iii) u has a uniform distribution with support $[0,5]$, (iv) the measurement noise w has a normal distribution with zero mean and unit variance, and (v) the enumeration algorithm (i.e., Algorithm 1) was applied with $\epsilon = 0$, $\delta = 0.01$, and $\Theta = [0,2]$.

The results averaged over 100 repetitions of sampling $n \in \{10, 30, 50, 100, 300, 500, 1000\}$ data points and then estimating the parameter θ are summarized in Table 2.2, and these

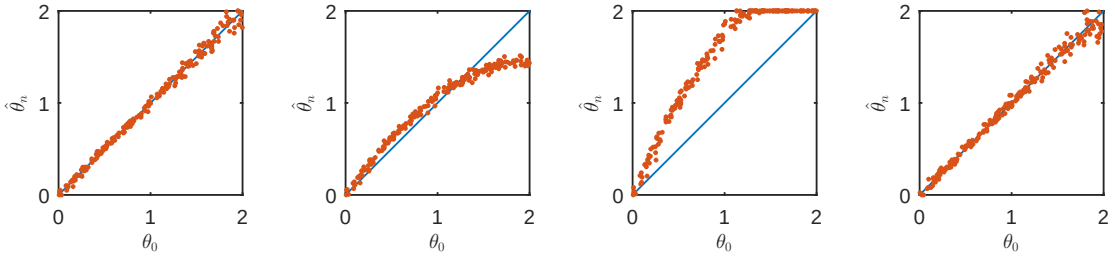

(a) ENA (b) KKA (c) VIA (d) SPA

Figure 2.2: Scatter plot comparing estimated parameter $\hat{\theta}_n$ versus true parameter θ_0 as computed by ENA, KKA, VIA and SPA algorithms at $n = 10,000$ when the data and model are both FOP-B.

Table 2.2: Normalized prediction error $Q(\hat{\theta}_n) - \text{var}(w)$ of enumeration algorithm (ENA) and benchmark algorithms (KKA and VIA) on two synthetic instances (n increasing, $p = 1$).

	n	10	30	50	100	300	500	1000
Data: FOP-C Model: FOP-B	ENA	0.0216	0.0184	0.0162	0.0150	0.0065	0.0046	0.0017
	KKA	0.0168	0.0124	0.0128	0.0151	0.0150	0.0150	0.0132
	VIA	0.0249	0.0185	0.0196	0.0149	0.0089	0.0072	0.0042
Data: SQR-1 Model: FOP-B	ENA	0.0294	0.0217	0.0152	0.0110	0.0073	0.0041	0.0024
	KKA	0.0394	0.0389	0.0398	0.0440	0.0504	0.0525	0.0518
	VIA	0.0343	0.0287	0.0243	0.0187	0.0122	0.0084	0.0072

results are normalized by subtracting var(w). The reason for this normalization is that the prediction error $\mathbb{E}((y - \xi(u))^2))$ of the prediction $\xi(u)$ of the true model (either FOP-C or SQR-M, respectively) is var(w) because $y = \xi(u) + w$ here. The enumeration algorithm has lower prediction error because it is risk consistent, whereas KKA and VIA are not risk consistent.

2.6.2 Semiparametric algorithm with synthetic data

We now examine the performance of the semiparametric algorithm (Algorithm 2) in four sets of experiments. In the first set of experiments, we generate data using a given FOP and then use the same equations in SAA-IOP. These experiments are situations where the model whose parameters are being identified exactly matches the model that generates the data. As a result, this setting consists of situations where **IC** is satisfied. We consider three different formulations for FOP. The first example is where: (i) FOP-D is $\min\{x'x - (\theta + u)'x \mid x \in [0,1]^p\}$, (ii) u has a uniform distribution with support $[0,2]^p$, (iii) the measurement noise w has a jointly Gaussian distribution with zero mean and identity covariance, (iv) the data is generated with $p = 10$ and $\theta_0 \in \mathbb{R}^p$ such that $\theta_{0k} = \frac{1}{2}$ for all $k \in [p]$, and (v) the semiparametric algorithm (i.e., Algorithm 2) was applied with γ, σ chosen using cross-validation (Hastie et al., 2009) and $\Theta = [0,2]$. The second example is where: (i) FOP-E

is

$$\min\left\{-\sum_{k=1}^{p}\theta_k\cdot\log(x_k+u_k)-\log(x_{p+1}+u_{p+1})\mid x_k\geq 0,\ \sum_{k=1}^{p+1}x_k=1\right\},\tag{2.19}$$

(ii) u has a uniform distribution with support $[1,2]^{p+1}$, (iii) the measurement noise w has a jointly Gaussian distribution with zero mean and identity covariance, (iv) the data is generated with $p=10$ and $\theta_0\in\mathbb{R}^p$ such that $\theta_{0k}=1$ for all $k\in[p]$, and (v) a modified version of the seimparametric algorithm (i.e., Algorithm 2) was applied with γ,σ chosen using cross-validation and $\Theta=[\frac{1}{2},2]$. The modification to Algorithm 2 is that we calculate $\tilde{x}_i=\min_x\{\|\overline{x}_i-x\|\mid x_k\geq 0\}$ and then compute $\hat{\theta}_n$ using SP–IOP–RISK–SAA with the $\tilde{x}_i$ replacing the $\overline{x}_i$. The $\tilde{x}_i$ are the projection of the $\overline{x}_i$ onto the nonnegative orthant, and it turns out this projection does not affect our theoretical results. In particular, a short proof using the continuous mapping theorem (van der Vaart, 2000) and the boundedness of the feasible set in **R1** gives that $\max_{i\in[n]}\|\tilde{x}_i-\mathcal{S}(u_i,\theta_0)\|\xrightarrow{p}0$. The projection is needed for this particular example because otherwise the inverse formulation would contain logarithms of negative numbers, which are complex-valued. More generally, a projection of $\overline{x}_i$ onto the feasible set of FOP will not affect our theoretical results, and can be added as a step in our semiparametric algorithm.

In Table 2.3, we present estimation results for the first and second examples, averaged over 100 repetitions for each value of $n\in\{10,30,50,100,300,500,1000\}$. We label the semiparametric algorithm (i.e., Algorithm 2) as SPA in the table. These results display estimation consistency of the semiparametric algorithm since it has lower estimation error as the data increases. To further illustrate estimation consistency, we conducted an experiment with the two situations above where the data was generated with $p=1$ and a θ_0 that was randomly chosen from a uniform distribution with support $[0,1]$ and $[\frac{1}{2},2]$ for the first and second situations, respectively. A plot comparing the estimates $\hat{\theta}_n$ to the true parameter θ_0 for the first situation when $n=1,000$ is shown in Figure 2.2, and a plot comparing the estimates $\hat{\theta}_n$ to the true parameter θ_0 for the second situation when $n=1,000$ is shown in Figure 2.3. Consistent estimates should line up along the diagonal, and hence these plots demonstrate the estimation consistency (inconsistency) of the semiparametric algorithm (KKA and VIA). It is worth comparing the results of the semiparametric and enumeration algorithms. As mentioned above, the semiparametric algorithm will generally have higher estimation error than the enumeration algorithm – this can be observed in these plots because the semiparametric algorithm estimates have a larger variation about the diagonal than the estimates of the enumeration algorithm.

In the second set of experiments, we generate data using a given model that is different than the FOP used to formulate SAA-IOP. In other words, this set of experiments are situations where the model whose parameters are being identified does not match the model that generates the data. As a result, this setting consists of situations where **IC** is *not* satisfied. The first setting is where: (i) the data is generated by FOP-C which is $\min\{\frac{3}{2}\cdot x'x-(1+u)'x\mid x\in[0,1]^{10}\}$, (ii) the model estimated by IOP-SAA is FOP-D with $p=10$, (iii) u has a uniform distribution with support $[0,5]^{10}$, (iv) the measurement noise w has a jointly Gaussian distribution with zero mean and identity covariance, and (v) the semiparametric algorithm (i.e., Algorithm 2) was applied with γ,σ chosen using cross-validation (Hastie et al., 2009) and $\Theta=[0,2]$. The second setting is where: (i) the data is generated

Table 2.3: Estimation error $\|\hat{\theta}_n - \theta_0\|$ of semiparametric algorithm (SPA) and benchmark algorithms (KKA and VIA) on two synthetic instances (n increasing, $p = 10$).

	n	10	30	50	100	300	500	1000
Data: FOP-D Model: FOP-D	SPA	2.4618	1.7025	1.2543	0.8535	0.4754	0.3750	0.2573
	KKA	2.2569	1.5513	1.2229	0.9281	0.6107	0.5435	0.4447
	VIA	3.3829	3.2603	3.1937	3.1501	3.0292	3.0324	2.9208
Data: FOP-E Model: FOP-E	SPA	0.9189	0.7982	0.7500	0.7487	0.6639	0.6070	0.5783
	KKA	1.6687	1.5850	1.5813	1.5865	1.5828	1.5806	1.5811
	VIA	1.9299	1.6781	1.6826	1.6132	1.6001	1.5973	1.5843

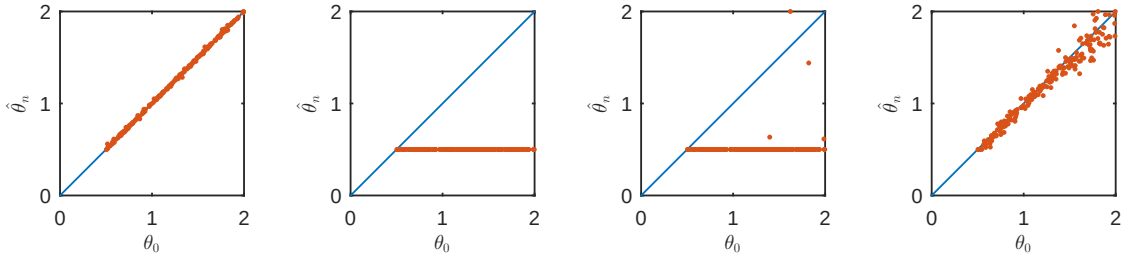

(a) ENA (b) KKA (c) VIA (d) SPA

Figure 2.3: Scatter plot comparing estimated parameter $\hat{\theta}_n$ versus true parameter θ_0 as computed by different algorithms at $n = 1,000$ when the data and model are both FOP-E.

by the statistical model SQR-P given by $y_i = \min\{\max\{\sqrt{u_i}, 0\}, 1\} + w_i$, (ii) the model estimated by IOP-SAA is FOP-D with $p = 10$, (iii) u has a uniform distribution with support $[0, 5]^{10}$, (iv) the measurement noise w has a jointly Gaussian distribution with zero mean and identity covariance, and (v) the semiparametric algorithm (i.e., Algorithm 2) was applied with γ, σ chosen using cross-validation (Hastie et al., 2009) and $\Theta = [0, 2]$. The results averaged over 100 repetitions of sampling $n \in \{10, 30, 50, 100, 300, 500, 1000\}$ data points and then estimating the parameter θ are summarized in Table 2.4, and these results are normalized by subtracting $\mathbb{E}(w'w)$. The reason for this normalization is that the prediction error $\mathbb{E}(\|y - \xi(u)\|^2))$ of the prediction $\xi(u)$ of the true model (either FOP-C or SQR-M, respectively) is $\mathbb{E}(w'w)$ because $y = \xi(u) + w$ here. The enumeration algorithm has lower prediction error because it is risk consistent, whereas KKA and VIA are not risk consistent.

In the third set of experiments, we generate data using the previous four settings. The difference in this set of experiments is that we fix $n = 1,000$ and vary $p \in \{1, 3, 5, 10, 30\}$. The results when the data/model are given by FOP-D/FOP-D and FOP-E/FOP-E, averaged over 100 repetitions and then estimating the parameter θ, are summarized in Table 2.5. These results show that the semiparametric algorithm has lower estimation error than KKA and VIA on these examples. The results when the data/model are given by FOP-C/FOP-B and SQR-M/FOP-B, averaged over 100 repetitions and then estimating the parameter θ, are summarized in Table 2.6. These results show that the semiparametric algorithm has lower prediction error than KKA and VIA on these examples.

Table 2.4: Normalized prediction error $Q(\hat{\theta}_n) - \mathbb{E}(w'w)$ of semiparametric algorithm (SPA) and benchmark algorithms (KKA and VIA) on two synthetic instances (n increasing, $p = 10$).

		n	10	30	50	100	300	500	1000
Data:	FOP-C	SPA	0.2319	0.1972	0.1744	0.1501	0.1029	0.0844	0.0529
Model:	FOP-D	KKA	0.1584	0.1308	0.1314	0.1349	0.1452	0.1497	0.1481
		VIA	0.3438	0.3407	0.3360	0.3205	0.2950	0.2816	0.2811
Data:	SQR-M	SPA	0.4180	0.3497	0.3195	0.2470	0.1572	0.0998	0.0658
Model:	FOP-D	KKA	0.3645	0.3885	0.3987	0.4537	0.5115	0.5114	0.5214
		VIA	0.3468	0.2784	0.2737	0.2524	0.2405	0.2458	0.2599

Table 2.5: Estimation error $\|\hat{\theta}_n - \theta_0\|$ of semiparametric algorithm (SPA) and benchmark algorithms (KKA and VIA) on two synthetic instances ($n = 1,000$, p increasing).

		p	1	3	5	10	30
Data:	FOP-D	SPA	0.0601	0.1464	0.1907	0.2794	0.4701
Model:	FOP-D	KKA	0.1178	0.2349	0.3038	0.4619	0.7978
		VIA	0.4943	1.2254	1.8099	2.9522	5.7737
Data:	FOP-E	SPA	0.0251	0.1258	0.2571	0.5890	0.5576
Model:	FOP-E	KKA	0.5000	0.8660	1.1174	1.5804	2.7377
		VIA	0.5000	0.8691	1.1231	1.5966	2.7628

Table 2.6: Normalized prediction error $Q(\hat{\theta}_n) - \mathbb{E}(w'w)$ of semiparametric algorithm (SPA) and benchmark algorithms (KKA and VIA) on two synthetic instances ($n = 1,000$, p increasing).

		p	1	3	5	10	30
Data:	FOP-C	SPA	0.0064	0.0171	0.0403	0.0628	0.2048
Model:	FOP-D	KKA	0.0538	0.1553	0.2619	0.5252	1.5712
		VIA	0.0078	0.0175	0.0745	0.2602	0.9654
Data:	SQR-M	SPA	0.0056	0.0194	0.0319	0.0606	0.1568
Model:	FOP-D	KKA	0.0148	0.0471	0.0761	0.1523	0.4394
		VIA	0.0055	0.0273	0.0821	0.2848	1.2896

Table 2.7: Estimation error $\|\hat{\theta}_n - \theta_0\|$ of semiparametric algorithm (SPA) and benchmark algorithms (KKA and VIA) on synthetic instance ($n = 1,000$, p increasing).

	p	5	10	20	50	100
Data: FOP-F Model: FOP-F	SPA	0.5535	0.8530	1.1522	2.0020	2.7205
	KKA	2.1753	4.6199	8.4599	12.4102	17.8112
	VIA	1.1825	1.8689	3.7320	5.9003	8.1874

Table 2.8: Normalized prediction error $Q(\hat{\theta}_n) - \mathbb{E}(w'w)$ of semiparametric algorithm (SPA) and benchmark algorithms (KKA and VIA) on synthetic instance ($n = 1,000$, p increasing).

	p	5	10	20	50	100
Data: FOP-F Model: FOP-F	SPA	0.2539	0.5117	0.9307	2.9423	5.8644
	KKA	8.2329	22.5149	60.4496	145.2210	302.1124
	VIA	2.3475	3.7495	10.8325	26.5713	48.3217

2.6.3 High-dimensional nonlinear forward problem with stochastic constraints

We now consider a setting where FOP is high dimensional, contains a logarithmic objective, and has an exponential stochastic constraint (i.e. the constraint depends on u). Specifically, we consider the following setting: (i) FOP-F is

$$\min \left\{ -\sum_{k=1}^{p} \theta_k \cdot u_k^{(1)} \cdot \log(x_k) \mid \tfrac{1}{p} \sum_{k=1}^{p} e^{x_k + u_k^{(1)}} - u_k^{(2)} \leq 0,\ x_k \geq 0 \right\}, \tag{2.20}$$

(ii) $u^{(1)}$ has a uniform distribution with support $[1,2]^p$ and $u^{(2)}$ has a uniform distribution with support $[50,100]^p$, (iii) the measurement noise w has a jointly Gaussian distribution with zero mean and identity covariance, (iv) the data is generated with $\theta_0 \in \mathbb{R}^p_+$ such that $\sum_{k=1}^{p} \theta_{0k} = p$, and (v) a modified version of the seimparametric algorithm (i.e., Algorithm 2) is applied where $\Theta = \{\theta \in \mathbb{R}^p_+ | \sum_{k=1}^{p} \theta_k = p\}$ and γ, σ is selected using cross-validation. We set $n = 1,000$ and repeat the sampling and estimation procedure 100 times for each value of $p \in \{5, 10, 20, 50, 100\}$. The average estimation and prediction errors are summarized in Tables 2.7 and 2.8, respectively, which show that the semiparametric algorithm is competitive with existing methods in this setting as well. Note that the magnitude of the errors is expected to increase with p, since we do not normalize the error for the number of parameters being estimated.

2.6.4 Empirical data: Estimating an energy-comfort utility function

We next apply our inverse optimization framework to the problem of estimating a utility function that describes the tradeoff made between occupant comfort and the amount of energy consumption when setting a thermostat temperature setpoint for air-conditioning. The data we use is collected from Sutardja Dai Hall on the Berkeley campus, which was used as part of the BRITE-S testbed in our past experiments (Aswani et al., 2012a,b,c) concerning robust learning-based optimization (Aswani et al., 2013) of heating, ventilation, and air-conditioning (HVAC) systems. Specifically, this building is equipped with a commercial

Table 2.9: Prediction error $Q(\hat{\theta}_n)$ of enumeration algorithm (ENA) and benchmark algorithms (KKA and VIA) on temperature preference dataset. (n increasing, p fixed).

	n	10	30	50	100	300	500	1000
Data: SDH-E Model: FOP-S	ENA	1.3656	1.3308	1.3255	1.3169	1.3112	1.3099	1.3090
	KKA	2.2439	2.2528	2.2508	2.2351	2.2225	2.2220	2.2200
	VIA	2.2975	2.2538	2.2472	2.2277	2.2163	2.2166	2.2138

web application (Building Robotics, 2016) that allows occupants to change the thermostat temperature setpoints in real-time, and so the setpoints are changed throughout the year by occupants in response to factors like the outside weather.

When a room is being cooled, a lower temperature setpoint requires increased energy consumption since the air-conditioner must provide more cold air; however, the purpose of air-conditioning is to improve comfort by lowering the room temperature. And so individuals must tradeoff comfort and energy consumption when choosing the setpoint. A simplified utility function model (expressed as minimization of the negative of the utility function) that captures this tradeoff is FOP-S:

$$\min_x \{\theta_1 \cdot (x - 76)^2 + (x - \theta_2 - u)^2 \mid x \in [70, 76]\}, \tag{2.21}$$

where $x \in \mathbb{R}$ is the thermostat temperature setpoint in units of degrees Fahrenheit (°F), and $u \in \mathbb{R}$ is the current outside temperaure in degrees Fahrenheit (°F). The term $(x - \theta_2 - u)^2$ indicates a preference for a temperature setpoint that is a fixed amount θ_2 above the outside temperature u (i.e., the preferred temperature is $\theta_2 + u$), and the reason for this term is that individuals prefer a higher indoor temperature as the outside temperature increases (ASHRAE, 2013). The term $\theta_1 \cdot (x-76)^2$ indicates a preference for a higher setpoint because of energy considerations, and the number 76 is used because 76°F–78°F is a relatively high setpoint temperature that is often recommended for saving energy. The parameter θ_1 quantifies the tradeoff between the preference for a higher setpoint to save energy versus the desired indoor temperature $\theta_2 + u$. Lastly, the constraints $x \in [70, 76]$ indicate observed setpoint limits.

The results averaged over 100 repetitions of sampling $n \in \{10, 30, 50, 100, 300, 500, 1000\}$ data points and then estimating the parameters θ are summarized in Table 2.9. The data set (which we label SDH-E in the table) used consists of outside temperature measurements (i.e., the u variable) and the chosen temperature set point (i.e., the x variable) of a single thermostat in Sutardja Dai Hall. In each repetition, the full data set was randomly split into a 1,000 point *training* data set and a 14,500 point *testing* data set. The n data points were randomly chosen from the training data set, and the prediction error of the estimated parameters were computed using the testing data set. To evaluate the statistical significance of the computed results, a bootstrap hypothesis test (Efron and Tibshirani, 1994a) was conducted. The computed p-value was less than 0.01, which indicates that the improved performance of the enumeration algorithm is statistically significant.

Chapter 3

Data-Driven Incentive Design via Inverse Optimization

3.1 Background and motivation

The United States spends more on healthcare than any other high-income nation in the world, both on a per capita basis and as a share of its GDP. In 2014, total healthcare spending in the U.S. exceeded 3 trillion dollars, equivalent to 17% of the U.S. GDP (OECD, 2015). Approximately 20% ($600b) of U.S. healthcare spending is through Medicare, the United States' federally-funded health insurance program for seniors and other qualifying individuals (CMS, 2016b). An additional 17% of healthcare spending is through Medicaid, which is targeted at low-income individuals and those with disabilities. Combined, government programs therefore account for 37% of all healthcare spending. Despite enormous spending on healthcare, the US underperforms its international peers on many indicators of health quality (Starfield, 2000, McGlynn et al., 2003), which suggests that high costs may be a consequence of *inefficiencies* in the healthcare system (Berwick and Hackbarth, 2012, Garber and Skinner, 2008, Wennberg et al., 2002). Moreover, although costs are already high, total healthcare spending – including Medicare – is projected to continue to climb at a rate of 5% per year over the next decade, and outpace GDP growth by over 1% (CMS, 2016b).

The growing cost of healthcare presents a significant challenge for Medicare. The issue is further complicated by a misalignment of incentives between Medicare and providers: while Medicare bears the cost of healthcare delivery for its beneficiaries, providers may be disincentivized from delivering efficient, high quality care under existing payment models (Rosenthal et al., 2004). For example, under traditional fee-for-service payment – where providers receive payments in proportion to the volume and intensity of services provided – it may be profitable for a provider to overutilize diagnostic imaging (such as CT scans) to generate additional Medicare payments, even if doing so does not necessarily improve patient outcomes (Hendee et al., 2010). Due to the potential impact on total Medicare spending, correcting adverse incentives in healthcare delivery has become one of the central issues of healthcare reform (Milgate and Cheng, 2006, Miller, 2009, Wilensky, 2013). As a consequence, Medicare payment models and provider incentive programs have recently received significant attention within the operations management community, with several studies focusing on analyzing and improving the efficacy of specific Medicare programs (Bastani et al., 2016,

Adida et al., 2016, Guo et al., 2016, Zhang et al., 2016a, Andritsos and Tang, 2015, Gupta and Mehrotra, 2015, Lee and Zenios, 2012, Ata et al., 2012).

The Medicare Shared Savings Program is a voluntary program that offers providers bonus payments for reducing the cost of providing care for Medicare beneficiaries, subject to satisfying certain quality standards. The MSSP has attracted significant attention in the health policy community due to its potential impact on total spending, and, as of 2017, accounts for one-sixth of the Medicare population. However, the financial performance of providers enrolled in the MSSP indicate mixed results. To that end, our focus in this chapter is to propose an alternate contract for the MSSP, and show that our proposed contract can lead to improved outcomes for both Medicare and participating providers.

The goal of the MSSP is to correct the misalignment of incentives between Medicare and providers by making it financially viable for providers to improve the efficiency of healthcare delivery. To participate in the MSSP, a group of Medicare providers are required to form an a cooperative entity referred to as an Accountable Care Organization (ACO). In contrast to the common practice of disaggregated providers treating patients independently, ACOs represent a shift towards an integrated care model wherein a group of Medicare providers coordinate care for a well-defined beneficiary population, and are held jointly responsible for the quality of care delivered (Berwick, 2011, Crosson, 2011). The main focus of the ACO model of healthcare delivery is thus to improve the overall efficiency of providing healthcare through better coordination (e.g., by avoiding duplication of health services), with the aim of reducing healthcare spending while maintaining a high standard of care. As an example, Akira Health, Inc. is an ACO based in Northern California that consists of 31 independent primary care providers (e.g. individual physicians or clinics), which together deliver care to 8,900 Medicare beneficiaries and account for $135 million in Medicare spending (Akira Health, Inc., 2017).

The defining feature of the MSSP is the establishment of *financial benchmarks* for each ACO, which are calculated based on the ACO's historical spending and are risk-adjusted based on attributes of the ACO's beneficiary population (Federal Register, 2011). An ACO whose annual Medicare spending is less than its financial benchmark is eligible for a *shared savings* payment, which is a bonus payment made to the ACO in addition to the usual Medicare fee-for-service payment. The shared savings payment is proportional to the savings generated by the ACO (i.e., the difference between the benchmark and actual spending), so to encourage the ACO to maximize savings. As of 2016, the Shared Savings Program has enrolled 477 ACOs, which together serve a total of 9 million Medicare beneficiaries, equivalent to one-sixth of the total Medicare population (CMS, 2016b).[1]

Despite significant interest from providers, the MSSP faces two notable challenges. First, financial data released by the Center for Medicare and Medicaid Services (CMS) suggests that ACOs have struggled to cut costs: out of the 392 ACOs participating in the MSSP as

[1]Note that the MSSP is distinct (but not mutually exclusive) from the Bundled Payments for Care Improvement (BPCI) Initiative, another recently established Medicare program that incentivizes providers to reduce the cost of healthcare delivery (CMS, 2017). The BPCI offers Medicare providers a single reimbursement for "bundles" of healthcare services received by a beneficiary during a single episode of care, in lieu of reimbursing the provider for each individual service provided. In contrast to the MSSP, participation in the BPCI does not require the formation of an ACO, or include financial bonuses in the form of shared savings payments. Unlike the MSSP, the BPCI is relatively well-studied in the operations management literature (Gupta and Mehrotra, 2015, Adida et al., 2016, Guo et al., 2016).

of 2015, only one-third generated enough savings to qualify for a shared savings payment (CMS, 2016e). As a result, the reduction in total Medicare spending has been modest. In 2015, the total savings across all ACOs was approximately $430 million, which represents a 0.6% decrease in Medicare spending (CMS, 2016e). Second, the low success rate among ACOs has made continued participation in the MSSP unattractive for many providers. A recent survey found that two-thirds of ACOs are uncertain if they will continue participating in the Shared Savings Program, with less than 10% certain that they will remain enrolled (National Association of ACOs, 2014). Since the success of the MSSP depends on voluntary participation by ACOs and a reduction in Medicare spending, these challenges raise serious questions regarding the sustainability of the MSSP in its current form. Moreover, a major barrier that ACOs face with respect to cutting healthcare costs is the significant *investment* that must be made in order to improve the efficiency of healthcare delivery (Haywood and Kosel, 2011). For example, an ACO may need to invest in new information technology to better coordinate care for its patients (Moore et al., 2011), or invest to increase the quality of care delivered so that costly readmissions are minimized (Anderson and Steinberg, 1984). ACOs thus face a delicate balancing act: in order to cut costs and receive shared savings payments from Medicare, they must also *increase* spending to achieve the necessary efficiency gains.

In this chapter, we take a predictive analytics approach to redesigning the MSSP contract, with the goal of addressing the challenges currently faced by the MSSP. Specifically, we propose a new type of contract that includes a *performance-based subsidy* to partially reimburse the investment of an ACO if it successfully reduces spending. The intention of the subsidy is twofold: First, to encourage ACOs to generate additional savings by offsetting the cost of efficiency improvements, and second, to boost total payments so to make the MSSP more attractive to current and prospective ACOs. Our analytical and empirical results suggest that the proposed subsidy scheme achieves both of these desirable outcomes.

Our analysis unfolds in two parts. First, we place the MSSP within a principal–agent framework and consider properties of the optimal contract. In the model, the ACO (agent) has the ability to make an investment to reduce spending on healthcare delivery, and in turn earn bonus payments that depend on both the savings generated and parameters of the MSSP contract. The space of feasible contracts that Medicare may select from is defined by two parameters: a *shared savings rate*, which is the fraction of savings that the ACO receives as bonus payment, and a *subsidy rate*, which is the fraction of the ACO's investment that is reimbursed by Medicare. The solution to Medicare's optimal contracting problem is not obvious. A generous contract may provide a strong incentive for the ACO to reduce spending, but will also increase Medicare's total payments to the ACO. Conversely, if the bonus payment offered to the ACO is too low, then the ACO may be insufficiently incentivized to generate any appreciable savings.

In the second part of our analysis, we design a new contract guided by financial performance data from 392 ACOs currently enrolled in the MSSP. We propose an estimator to infer parameters of the principal-agent model from the ACO data. In contrast to previous work on data-driven incentive design, which uses approaches like linear or logistic regression (e.g. Lee and Zenios (2012), Yamin and Gavious (2013)), our estimation procedure is in the spirit of *inverse optimization*, which refers to the inference of optimization model parameters from noisy solution data (Aswani et al., 2018b, Bertsimas et al., 2015). Our approach is also

distinguished from previous papers in that the key model primitive that we aim to estimate is itself a probability distribution. Using the estimated principal-agent model, we then solve for the optimal MSSP contract, which we formulate as a pure-integer optimization problem, under both the current (non-subsidized) and proposed (subsidized) contracts. Finally, we estimate the potential improvement in savings due to the subsidy by simulating ACO performance under both contracts. The remainder of the paper is organized as follows. In Section 3.2, we review related literature. In Section 3.3, we formulate the principal-agent model for the MSSP, provide an overview of the existing contract, and propose our alternate contract. In Section 3.4, we formulate Medicare's optimal contracting problem and present our key analytical results. In Section 3.5, we present an inverse-optimization based estimator for the principal-agent model and outline our empirical method. In Section 3.6, we present results from the structural estimation and discuss their policy implications with respect to the MSSP. In Section 3.7, we consider extensions to the model and show that our main findings persist under several robustness checks.

3.2 Related literature

This chapter builds on a recent and growing body of work on healthcare contracts in the operations management literature. We also contribute to the health policy literature on accountable care organizations and the broader literature on incentive design.

Healthcare contracts. Incentive problems in healthcare delivery have received significant attention in the operations management literature recently, in part due to the focus on healthcare reform in the United States. A large share of this work has focused on principal-agent settings where the principal (e.g. Medicare) is required to design a payment model for an agent (e.g. provider) that yields socially beneficial outcomes, such as improvements to patient health or a reduction in healthcare spending. One of the first papers in the operations management literature to consider a healthcare contracting problem is by So and Tang (2000), who consider a setting where the principal reimburses an agent for drugs prescribed to patients. The authors focus on analyzing the agent's response to changes in a reimbursement policy that is tied to patient outcomes. Fuloria and Zenios (2001) consider a general problem where an agent determines the intensity of treatment for a patient and the principal reimburses the agent for the services provided. Jiang et al. (2012) consider an optimal contracting problem in a general healthcare setting where the agent's decision is to allocate outpatient service capacity to different groups of patients, and the principal wishes to minimize service cost subject to constraints on agent performance. A common conclusion in all three of these papers is that linking provider reimbursement to patient health can lead to improved health outcomes. The contract we consider in this work is different in that it depends primarily on the financial performance of the provider instead of the quality of healthcare delivered.

Several previous works have analyzed specific Medicare programs. Ata et al. (2013) consider Medicare's payment policies for hospice care. They highlight adverse incentives in the existing policies and propose an alternative payment model which corrects the misalignment of incentives. Gupta and Mehrotra (2015) formulate a game theoretic model for the Bundled Payments for Care Improvement Initiative, which is a new payment model in which Medicare providers receive a single payment for a collection of services provided to a beneficiary. Adida et al. (2016) and Guo et al. (2016) also analyze bundled payment models, and compare

their performance with traditional fee-for-service payment. Zhang et al. (2016a) propose a game theoretic model to study the behavior of hospitals under the recently created Hospital Readmissions Reduction Program, which is a mandatory program that penalizes hospitals that do not reduce readmissions below target levels. Their main result is a set of conditions under which a hospital would rather pay penalties than reduce readmissions. Andritsos and Tang (2015) compare the effect of different Medicare payment models (e.g. fee-for-service vs performance-based payment) on hospital readmissions in a setting where service is co-produced by both the provider and patient. Bastani et al. (2016) propose a general principal-agent framework for pay-for-performance contracts, and examine three Medicare programs as special instances. Jiang et al. (2016) and Savva et al. (2016) also consider the realignment of provider incentives in a more general setting, and focus on the role of competition in reducing costs. To the best of our knowledge, the only existing work to consider the MSSP is by Zhang et al. (2016b), who analyze the impact of the MSSP on the use of computed tomography. This work is different in that we consider the impact of investment subsidies on the financial performance of ACOs, and use observational data from the MSSP to estimate our model.

We are aware of two other papers which take a data-driven approach to designing incentives in a healthcare setting. The first is by Lee and Zenios (2012), who consider the problem of designing a payment model for Medicare's End Stage Renal Disease Program (ESRD). A key methodological distinction with our work is that the formulation of the agent problem in Lee and Zenios (2012) lends itself naturally to the use of linear regression for estimating the model parameters. By contrast, the agent problem in this work is a more general optimization problem and has no closed form solution, which makes the use of linear regression nonviable. Instead, we take a maximum likelihood estimation approach, which results in requiring us to solve an inverse optimization problem (Aswani et al., 2018b). A second distinction is that Lee and Zenios (2012) consider a linear contract, whereas the current MSSP contract that we analyze is piecewise linear and discontinuous in the agent's output. The second related work is by Yamin and Gavious (2013), who consider incentives offered to patients to encourage them to receive influenza vaccinations. The authors formulate the vaccine problem within a game theoretic framework and use logistic regression to estimate the size of the incentive required to optimally vaccinate the population, using phone survey data. The context in the vaccine problem is markedly different than ours since the incentive is offered directly to patients, rather than healthcare providers, and the payment is provided as a lump-sum, rather than being a function of agent output or effort.

Health policy. While the Medicare Shared Savings Program is relatively recent (the first cohort of ACOs enrolled in 2012), it has received significant attention in the health policy literature due its potential impact on healthcare spending. However, to date, quantitative analyses of ACO performance and the MSSP has been limited. McWilliams et al. (2016) analyze early ACO performance data and find that ACOs achieve minimal savings in their first year, suggesting a transition phase for ACOs once they enroll in the MSSP. They also find that ACOs consisting of independent primary care groups tend to save more than those integrated with hospitals. The authors carry out a similar analysis in McWilliams et al. (2015) and McWilliams et al. (2013). Eddy and Shah (2012) develop a simulation model for ACO performance within the MSSP, and find that the existing rules of the MSSP offer little incentive to ACOs to improve the quality of care delivered. Liu and Wu (2014) per-

form a simulation study that considers ACOs and patients as individual agents, and focuses specifically on congential heart failure. There has also been significant qualitative discussion around the rules of ACO formation (Fisher et al., 2012, Lieberman and Bertko, 2011) and the MSSP benchmarking methodology (Douven et al., 2015, Chernew et al., 2014). This work contributes to the literature on ACOs by analyzing the MSSP from a modeling perspective and by being the first to examine the potential impact of ACO subsidies.

Incentive design. This chapter builds on a rich and extensive literature on incentive design and principal-agent problems. For an overview of foundational work in the economics literature, we refer the reader to work by Mirrlees (1974), Hölmstrom (1979), Grossman and Hart (1983) and Hart and Holmström (1986). An overview of principal-agent problems is given by Gibbons (1998) and Laffont and Martimort (2009). Incentive design problems have also recently received significant attention in the operations management literature. Plambeck and Zenios (2000) propose a general framework for dynamic principal-agent problems based on a Markov decision process. Incentive problems have been considered in wide variety of contexts in addition to healthcare, including software development (Whang, 1992), finance (Grinblatt and Titman, 1989, Raghu et al., 2003), sales (Chen, 2000, DeHoratius and Raman, 2007, Khanjari et al., 2013), project management (Chen et al., 2015), manufacturing (Balasubramanian and Bhardwaj, 2004) and supply chain management (Khouja and Zhou, 2010, Guajardo et al., 2012, Lariviere, 2015, Chen and Lee, 2016). There is also a growing literature on the use of subsidies by a central planner in attaining socially desirable outcomes. Most of this work has focused on the use of subsidies to encourage product adoption, such as influenza vaccines (Chick et al., 2008, Arifoglu et al., 2012, Mamani et al., 2013), malaria drugs (Taylor and Xiao, 2014, Levi et al., 2016), and renewable energy technologies (Ata et al., 2012, Chemama et al., 2014, Cohen et al., 2015a,b). We consider subsidies in a slightly different context, where they are used to offset the cost of agent effort, rather than reduce the price of a product.

Lastly, we note here that the work in this chapter coincides with the paper by Aswani et al. (2018a).

3.3 Model

We begin by developing the principal-agent model for the MSSP, which we formulate as a single-period sequential game between Medicare and a single ACO. The ACO provides healthcare to a beneficiary population at a cost that is entirely incurred by Medicare. The interaction proceeds in four steps. First, Medicare selects a contract that depends on the ACO's savings and investment. The ACO then observes the contract and makes an investment (i.e. exerts effort) to reduce spending. Next, the actual savings generated by the ACO is realized. Finally, Medicare observes the actual savings and investment and pays the ACO according to the selected contract.

3.3.1 Preliminaries

The ACO is defined by two attributes: μ and θ. We refer to μ as the ACO's *benchmark*, which is known to Medicare, and θ as the ACOs *type*, which represents the ACO's private information and is unknown to Medicare. The benchmark serves as the primary reference

point for determining whether the ACO has generated savings (if spending is less than μ) or losses (if spending is greater than μ). The benchmark is calculated based on the ACO's historical spending in the years prior to joining the MSSP, and is inflation-adjusted to serve as an estimate of the cost of providing healthcare to the beneficiary population (Federal Register, 2011). The type parameter θ governs the ACO's ability to generate savings. Let the interval $\Theta = [\underline{\theta}, \bar{\theta}]$ denote a continuum of possible ACO types. To reflect uncertainty in the ACO's type, let θ be a random variable supported on Θ, where $F(\cdot|\mu)$ and $f(\cdot|\mu)$ are the distribution and density, respectively, of an ACO with benchmark μ. We assume that F and f are known to Medicare. Note that the type distribution depends on the ACO's benchmark, meaning the ACO's ability to generate savings may depend on its historical Medicare spending. We will often suppress dependence on μ in the notation. A type θ ACO can reduce spending by x by investing $c(x, \theta)$ into efficiency improvements, where $x \in [0, \bar{x}]$. The upper-bound $\bar{x}$ is included for technical purposes and is without loss of generality. In general, if θ is large, we say the ACO is effective at generating savings and is a "high-type" ACO; conversely, if θ is small, then the ACO is ineffective at generating savings and is "low-type". This notion is formalized in Assumption 3.1.

Assumption 3.1. *Let the following assumptions hold.*

1. $c(0, \theta) = 0$ *for all* θ,
2. $c(x, \theta)$ *is strictly convex and increasing in* x,
3. $c(x, \theta)$ *and* $(\partial/\partial x)c(x, \theta)$ *are decreasing in* θ, *and* $\lim_{\theta \to 0} c(x, \theta) = \infty$ *for any* x.

The assumption that $c(x, \theta)$ is strictly convex and increasing in x implies diminishing returns on investment, since the marginal cost of generating an additional unit of savings increases with x. The assumption that $c(x, \theta)$ is decreasing in θ implies that high-type ACOs are more efficient at generating savings than low-type ACOs. Similarly, the assumption that $(\partial/\partial x)c(x, \theta)$ is decreasing in θ implies that the rate at which the return in investment diminishes is lower for high-type agents, which is consistent with the notion that high-type agents are more effective at generating savings. Examples of functional forms that satisfy Assumption 3.1 are x^2/θ and $(x/\theta)\log(x + 1)$. Due to variations in the precise healthcare needs of the population, there may be uncertainty in the exact cost of delivering care. We account for this uncertainty through a random shock to the ACO's spending, denoted by ξ. Let $G(\cdot)$ and $g(\cdot)$ be the distribution and density of ξ, respectively. We now impose the following conditions on the shock.

Assumption 3.2. *Let the following assumption hold.*

i) $\mathbb{E}[\xi] = 0$, $\mathbb{E}[\xi^2] = \sigma^2 < \infty$.

ii) The shock density g *is continuous, almost everywhere differentiable, unimodal and symmetric around 0.*

iii) There exists constants $\bar{g}$, $\bar{g}' < \infty$ *such that* $g(\xi) \leq \bar{g}$, $g'_-(\xi) \leq \bar{g}'$, *and* $g'_+(\xi) \leq \bar{g}'$ *for all* ξ, *where* $g'_-(\xi)$ *and* $g'_+(\xi)$ *are left and right derivatives of* g, *respectively.*

Assumption 3.2 is fairly mild and admits a large class of probability distributions, such as the Gaussian, Laplace and logistic distributions. Statement *iii*) simply means that the density and its derivative are bounded, and is stated with respect to the left and right derivatives of g to permit distributions that are not differentiable everywhere (e.g. the Laplace density). Note that the shock distribution is independent of the ACO's benchmark. This assumption allows us to avoid overfitting by reducing the number of parameters in our estimation procedure (discussed in Section 3.5). We note that we obtain a strong model fit despite this simplifying assumption.

We can now define the ACO's actual spending on delivering healthcare, or the *delivery cost*, as $D = \mu - x - \xi$. It may be convenient to interpret D as demand for healthcare that must be satisfied by the ACO. Since in our setting, demand represents the healthcare needs of the patient population, we model it as being exogenous to costs. Note that since $\mathbb{E}[\xi] = 0$, the expected delivery cost if the ACO makes no investment (implying $x = 0$ by Assumption 3.1) is given by $\mathbb{E}[D] = \mu$. We can now define the *savings* generated by the ACO as the difference between the benchmark and actual delivery cost as $y = \mu - D$. Since $D = \mu - x - \xi$, we can write the savings equivalently as

$$y = x + \xi.$$

Therefore, the actual savings generated by an ACO is equal to its reduction in spending plus an exogenous shock. Let $H(\cdot|x)$ and $h(\cdot|x)$ denote the distribution and density of savings, respectively, given the ACO's action x. Note that the actual savings y is random due to the shock, and that $\mathbb{E}[y] = x$. We also assume throughout that $\mathbb{P}[\mu - \bar{x} - \xi < 0] = 0$, meaning the demand is always non-negative. This assumption can be ensured by truncating the distribution of the shock variable ξ (and holds trivially in practice, given that healthcare costs cannot be driven to 0). Next, we define the contract that determines the ACO's shared savings payment, which depends on the realized savings, y.

3.3.2 Baseline contract: Shared savings

In the MSSP, an ACO enters a "one-sided" contract for an initial agreement period (usually three years), after which it is transitioned to a "two-sided" contract. The one-sided contract is risk-free for the ACO: it receives payments from Medicare for generating savings, but is not penalized if spending exceeds the benchmark. The one-sided contract depends on three parameters: the *shared savings rate* $\alpha \in \mathcal{A}$, a minimum savings threshold h, and a fixed savings cap, C_u. The shared savings payment received by the ACO in the one-sided contract is then given by

$$r_+(y, \alpha) = \begin{cases} 0, & y \in (-\infty, h], \\ \alpha y, & y \in (h, C_u], \\ \alpha C_u, & y \in (C_u, \infty). \end{cases} \tag{3.1}$$

The shared savings rate α represents the fraction of savings that the ACO receives as payment if the realized savings y is positive. The threshold h is the minimum savings required for the ACO to receive a payment from Medicare. The minimum savings threshold accounts for natural variation in healthcare costs by ensuring that any observed savings are "real", i.e., due to ACO effort and not chance (Federal Register, 2011). The savings cap C_u reduces the

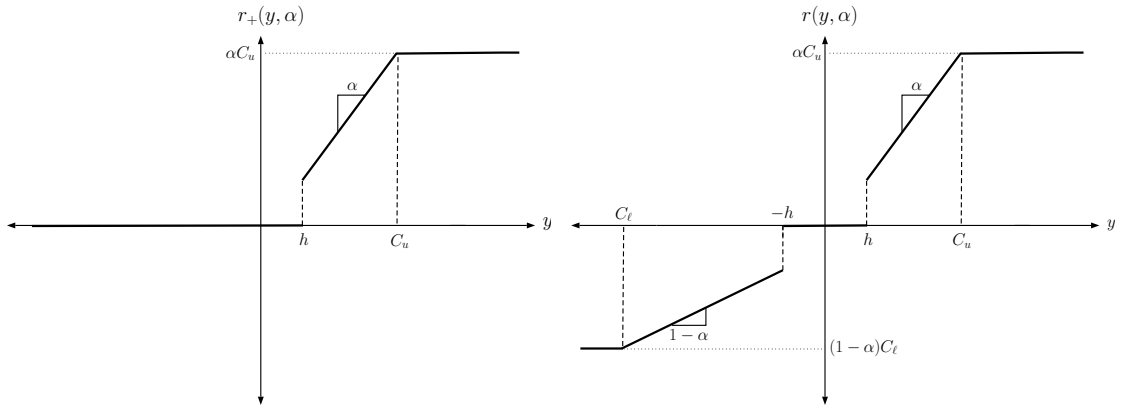

(a) One-sided contract.

(b) Two-sided contract.

Figure 3.1: Shared savings payment functions under one-sided and two-sided contracts.

risk to Medicare due to the shock and the information assymmetry, by protecting Medicare from making an excessively large shared savings payment to the ACO.

In the two-sided contract, both savings and losses are shared with Medicare. In addition to receiving a shared savings payment for generating positive savings, the ACO pays a penalty of $(1-\alpha)y$ to Medicare if $y < -h$, that is, if the excess spending above the benchmark is greater than the threshold h. Similar to the one-sided case, the maximum penalty that can be paid by the ACO is capped at $(1-\alpha)C_\ell$. Writing the penalty terms as negative payments, the payment received by the ACO in the two-sided contract is then given by

$$r(y,\alpha) = \begin{cases} (1-\alpha)C_\ell, & y \in (-\infty, C_\ell], \\ (1-\alpha)y, & y \in (C_\ell, -h], \\ 0, & y \in (-h, h], \\ \alpha y, & y \in (h, C_u], \\ \alpha C_u, & y \in (C_u, \infty). \end{cases} \tag{3.2}$$

Figure 3.1 shows a schematic of the one-sided and two-sided contracts of the MSSP. We largely focus on the two-sided contract in this work, since it is the intended long-term contract of the MSSP. The one-sided contract is most relevant in Section 3.5, since the financial data that we use for estimation corresponds to ACO performance under the one-sided contract.

3.3.3 Proposed contract: Shared savings + performance-based subsidy

The MSSP in its current form is entirely performance-based, in that it rewards the ACO for generating savings but is otherwise agnostic to the ACO's investment. The presence of the random shock under this payment model is risky for the ACO, since it may be penalized for exceeding the benchmark, despite investing in efficiency improvements. Here, we define a new contract which partially mitigates the ACO's risk by providing a subsidy for the investment

in addition to the the existing shared savings payment. In this sense, our proposed contract is both effort-based and performance-based. Further, as discussed in Section 1, our inclusion of a subsidy component in the MSSP contract is related to evidence that suggests that the investments required by ACOs to improve efficiency can be a barrier to the generation of savings (CMS, 2016c). Moreover, our approach of contracting on both the agent's action as well as its performance is inspired by a well-established body of literature that suggests that contracting on both action and outcome can improve outcomes (Shavell, 1979a,b, Harris and Raviv, 1979, Hölmstrom, 1979). As we discuss later in this section, the key consequence of our proposed is that, under reasonable conditions, it can dominate the existing contract by generating a higher payoff for *both* Medicare and the ACO. Under our proposed contract, Medicare observes the ACO's investment and provides a subsidy proportional to the investment, in addition to the shared savings payment. Let $\beta \in \mathcal{B}$ be the *subsidy rate* selected by Medicare. The subsidy payment is then given by

$$s(y, \beta) = \begin{cases} 0, & y \in (-\infty, -h), \\ \beta c(x, \theta), & y \in (h, \infty). \end{cases} \tag{3.3}$$

Note that like the shared savings payment, the subsidy payment is also performance-based, in the sense that the ACO only receives it if the savings is positive and exceeds the minimum savings threshold h. Note that since the subsidy payment $\beta c(x, \theta)$ depends on the private information θ, which is a priori unknown to Medicare. Following standard mechanism design theory, in Section 3.4.2 we will formulate Medicare's optimal contracting problem in a manner that incentivizes ACOs to truthfully report their type θ, which enables Medicare to contract directly on the ACO's investment.

We also highlight here that the subsidy scheme presented in this work is different from the the ACO Investment Model that was recently created by the CMS (CMS, 2016c). The most significant difference is that the subsidy payments offered to ACOs through the AIM are later deducted from shared savings payments, which effectively makes the AIM a loan program. We instead consider an entirely separate incentive that exists in addition to the shared savings payment. Further, the AIM is targeted at a subset of ACOs that meet a well-defined criteria (e.g. those that are in rural areas and exclude hospitals), whereas we consider a more general program that can apply to any ACO.

3.3.4 Discussion of modeling assumptions

Before analyzing Medicare's contracting problem, we first discuss some of the key assumptions of our model.

Quality. One might reasonably expect that an ACO could decrease the quality of healthcare delivered (e.g., by cutting services) as a way to reduce spending and generate bonus payments within the MSSP. However, we assume throughout our analysis that quality of care is fixed, and that the ACO only decides on the size of the investment. In other words, our model does not permit the ACO to decrease quality of care to generate savings. This assumption is supported by ACO performance data and the existing regulations of the MSSP. Data released by the CMS shows that quality has actually *improved* under the MSSP – for example, the average ACO quality score increased by approximately 15% from 2014 to 2015 (CMS, 2016a). Moreover, the regulations of the MSSP require Medicare to verify that any

savings generated by an ACO are not due to quality reductions. Medicare upholds quality through close monitoring of ACOs, e.g., to ensure that they do not avoid at-risk patients or underuse health care services (CMS, 2016d), and failure to meet quality standards or comply with Medicare monitoring may jeopardize an ACOs participation in the MSSP (Federal Register, 2011). As a consequence, ACOs have neither the incentive nor the ability to use quality as a lever to generate shared savings payments. We therefore focus our analysis on the ACO's investment behavior.

Benchmark independence and ACO size. In our model the ACO's optimal savings $x(\theta)$ depends explicitly on α, β and θ, but does not depend on the benchmark μ when θ is held fixed. To capture the effect of ACO benchmark on the savings generated, we allow the distribution over the parameter θ to depend on the ACO's benchmark in our empirical analysis in Section 3.5. We also do not explicitly incorporate ACO size (i.e. number of assigned beneficiaries) into the model, and instead perform our analysis at the beneficiary level. This normalization does not change the analytical results presented in this section, but is useful for simplifying the estimation in Section 3.5. A cursory look at the data suggests a weak relationship between ACO size and savings (a correlation coefficient of $\rho = 0.08$), which supports this normalization. While we normalize for beneficiaries for estimation purposes, we explicitly incorporate ACO size when estimating ACO performance under the proposed contract in Section 3.5.

Single ACO. Although the MSSP includes many participating ACOs, for our analytical results we consider the interaction between Medicare and a single ACO. This is because an ACO's payment depends only its own savings, benchmark and investment, and does not depend on the performance of other ACOs. Moreover, ACOs are prohibited by legislation from colluding or sharing information with each other (Federal Register, 2011). As a consequence, the key insights can be obtained by focusing on the contracting problem between Medicare and one ACO.

Single period. In practice, the financial benchmark for ACOs are updated annually, based on the ACO's most recent three years of spending and the growth rate in national healthcare spending. The MSSP might therefore be viewed as a multi-period problem with dynamically updating spending benchmarks. However, since our aim in this work is to analyze the potential impact of investment subsidies, it suffices to restrict our attention to a single period model. We note here that additional insights may be gained by considering a multi-period model for the MSSP. We take a first step toward this extension in Appendix 3.7.1.

3.4 Model analysis

In this section, we consider the ACO's investment behavior and Medicare's optimal contracting problem under the existing and proposed MSSP contracts. We examine the ACO's problem in Section 4.1 and Medicare's contracting problem in Section 4.2. In Section 4.3, we present analytical results regarding the optimal contract structure.

3.4.1 ACO's Investment Problem

We begin our analysis by consdering the investment problem faced by the ACO. Medicare providers typically generate profit from delivering healthcare to beneficiaries. Let the profit

associated with spending D be γD, where $\gamma > 0$ is the profit margin. The expected service-related profit is then given by $\mathbb{E}[\gamma D] = \gamma(\mu - x)$. Since μ is a constant, we focus on the ACO's profit loss due to a reduction in spending, which is simply γx.[2] The ACO's *expected payoff* associated with a savings level of x can now be written as

$$u(x, \alpha, \beta, \theta) = \int_{-\infty}^{\infty} (r(y, \alpha) + s(y, \beta))\, h(y|x) dy - \gamma x - c(x, \theta) \tag{3.4}$$

The payoff function $u(x, \alpha, \beta, \theta)$ is the sum of the expected shared savings and subsidy payments from Medicare, minus the profit loss and investment. We assume that the ACO wishes to maximize expected payoff. The ACO's *optimal savings* is then given by

$$x(\theta) = \underset{x \in [0, \bar{x}]}{\operatorname{argmax}} \; u(x, \alpha, \beta, \theta) \tag{3.5}$$

where $\bar{x}$ is the maximum savings the ACO can achieve (e.g. $\bar{x}$ can be trivially set to μ). In general, x depends on α, β and σ as well, although we suppress dependence on these parameters for conciseness when it is clear from context. For convenience, we also define $U(\alpha, \beta, \theta) = u(x(\theta), \alpha, \beta, \theta)$ to be the ACO's optimal profit. We assume that the ACO participates in the MSSP if $U(\alpha, \beta, \theta) \geq 0$.[3] The model naturally generalizes to account for a fixed cost associated with enrolling in the MSSP, which can be represented by subtracting a constant from $u(x, \alpha, \beta, \theta)$. Including this fixed cost does not fundamentally alter our results, so we assume it to be 0. We now turn our attention to analyzing the ACO's behavior, starting with the following assumption.

Assumption 3.3. *The inequality*

$$\bar{\alpha}(h\bar{g}' + \bar{g}) + 2\bar{\beta}\bar{g}(\partial/\partial x)c(x, \theta) \leq (1 - \bar{\beta})(\partial^2/\partial x^2)c(x, \theta).$$

holds for all $x \in [0, \bar{x}]$ *and* $\theta \in \Theta$.

We now make the following observation.

Lemma 3.1. *Under Assumption 3.3, the ACO payoff function* $u(x, \alpha, \beta, \theta)$ *is strictly concave over* $[0, \bar{x}]$.

Assumption 3.3 is a technical condition that faciliates our analysis by guaranteeing that the ACO's payoff function has a unique maximizer, and can be shown to hold under parameter values from our data set (see Appendix 3.7.2). We emphasize that strict concavity in the ACO's payoff function is not crucial, and our results remain valid for any payoff function with a unique maximizer (i.e. non-convex unimodal functions). We now define a quantity that will be useful in the remainder of this section.

[2] To keep the the model tractable and amenable to estimation, we assume that the provider's service-related profit decreases with its cost reduction efforts, but the profit margin itself remains constant at γ. This is supported by the notion that the MSSP aims, in part, to generate savings by reducing the total number of healthcare services provided, meaning the average per-service profit margin will not necessarily be impacted by the provider's operational decisions.

[3] In general, we assume that the ACO participates if its payoff is non-negative. In our analysis, however, we shall use a stronger constraint to guarantee that the ACO's payoff is no less than it would be under the existing MSSP contract (i.e. strictly positive). We impose this stronger condition to restrict attention to contracts that improve Medicare savings without leaving the ACO worse off.

Definition 3.1. *Define* $\theta_{\alpha,\beta} = \inf\{\theta \in \Theta | x(\theta) > 0$ *under* $\alpha, \beta\}$.

The threshold $\theta_{\alpha,\beta}$ represents the lowest type ACO that would generate a strictly positive savings given the shared savings rate α and subsidy rate β (note that $\theta_{\alpha,0}$ may not exist if α is is very small; in general, we assume that the minimum shared savings rate in $\mathcal{A}$ is large enough such that $\theta_{\alpha,0}$ exists for all $\alpha \in \mathcal{A}$). We now introduce two lemmas that characterize the ACO's behavior and will be useful for our main results in the next section. The first of these relates the savings generated by an ACO to its type.

Lemma 3.2. *For any* $\alpha \in \mathcal{A}$ *and* $\beta \in \mathcal{B}$*, the ACO's optimal savings* $x(\theta)$ *is non-decreasing over* Θ *and strictly increasing over* $[\theta_{\alpha,\beta}, \bar{\theta}]$.

Lemma 3.2 states that an ACO with high investment efficacy will generate higher savings. This is not a surprising result – since it is less costly for a high-type ACO to achieve a fixed savings level x, one might expect that it would be optimal for a high-type ACO to save more than a low-type ACO. The parameter $\theta_{\alpha,\beta}$ is the threshold beyond which ACO savings are guaranteed to be strictly increasing in type. This threshold exists because for low-type ACOs where $x = 0$, a small increase in investment efficacy may not be sufficient to make it profitable for the ACO to generate positive savings, and thus x may remain at 0. However, if $x > 0$, then even slight improvements in the ACO's investment efficacy will lead to higher savings. The next result shows that offering a subsidy can boost ACO payoff while also incentivizing the ACO to generate higher savings.

Lemma 3.3. *For any* $\alpha \in \mathcal{A}$ *and* $\theta \in \Theta$*, the ACO's optimal savings* $x(\theta)$ *and payoff* $U(\alpha, \beta, \theta)$ *are non-decreasing in* β*. For any* α *and* $\theta \in [\theta_{\alpha,0}, \bar{\theta}]$*, the savings* $x(\theta)$ *and* $U(\alpha, \beta, \theta)$ *are strictly increasing in* β*.*

Intuitively, the subsidy effectively increases the ACO's investment efficacy, since it makes it less costly for the ACO to generate a fixed saving of x. As a consequence, the ACO's investment and optimal savings level increases with β. Further, the ACO's optimal payoff $U(\alpha, \beta, \theta)$ also increases with β. This suggests that introducing a subsidy may help improve participation by boosting total ACO payments.

3.4.2 Medicare's optimal contract

For a given ACO type θ, Medicare's total expected spending is given by the sum of the delivery cost and the shared savings and subsidy payments made to the ACO. Since the expected delivery cost is $\mathbb{E}[D] = \mu - x$, and μ is a constant, we can formulate Medicare's problem equivalently as one of maximizing savings. We thus write *Medicare's savings* as the ACO's savings minus the payments made to the ACO:

$$v(x, \alpha, \beta, \theta) = x - \int_{-\infty}^{\infty} (r(y, \alpha) + s(y, \beta))\, h(y|x) dy. \tag{3.6}$$

For convenience, let $V(\alpha, \beta, \theta) = v(x(\theta), \alpha, \beta, \theta)$ be Medicare's savings under the ACO's optimal savings level $x(\theta)$. We consider the contracting problem in two settings: a full information case, where the ACO's type is known to Medicare, and a partial information case, where type is unknown. We refer to the optimal contracts obtained under the full and

partial information cases as the *first-best* and *second-best* contracts, respectively. In practice, the first-best contract cannot be implemented due to the information assymetry between Medicare and the ACO. Note that the full information case represents the case where the ACO's optimal investment is observable, since the ACO's optimal investment level under a given contract is encoded by the type parameter θ. Similarly, the partial information case where θ is private information represents the case where the ACO's optimal investment level is unobservable by Medicare. In general, we will assume that the partial information case represents the reality of the MSSP, where the ACO's investment is not known to Medicare, but can be elicited by designing the contract appropriately.

Full information (First-best). In the full information case, the optimal contract may depend on the ACO's type, since it is fully known by Medicare. For a known ACO type θ, the first-best contract to offer a type θ ACO is given by the solution to

$$\begin{aligned}
\underset{\alpha,\beta}{\text{maximize}} \quad & \int_{\Theta} \left(x(\theta) - \int_{-\infty}^{\infty} \left(r(y,\alpha) + s(y,\beta)\right) h(y|x(\theta))dy \right) f(\theta)d\theta && \text{(3.7a)} \\
\text{subject to} \quad & \int_{-\infty}^{\infty} \left(r(y,\alpha) + s(y,\beta)\right) h(y|x(\theta))dy - \gamma x(\theta) - c(x(\theta),\theta) \geq \bar{u}(\theta), && \text{(3.7b)} \\
& \alpha \in \mathcal{A}, && \text{(3.7c)} \\
& \beta \in \mathcal{B}. && \text{(3.7d)}
\end{aligned}$$

The constraint (3.7b) is a participation constraint that ensures the ACO receives a payoff of at least $\bar{u}(\theta)$. Participation contraints are standard in the mechanism design literature, and are also referred to as *individual rationality* constraints (e.g. see Laffont and Martimort (2009) and Borgers et al. (2015)). This constraint is particularly important in the context of the MSSP because voluntary participation by a large number of ACOs is crucial for the success of the program (Rosenthal et al., 2011). Since our focus is on the impact of the subsidy payment on improving payoffs for both Medicare and ACOs, in our analysis we set $\bar{u}(\theta)$ equal to the expected payoff that a type θ ACO receives under the existing contract. This allows us to restrict attention to subsidy contracts that leaves both Medicare and the ACO better off.

Partial information (Second-best). We now consider the contracting problem where the ACO's type is unknown by Medicare. In particular, we assume that the the parameter θ is only known to be drawn from the distribution $f(\theta)$. In this setting, Medicare's contracting problem is to maximize savings by designing a menu of contracts $(\alpha(\theta), \beta(\theta))$, $\theta \in \Theta$, where $\alpha(\theta) \in \mathcal{A}$ and $\beta(\theta) \in \mathcal{B}$. An ACO that reports its type to be θ then receives a contract parameterized by $\alpha(\theta)$ and $\beta(\theta)$. The second-best optimal contract is given by the solution to

$$\begin{aligned}
\underset{\alpha(\theta),\beta(\theta)}{\text{maximize}} \quad & \int_{\theta\in\Theta} V(\alpha(\theta),\beta(\theta),\theta) f(\theta)d\theta && \text{(3.8a)} \\
\text{subject to} \quad & U(\alpha(\theta),\beta(\theta),\theta) \geq \bar{u}(\theta), \quad \theta \in \Theta, && \text{(3.8b)} \\
& U(\alpha(\theta),\beta(\theta),\theta) \geq U(\alpha(\theta),\beta(\theta),\theta') \quad \theta,\theta' \in \Theta, && \text{(3.8c)} \\
& \alpha(\theta) \in \mathcal{A}, \quad \theta \in \Theta, && \text{(3.8d)} \\
& \beta(\theta) \in \mathcal{B}, \quad \theta \in \Theta. && \text{(3.8e)}
\end{aligned}$$

Let $(\alpha^*(\theta), \beta^*(\theta))$, $\theta \in \Theta$ denote optimal contract parameters attained at a solution to (3.8). The objective function (3.8a) represents Medicare's expected savings, where the expectation is taken over both the shock and type. As in the full information problem (3.7), we include a participation constraint to ensure that the ACO's expected payoff is at least $\bar{u}(\theta)$. Constraints (3.8c) are *incentive compatibility* constraints which enable Medicare to accurately elicit the ACO's private information θ (Laffont and Martimort, 2009). The interpretation of constraint (3.8c) is that the payoff that the ACO enjoys from truthfully reporting its type θ to Medicare must be at least as high as the payoff it receives from declaring any other type θ'. These constraints ensure that the feasible set of contracts is restricted to those that allow the ACO's type θ, and thus its investment $c(x(\theta), \theta)$, to effectively be observable by Medicare. Note also that Medicare's optimal contracting problems (3.7) and (3.8) are bi-level programs, since the constraint sets depends on the optimal solution of the ACO problem $x(\theta)$ (see Dempe (2002) for an overview of bi-level programming). Bi-level programs are typically non-convex, and so uniqueness of $\alpha^*(\theta)$ and $\beta^*(\theta)$ is generally not guaranteed. We now present three results related to the structure of the optimal contract. First, let the following assumption hold in the remainder of this section.

Assumption 3.4. *The inequality* $(\partial/\partial x)R(x, \alpha) < 1$ *holds for all* $x \in [0, \bar{x}]$.

This assumption is relatively mild and required for technical purposes. Intuitively, we would expect Assumption 3.4 to hold at larger values of x, because the expected shared savings payment is sub-linear at higher savings levels (due to $\alpha < 1$). However, there can exist values of h and σ such that Assumption 3.4 is violated for small x, due to the discontinuity of the shared savings payment at $y = h$. This assumption is validated by our dataset (see Appendix 3.7.2). Our first result regarding the optimal contract focuses on low-type ACOs.

Proposition 3.1. *There exists* $\psi > 0$ *and* $\theta_0 > 0$ *such that if* $C_u \geq \psi$ *and* $|C_\ell| \geq \psi$, *then* $\alpha^*(\theta) \geq 1/2$ *and* $\beta^*(\theta) = 0$ *for all* $\theta \leq \theta_0$.

The condition that C_u and $|C_\ell|$ be sufficiently large simplifies the analysis but is not restrictive.[4] To further illustrate Proposition 3.1, recall that the optimal savings $x(\theta)$ of a very low-type ACO will be small (Lemma 3.2). As a consequence, the shared savings rate α must be sufficiently large to guarantee the ACO a non-negative payoff. From the ACO performance data, we observed that the average quality-adjusted shared savings rate that ACOs received in 2015 was 0.55, with approximately 15% of ACOs receiving a shared savings rate of less than 0.5. In light of Proposition 3.1, this suggests that a large share of ACOs may be at risk of dropping out of the MSSP once they transition to the two-sided contract. The next result characterizes the dependence of optimal contract parameters $\alpha^*(\theta)$ and $\beta^*(\theta)$ on θ.

Proposition 3.2. *Suppose* $\Theta = \{\theta_L, \theta_H\}$, *where* $\theta_H > \theta_L > 0$. *There exists* $\delta > 0$ *such that if* $\theta_H - \theta_L \geq \delta$, *then* $\alpha^*(\theta_L) \geq \alpha^*(\theta_H)$ *and* $\beta^*(\theta_L) \leq \beta^*(\theta_H)$.

Proposition 3.2 states it is optimal for Medicare to offer a high-type ACO a higher subsidy rate and a lower shared savings rate compared to a low-type ACO. The condition that the

[4] In practice, the parameters C_u and $|C_\ell|$ are large relative to typical ACO savings. Based on existing MSSP guidelines, C_ℓ and C_u are set at 15% and 20% of the ACO's benchmark, respectively (Federal Register, 2011), whereas data released by the CMS shows that the majority of ACO savings and losses are within 5% of the benchmark.

ACO types be sufficiently far apart (as represented by the condition that $\theta_H - \theta_L \geq \delta$ for some $\delta > 0$) is included because the generality of $g(\xi)$ and $c(x, \theta)$, in addition to the absence of a closed form expression for $x(\theta)$, precludes obtaining stronger monotonicity results regarding $\alpha^*(\theta)$ and $\beta^*(\theta)$. To see the intuition behind Proposition 3.2, note that a very high-type ACO can generate a large savings with a relatively small investment. Note also that the resulting subsidy payment associated with a subsidy rate of β is proportional to the ACO's investment. As a result, it is worthwhile for Medicare to provide a large subsidy rate to a high-type ACO, since the actual subsidy payment will be small relative to the increase in savings generated by the ACO. The converse is true for low-type ACOS. Our final result in this section compares the outcomes of the optimal non-subsidized[5] and subsidized contracts.

Proposition 3.3. *Let $U_0^*(\theta)$ and $U_s^*(\theta)$ be a type θ ACO's payoff under the optimal non-subsidized and subsidized contracts, respectively. Let $V_0^*(\theta)$ and $V_s^*(\theta)$ be the associated Medicare savings. Then there exists $\theta_s > 0$ such that $U_s^*(\theta) > U_0^*(\theta)$ and $V_s^*(\theta) > V_0^*(\theta)$ for all $\theta \geq \theta_s$.*

Proposition 3.3 states that for sufficiently high-type ACOs, the subsidized contract dominates the existing shared savings only contract, in the sense that it produces a strictly higher payoff for both Medicare and the ACO. The condition that the ACO type is at least θ_s is necessary to ensure the ACO is effective enough at generating savings for it to be worthwhile for Medicare to provide an additional payment in the form of a subsidy. Conversely, if the ACO type is too low, then it may be the case that the additional savings generated by the ACO are less than the subsidy payment itself, which would make it suboptimal for Medicare to offer a subsidy of any size. From our numerical results, we find that this type threshold θ_s is relatively low (see Section 4.6). To illustrate the need for Assumption 3.4 in Proposition 3.3, recall that increasing β from 0 to a positive value increases the ACO savings x (Lemma 3.3), and thus increases the expected shared savings payment as well. Now consider the case where σ is very close to 0, so that $y \approx x$. Since $r(y, \alpha)$ contains a jump discontinuity at $y = h$, a small increase in x may lead to a relatively large increase in $r(y, \alpha)$ if x is close to h before the subsidy. Assumption 3.4 ensures that the increase in the expected shared savings payment is not greater than the increase in x itself. The key implication of Proposition 3.3 is that Medicare can generate additional savings using the subsidy *without* jeopardizing ACO participation through lower payments. In the remainder of the paper, we outline an empirical approach for estimating Medicare's potential savings under the proposed subsidized contract.

3.5 Estimation

Typically, in the mechanism design literature, the distribution over agent types is common knowledge to the principal. In practice, however, this distribution is unlikely to be known *a priori*. Our aim in this section is to use ACO financial data made available by the CMS to estimate the type density $f(\theta)$, which is an important input to identifying the optimal MSSP contract.[6] We also estimate the variance of the shock, σ^2 from the data. We then

[5] By "optimal non-subsidized contract", we mean the existing MSSP structure under the optimal shared savings parameters $\alpha(\theta)$, where $\beta(\theta) = 0$ for all $\theta \in \Theta$.

[6] Note that because we have only a single observation for each ACO, we cannot estimate the type parameter θ for each individual ACO in the dataset. We therefore focus on estimating the distribution over ACO types

use these estimates to solve for the optimal parameters $\alpha(\theta)$ and $\beta(\theta)$ under the proposed contract. Lastly, we estimate the performance of ACOs in the MSSP under the existing and proposed contracts, which allows us to estimate the improvement in Medicare savings that might result from introducing a performance-based subsidy to the MSSP.

3.5.1 Data

We use a dataset made publicly available by the Centers for Medicare and Medicaid Services (CMS, 2016e). The dataset contains the number of Medicare beneficiaries, benchmark expenditures, actual expenditures, and quality score for 392 ACOs participating within the Medicare Shared Savings program in 2015. A summary of the dataset is given in Tables 1 and 2. The ACOs in our dataset represent 7.27 million Medicare beneficiaries across the United States. In 2015, this group of ACOs were assigned an aggregate spending benchmark of \$73.30 billion, and had an actual spending of \$72.87 billion, representing a \$430 million net reduction in total Medicare spending. Although the total savings was \$430 million, the total shared savings payments earned by ACOs was \$645 million, resulting in a net loss of \$215 million for Medicare. This loss is due to the fact that as of 2015, all participating ACOs were under an initial three-year agreement in which the one-sided contract was in effect, meaning ACOs were rewarded for generating savings but not penalized for exceeding their assigned benchmarks. Note that the mean per-beneficiary savings by an ACO in the study cohort was \$101, which is approximately 1% of the mean benchmark. The standard deviation of the per-beneficiary ACO savings was \$680, suggesting substantial variation in ACO performance. Figure 3.2 illustrates the distribution in the number of beneficiaries, spending benchmarks and savings for all 392 ACOs.

Let μ_i, b_i, y_i and θ_i represent the benchmark, number of beneficiaries, actual savings and type of the i^{th} ACO, respectively. In practice, μ_i, b_i and y_i are directly observed by Medicare, but θ_i represents private information. To maximize Medicare's savings, we would ideally estimate $\theta_i, \ldots, \theta_n$ from the data, and then solve problem (3.7) to identify the first-best contract to offer to each of the n ACOs, based on their types. However, this approach is not viable since limited data is available for each ACO (in many cases only a single year). As a consequence, the type parameters $\theta_i, \ldots, \theta_n$ cannot be estimated directly. Under appropriate assumptions, however, we may instead estimate the distribution over ACO types, $f(\theta)$, which equips us to identify the optimal second-best contract. For our numerical experiments, we select a random subset of 275 ACOs (70% of the data) to estimate the model, and validate the resulting model fit against an out-of-sample dataset containing the performance of the remaining 117 ACOs (30% of the data).

3.5.2 Model parameterization

We have so far imposed minimal assumptions on the ACO type and spending shock distributions. Our goal now is to estimate these distributions from the ACO performance data in a manner that is both tractable and captures the dependence of an ACO's type on its benchmark. Here, we outline a parameterization of our model to be used in the estimation.

in aggregate. However, it may be possible to estimate the exact type parameter for each ACO given multiple observations of the same ACO over several years. We note this as a potentially fruitful direction for future analyses of the MSSP.

Table 3.1: Performance of ACOs in Medicare Shared Savings Program in 2015 under one-sided contract.

Variable	Value
ACOs	392
Total Medicare beneficiaries	7.27 million
Total benchmark expenditures	$73.30 billion
Total actual expenditures	$72.87 billion
Total savings	$430 million
ACO shared savings payments	$645 million
Medicare savings	−$215 million

Table 3.2: ACO summary statistics.

	Mean (SD)	Median	IQR	Range
Beneficiaries (per ACO)	18,547 (18,508)	12,545	7,954 – 21,286	513 – 149,633
Benchmark ($ per beneficiary)	10,403 (2,360)	9,863	8,827 – 11,357	5,548 – 22,777
ACO savings ($ per beneficiary)	101 (680)	13	−252 – 394	−3,136 – 2,586

The model primitives to be specified are the ACO conditional type distribution $f(\theta|\mu)$, the shock distribution $G(\xi)$, and the ACO's investment function $c(x, \theta)$. We assume ACO types are distributed on the type interval $\Theta = [1, 1 \times 10^4]$ according to a mixture of (truncated) exponential distributions. Each mixture component corresponds to an ACO benchmark group (e.g. 'low' or 'high' benchmark), which allows us to capture the dependence of ACO type on benchmark in a semi-parametric manner. Although our approach can be extended to alternate choices for the type distribution, we choose the exponential distribution based on the observation that most ACOs generate minimal savings, which suggests that ACO types are concentrated at lower values of the type parameter. Moreover, use of the exponential distribution requires us to estimate only a single parameter for each mixture component, which limits the complexity of the model and prevents overfitting. Let m be the number of mixture components, indexed by j. Then let $\mathcal{M}_1, \ldots, \mathcal{M}_m$, be a disjoint partitioning of the positive real line, where each interval $\mathcal{M}_j$ represents a set of ACO benchmarks. With a slight abuse of notation, let $f(\theta|\lambda_j)$ denote the exponential distribution with parameter λ_j. We assume that the i^{th} ACO's type is drawn from $f(\theta|\lambda_j)$ if the ACO's benchmark μ_i belongs to $\mathcal{M}_j$. For the shock distribution, we assume that the variable ξ is distributed according to the zero-mean Laplace distribution $G(\xi|\sigma)$, where σ is the standard deviation. Note that this specification of the shock distribution satisfies Assumption 3.2. Based on this parameterization, the quantities to be estimated are the m shape parameters, $\lambda_1, \ldots, \lambda_m$, and the shock parameter σ. Note that we are not required to determine which mixture component to assign each observation to, since the assignment depends only on the benchmark μ_i, which is known from the data. Lastly, for the investment function, we set $c(x, \theta) = x^2/\theta$, which satisfies Assumption 3.1.

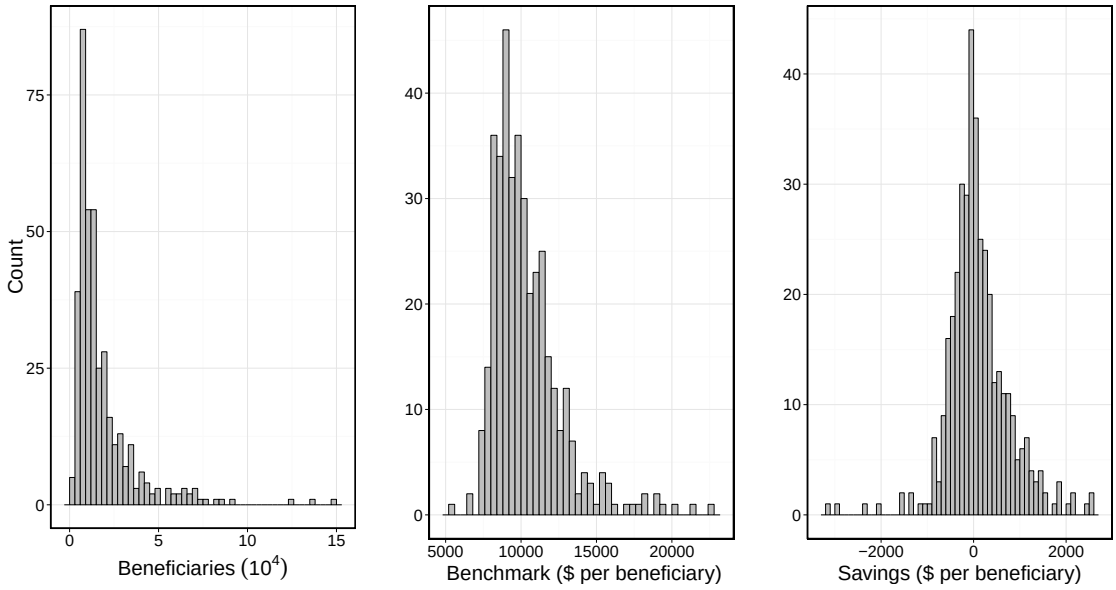

Figure 3.2: Distribution of number of Medicare beneficiaries, spending benchmarks and savings for 392 ACOs participating in the MSSP in 2015.

3.5.3 Model idenfiability and maximum likelihood estimation

The parameters $(\lambda_1, \ldots, \lambda_m, \sigma)$ can show to be statistically identifiable under an additional condition that is relatively mild. This condition is formalized in Proposition 3.4.

Proposition 3.4 (Identification)**.** *If the condition*

$$\int_{\Theta} e^{-\sqrt{2}x(\theta)/\sigma} \left(f(\theta|\lambda) - f(\theta|\tilde{\lambda}) \right) \neq 0 \tag{3.9}$$

holds for all $\sigma \in \Sigma$*, and all* $\lambda, \tilde{\lambda} \in \Lambda$ *such that* $\lambda \neq \tilde{\lambda}$*, then the model parameters* $(\lambda_1, \ldots, \lambda_m, \sigma)$ *are identifiable.*

Corollary 3.1. *If* Λ *and* Σ *are finite sets and* $\bar{\theta}$ *is sufficiently large, then the identification condition* (3.9) *holds.*

The identification condition (3.9) ensures that there cannot exist multiple values of the shape parameter λ that give rise to the same savings distribution, for a given benchmark group. Identifiability of the model can be guaranteed by placing appropriate assumptions on Λ, Σ and Θ. Corollary 3.1 provides one such set of assumptions. Note that the conditions in Corollary 3.1 are mild, since one can construct finite sets Λ and Σ by discretizing a continuous parameter space to an arbitrarily fine degree. Additionally, assuming that $\bar{\theta}$ is very large does not materially change the model or estimation approach, since it simply expands the space of possible ACO types.

We now outline a maximum likelihood estimation approach for the parameters $(\lambda_1, \ldots, \lambda_m, \sigma)$. Obtaining maximum likelihood estimates for our principal-agent model requires us to solve an

inverse optimization problem, which refers to the estimation of optimization model parameters from (potentially noisy) solution data (Aswani et al., 2018b, Bertsimas et al., 2015, Ahuja and Orlin, 2001b). The MLE problem takes the form of an inverse optimization problem since the observed savings data represent noisy observations of the ACO's optimal decision (recall $y = x + \xi$). We first require the following assumption.

Assumption 3.5. *The data $(\mu_i, b_i, y_i, \theta_i)$, $i = 1, \ldots, n$, are drawn independently from a common distribution.*

Assumption 3.5 states that the benchmark, number of beneficiaries, type and savings of an ACO is independent of other ACOs. This assumption is reasonable given that each ACO operates independently, and is necessary for tractability of the estimator. Note that Assumption 3.5 permits dependence between the attributes of a given ACO. We now formalize the estimation problem. Let $\Lambda_1, \ldots, \Lambda_m$ and Σ denote the parameter sets for $\lambda_1, \ldots, \lambda_m$ and σ, respectively.

Proposition 3.5. *The maximum likelihood estimate of $(\lambda_1, \ldots, \lambda_m, \sigma)$ is given by*

$$(\hat{\lambda}, \hat{\sigma}) = \underset{\lambda \in \Lambda, \sigma \in \Sigma}{\operatorname{argmax}} \sum_{i=1}^{n} \log \left(\int_{\Theta} g\left(y_i - x(\theta) | \sigma, \theta\right) f(\theta | \lambda(\mu_i)) d\theta \right)$$

$$x(\theta) = \underset{x \geq 0}{\operatorname{argmax}} \int_{-\infty}^{\infty} r_+(y, \alpha) h(y|x) dy - \gamma x - c(x, \theta)$$

where $\lambda(\mu_i) = \lambda_j$ if $\mu_i \in \mathcal{M}_j$.

In Proposition 3.5, $x(\theta)$ refers to the optimal savings of a type θ under the contract that generated the data. Note that since our dataset contains the performance of ACOs under the one-sided contract, we use the one-sided shared savings function r_+ (given in (3.1)) in the ACO's problem in Proposition 3.5. If the dataset represented ACO performance under the two-sided contract, then we would simply replace r_+ with the two-sided payment function r, given in (3.2).[7]

Maximum likelihood estimators are often difficult to solve to global optimality due to the likelihood function being non-convex. As a result, estimation approaches which find local maxima of the likelihood function (e.g. expectation maximization) are typically used (Hastie et al., 2005). However, since the number of parameters is relatively small in our setting (assuming m is not too large), we obtain an approximately optimal solution to the MLE problem as follows. Let Λ, Σ and Θ be discrete sets. First, we numerically solve the ACO's problem to obtain $x(\theta)$ for each $(t, \sigma) \in \Theta \times \Sigma$ (note that the ACO's problem depends on σ as well as θ). Then, for each $\sigma \in \Sigma$, we evaluate the likelihood function given in Proposition 3.5 for each $(\lambda_1, \ldots, \lambda_m) \in \Lambda_1 \times \ldots \times \Lambda_m$, and select the parameter vector $(\lambda_1, \ldots, \lambda_m)$ with the largest likelihood.

[7]Note that we estimate the type distribution $f(\theta)$ by using the normalized, per-beneficiary savings of each ACO. In other words, to reduce model complexity we do not explicitly account for variation in ACO size in the estimator. We instead capture variation in ACO size when simulating contract performance, by sampling the number of beneficiaries along with other ACO attributes in the bootstrap.

3.5.4 Specification of parameters

Mixture distribution. Note that if the ACO type distribution is defined by a large number of mixture components, then the resulting model may overfit to the training data and thus not be generalizable outside of the sample. Thus, two additional modeling decisions that remain are the number of benchmark groups (i.e., mixture components), and the range of ACO benchmarks covered by each group. We tuned these additional parameters by using k-means clustering to identify the endpoints of the intervals corresponding to each of k benchmark groups, and by using 10-fold cross validation (on the training data) to identify the optimal number of benchmark groups. For the cross validation step, we fit the model specified above using maximum likelihood estimation (described in Section 4.5.2) and evaluated model performance by computing the likelihood function for the validation set. We performed the cross validation by varying the number of benchmark groups from one and six, and found the optimal number of mixture components to be three. The associated benchmark groups implied by the clustering were $\mathcal{M}_1 = [0, 1.03 \times 10^4]$, $\mathcal{M}_2 = (1.03 \times 10^4, 1.43 \times 10^4]$ and $\mathcal{M}_3 = (1.43 \times 10^4, \infty)$.

ACO profit function. To solve the MLE problem, we must also specify the parameters in the ACO's profit function, $u(x, \alpha, \beta, \theta)$. We set $\alpha = 0.46$ to represent the "effective" shared savings rate, which we observed from the ACO data to be the average shared savings rate that ACOs received after quality adjustments (a maximum rate of 0.5 multiplied by the average quality score of 92%). Although the quality scores vary from one ACO to the next, we assume a fixed score of 92% across all ACOs to reduce model complexity during estimation. We find that we obtain a strong model fit despite this simplifying assumption. We explicitly incorporate quality scores in our estimate of ACO performance under the optimal contract by sampling each ACOs quality in the bootstrapping procedure, where it is used to adjust the nominal shared savings rate.

The remaining parameters to be defined are the payment cap C_u, the penalty cap C_ℓ, the minimum savings and losses threshold h, and the ACO's Medicare profit margin γ. We set $C_u = 2,000$ and $C_\ell = -1,500$ based on MSSP guidelines (Federal Register, 2011). We set $h = 200$, since the minimum savings and losses threshold is mandated to be 2% of the ACO's benchmark, and the average benchmark in the dataset was $\$10,403$ per beneficiary. Lastly, we set the ACO's Medicare profit margin as $\gamma = 0.03$, based on a recent report by the Medicare Payment Advisory Commission (Miller, 2013).

3.5.5 Estimation of contract performance

Using the parameter estimates obtained from Section 3.5, we solve for the optimal second-best subsidized contract, which we label II-S. The space of possible shared savings and subsidy rates is given by $\mathcal{A} = \{0, 0.05, \ldots, 0.9\}$ and $\mathcal{B} = \{0, 0.05, \ldots, 0.9\}$, respectively. By also letting the type space be discrete, i.e., by setting $\Theta = \{1, 100, 200, \ldots, 5000\}$, we can formulate and solve a discrete counterpart to the contracting problem (3.8) as an integer optimization model. The formulation is given in Appendix 3.5.7. Given the importance of ACO participation, our focus is on identifying a contract that improves both Medicare savings and the ACO's payoff. Therefore, for each $\theta \in \Theta$, we set $\bar{u}(\theta)$ equal to the expected ACO payoff under the existing two-sided contract. This restricts the set of feasible contracts in problems (3.8) to those that produce at least as high of a payoff for the ACO as the current MSSP contract. To maintain

our focus on the impact of the subsidy on the performance of the MSSP, we do not adjust the minimum savings threshold h or the savings and losses cap C_u and C_ℓ as part of the contract optimization. Moreover, since our aim is to examine the effect of introducing a subsidy to the MSSP as it exists currently, we keep the parameters h, C_u and C_ℓ at the values specified by the existing MSSP regulations.

Using the estimated model parameters, we simulate ACO performance under the existing MSSP contract and our proposed contract using a standard non-parametric bootstrapping procedure (see Efron and Tibshirani (1994b) for a comprehensive overview of bootstrapping techniques). In particular, for each ACO we sample its benchmark spending, actual spending, quality score, and the number of beneficiaries, which allows us to estimate the total savings under each contract. We also compute 95% confidence intervals for the savings under each contract, and perform bootstrap hypothesis tests to assess the statistical significance of the estimated improvement in contract performance due to the investment subsidy.

3.5.6 Non-parametric contract

As a point of reference, we also consider a non-parametric contract that does not involve separate shared savings or subsidy payments, but instead assigns a single payment to each ACO type. This contract can be viewed as the most flexible possible contract for the MSSP, since it involves constructing general functions which, for each $\theta \in \Theta$, map a payment $\rho(x)$ to each savings level x. By the well-known revelation principle (Myerson, 1981) it suffices to restrict attention to designing a menu of contracts $(\rho(\theta), x(\theta))$, $\theta \in \Theta$. Under incentive compatibility, an ACO that reports type θ finds it optimal to generate savings $x(\theta)$, for which it receives payment $\rho(\theta)$. Therefore, the ACO's profit is given by $\rho(\theta) - \gamma x(\theta) - c(x(\theta), \theta)$. Medicare's savings associated with a type θ ACO is then given by $x(\theta) + \xi - \rho(\theta)$, which, in expectation, is simply $x(\theta) - \rho(\theta)$. As with the subsidized contract, Medicare's optimal contracting problem can be formulated as an integer optimization problem (see Appendix 3.5.7). The structure of this contracting problem, which designs a menu of reward-action pairs over all possible agent types is standard in the mechanism design literature (Laffont and Martimort, 2009). For optimization purposes, we assume that for each θ, Medicare selects $\rho(\theta)$ from the set $\mathcal{R} = \{0, 20, 40, \ldots, 2000\}$ and $x(\theta)$ from the set $\mathcal{X} = \{0, 20, 40, \ldots, 2000\}$.

3.5.7 Integer optimization models for Medicare's contracting problem

This section presents the formulations for Medicare's optimal contracting problem. We consider a discrete counterpart to the contracting problem given in (3.8), which allows us to formulate the contracting problem as an integer optimization model. Let $\{\alpha_1, \alpha_2, \ldots, \alpha_{|\mathcal{I}|}\}$ be the set of shared savings rates, $\{\beta_1, \beta_2, \ldots, \beta_{|\mathcal{J}|}\}$ be the set of subsidy-rates, and $\{\theta_1, \theta_2, \ldots, \theta_{|\mathcal{K}|}\}$ be the set of ACO types. Let $u(\alpha_i, \beta_j, \theta_k)$ be the optimal payoff of a type θ_k ACO under contract parameters α_i and β_j, which is a parameter that can be computed numerically in advance. Similarly, let $v(\alpha_i, \beta_j, \theta_k)$ be Medicare's savings under parameters $(\alpha_i, \beta_j, \theta_k)$. As in formulation (3.8), let $\bar{u}(\theta_k)$ be the minimum pay-off for a type θ_k ACO. Let $p(\theta_k)$ be the probability that an ACO is type θ_k, where $p(\theta_k)$ can be computed by discretizing the type distribution estimated in Section 3.5. The key decision variable in the integer optimization model is the binary variable $z(\alpha_i, \beta_j, \theta_k)$, which is equal to 1 if the contract parameters α_i and β_j are mapped to a type θ_k ACO, and 0 otherwise. The optimal contracting problem

can then be formulated as

$$\underset{\mathbf{z}}{\text{maximize}} \quad \sum_{i\in\mathcal{I}}\sum_{j\in\mathcal{J}}\sum_{k\in\mathcal{K}} v(\alpha_i,\beta_j,\theta_k)\cdot z(\alpha_i,\beta_j,\theta_k)\cdot p(\theta_k) \tag{3.10a}$$

$$\text{subject to} \quad \sum_{i\in\mathcal{I}}\sum_{j\in\mathcal{J}} u(\alpha_i,\beta_j,\theta_k)\cdot z(\alpha_i,\beta_j,\theta_k) \geq \bar{u}(\theta_k), \quad k\in\mathcal{K}, \tag{3.10b}$$

$$\sum_{i\in\mathcal{I}}\sum_{j\in\mathcal{J}} u(\alpha_i,\beta_j,\theta_k)\cdot z(\alpha_i,\beta_j,\theta_k) \geq \sum_{i\in\mathcal{I}}\sum_{j\in\mathcal{J}} u(\alpha_i,\beta_j,\theta_k)\cdot z(\alpha_i,\beta_j,\theta_{k'}), \quad k,k'\in\mathcal{K}, \tag{3.10c}$$

$$\sum_{i\in\mathcal{I}}\sum_{j\in\mathcal{J}} z(\alpha_i,\beta_j,\theta_k) = 1, \quad k\in\mathcal{K}, \tag{3.10d}$$

$$z(\alpha_i,\beta_j,\theta_k)\in\{0,1\}, \quad i\in\mathcal{I},\ j\in\mathcal{J},\ k\in\mathcal{K}. \tag{3.10e}$$

The optimal contract in the non-parametric setting can be formulated in a similar manner. Let $\{\rho_1,\rho_2,\ldots,\rho_{|\mathcal{L}|}\}$ be the set of possible payments, $\{x_1,x_2,\ldots,x_{|\mathcal{M}|}\}$ be the set of ACO savings levels, and $\{\theta_1,\theta_2,\ldots,\theta_{|\mathcal{K}|}\}$ be the set of ACO types. By the revelation principle (Myerson, 1981), it suffices to restrict attention to incentive-compatible contracts that map a payment ρ and savings x to each ACO type θ. The decision variable in the non-parametric formulation is thus $w(\rho_l,x_m,\theta_k)$ which is equal to 1 if the payment-savings pair (ρ_l,x_m) is assigned to a type θ_k ACO. The integer optimization model for the optimal non-parametric contract is then given by

$$\underset{\mathbf{w}}{\text{maximize}} \quad \sum_{l\in\mathcal{L}}\sum_{m\in\mathcal{M}}\sum_{k\in\mathcal{K}} v(\rho_l,x_m,\theta_k)\cdot w(\rho_l,x_m,\theta_k)\cdot p(\theta_k)$$

$$\text{subject to} \quad \sum_{l\in\mathcal{L}}\sum_{m\in\mathcal{M}} u(\rho_l,x_m,\theta_k)\cdot z(\rho_l,x_m,\theta_k) \geq \bar{u}(\theta_k), \quad k\in\mathcal{K},$$

$$\sum_{l\in\mathcal{L}}\sum_{m\in\mathcal{M}} u(\rho_l,x_m,\theta_k)\cdot w(\rho_l,x_m,\theta_k) \geq \sum_{l\in\mathcal{L}}\sum_{m\in\mathcal{M}} u(\rho_l,x_m,\theta_k)\cdot z(\rho_l,x_m,\theta_{k'}), \quad k,k'\in\mathcal{K},$$

$$\sum_{l\in\mathcal{L}}\sum_{m\in\mathcal{M}} w(\rho_l,x_m,\theta_k) = 1, \quad k\in\mathcal{K},$$

$$w(\rho_l,x_m,\theta_k)\in\{0,1\}, \quad l\in\mathcal{L},\ m\in\mathcal{M},\ k\in\mathcal{K}.$$

Both formulations can be solved using off-the-shelf optimization solvers (e.g. Gurobi). Note that the optimal subsidized and non-parametric contracts can be recovered from the optimal assignment vectors $\mathbf{z}^*$ and $\mathbf{w}^*$, respectively.

3.6 Results and discussion

For validation purposes, we estimate the model parameters using data from a random subset of 275 ACOs (70% of the data), and validate it against an out-of-sample dataset containing the performance of the remaining 117 ACOs (30% of the data). Table 3.3 presents the parameter estimates for each of the three benchmark groups, where $\hat{\lambda}$ is the estimated shape parameter of the exponential distribution, and $\hat{\sigma}$ is the estimated standard deviation of the

Table 3.3: Maximum likelihood parameter estimates for three benchmark groups.

Benchmark	# of ACOs	$\hat{\lambda}$	$\mathbb{E}[\theta]$	$\hat{\sigma}$	$\mathbb{E}[x(\theta)]$
$< \$10,300$	228	1	1	550	\$1
$\$10,300 - \$14,300$	140	1.1×10^{-3}	900	550	\$140
$> \$14,300$	24	6.7×10^{-4}	1500	550	\$260

random shock. Note that $\hat{\sigma}$ is the same for all three benchmark groups, since we estimate only a single shock distribution to avoid overfitting. The estimated mean of the exponential distributions is given by $1/\hat{\lambda}$, which we may interpret as the average ACO type for the given benchmark category. Since the parameter estimates in Table 3.3 are difficult to interpret directly, we also report the expected optimal savings, $\mathbb{E}_t[x(\theta)]$, under the MSSP's current one-sided contract.

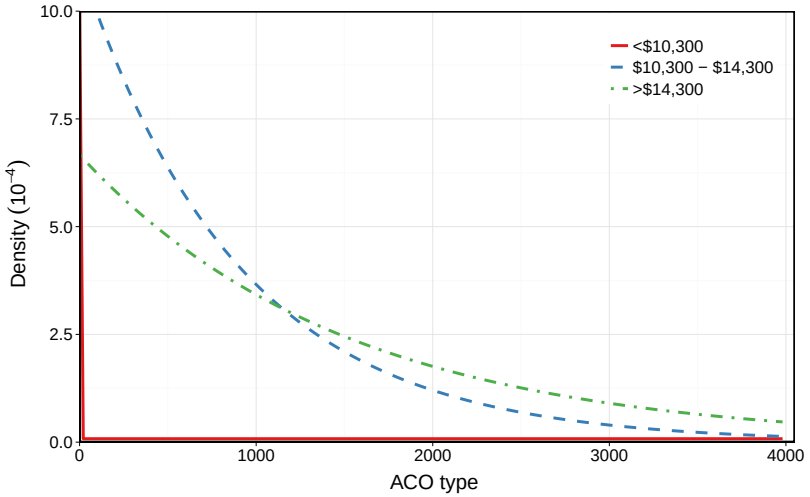

Figure 3.3: Estimated ACO type distributions for low (<\$10,300), intermediate (\$10,300-\$14,300) and high (>\$14,300) benchmark groups.

To validate our model, we simulated the savings that would be generated under the current contract using the parameter estimates given in Table 3.3. Figure 3.4 shows the empirical savings distribution for both the in-sample and out-of-sample datasets and the simulated CDF from our fitted model (with 10,000 samples). We performed a Kolmogorov-Smirnov test to assess the goodness-of-fit of our model with respect to the empirical savings data. The KS-statistic (i.e. the maximum vertical distance between the empirical and simulated CDFs) on the out-of-sample dataset was 0.076, with an associated p-value of $p > 0.2$. This p-value implies a *failure* to reject the null hypothesis that the two CDFs were generated under different statistical models at a confidence level of 80%, suggesting a strong model fit.

Our results show that, in general, ACOs with high benchmarks are more likely to be high-type, and thus more likely to generate positive savings. As shown in Table 3.3, we estimated

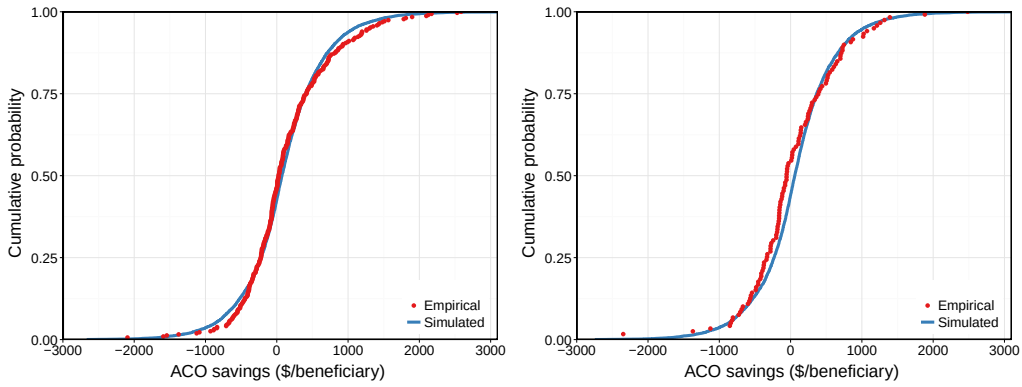

(a) In-sample.

(b) Out-of-sample.

Figure 3.4: Empirical and simulated savings under one-sided contract, for in-sample and out-of-sample data.

the expected per-beneficiary savings of low ($< \$10,300$), intermediate ($\$10,300 - \$14,300$) and high ($> \$14,300$) benchmark ACOs to be \$1, \$140 and \$260, respectively. Figure 3.3 provides an accompanying visualization which shows the estimated type distribution for each benchmark group. Observe that the type distribution for low benchmark ACOs is concentrated near $\theta = 0$, whereas the distributions for the intermediate and high benchmark groups place mass across a range of values of θ. As one might expect, the parameter estimate $\hat{\lambda}_j$ for the intermediate benchmark group falls in between the low and high benchmark values. We numerically found the type threshold θ_s discussed in Section 3.4.2 to be approximately 80 for our choice of the ACO investment function. As can be seen from Figure 3.3, ACOs in the intermediate and high benchmark groups have a high probability of being above this type threshold, suggesting that the subsidy may be useful in boosting savings from these groups. In contrast, since low benchmark ACOs are low-type, the subsidy may not be effective in improving savings from low benchmark ACOs.

The observation that high benchmark ACOs tend to save more may be a consequence of the benchmark calculation methodology that is currently employed by the MSSP. As discussed previously, an ACO's benchmark is largely determined by its historical spending. Since there are multiple factors that may contribute to an ACO's historical spending, such as regional variations in healthcare costs and the exact needs of the beneficiary population, it is difficult to pinpoint exactly why a particular ACO has high historical spending. However, our results suggest that ACOs with high benchmark expenditures may have an easier time reducing spending because they were historically cost inefficient, and thus have "more room" for improvement. By contrast, an ACO with a low benchmark may be less able to generate savings because it has historically been cost efficient, and thus additional reductions in spending are more difficult to attain. Based on this result, we propose that it may be beneficial for the CMS to revisit the benchmark calculation methodology, so that ACOs with low benchmarks are not inadvertently penalized for good historical performance.

Figure 3.5 shows the optimal contract parameters of the first-best and second-best subsidized contracts. Note that the right panel of Figure 3.5 exhibits the behavior predicted by Proposition 3.1, where low type ACOs are assigned a shared-savings rate where $\alpha(\theta) \geq 1/2$.

Additionally, the results shown in Figure 3.5 align with Proposition 3.2, where high-type ACOs are assigned relatively higher subsidy rates and lower shared savings rates, with the converse holding true for low-type ACOs.

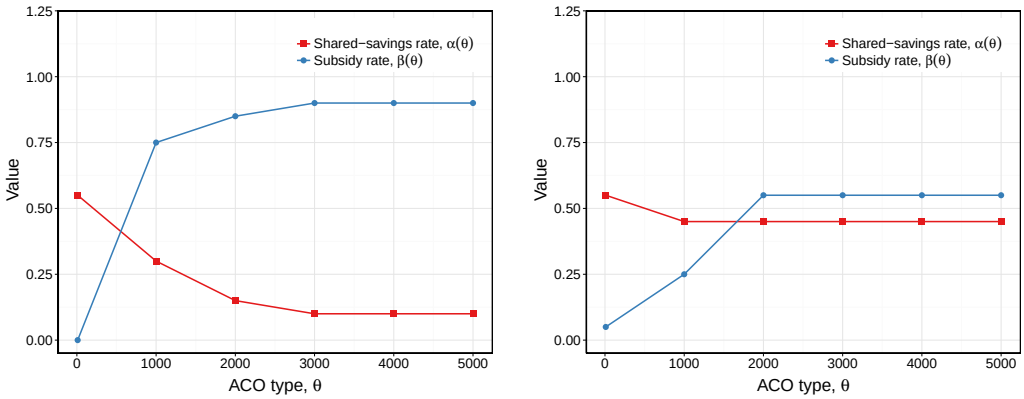

(a) Full information. (b) Partial information.

Figure 3.5: Optimal contract parameters for subsidized MSSP in full and partial information settings.

Table 3.4 shows the estimated ACO payoff, Medicare savings, and total welfare under the optimal second-best subsidized contract. For comparison purposes, we simulate the performance of the 392 ACOs once they enter the current two-sided contract of the MSSP, and report the results under the 'Baseline' column. We also report the difference between the baseline and optimal contracts, denoted by Δ, as well as a one-sided p-value to test the significance of the difference (see Efron and Tibshirani (1994b) for an overview of hypothesis testing with the bootstrap). We also report the 95% confidence intervals for the payoff estimates under the baseline contract, subsidized contract, and their relative performance gap Δ. We note here that although we held quality constant throughout our analysis thus far, for simulation purposes we also sample quality scores for each ACO in the bootstrap, and calculate the savings based on the quality-adjusted value of the shared savings parameter α.

Table 3.4: Bootstrap estimates for subsidized contract (in millions).

	Baseline		II-S		Δ		
	Mean	95% C.I.	Mean	95% C.I.	Mean	95% C.I.	p
ACOs	\$282	(\$188, \$382)	\$316	(\$216, \$423)	\$34	(\$22, \$49)	<0.01
Medicare	\$146	(\$39, \$260)	\$207	(\$97, \$328)	\$62	(\$48, \$79)	<0.01
Total	\$427	(\$234, \$642)	\$523	(\$320, \$741)	\$96	(\$78, \$118)	<0.01

In the subsidized contract (II-S), the ACOs and Medicare both experience an improvement compared to the baseline. This finding aligns with the result in Proposition 3.3. Specifically, ACO payoff increases by 12% (\$282 to \$316 million) and Medicare savings increases by 43% (\$146 to \$207 million), relative to their baseline levels. Note also that the increase in total welfare (the sum of ACO payoff and Medicare savings) is 22% under the subsidized contract.

As an intermediate step in our analysis, we also solved for the optimal non-subsidized contract where β is held fixed at 0 and the contracting problem is solved over $\alpha(\theta)$ only to identify the optimal set of shared savings rates. The improvement in Medicare savings from optimizing over the shared savings rate alone was found to be negligible (results omitted). This result suggests that in the absence of a subsidy, Medicare cannot boost savings by adjusting the shared savings rate without compromising ACO payoff, and therefore their participation in the MSSP. While it has been suggested that the MSSP should increase the shared savings rate to strengthen incentives for ACOs (Chernew et al., 2014), our analysis suggests that doing so will only further reduce Medicare savings, due to the associated increase in ACO payments.

In Section 3.4.2, we showed that subsidizing the investments made by ACOs to reduce healthcare spending can improve both Medicare savings and ACO payoff. This result is validated by our empirical analysis and the findings in Table 3.4. Further, we note that the vast majority of ACOs in the MSSP remain under the initial one-sided contract, and are therefore not penalized for accruing losses. However, out of the 392 ACOs in our dataset, 28 attained quality scores that would result in an effective shared savings rate that is less than 0.5 in the two-sided contract. In light of Proposition 3.1, these ACOs are particularly at risk of dropping out of the MSSP once they transition to the two-sided contract. Our results suggest that subsidizing ACO investments may be useful in mitigating this risk.

Table 3.5: Bootstrap estimates for non-parametric contract (in millions).

	Baseline		II-NP		Δ		
	Mean	95% C.I.	Mean	95% C.I.	Mean	95% C.I.	p
ACOs	\$282	(\$188, \$382)	\$301	(\$265, \$346)	\$19	(−\$73, \$109)	>0.1
Medicare	\$146	(\$39, \$260)	\$237	(\$92, \$381)	\$91	(\$30, \$152)	<0.01
Total	\$427	(\$234, \$642)	\$538	(\$325, \$745)	\$110	(\$82, \$122)	<0.01

For comparison purposes, Table 3.5 shows the estimated performance of the optimal non-parametric contract. The optimal contract parameters are shown in Figure 3.6. Due to its generality, this contract represents the theoretical best possible performance of any MSSP contract (given the estimated type distribution). Under the non-parametric contract, Medicare's total savings increases from \$146 to \$237 million, which represents a 63% increase in total savings compared to the baseline contract. The non-parametric contract also improves upon the subsidized contract (\$237 vs \$207 million). This is unsurprising given the additional flexibility provided by the non-parametric contract. However, it is worth noting that a majority of the improvement potential associated with the non-parametric contract (\$91 over baseline) is captured by the subsidized contract (\$62 million over baseline). This result suggests that a subsidized contract performs favorably in comparison to the theoretically best non-parametric contract.

It is unlikely that there is a single solution to the challenges currently faced by the MSSP. More likely, a multi-faceted restructuring of the existing contract will be required to boost Medicare savings and increase the appeal of the MSSP to existing and prospective ACOs. Our analytical and empirical results suggest that ACO investment subsidies have the potential to play an important role in this restructuring. Moreover, as the MSSP continues and more

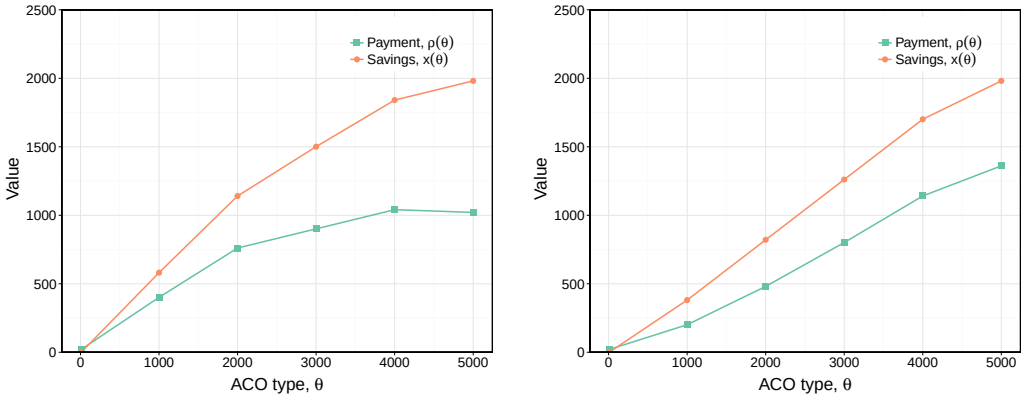

(a) Full information.

(b) Partial information.

Figure 3.6: Optimal contract parameters for non-parametric contract in full and partial information settings.

data is collected, a clearer picture of ACO performance within the MSSP will emerge. We hope the empirical approach developed in this work can serve as a springboard for future analyses of the MSSP.

3.7 Extensions and sensitivity tests

In this section we consider three extensions to our main model. First, we repeat the empirical analysis (parameter estimation, contract optimization and bootstrap simulation) under two alternate parametric forms for the ACO's investment function $c(x, \theta)$, and compare these results with our findings from Section 4.6. Second, we numerically analyze a setting where there is inaccuracy in the financial benchmark, i.e., Medicare observes a noisy signal of μ instead of μ directly. Third, we consider an alternate contract that only depends on the ACO's benchmark μ, instead of θ.

Table 3.6 presents the estimated savings under the baseline and optimal contracts for three different investment functions, including $c(x, \theta) = x^2/\theta$, which is used in our main analysis in Sections 3.5. Under the current two-sided contract, we estimate the total Medicare savings to range from \$135 to \$146 million. Our estimates of Medicare's savings under the optimal subsidy contract range from \$187 to \$303 million. Given that total spending by this group of ACOs was approximately \$73 billion, these results suggest that our estimates of Medicare savings under the optimal contract are reasonably robust to the choice of the investment function.

Table 3.6: Bootstrap estimates for subsidized contract under alternate investment functions (in millions).

		Baseline		II-S		Δ		
$c(x,\theta)$		Mean	95% C.I.	Mean	95% C.I.	Mean	95% C.I.	p
x^2/θ	ACOs	\$282	(\$188, \$382)	\$316	(\$216, \$423)	\$34	(\$22, \$49)	<0.01
	Medicare	\$146	(\$39, \$260)	\$207	(\$97, \$328)	\$62	(\$48, \$79)	<0.01
	Total	\$427	(\$234, \$642)	\$523	(\$320, \$741)	\$96	(\$78, \$118)	<0.01
$(x^2+x)/\theta$	ACOs	\$256	(\$162,\$353)	\$284	(\$188, \$389)	\$28	(\$18, \$42)	<0.01
	Medicare	\$137	(\$31, \$249)	\$187	(\$79, \$303)	\$50	(\$38, \$63)	<0.01
	Total	\$393	(\$200, \$598)	\$449	(\$250, \$660)	\$78	(\$63, \$94)	<0.01
$(x/\theta)\log(x+1)$	ACOs	\$221	(\$127, \$326)	\$249	(\$137, \$355)	\$27	(\$10, \$47)	<0.01
	Medicare	\$135	(\$8, \$272)	\$383	(\$228, \$535)	\$248	(\$194, \$321)	<0.01
	Total	\$356	(\$150, \$591)	\$632	(\$392, \$877)	\$276	(\$214, \$354)	<0.01

In general, we have assumed that the benchmark μ represents the ACO's expected spending on healthcare delivery in the status-quo, that is, with no ACO investment. In practice, the ACO's expected spending under no investment and the benchmark assigned by Medicare may be different. We numerically consider a setting where the true benchmark for a given ACO is μ, which is known to the ACO, but Medicare observes it to be $\mu + \omega$, where ω is a random noise term. We set ω to be a zero-mean uniform random variable with support $[-0.05\mu, 0.05\mu]$. In other words, we consider a setting where Medicare's observation of each ACO's benchmark may be up to 5% higher or lower than the true benchmark. The results in Table 3.7 indicate that inaccurate estimates of the benchmark can decrease Medicare's savings under both the baseline and subsidized contracts. However, the improvement potential associated with the performance-based subsidy remains fairly constant at \$60 million in both cases. The ACO's total payoff increases when Medicare's observation of the benchmark is noisy, which is unsurprising given the informational advantage that ACOs enjoy in this setting.

Lastly, we consider a contract that depends only on the ACO's benchmark, instead of its type parameter. This contract may be considered be more practical than a menu of contracts over ACO types, given that the benchmark is a tangible and observable attribute of each ACO. Specifically, we consider a contract that assigns a single shared savings rate α and subsidy rate β to each of the three benchmark clusters from Section 4.6. Further, we assume in this setting that Medicare observes a noisy signal of the ACO's investment. We posit that observing the ACO's investment can be feasible in practice since Medicare already extensively monitors ACOs within the MSSP, and has a legislated mandate to audit Medicare providers in general. To reflect the incentive that ACOs have to inflate their reported investments (so to earn a larger subsidy), we assume Medicare observes the ACO's investment to be $(1+\eta)c(x, \theta)$, where η is an exponential random variable with a mean of 0.1. In other words, we assume Medicare observation of the true investment is always inflated, with an average inflation of 10% above the true investment.

Figure 3.7 shows the optimal contract parameters for the benchmark-based contract. The optimal shared savings rate for all benchmark groups was found to be 0.55. This is a consequence of the the contract being unable to distinguish between ACO types combined with Proposition 3.1. Note also that the subsidy rate is $\beta = 0$ for low benchmark ACOs, since low benchmark ACOs are all concentrated as low-type ACOs, and thus ineffective at generating savings (cf. Table 3.3). As a consequence, it is optimal for Medicare to not offer the investment subsidy to low benchmark ACOs. By contrast, the optimal subsidy rate for intermediate and high benchmark ACOs is $\beta = 0.3$. Table 3.8 shows the simulated savings under the regular and benchmark-based contracts. Interestingly, the benchmark-based contract achieves nearly the same level of Medicare savings as the full type-based contract (\$187 vs \$207). This is likely due to the fact that the optimal contract parameters in the type-based contract are relatively insensitive to the ACO type, as shown in Figure 3.5b).

3.7.1 ACO savings over multiple periods

From the perspective of Medicare, it may be ideal for the ACO to behave myopically, that is, generate as large of a savings as possible each year. However, since the existing MSSP benchmarking mechanism lowers the benchmark based on the ACO's savings in the previous

year, an ACO might find it optimal to strategically curtail investment (and generate lower savings) in earlier years, in order to receive more favorable benchmarks in later years. In this section, we address the multiple-period setting where the ACO may behave in this strategic manner. We briefly discuss relevant literature and potential approaches to modifying the benchmarking methodology to mitigate the investment distortions that might arise from such strategic behavior. We also present a result that shows that for a forward looking ACO, the deviation in savings from the myopic strategy vanishes over time under the existing benchmarking mechanism.

The notion that an agent might find it optimal to strategically delay effort in multi-period incentive problems is well documented in the operations literature. Chen (2000) consider a salesforce compensation contract where agents receive payments based on an annual sales quota system. The author shows how the "sales-hockey stick" phenomenon, where agent effort is delayed to the final period of the horizon, can arise in this setting. Sohoni et al. (2010) also consider threshold-based contracts, and show that adjusting the sales threshold based on a relevant market signal can mitigate the sales-hockey stick phenomenon. In a similar vein, Besbes et al. (2016) consider how the presence of debt in a dynamic pricing setting can lead to price distortions and efficiency losses.

One approach to regulating the strategic underinvestment of a forward-looking ACO might be to set the financial benchmarks of an ACO according to the spending patterns of comparable ACOs, rather than setting benchmarks for each ACO individually based on their historical spending. In a seminal paper, Shleifer (1985) discusses how regulated local monopolies have little incentive to reduce costs in an environment where the regulator sets prices according to the firm's costs. The author proposes a mechanism known as *yardstick competition* whereby the regulator sets prices for a firm according to the costs of identical firms. This mechanism has the effect of inducing competition between firms that serve different markets. With respect to the MSSP, incorporating elements of yardstick competition when setting the financial benchmarks of ACOs may be useful in disincentivizing strategic underinvestment. However, while a purely yardstick competition approach to setting the ACO benchmarks might dampen underinvestment, due to the voluntary nature of the MSSP contract it may also jeopardize ACO participation if some ACOs find the benchmarks to be unattainable, which is an undesirable outcome for Medicare. We therefore posit that a benchmarking methodology that is based on a blend of the ACO's own historical spending and yardstick competition might be a more fruitful approach to mitigating strategic delay of investment, while also encouraging ACO participation, and underline this as a direction for future work. We also note that yardstick competition for controlling costs is not new to Medicare. For example, the Medicare Prospective Payment System (PPS) reimburses providers based on the costs of comparable hospitals (Fetter, 1991). This program was created precisely in response to the adverse incentives that arise from purely cost-based reimbursement systems. For a more recent example, Savva et al. (2016) consider a modification of traditional cost-based yardstick competition in the context of a queueing system, motivated by the problem of reducing wait times in emergency departments.

It is also worth noting that like many of the other recent innovations in Medicare provider payment systems, the MSSP was also created to realign incentives and reduce provider spending. In contrast to other Medicare programs, however (e.g., the PPS), the MSSP emphasizes self-improvement of each ACO independently, rather than competition between ACOs. While

the existing MSSP benchmarking mechanism may allow ACOs to strategically delay investment, and can potentially be improved upon, our analysis suggests that a revised contract with performance-based subsidies can still lead to overall cost improvements. In the remainder of this section, we consider a simple multi-period model for the MSSP in which the ACO's benchmark is updated based on its savings in the previous year. We find that the ACOs distortion in its optimal savings level relative to a myopic policy diminishes over time.

Consider a single ACO with type θ that faces a multiple-period savings problem with an infinite horizon, where $t = 0, 1, 2, \ldots,$ indexes each period. Let α and β be fixed throughout, which reflects the fact that the MSSP uses the same shared savings formula each year. We suppress θ, α and β in the notation. Let $\mu_0, \mu_1, \mu_2, \ldots,$ be the benchmarks in each period, and let $x_0, x_1, x_2, \ldots,$ be the ACO's decision variables, which represents the ACO's reduction in spending. We emphasize that x_t represents the current-period savings relative to the relevant benchmark μ_t. Define $s_t = \sum_{i=0}^{t} x_i$, so that $s_0, s_1, s_2, \ldots,$ is the cumulative savings up to and including period t. We assume that the cost of generating savings x_t in period t is $c(s_{t-1} + x_t)$, where $c(\cdot)$ is the single period cost function described in Assumption 3.1. In other words, the investment required in period t to reduce spending down to a level of $\mu_t - x_t$ is measured with respect to the expected spending in the baseline "do-nothing" case, which is given by μ_0. We note here that if the investment required to achieve a savings of x_t were $c(x_t)$, then it can be shown that the ACO would simply behave myopically each year.

In the MSSP, the benchmarks are updated based on the previous period's benchmark and the savings generated by the ACO. The benchmark only decreases if the ACO generates positive savings, i.e., the benchmark is never increased when ACO spending exceeds it (Federal Register, 2011). To reflect the MSSP's current practice, for $t \geq 2$ let $\mu_t = \mu_{t-1} - x_{t-1}$. The initial benchmark, $\mu_0 = \mu$, is determined based on the ACO's historical spending (Federal Register, 2011). Note that s_t is then given by $s_t = \mu_0 - \mu_t + x_t$. Next, let $y_t = x_t + \varepsilon$ be the realized savings, as in the single period model. The ACO's total expected payment in period t is then given by $P_t(x_t) = \int_{-\infty}^{\infty} [r(y_t, \alpha) + s(y_t, \beta)]\, h(y_t|x_t) dy_t$, and the ACO's profit in period t is then $u_t(x_t) = P_t(x_t) - \gamma x_t - c(s_{t-1} + x_t)$. Let $\delta < 1$ be the discounting factor. The ACO's multiple-period optimization problem is then given by

$$\begin{aligned}
\underset{\mathbf{x}}{\text{maximize}} \quad & \sum_{t=0}^{\infty} \delta^t \left(P_t(x_t) - \gamma x_t - c(s_{t-1} + x_t) \right) \\
\text{subject to} \quad & s_t = \mu_0 - \mu_t + x_t, \quad t = 1, 2, \ldots, \\
& \mu_t = \mu_{t-1} - x_t, \quad t = 1, 2, \ldots, \\
& \mu_0 = \mu, \\
& s_0 = 0, \\
& x_t \in [0, \bar{x}], \quad t = 0, 1, 2, \ldots,
\end{aligned}$$

Let the optimal solution to the above problem be $(\tilde{x}_0, \tilde{x}_1, \tilde{x}_2, \ldots,)$. If the ACO were to reduce spending myopically, then in each period t it would simply solve the single period problem

$$\begin{aligned}
\underset{x_t}{\text{maximize}} \quad & P_t(x_t) - \gamma x_t - c(s_{t-1} + x_t) \\
\text{subject to} \quad & x_t \in [0, \bar{x}],
\end{aligned}$$

where $s_{t-1} = \sum_{i=0}^{t-1} x_i$. Let the optimal solution to the myopic problem be $(x_0^*, x_1^*, x_2^*, \ldots)$. We can now interpret $|\tilde{x}_t - x_t^*|$ as the distortion due to the ACOs strategic saving. The following result shows that the distortion

Proposition 3.6. *In the multiple-period setting, the difference in ACO savings under myopic and forward-looking behavior vanishes,* $\lim_{t\to\infty} |\tilde{x}_t - x_t^*| = 0$.

Proposition 3.6 is a straightforward result – it shows that the deviation from the myopic behavior in a multi-period setting must vanish because the ACO savings itself under both myopic and forward-looking behavior goes to 0. This result follows from the fact that the monotonic reduction in the financial benchmarks within the MSSP makes it increasing costly for ACOs to receive bonus payments. Therefore, while the existing benchmarking mechanism may allow ACOs to strategically delay investment, this behavior is unlikely to persist in the long term.

3.7.2 Data-based validation of Assumptions 3.3 and 3.4

Here we show that Assumptions 3.3 and 3.4 are validated by the dataset. First we provide values of the model parameters based on the data such that Assumption 3.3 holds. Specifically, we wish to show that $\bar{\alpha}(h\bar{g}' + \bar{g}) + 2\bar{\beta}\bar{g}(\partial/\partial x)c(x,\theta) \leq (1-\bar{\beta})(\partial^2/\partial x^2)c(x,\theta)$ holds for all $x \in [0, \bar{x}]$ and $\theta \in \Theta$. We shall consider the model and parameter estimates from Sections 3.5 and 4.6. Since we use $c(x,\theta) = x^2/\theta$ for the ACO investment function, we have $(\partial/\partial x)c(x,\theta) = 2x/\theta$ and $(\partial^2/\partial x^2)c(x,\theta) = 2/\theta$. Let $\bar{\alpha} = 0.55$, $\bar{\beta} = 0.05$, $\bar{x} = 1,000$, $\underline{\theta} = 1$ and $\bar{\theta} = 2,000$. Based on the guidelines of the MSSP, the parameter h is approximately 200 (2% of the benchmark). Since the shock density g is assumed to be Laplace distributed with an estimated standard deviation of $\hat{\sigma} = 550$, this corresponds to $\bar{g} = 1/(\sqrt{2}\hat{\sigma}) = 0.0013$ and $\bar{g}' = (1/\hat{\sigma}^2) = 2.8 \times 10^{-6}$. Using these parameter values, we can numerically verify that the inequality holds for all $x \in [0, \bar{x}]$ and $\theta \in \Theta$. Next, for Assumption 3.4, we have

$$\begin{aligned}
(\partial/\partial x)R(x,\alpha) &= (1-\alpha)\left(hg(-h-x) + G(-h-x) - G(C_\ell - x)\right) \\
&\quad + \alpha\left(hg(h-x) + G(C_u - x) - G(h-x)\right) \\
&< (1-\alpha)\left(hg(-h-x) + 1/2\right) + \alpha\left(hg(h-x) + 1\right) \\
&\leq (1-\alpha)\left(hg(-h) + 1/2\right) + \alpha\left(h\bar{g} + 1\right)
\end{aligned}$$

The first line is given by the expression for $(\partial/\partial x)R(x,\alpha)$ given in Lemma A.3. The second line follows from noting that $G(-h-x) - G(C_\ell - x) \leq 1/2$ (since $-h-x < 0$) and $G(C_u - x) - G(h-x) \leq 1$. The third line follows from the fact that $\bar{g} \geq g(h-x)$ and $g(-h) \geq g(-h-x)$ for $x \geq 0$. For Assumption 4 to hold, we require $(1-\alpha)\left(hg(-h) + 1/2\right) + \alpha\left(h\bar{g} + 1\right) < 1$ for all $\alpha \in \mathcal{A}$, which holds if $\bar{\alpha} \leq (1 - 2hg(-h))/(1 + 2h(\bar{g} - g(-h)))$. The inequality can be shown to hold for $\bar{\alpha} = 0.55$, $h = 200$ and $\hat{\sigma} = 550$, where $\bar{g} = 1/(\sqrt{2}\hat{\sigma})$ and $g(-h) = 1/(\sqrt{2}\hat{\sigma})e^{\sqrt{2}|-h|/\hat{\sigma}}$.

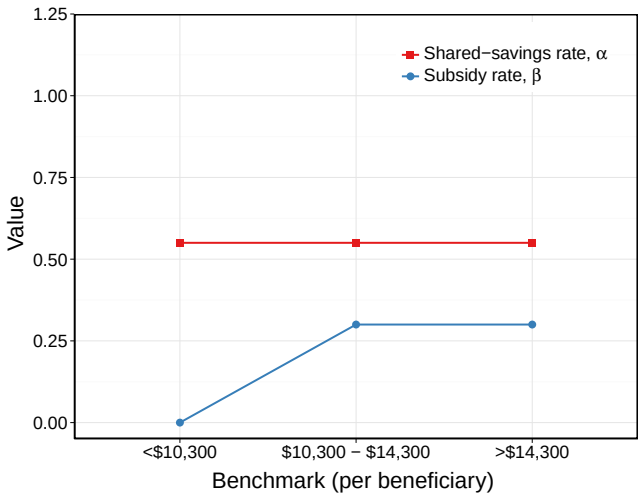

Figure 3.7: Optimal contract parameters in benchmark-only contract.

Table 3.7: Bootstrap estimates for subsidized contract under noisy benchmark observations (in millions).

		Baseline		II-S		Δ		
μ-Error		Mean	95% C.I.	Mean	95% C.I.	Mean	95% C.I.	p
0%	ACOs	\$282	(\$188, \$382)	\$316	(\$216, \$423)	\$34	(\$22, \$49)	<0.01
	Medicare	\$146	(\$39, \$260)	\$207	(\$97, \$328)	\$62	(\$48, \$79)	<0.01
	Total	\$427	(\$234, \$642)	\$523	(\$320, \$741)	\$96	(\$78, \$118)	<0.01
5%	ACO	\$309	(\$194, \$421)	\$337	(\$218, \$452)	\$27	(\$15, \$41)	<0.01
	Medicare	\$111	(−\$26 , \$219)	\$172	(\$31, \$292)	\$60	(\$44, \$82)	<0.01
	Total	\$420	(\$169, \$639)	\$508	(\$251, \$723)	\$88	(\$68, \$111)	<0.01

Table 3.8: Bootstrap estimates with partial observability of ACO investment (in millions)

		Baseline		II-S		Δ		
Contract		Mean	95% C.I.	Mean	95% C.I.	Mean	95% C.I.	p
θ-Based	ACOs	\$282	(\$188, \$382)	\$316	(\$216, \$423)	\$34	(\$22, \$49)	<0.01
	Medicare	\$146	(\$39, \$260)	\$207	(\$97, \$328)	\$62	(\$48, \$79)	<0.01
	Total	\$427	(\$234, \$642)	\$523	(\$320, \$741)	\$96	(\$78, \$118)	<0.01
μ-Based	ACOs	\$278	(\$182, \$375)	\$332	(\$232, \$437)	\$54	(\$40, \$73)	<0.01
	Medicare	\$127	(\$17, \$233)	\$187	(\$72, \$301)	\$60	(\$45, \$79)	<0.01
	Total	\$406	(\$202, \$602)	\$519	(\$308, \$734)	\$124	(\$99, \$158)	<0.01

Chapter 4

A Distributionally Robust Optimization Approach to Facility Location

4.1 Background and motivation

Sudden cardiac arrest is a leading cause of death in North America, and is responsible for over 400,000 deaths each year (Heart and Stroke Foundation, 2013, Mozaffarian et al., 2016). Chances of survival have been estimated to decrease by as much as 10 percent with each minute of delay in treatment (Larsen et al., 1993), making treatment of cardiac arrest extremely time sensitive. Currently, the probability of survival from an out-of-hospital cardiac arrest is low, with only 5-10% of out-of-hospital cardiac arrest victims surviving to hospital discharge (Weisfeldt et al., 2010).

Treatment of cardiac arrest involves cardiopulmonary resuscitation (CPR) and electrical shocks – known as defibrillation – from an automated external defibrillator (AED). AEDs are portable devices that can automatically assess heart rhythms and perform defibrillation if necessary, making their use by non-professionals (lay responders) in a cardiac arrest emergency viable (Gundry et al., 1999). Treatment with a defibrillator following cardiac arrest is vital: the probability of survival from cardiac arrest has been estimated to be as high as 40-70% if defibrillation is administered within three minutes of the victim's collapse (Valenzuela et al., 2000, Page et al., 2000, Caffrey et al., 2002). As a consequence, there is growing interest in the development of public access defibrillation (PAD) programs, which are organized efforts to place defibrillators in public areas, such as malls, coffee shops, restaurants, and subway stations, with the hopes of improving easy access to defibrillators in the event of a cardiac arrest emergency. With appropriate positioning, public AEDs can enable lay responders to treat cardiac arrest victims while awaiting the arrival of professional emergency medical responders. Indeed, PAD programs have been shown to decrease the time to treatment and improve survival outcomes (PAD Trial Investigators, 2004), and it is estimated that their widespread implementation may save between 2,000 and 4,000 lives annually in the United States (Hazinski et al., 2005). Thus, the strategic positioning of public AEDs has a crucial role to play in strengthening the overall response to cardiac arrest.

Despite evidence that supports improved chances of survival with PAD programs, AEDs

are used infrequently. For example, only 8% of public location cardiac arrests in Toronto, Canada involve the use of an AED by a bystander (Sun et al., 2016). Although the reasons for this low usage rate are multi-faceted, it is clear that AEDs that are poorly positioned may be located too far from the locations of cardiac arrest emergencies for them to be accessed and used by lay responders. Recently, several mobile phone applications have been developed that aim to improve AED usage during public location cardiac arrests. These applications notify volunteer lay responders of the location of nearby AEDs so that they can be retrieved and used quickly (e.g. PulsePoint (2015)). In order to be effective, these emerging technologies – and public access defibrillation in general – rely critically on AEDs being placed strategically in public locations. The question of where to place public access AEDs remains relatively open in the medical literature (Portner et al., 2004, Folke et al., 2009).

In this chapter, we present a data-driven optimization model for deploying AEDs in public locations. We consider a setting where cardiac arrests occur continuously over a bounded service area. Our model addresses a key challenge in the AED deployment problem, which is the uncertainty in the locations of future cardiac arrests during the deployment of AEDs. Since the chances of survival from cardiac arrest diminish rapidly with each passing minute, it is critical that witnesses of cardiac arrest are able to quickly retrieve and use nearby AEDs when needed. Due to this extreme time sensitivity, protecting against cardiac arrest location uncertainty can help improve survival outcomes by shortening the distance (and thus the response time) to the nearest AED during a cardiac arrest emergency.

Although cardiac arrests are not uncommon within the general population, they are extremely rare events for any *particular* building or geographic location. Thus, any estimate of the risk of cardiac arrest at a specific location is likely to be subject to considerable uncertainty. To account for this uncertainty while still leveraging historical data, we take a distributionally robust optimization approach in which service is optimized for the worst-case spatial distribution of cardiac arrests. Our approach involves discretizing the continuous service area into a large set of scenarios, where the probability vector that describes the probability of the next cardiac arrest arriving at each scenario is only known to reside within a polyhedral uncertainty set. We construct the uncertainty set by representing the service area as the union of a set of *uncertainty regions*, where only the aggregate probability of the next cardiac arrest arriving in each uncertainty region (i.e., group of scenarios) is known. An important consequence of modeling uncertainty in this manner is that it induces sparsity in the worst-case distribution (with respect to the number of scenarios), which we show can be exploited through an efficient row-and-column algorithm to obtain optimal solutions to the robust problem.

For a given AED deployment and a set of cardiac arrest locations, one can construct a *distance distribution* which describes the histogram of distances between the cardiac arrests and the nearest AED in the given deployment. Inspired by the use of quantile-based targets in emergency medical services (Pell et al., 2001, Pons and Markovchick, 2002), we use a conditional value-at-risk (CVaR) objective function to directly optimize the tail of the worst-case distance distribution, which allows us to mitigate the risk of unacceptably long distances between cardiac arrest locations and their nearest AEDs. An additional benefit of a CVaR-based objective function is that it permits more flexibility with respect to selecting a performance metric to optimize, e.g., we can minimize the mean or tail of the worst-case distribution by using an appropriate parameterization of the model.

4.2 Related literature

This chapter primarily contributes to the operations research literature on facility location under uncertainty and the medical literature on cardiac arrest and public access defibrillation. We also briefly review the connection between our model and the general literature on optimization under uncertainty.

Connection to Siddiq (2013). This chapter expands significantly on the work of Siddiq (2013). That work considers robust analogues to the p-median and p-center problems in a setting where each demand point is only known to lie within a planar uncertainty region. By contrast, the model presented in this chapter considers a continuous demand distribution in the plane and uses polyhedral uncertainty sets to model distributional uncertainty. Indeed, it can be shown that the model presented in Siddiq (2013) is a special case of the modeling approach taken in this chapter. Siddiq (2013) also proposes a decomposition algorithm for solving a two-stage robust optimization problem. However, that approach taken in that work involves a simple subproblem that can be solved using a sort, whereas the algorithm in this chapter contains a more general linear programming subproblem. In addition, the algorithm in Siddiq (2013) does not require the innermost minimization problem of the two-stage robust optimization model to be dualized due to the special structure of the problem considered.

Public access defibrillation. Most of the literature in the medical community on AED location has focused on identifying building types that have a high incidence of cardiac arrest, such as shopping malls (Becker et al., 1998, Engdahl and Herlitz, 2005, Brooks et al., 2013). However, relying on building type alone to guide AED locations has several drawbacks. First, it does not account for cardiac arrests that occur outdoors, since these events cannot be assigned to a building type. Second, a high risk building type in one city may not be high risk in another. Lastly, identifying high risk buildings does not address the question of how to optimally deploy a limited number of defibrillators. By contrast, our proposed approach is guided by historical cardiac arrest data and is agnostic to building type, which makes it generalizable to any city.

There is a recent but limited body of work on optimization-based approaches to placing AEDs in public locations (Siddiq et al., 2013, Chan et al., 2013, 2014, Sun et al., 2016). This work differs from and extends this literature in a few substantive ways. First, previous approaches have solely optimized over exact historical cardiac arrest locations. While these approaches serve as a good first step toward improving public access defibrillation, they do not explicitly account for uncertainty in future cardiac arrest locations. This work is the first to explicitly incorporate the uncertainty in cardiac arrest locations into the optimization model. We find that accounting for cardiac arrest location uncertainty leads to improved outcomes based on several performance metrics, compared to a baseline method of optimizing over historical locations only. Second, existing AED models focus on maximizing cardiac arrest *coverage*, which is the number of cardiac arrests that are within some fixed distance (e.g., 100m) of an AED. While this approach can be effective for the *median* cardiac arrest patient, it can have the adverse effect of neglecting cardiac arrest victims at the tail of the distance distribution. By contrast, our modeling approach involves a risk-based objective function (inspired by ambulance response time guidelines), which allows us to mitigate the risk of unacceptably large distances between cardiac arrest victims and the nearest AED. Lastly, existing work on AED location has focused on identifying general hotspots for cardiac arrest within a large service area (e.g., Chan et al. (2014) identify historical clusters of cardiac

arrests within a city). We complement this work by taking a more tactical approach to the AED deployment problem, by focusing on the precise placement of AEDs within a small service area consisting of several city blocks.

Facility location under uncertainty. Facility location problems under uncertainty have received considerable attention, particularly with respect to uncertainty in demand node weights and edge lengths (Owen and Daskin, 1998, Snyder, 2006, Baron et al., 2011), and more recently in the risk of service disruptions at the facilities (Lim et al., 2010, Cui et al., 2010, Shen et al., 2011). Our model differs from the existing literature in that we consider the arrival of demand points over a continuous, planar service area, rather than being restricted to a small set of locations. As a result, we do not aggregate demand into a small set of known locations, which is a common approach in facility location models related to emergency response (Cho et al., 2014, Erkut et al., 2008, Brotcorne et al., 2003) and in general (Francis et al., 2000). Previous work on demand location uncertainty in facility location problems is limited and has generally focused on the placement of a single facility (Cooper, 1978, Drezner, 1989, Averbakh and Bereg, 2005).

This work is most closely related to the paper by Baron et al. (2011), who also present a robust optimization model for facility location. In their paper, demand is again restricted to a set of known locations, but the vector of demands over all locations is only known to live within a given uncertainty set. In contrast, we model total demand as being known (i.e., normalized to 1), and consider uncertainty in how this demand is distributed in space.

Our work is also related to facility interdiction problems, which focus on protecting against potential disruptions at facilities. Facility interdiction models, which, like ours, take the form of min-max optimization problems, were first considered by Church and Scaparra (2007) and later by Scaparra and Church (2008), Liberatore et al. (2011) and Losada et al. (2012). With respect to CVaR, the only other work to use it in a facility location context is the *mean-excess* model presented in Chen et al. (2006), which focuses on demand level uncertainty.

Optimization under uncertainty. Our main formulation takes the form of a robust optimization problem in which a large set of scenario probabilities are uncertain and only known to reside within a given polyhedron. In the special case where there is no uncertainty (i.e. the uncertainty set is a singleton representing the true distribution), our formulation simplifies to an extensive form stochastic program with a large set of scenarios (Birge and Louveaux, 2011). Other papers which model uncertainty in probabilities using polyhedral uncertainty sets include Chan and Mišić (2013), who consider radiation therapy with uncertainty in patient breathing patterns, and Farias et al. (2013), who consider uncertainty in consumer choice probabilities. Ben-Tal et al. (2013) present a general framework which uses ϕ-divergences to model distributional uncertainty sets, and obtain polyhedral sets as a special case. For a review of the broader robust optimization literature, we refer the reader to Ben-Tal et al. (2009) and Bertsimas et al. (2011).

Since the source of uncertainty in our model stems from an unknown discrete probability distribution, our model can also be interpreted as an instance of distributionally robust optimization, which refers to optimization problems in which only a partial description of a probability distribution is available (Scarf et al., 1958, Dupačová, 1980, Birge and Wets, 1987, Shapiro and Kleywegt, 2002, Delage and Ye, 2010, Goh and Sim, 2010, Xu et al., 2012). Recently, Wiesemann et al. (2014) proposed a canonical framework for distributionally robust convex optimization, and show that it subsumes well-known robust optimization models. A

common approach to modeling distributional uncertainty sets is to impose constraints on the moments of the distribution (e.g., Ghaoui et al. (2003), Delage and Ye (2010)). Since we instead model distributional uncertainty using polyhedral sets, from a technical perspective our model is more aligned with the literature on robust linear optimization (e.g., Bertsimas and Sim (2004)).

Due to our use of a CVaR objective function, our main formulation takes the form of a two-stage (min-max-min) robust optimization model, where the inner minimization problem is used to compute CVaR. Gabrel et al. (2014) and Zeng and Zhao (2013) also present two-stage models for facility location problems, where the focus is on uncertainty in demand level rather than location. Other settings in which two-stage robust optimization models have been considered include network flows (Atamturk and Zhang, 2007, Ordonez and Zhao, 2007), power systems (Bertsimas et al., 2015, Zhao and Zeng, 2012), and military applications, including the defense of critical infrastructure (Brown et al., 2006, Alderson et al., 2011) and network interdiction (Brown et al., 2009).

Lastly, our modeling approach has some similarities with work by Carlsson and Delage (2013), in that they also consider a distributionally robust optimization problem in a spatial setting where demand points arrive in a continuous service area. Assuming the mean and covariance of the demand distribution is known, they focus on a vehicle routing problem where the goal is to partition the service area into a set of subregions such that the worst-case load for any vehicle across all subregions is minimized. By comparison, we assume that a fixed set of subregions is given, and the only distributional information available is the probability of an arrival in each subregion.

We note here that the work in this chapter coincides with the paper by Chan et al. (2018).

4.3 Model

Let $\mathcal{A} \subset \mathbb{R}^2$ represent the continuous service area over which cardiac arrests arrive. Let $\mathcal{I}$ be an index set for m candidate sites, and let $\mathbf{y} \in \{0,1\}^m$ be a binary vector, where $y_i = 1$ indicates the presence of an AED in the i^{th} candidate site. Let the function $d : \mathcal{A} \times \mathcal{A} \to \mathbb{R}_+$ measure the distance between two locations in $\mathcal{A}$. With a slight abuse of notation, let $\mathbf{y}(a) \in \mathcal{A}$ represent the nearest facility to a point $a \in \mathcal{A}$, so that $d(a, \mathbf{y}(a))$ is the distance between a point $a \in \mathcal{A}$ and its nearest facility.

Let ξ be a random vector representing the location of the next cardiac arrest event within $\mathcal{A}$, with distribution μ. The distribution μ can be interpretted as describing the risk of a cardiac arrest event occurring over $\mathcal{A}$. Without loss of generality, we assume $P_\mu(\xi \in \mathcal{A}) = 1$. For a set of AED locations $\mathbf{y}$, $d(\xi, \mathbf{y}(\xi))$ represents the distance between the next demand arrival and its nearest facility. Since $d(\xi, \mathbf{y}(\xi))$ is itself random due to ξ, we shall refer to the distribution of $d(\xi, \mathbf{y}(\xi))$ as the *distance distribution* induced by $\mathbf{y}$. We can think of the distribution of $d(\xi, \mathbf{y}(\xi))$ as fully describing the performance of the AED deployment.

As discussed in Section 1, we wish to place AEDs in a manner that mitigates the risk of large distances between cardiac arrest victims and the nearest AED. Restated, our goal is to identify an AED deployment $\mathbf{y}$ that controls the right-tail of the distance distribution, $d(\xi, \mathbf{y}(\xi)$. Since placing an AED in every possible candidate site is prohibitively expensive, we assume throughout that up to P AEDs are available for deployment. If the distribution μ is known, then the problem of minimizing the CVaR of the distance distribution is given

by

$$\underset{\mathbf{y}\in\mathbf{Y}}{\text{minimize}} \;\; \text{CVaR}_{\mu}[d(\xi, \mathbf{y}(\xi))] \tag{4.1}$$

where

$$\mathbf{Y} = \left\{ \mathbf{y} \in \{0,1\}^m \;\middle|\; \sum_{i\in\mathcal{I}} y_i = P \right\}.$$

CVaR is closely related to the value-at-risk (VaR) measure, both of which have origins in the finance literature (see Rockafellar and Uryasev (2000, 2002)). For a given loss distribution, β-VaR represents the smallest value α such that the loss does not exceed α with probability $1-\beta$. By contrast, β-CVaR represents the expected loss conditional on the loss exceeding β-VaR.

In practice, the distribution of cardiac arrest locations is unlikely to be known precisely. Since – for a given location – cardiac arrests are low probability events, optimizing AED locations solely with respect to limited historical data (effectively a sample average approximation) may result in a poor deployment of AEDs. Moreover, any estimate of μ from historical data is also likely to be subject to uncertainty, especially if a limited amount of data is available.

To account for this uncertainty, we instead assume that the distribution μ is only known to belong to an uncertainty set, $\mathcal{U}$. A key question at this juncture is how to structure the set $\mathcal{U}$. To obtain a formulation that is both tractable and effective at capturing uncertainty in μ, we take the following approach. Suppose the service area $\mathcal{A}$ can be divided into (possibly overlapping) *uncertainty regions* $\mathcal{A}_1, \mathcal{A}_2, \ldots, \mathcal{A}_{|\mathcal{J}|}$, where $\mathcal{A} = \mathcal{A}_1 \cup \mathcal{A}_2 \cup \ldots \cup \mathcal{A}_{|\mathcal{J}|}$. Suppose now that the only available information about the distribution μ is the probability of the next cardiac arrest occuring in each of the regions $\mathcal{A}_1, \ldots, \mathcal{A}_{|\mathcal{J}|}$, which we denote by $\lambda_1, \ldots, \lambda_{|\mathcal{J}|}$. The worst-case CVaR can now be minimized by solving the problem

$$\min_{\mathbf{y}\in\mathbf{Y}} \max_{\bar{}\in\mathcal{U}} \;\; \text{CVaR}_{\mu}[d(\xi, \mathbf{y}(\xi))], \tag{4.2}$$

where

$$\mathcal{U} = \{\mu \mid \mathbb{P}_{\mu}(\xi \in \mathcal{A}_1) = \lambda_1, \ldots, \mathbb{P}_{\mu}(\xi \in \mathcal{A}_{|\mathcal{J}|}) = \lambda_{|\mathcal{J}|}\},$$

is the distributional uncertainty set. Note that since the uncertainty regions are permitted to overlap, the probabilities $\lambda_1, \ldots, \lambda_{|\mathcal{J}|}$ may sum to a value greater than 1. We assume throughout that the uncertainty regions are given. In practice, the design of the uncertainty regions $\mathcal{A}_j$ should be guided by several considerations such as tractability, availability of historical data, and model interpretability (we discuss these considerations in greater detail in Section 4.5).

There are several reasons behind our choice of the uncertainty set $\mathcal{U}$. First, we posit that obtaining estimates of the cardiac arrest risk at an aggregate (i.e., uncertainty region) level is less onerous than estimating the distribution μ directly, especially if data on historical cardiac arrest locations is limited. Estimating the parameters $\lambda_1, \ldots, \lambda_{|\mathcal{J}|}$ in lieu of μ itself allows us to partially capture the underlying distribution while accounting for uncertainty. This approach is well-aligned with the notion that aggregate forecasts tend to be more reliable

than point forecasts when dealing with uncertainty (Simchi-Levi et al., 2004, Sheffi, 2005, Bertsimas and Thiele, 2006, Nahmias and Cheng, 2009). In addition, our choice for the structure of $\mathcal{U}$ lends itself to a tractable mathematical programming formulation of (4.2). The key to this tractability is the observation that for any $\mathbf{y}$, the worst-case distribution in $\mathcal{U}$ is supported on a relatively small number of locations in $\mathcal{A}$. This sparsity can in turn be exploited through a decomposition algorithm, which allows us to solve the robust optimization problem efficiently (see Section 4). Our choice of $\mathcal{U}$ also provides a natural generalization of classical facility location models, which have a rich and extensive history. In particular, classical models typically assume demand is concentrated at n weighted demand points (e.g., a weight of λ_j for demand point $j = 1, \ldots, n$), which is recovered from our model by simply shrinking each $\mathcal{A}_j$ to a singleton. Thus, a natural interpretation of our distributionally robust model is as an extension of classical models where each demand point lives in a region $\mathcal{A}_j$, wherein the true demand location will be realized after facility siting. We discuss the connection between our model and the classical location literature in more detail in Section 4.3.4. In particular, we show that our model unifies the classical p-median and p-center problems, as well as their robust analogues.

We note here that our choice of uncertainty set bears similarities to the uncertainty set discussed in the paper by Wiesemann et al. (2014), who present a canonical modeling approach for distributionally robust optimization problems. The authors formulate the uncertainty set using constraints on the probability that a random vector is realized within various convex sets. They show that within their framework, simply checking the whether the distributional uncertainty set is empty can be a strongly NP-hard problem, unless the constituent convex sets satisfy a particular nesting condition. This hardness follows from the generality of the uncertainty set, which can be shown to subsume integer programming, in the absence of the nesting condition. Our setting is considerably more structured. By discretizing the service area into a large number of scenarios, we obtain a representation of $\mathcal{U}$ that is polyhedral. As a consequence, the worst-case distribution can be identified by solving a linear program.

In the remainder of this section, we show how our discretization approach allows us to formulate the optimization problems (4.1) and (4.2) as tractable mathematical programs. We then propose a bound on the error introduced by the discretization.

4.3.1 Known demand distribution

For ease of exposition, we first consider the case where μ is known. A standard approach for formulating stochastic programs is to discretize the outcome space into a finite set of *scenarios* and to assign a probability to each scenario. Since in our setting, ξ is the random location of the next cardiac arrest event, each scenario (i.e., realization of ξ) has a one-to-one correspondence with a particular geographic location in $\mathcal{A}$. Accordingly, let $\Xi \subset \mathcal{A}$ be a discrete set of locations that approximate the continuous service area $\mathcal{A}$, so that Ξ is a discrete approximation of the outcome space of ξ, and each element of Ξ corresponds to a distinct location in $\mathcal{A}$. Letting $\mathcal{K}$ be an index set for the scenarios in Ξ, we can write $\Xi = \left\{\xi_1, \xi_2, \ldots, \xi_{|\mathcal{K}|}\right\}$ to represent the set of all possible realizations of ξ. Let u_k be the probability that the next cardiac arrest will occur at location ξ_k. Since in this discrete setting the possible cardiac arrest locations are restricted to Ξ, we have $\sum_{k\in\mathcal{K}} u_k = 1$.

Let the parameter d_{ik} represent the distance between candidate site i and scenario ξ_k based on the distance function $d(\cdot, \cdot)$. Let z_{ik} be an assignment variable that is equal to 1 if

a cardiac arrest realized at location ξ_k is assigned to an AED at location i. The feasible set of assignments is given by

$$\mathbf{Z}(\mathbf{y}) = \left\{ z_{ik} \geq 0,\ i \in \mathcal{I}, k \in \mathcal{K} \;\middle|\; \sum_{i \in \mathcal{I}} z_{ik} = 1, k \in \mathcal{K};\ z_{ik} \leq y_i,\ i \in \mathcal{I}, k \in \mathcal{K} \right\}.$$

Note that based on the discretization, we have $d(\xi_k, \mathbf{y}(\xi_k)) = \min_{\mathbf{z} \in \mathbf{Z}(\mathbf{y})} \sum_{i \in \mathcal{I}} d_{ik} z_{ik}$. Hence, for fixed $\mathbf{y}$, $\text{CVaR}_\mu[d(\xi, \mathbf{y}(\xi))]$ can be expressed as (Rockafellar and Uryasev, 2000, 2002):

$$\text{CVaR}_\mu[d(\xi, \mathbf{y}(\xi))] \approx \underset{\mathbf{z}, \alpha}{\text{minimize}}\ \ \alpha + \frac{1}{(1-\beta)} \sum_{k \in \mathcal{K}} u_k \max\left\{ \sum_{i \in \mathcal{I}} d_{ik} z_{ik} - \alpha, 0 \right\}. \tag{4.3}$$

Strictly speaking, the quantity in (4.3) is known as β-CVaR$^-$, which is an approximation of β-CVaR. This is a commonly used approximation for the case of discrete distributions – see Rockafellar and Uryasev (2000) for further details. The nominal optimization problem (4.1) for a known distribution can now be written as

$$\begin{aligned} \underset{\mathbf{y}, \mathbf{Z}, \alpha}{\text{minimize}}\ \ & \alpha + \frac{1}{(1-\beta)} \sum_{k \in \mathcal{K}} u_k \max\left\{ \sum_{i \in \mathcal{I}} d_{ik} z_{ik} - \alpha, 0 \right\} \\ \text{subject to}\ \ & \mathbf{z} \in \mathbf{Z}(\mathbf{y}),\ \mathbf{y} \in \mathbf{Y}, \\ & \alpha \geq 0, \end{aligned} \tag{N-AED}$$

which is a mixed-integer linear program after linearizing the maximization term in the objective.

4.3.2 Distributionally robust model for unknown demand distribution

We now extend formulation N-AED to accomodate uncertainty in the distribution μ. Given the uncertainty regions $\mathcal{A}_1, \ldots, \mathcal{A}_{|\mathcal{J}|}$, we construct scenario sets $\Xi_1, \Xi_2, \ldots, \Xi_{|\mathcal{J}|}$, where

$$\Xi_j = \{\xi_k, k \in \mathcal{K} \mid \xi_k \in \mathcal{A}_j\}$$

and $\bigcup_{j=1}^n \Xi_j = \Xi$. Let $\mathcal{K}_j$ index the scenarios in Ξ_j, and let $\mathcal{J}$ index the set of uncertainty regions. Since in the discretized setting we assume a cardiac arrest arriving in $\mathcal{A}_j$ can only be realized at one of the locations in Ξ_j, we have $P_\mu(\xi \in \Xi_j) = P_\mu(\xi \in \mathcal{A}_j) = \lambda_j$. The worst-case CVaR of the distance distribution given a feasible $\mathbf{y}, \mathbf{z}$ is then given by optimal value of the

following max-min problem:

$$\begin{aligned} \max_{\mathbf{u}} \min_{\alpha} \ & \alpha + \frac{1}{(1-\beta)} \sum_{k \in \mathcal{K}} u_k \max \left\{ \sum_{i \in \mathcal{I}} d_{ik} z_{ik} - \alpha, 0 \right\} \\ \text{subject to } & \sum_{k \in \mathcal{K}} u_k = 1, \\ & \sum_{k \in \mathcal{K}_j} u_k = \lambda_j, \quad j \in \mathcal{J}, \\ & u_k \geq 0, \quad k \in \mathcal{K}, \\ & \alpha \geq 0. \end{aligned} \tag{4.4}$$

For conciseness, define

$$\mathbf{U} = \left\{ \mathbf{u} \in \mathbb{R}_+^{|\mathcal{K}|} \;\middle|\; \sum_{k \in \mathcal{K}} u_k = 1, \sum_{k \in \mathcal{K}_j} u_k = \lambda_j, \; j \in \mathcal{J} \right\}.$$

Note that the set $\mathbf{U}$ may be empty if the parameters $\lambda_1, \ldots, \lambda_{|\mathcal{J}|}$ are selected arbitrarily (e.g., consider two regions $\mathcal{A}_1$, $\mathcal{A}_2$ where $\mathcal{A}_1 \cap \mathcal{A}_2 = \emptyset$ but $\lambda_1 = \lambda_2 = 1$). However, if λ is estimated appropriately from historical cardiac arrest data, then nonemptiness of $\mathbf{U}$ is guaranteed (see Proposition 4.3). We discuss the design of uncertainty regions and the estimation of λ further in Section 4.5.2. Note also that our discretization scheme allows us to model generally shaped uncertainty regions. This flexibility allows physical obstacles to be incorporated into the uncertainty regions, such as certain private buildings where public cardiac arrests cannot occur.

Minimizing the worst-case CVaR over $\mathbf{y}, \mathbf{z}$, we now arrive at the following two-stage robust optimization problem:

$$\begin{aligned} \min_{\mathbf{y},\mathbf{z}} \max_{\mathbf{u}} \min_{\alpha} \ & \alpha + \frac{1}{(1-\beta)} \sum_{k \in \mathcal{K}} u_k \max \left\{ \sum_{i \in \mathcal{I}} d_{ik} z_{ik} - \alpha, 0 \right\} \\ \text{subject to } & \mathbf{u} \in \mathbf{U}, \mathbf{z} \in \mathbf{Z}(\mathbf{y}), \mathbf{y} \in \mathbf{Y}, \\ & \alpha \geq 0. \end{aligned} \tag{R-AED}$$

With respect to the size of R-AED, in practice we expect $\mathcal{K}$ to be large (on the order of 10^3 or larger), so that our discretization procedure reasonably approximates the underlying continuous problem. For example, in our case study (Sections 5 and 6), we take $|\mathcal{K}| = 2,600$. Similarly, the number of candidate sites $|\mathcal{I}|$ in our case study is 120, and the number of uncertainty regions $|\mathcal{J}|$ is 15.

It is worth discussing how formulation R-AED compares to other robust location models with demand uncertainty. The model closest to ours is the one proposed by Baron et al. (2011), who consider demand-level uncertainty at each of a fixed set of demand locations. Our model can also be viewed as modeling demand-level uncertainty, although there are some key differences with Baron et al. (2011). In formulation R-AED, the discrete uncertainty set $\mathbf{U}$ forces the aggregate demand from the locations $\mathcal{K}_j$ to be λ_j, and normalizes the total

demand over all locations in $\mathcal{K}$ to be 1. This is in contrast to the approach of Baron et al. (2011), who model the uncertainty by allowing the vector of demands over all locations to reside within a box or ellipsoidal uncertainty set. This difference in how demand uncertainty is modeled turns out to be consequential. Based on our choice of the (polyhedral) uncertainty set, the worst-case distribution in $\mathbf{U}$ will generally be supported on a relatively small number of locations, compared to the total number of scenarios. This means that, in the worst-case distribution, the demand u_k in the vast majority of locations in $\mathcal{K}$ will be 0. Moreover, increasing the size of the scenario $\mathcal{K}$ has no effect on the number of locations that support the worst-case distribution in $\mathbf{U}$. This sparsity and independence from the total number of scenarios in $\mathcal{K}$ is exploited by our row-and-column generation algorithm to obtain solutions to R-AED (Section 4). By contrast, in Baron et al. (2011), the uncertain demand at each location typically has a non-zero lower bound, which leads to all locations having non-zero demand, and thus precludes the application of our decomposition technique.

In the most general case, formulation R-AED can be interpreted as a facility location model with a single demand point (i.e., the location of the "next" cardiac arrest) and stochastic edge lengths, where d_{ik} gives the length of the edge between the demand point and candidate site i under scenario k. Under this interpretation, λ_j represents the probability that one of the scenarios in the set $\mathcal{K}_j$ occurs. Our model differs from the existing literature on stochastic edge lengths since we assume that the vector of probabilities $\mathbf{u}$ is itself uncertain, and only known to reside within a polyhedral uncertainty set.

Formulation R-AED can also be shown to unify the well known p-median and p-center location problems, as well as their robust analogues under demand location uncertainty. We formalize this connection in Section 4.3.4.

4.3.3 Discretization error bounds

The model R-AED is an approximation in the sense that it restricts the continuous demand distribution to be supported on a discrete set. As a result, the "true" worst-case CVaR in the continuous setting may be underestimated by the optimal value of formulation R-AED. In this section, we present bounds on this underestimation, which we refer to as the *discretization error*. The discretization error can be interpreted as the optimality gap associated with an optimal solution of formulation R-AED with respect to the underlying continuous problem (4.2).

To define and prove the discretization error bounds, we first construct a partition of the service area $\mathcal{A}$. The partitioning is required for technical purposes only, and can be applied to any general setting with overlapping uncertainty regions. Note that any configuration of the uncertainty regions $\mathcal{A}_j$ implies a partitioning of the full service area $\mathcal{A}$ into potentially more than n "subregions". For example, suppose $\mathcal{A} = \mathcal{A}_1 \cup \mathcal{A}_2$, where $\mathcal{A}_1 \cap \mathcal{A}_2 \neq \emptyset$. Then we consider three disjoint subregions: $\mathcal{A}'_1 = \mathcal{A}_1 \setminus \mathcal{A}_2$, $\mathcal{A}'_2 = \mathcal{A}_2 \setminus \mathcal{A}_1$, and $\mathcal{A}'_3 = \mathcal{A}_1 \cap \mathcal{A}_2$. In general, let $\mathcal{A}'_1, \ldots, \mathcal{A}'_{|\mathcal{R}|}$ represent these subregions, and let $\Xi'_1, \ldots, \Xi'_{|\mathcal{R}|}$ be the discrete counterparts to the $\mathcal{A}'_r$. Let $\mathcal{R}$ be the index set for the subregions and let $\mathcal{K}_r$ index the scenarios in subregion r. Since by definition the $\mathcal{A}'_r$ form a partition of $\mathcal{A}$, it follows that $\Xi'_r \cap \Xi'_{r'} = \emptyset$ for all $r \neq r'$. Note that in the most general case, we may have $|\mathcal{R}| = 2^{|\mathcal{J}|}$, if every uncertainty region partially overlaps with every other uncertainty region. However, this extreme case is unlikely to arise naturally in the context of AED deployment and other facility location problems. Note also that if there is no overlap among the $\mathcal{A}_j$, then $|\mathcal{R}| = |\mathcal{J}|$.

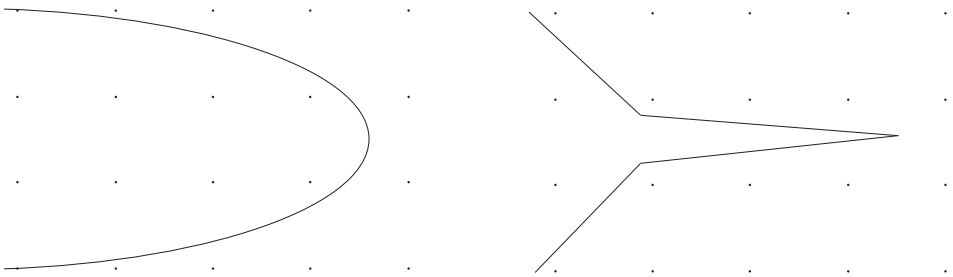

(a) Assumption 4.2 satisfied. (b) Assumption 4.2 violated.

Figure 4.1: Examples for Assumption 4.2.

To construct the bound, we first require the following assumptions.

Assumption 4.1. *$\Xi'_r = \mathcal{A}'_r \cap \mathcal{L}$ for all $r \in \mathcal{R}$, where $\mathcal{L}$ is a square lattice in the plane with a grid spacing of length σ.*

Assumption 4.2. *The shape of the subregion $\mathcal{A}'_r$ is such that for any point $a \in \mathcal{A}'_r \setminus \Xi'_r$, at least one of the four points in $\mathcal{L}$ that define the smallest square containing a is inside $\mathcal{A}'_r$.*

Assumption 4.1 states that the uncertainty regions are discretized using a square lattice, which simplifies the analysis considerably without loss of generality. Assumption 4.2 is meant to exclude pathologically shaped uncertainty regions relative to the size of the discretization (see Figure 4.1), and can be satisfied with a sufficiently granular discretization of the uncertainty region. We now define a parameter $\ell(\sigma)$ with the following property:

$$\ell(\sigma) \geq \max_{a \in \mathcal{A}'_r} \min_{\xi \in \Xi'_r} d(a, \xi), \qquad \text{for all } r \in \mathcal{R}. \tag{4.5}$$

Here, $\ell(\sigma)$ is an upper bound on the maximum distance between any point in $\mathcal{A}'_r$ and its closest scenario location in Ξ'_r. To make the expression in (4.5) concrete, we provide some examples of the parameter $\ell(\sigma)$ in Remark 1, depending on the choice of the distance metric.

Remark 4.1. Let Assumptions 4.1 and 4.2 hold. Let $\|\cdot\|_1$, $\|\cdot\|_2$ and $\|\cdot\|_\infty$ be the Euclidean, rectilinear and maximum metric on $\mathbb{R}^2$.

1. If $d(a, b) = \|a - b\|_1$, then $\ell(\sigma) = 2\sigma$.
2. If $d(a, b) = \|a - b\|_2$, then $\ell(\sigma) = \sqrt{2}\sigma$.
3. If $d(a, b) = \|a - b\|_\infty$, then $\ell(\sigma) = \sigma$.

Note that setting $\ell(\sigma) = 2\sigma$ is also valid for any distance function based on a p-norm with $p \geq 1$, since $\|\cdot\|_p$ is non-increasing in p. Let Z_D be the optimal value of formulation R-AED, Z_C be the optimal value of the associated continuous problem (4.2), and

$$\Delta := \frac{Z_C - Z_D}{Z_C}$$

be the discretization error. We now present the main bound.

Theorem 4.1. *For any $\varepsilon > 0$, if $\ell(\sigma) \leq \varepsilon(1-\beta)Z_D$, then $\Delta \leq \varepsilon$.*

Theorem 4.1 characterizes how finely the uncertainty regions must be discretized to achieve a certain discretization error. Since $\ell(\sigma)$ vanishes with σ, the discretization error can be made arbitrarily small by selecting an appropriate value for σ. Thus, we can obtain near-optimal solutions to the underlying continuous uncertainty problem by solving formulation R-AED with a sufficiently fine discretization (and thus a sufficiently large number of scenarios). We can also re-arrange the expression in Theorem 4.1 to calculate an upper bound on Z_C as a function of σ and Z_D. The upper bound $\bar{Z}_C$ is constructed as follows:

$$Z_C \leq \bar{Z}_C = Z_D(1+\varepsilon), \text{ where } \varepsilon = \frac{\ell(\sigma)}{Z_D(1-\beta)}. \tag{4.6}$$

In the next section, we outline a decomposition technique which enables us to compute a tight bound $\bar{Z}_C$ by solving instances of R-AED with an extremely large number of scenarios.

4.3.4 Unifying the p-median and p-center problems

The p-median and p-center problems are two of the most well-studied location models, having served as the foundation for a significant portion of the existing facility location literature (Owen and Daskin, 1998, Snyder, 2006, Melo et al., 2009). In this section, we discuss how a non-stochastic interpretation of R-AED specializes to robust variants of the p-median and p-center problems, depending on how β is selected. A straightforward corollary of this result is that if the $\mathcal{A}_j$ are all singletons, then formulation R-AED unifies the classical p-median and p-center problems (i.e. without any uncertainty).

We begin by defining robust variants of the p-median and p-center problems, and then prove an equivalence between these problems and formulation R-AED. Let $\mathcal{J}$ index a finite and fixed set of demand points. In the classical p-median problem, one wishes to site P facilities so to minimize the (weighted) total distance between all demand points and their nearest facilities. Suppose now that the j^{th} demand point is only known to reside within an uncertainty region $\mathcal{A}_j$. We formulate a generalization of the p-median problem that aims to minimize the *worst-case* total distance. Let each uncertainty region be discretized into a set of locations $\Xi_j \subset \mathcal{A}_j$, so that the demand point j is only known to be at one of the locations in Ξ_j. Let $\mathcal{K}_j$ index the locations in Ξ_j. Let $\mathbf{Z}(\mathbf{y})$ and $\mathbf{Y}$ retain their definitions from Section 4.3. Let λ_j be the usual notion of "weight" placed on demand point j. Now define $\mathbf{X} = \mathbf{X}_1 \times \ldots \times \mathbf{X}_{|\mathcal{J}|}$, where

$$\mathbf{X}_j = \left\{ x_{jk} \geq 0,\ k \in \mathcal{K} \;\middle|\; \sum_{k \in \mathcal{K}_j} x_{jk} = 1 \right\}, \quad j \in \mathcal{J}. \tag{4.7}$$

We can now write the *robust p-median problem* to minimize the worst-case total distance as

$$\min_{\mathbf{y},\mathbf{z}} \max_{\mathbf{x}} \sum_{i \in \mathcal{I}} \sum_{k \in \mathcal{K}_j} \sum_{j \in \mathcal{J}} \lambda_j d_{ik} z_{ik} x_{jk} \tag{4.8}$$
$$\text{subject to } \mathbf{x} \in \mathbf{X},\ \mathbf{z} \in \mathbf{Z}(\mathbf{y}),\ \mathbf{y} \in \mathbf{Y}.$$

In formulation (4.8), the set $\mathbf{X}$ can be interpreted as the decision space of an adversary which,

after observing the facility locations, realizes each demand point within its uncertainty region at a location that maximizes the distance to the nearest facility. Using the same definitions as above, we can also define the *robust p-center problem*, which minimizes the worst-case *maximum* distance between a demand point and its nearest facility:

$$\begin{aligned} \min_{\mathbf{y},\mathbf{z}} \max_{\mathbf{x}} \min_{t} \; & t \\ \text{subject to } & t \geq \sum_{i \in \mathcal{I}} \sum_{k \in \mathcal{K}_j} d_{ik} z_{ik} x_{jk}, \quad j \in \mathcal{J}, \\ & \mathbf{x} \in \mathbf{X},\ \mathbf{z} \in \mathbf{Z}(\mathbf{y}),\ \mathbf{y} \in \mathbf{Y}. \end{aligned} \tag{4.9}$$

We now show that the robust p-median and p-center problems are unified by formulation R-AED.

Proposition 4.1. *Assume that each scenario k belongs to exactly one of the sets $\mathcal{K}_1, \ldots, \mathcal{K}_{|\mathcal{J}|}$, and $\lambda_1 = \ldots = \lambda_{|\mathcal{J}|} = 1/n$. If $\beta = 0$, then formulation R-AED is equivalent to the robust p-median problem* (4.8) *with equally weighted demand points. If $\beta \geq (n-1)/n$, then formulation R-AED is equivalent to the non-weighted robust p-center problem* (4.9)*.*

Based on Proposition 4.1, we can interpret formulation R-AED as subsuming a spectrum of location models with the robust p-median problem at one end ($\beta = 0$) and the robust p-center problem at the other (β close to 1). While the specialization of CVaR to the mean and maximum of a general discrete distribution is well known, this explicit unification of the p-median and p-center problems via a CVaR objective does not appear to have been discussed previously in the facility location literature. We note also that the assumption in Proposition 4.1 that each k belong to exactly one $\mathcal{K}_j$ is needed to ensure the equivalence between the main formulation R-AED and the intermediate model (A.45) used in the proof. If this assumption is relaxed, then formulation (A.45) still serves as a unifying model for the p-median and p-center problems, although the equivalence with R-AED no longer holds. The *ordered median problem* (Kalcsics et al., 2002, Nickel and Puerto, 1999) also unifies median and center problems, but requires an ordering of the demands according to their distance to the sited facilities so that a set of appropriate parameters can be identified. By contrast, our approach involves the adjustment of a single parameter and requires no such ordering.

To the best of our knowledge, (4.8) and (4.9) are the first robust formulations of the general p-median and p-center problem with non-stochastic uncertainty in the demand location (i.e., in the spirit of robust optimization rather than scenario-based stochastic optimization). Previous work in robust facility location has considered only one facility (Averbakh and Berman, 2000, Averbakh, 2003) or uncertainty in demand level rather than location (Baron et al., 2011). As discussed in Section 3.2, the formulation in (Baron et al., 2011) is well equipped to handle uncertainty in the level of demand at each of a set of known demand points. The robust p-median and p-center formulations presented above may be more appropriate in applications where there is continuous uncertainty in future demand locations (e.g., emergency response or peer-to-peer ridesharing applications).

4.4 Row-and-column generation algorithm

It is well known that two-stage robust optimization problems with the structure of (4.11) can be solved using Benders-based decomposition algorithms (see, for example, Brown et al. (2006) Brown et al. (2009) Alderson et al. (2011)). These approaches decompose the formulation into a master problem, which is a relaxation of the original problem, and a subproblem, which generates constraints to be added to the master problem. The master and subproblems are solved iteratively until convergence to a provably optimal solution. However, when applied to (4.11), the number of constraints and decision variables of the (integer) master problem using these approaches depends on the number of scenarios $|\mathcal{K}|$. As a consequence, these decomposition approaches can fail to be tractable if $|\mathcal{K}|$ is even moderately sized. We overcome this issue by proposing an algorithm that decouples the size of the master problem from the number of scenarios. A key consequence of this decoupling is that the tractability of the resulting formulation no longer depends on the granularity of the discretization, which, when combined with the bounds in Section 4.3.3, allows us to approximate the continuous problem using an extremely large number of scenarios. A similar row-and-column generation algorithm for solving min-max optimization problems was independently proposed by Zeng and Zhao (2013). A key difference is that unlike the model in Zeng and Zhao (2013), which has an integer subproblem, our subproblem is a linear program.

We compare our proposed method with two other solution approaches that are natural candidates for solving formulation R-AED – a Benders-based row generation algorithm and a mixed integer programming (MIP) reformulation of R-AED.

4.4.1 Algorithm overview

We now provide an overview of the row-and-column generation algorithm. First, observe that the inner minimization problem of (4.4) is a linear program. Letting $\mathbf{p}$ denote the dual vector for this linear program, the dual polyhedron – which depends on the distribution $\mathbf{u}$ – can be written as

$$\mathbf{P}(\mathbf{u}) = \left\{ p_k \geq 0, k \in \mathcal{K} \;\middle|\; \sum_{k\in\mathcal{K}} p_k \leq 1,\ p_k \leq u_k/(1-\beta), k \in \mathcal{K} \right\}. \tag{4.10}$$

The optimization problem (4.4) can now be reformulated as

$$\begin{aligned} \min_{\mathbf{y},\mathbf{z}} \max_{\mathbf{u},\mathbf{p}} \ & \sum_{i\in\mathcal{I}} \sum_{k\in\mathcal{K}} d_{ik} z_{ik} p_k \\ \text{subject to } & \mathbf{p} \in \mathbf{P}(\mathbf{u}), \mathbf{u} \in \mathbf{U}, \\ & \mathbf{z} \in \mathbf{Z}(\mathbf{y}), \mathbf{y} \in \mathbf{Y}. \end{aligned} \tag{4.11}$$

By introducing the auxilliary variable t, we obtain the following epigraph representation of (4.11):

$$\underset{\mathbf{y},\mathbf{z},t}{\text{minimize}} \;\; t \tag{4.12a}$$

$$\text{subject to} \;\; t \geq \sum_{i\in\mathcal{I}} \sum_{k\in\mathcal{K}} d_{ik} z_{ik} p_k, \quad \mathbf{p} \in \mathbf{P}(\mathbf{u}), \mathbf{u} \in \mathbf{U}, \tag{4.12b}$$

$$\mathbf{z} \in \mathbf{Z}(\mathbf{y}), \mathbf{y} \in \mathbf{Y}. \tag{4.12c}$$

Note that formulation (4.12) is a semi-infinite optimization problem, since the constraint (4.12b) is enforced over all possible $\mathbf{p}$. To obtain a feasible solution to (4.12), we may instead solve a relaxation where the constraint (4.12b) is enforced over a finite subset of $\mathbf{P}$. Let this subset be $\mathbf{P}_{|\mathcal{S}|} = \{\mathbf{p}^1, \ldots, \mathbf{p}^{|\mathcal{S}|}\}$, and let $\mathcal{S}$ be an index set for the vectors in $\mathbf{P}_{|\mathcal{S}|}$. To obtain a relaxation of (4.12b) with a finite number of constraints, we replace the constraint set represented by (4.12b) with the following set of $|\mathcal{S}|$ constraints:

$$t \geq \sum_{i\in\mathcal{I}} \sum_{k\in\mathcal{K}} d_{ik} z_{ik} p_k^s, \quad s \in \mathcal{S}. \tag{4.13}$$

Although replacing the constraint (4.12b) with (4.13) in formulation (4.12) produces a mixed-integer linear program, the resulting formulation may still be intractably large, due to the dependence of the number of constraints and variables on the size of the scenario set, $\mathcal{K}$. Consider now the set

$$\mathcal{K}_+ = \left\{ k \in \mathcal{K} \;\middle|\; \sum_{s\in\mathcal{S}} p_k^s > 0 \right\}, \tag{4.14}$$

and note that constraint (4.13) can be written as

$$t \geq \sum_{i\in\mathcal{I}} \sum_{k\in\mathcal{K}_+} d_{ik} z_{ik} p_k^s + \sum_{i\in\mathcal{I}} \sum_{k\in\mathcal{K}\setminus\mathcal{K}_+} d_{ik} z_{ik} p_k^s, \quad s \in \mathcal{S}. \tag{4.15}$$

Since by definition we have $\sum_{k\in\mathcal{K}\setminus\mathcal{K}_+} p_k^s = 0$ for all $s \in \mathcal{S}$, we may drop the second term on the right hand side of (4.13), which further simplifies the constraint to

$$t \geq \sum_{i\in\mathcal{I}} \sum_{k\in\mathcal{K}_+} d_{ik} z_{ik} p_k^s, \quad s \in \mathcal{S}. \tag{4.16}$$

Observe that if $k \notin \mathcal{K}_+$, then the value of the corresponding assignment variable z_{ik} has no impact on the objective function, since its coefficient is 0 for all $s \in \mathcal{S}$. We can therefore remove all z_{ik} variables and the constraints $z_{ik} \leq y_i$ and $z_{ik} \geq 0$, for $k \in \mathcal{K} \setminus \mathcal{K}_+$. This leads

to the following relaxation of (4.12):

$$
\begin{aligned}
\underset{\mathbf{y},\mathbf{z},\mathbf{t}}{\text{minimize}} \;\; & t \\
\text{subject to} \;\; & t \geq \sum_{i\in\mathcal{I}} \sum_{k\in\mathcal{K}_+} d_{ik} z_{ik} p_k^s, \quad s \in \mathcal{S}, \\
& \sum_{i\in\mathcal{I}} y_i = P, \\
& \sum_{i\in\mathcal{I}} z_{ik} = 1, \quad k \in \mathcal{K}_+, \\
& z_{ik} \leq y_i, \quad i \in \mathcal{I},\ k \in \mathcal{K}_+, \\
& z_{ik} \geq 0, \quad i \in \mathcal{I},\ k \in \mathcal{K}_+, \\
& y_i \in \{0,1\}, \quad i \in \mathcal{I}.
\end{aligned}
\tag{R-AED-MP}
$$

The relaxed problem R-AED-MP yields a lower bound on the optimal value of the original problem (4.11). To tighten the relaxation, we can add a new vector $\mathbf{p}$ to $\mathbf{P}_{|\mathcal{S}|}$ and re-solve R-AED-MP, with the aim of obtaining an improved lower bound. Let $\bar{\mathbf{y}}$ be a deployment obtained at a solution of R-AED-MP. We then construct a feasible assignment vector $\bar{\mathbf{z}}$ that assigns each scenario to its nearest AED as follows. Let $i^*(k) = \text{argmin}_{\{i\in\mathcal{I}|\bar{y}_i=1\}}\{d_{ik}\}$, which represents the closest AED to location k, given the deployment $\bar{\mathbf{y}}$. Then for each $k \in \mathcal{K}$, set

$$
\bar{z}_{i^*(k),k} = 1 \text{ and } \bar{z}_{ik} = 0 \text{ for all } i \neq i^*(k). \tag{4.17}
$$

It is straightforward to show that the $\bar{\mathbf{z}}$ given by (4.17) is also optimal for R-AED-MP. To identify a new $\mathbf{p}$, we then solve the following subproblem, where $\bar{\mathbf{y}}, \bar{t}$ is obtained as the incumbent solution to R-AED-MP and $\bar{\mathbf{z}}$ is given by (4.17):

$$
\begin{aligned}
\underset{\mathbf{u},\mathbf{p}}{\text{maximize}} \;\; & \sum_{i\in\mathcal{I}} \sum_{k\in\mathcal{K}} d_{ik} \bar{z}_{ik} p_k \\
\text{subject to} \;\; & \mathbf{p} \in \mathbf{P}(\mathbf{u}), \mathbf{u} \in \mathbf{U}.
\end{aligned}
\tag{R-AED-SP}
$$

To understand why the construction of $\bar{\mathbf{z}}$ using (4.17) is necessary, note that since R-AED-MP minimizes CVaR, the optimal assignment variables $\mathbf{z}$ obtained at a solution to R-AED-MP may not assign every location k to its nearest AED, since only the locations in the tail of the distance distribution affect the optimal objective function value. As a consequence, solving R-AED-SP using the $\bar{\mathbf{z}}$ obtained directly from R-AED-MP may produce a solution to R-AED-SP that does not correspond to a worst-case distribution (with respect to the incumbent deployment $\mathbf{y}$). While in theory, the algorithm converges in finite time without reconstructing $\bar{\mathbf{z}}$ according to (4.17), in practice we found that including this additional assignment step results in fewer iterations and a significant speed improvement.

Observe that formulation R-AED-SP is a linear program. Now let $\mathbf{u}^*, \mathbf{p}^*$ be an optimal solution to R-AED-SP, and let Z_{SP} be the optimal value. If $\bar{t} \geq Z_{\text{SP}}$, then it follows that $\bar{t} \geq \sum_{i\in\mathcal{I}} \sum_{k\in\mathcal{K}} d_{ik} \bar{z}_{ik} p_k$ for all $\mathbf{p} \in \mathbf{P}(\mathbf{u}), \mathbf{u} \in \mathbf{U}$, and thus the incumbent solution is certifiably optimal to the original problem. If $\bar{t} < Z_{\text{SP}}$, then we introduce the variables z_{ik}, $k \in \mathcal{K}^*$ to the master problem, as well as the constraints $z_{ik} \geq 0$, $z_{ik} \leq y_i$, $k \in \mathcal{K}^*$ and $t \geq$

$\sum_{i\in\mathcal{I}}\sum_{k\in\mathcal{K}} d_{ik}\bar{z}_{ik}p_k^*$, where $\mathcal{K}^* = \{k \in \mathcal{K} \mid p_k^* > 0\}$. We then solve R-AED-MP again to obtain a new solution $\bar{\mathbf{y}}, \bar{\mathbf{z}}, t$, which begins a new iteration of the algorithm. Since there are a finite (but possibly large) number of rows and columns that can be generated, convergence after a finite number of iterations is guaranteed. The algorithm is summarized in Algorithm 3.

Algorithm 3: Row-and-column generation

1 **Initialize** Set $\mathcal{S} = \{0\}$, $s = 1$, and $\epsilon > 0$. Let $\mathcal{K}_+ = \emptyset$.

2 Solve R-AED-MP to obtain solution $(\bar{\mathbf{y}}, \bar{t})$ and objective value Z_{MP}.

3 Construct $\bar{\mathbf{z}}$ according to (4.17).

4 Solve R-AED-SP with fixed $\bar{\mathbf{z}}$ to obtain $\bar{\mathbf{p}}^s$ and objective value Z_{SP}. **If** $\frac{Z_{\text{SP}} - Z_{\text{MP}}}{Z_{\text{SP}}} \leq \epsilon$, terminate and return optimal deployment $\bar{\mathbf{y}}$ and worst-case CVaR $\bar{t}$. **else**

i) Construct index set

$$\mathcal{K}^* = \{k \in \mathcal{K} \setminus \mathcal{K}_+ \mid \bar{p}_k^s > 0\}$$

and set $\mathcal{K}_+ \leftarrow \mathcal{K}_+ \cup \mathcal{K}^*$.

ii) Add variables z_{ik}, $i \in \mathcal{I}$, $k \in \mathcal{K}^*$ to master problem R-AED-MP.

iii) Add constraints

$$\begin{aligned} t &\geq \sum_{i\in\mathcal{I}} \sum_{k\in\mathcal{K}_+} d_{ik} z_{ik} \bar{p}_k^s, \\ z_{ij}^k &\leq y_i, \quad i \in \mathcal{I},\ k \in \mathcal{K}^*, \\ \sum_{i\in\mathcal{I}} z_{ij}^k &= 1, \quad k \in \mathcal{K}^*, \end{aligned}$$

to master problem R-AED-MP.

iv) Set $\mathcal{S} \leftarrow \mathcal{S} \cup \{s\}$. Increment s and return to step 2.

The intuition behind the algorithm is as follows. At a given iteration, the master problem only includes the subset of scenarios that are assigned a non-zero probability mass in the worst case distribution for at least one of the previous iterations. Based on the incumbent solution, the subproblem then identifies the worst-case distribution using the entire scenario set. The scenarios which support the worst-case distribution (obtained by solving the subproblem) are then added to the master problem on an as needed basis, along with the relevant assignment variables and related constraints. Due to this decomposition, at any iteration of the algorithm the size of the master problem R-AED-MP depends only on the number of candidate sites, $|\mathcal{I}|$, and the size of the restricted scenario set, $\mathcal{K}_+$, which may be substantially smaller than $\mathcal{K}$. In fact, due to the structure of the uncertainty set $\mathbf{U}$, the size of the restricted scenario set $\mathcal{K}_+$ can be shown to depend only on the configuration of the uncertainty regions and the number of iterations completed by the algorithm, and is otherwise independent of the total number of scenarios in $\mathcal{K}$. This result is formalized in the following proposition.

Proposition 4.2. *Assume that for any k and k' such that $k \neq k'$, we have $d_{ik} \neq d_{ik'}$ for all $i \in \mathcal{I}$. Then the number of constraints and variables in the master problem R-AED-MP in the s^{th} iteration of Algorithm 1 is $O(|\mathcal{I}|(|\mathcal{J}|+1)s)$.*

The assumption that $k \neq k'$ implies $d_{ik} \neq d_{ik'}$ for all $i \in \mathcal{I}$ is mild, and ensures that for any candidate site i, there are no exact ties for the closest location in $\mathcal{K}$ (indeed, this assumption can be relaxed if we instead assume that the solution to R-AED-SP is always a basic feasible solution, which is the case if the simplex method is used to solve R-AED-SP). It is worth noting that the size of the master problem R-AED-MP may still depend indirectly on $\mathcal{K}$ if the total number of scenarios is small (i.e., if $|\mathcal{R}|s > |\mathcal{K}|$). Additionally, Proposition 1 only guarantees that the size of the master problem in any given iteration is independent of $\mathcal{K}$, but does not guarantee that the *total* number of iterations required will be independent of $\mathcal{K}$. However, our numerical results indicate that increasing the number of scenarios generally has a minimal impact on the performance of the row-and-column generation algorithm, because the worst-case distribution is generally sparse with respect to $\mathcal{K}$. Note that unlike the master problem, the subproblem R-AED-SP still depends on the size of $\mathcal{K}$. However, the subproblem is a linear program with $|\mathcal{K}|$ decision variables, and we found that it can be solved efficiently for practically sized problems, even if the set $\mathcal{K}$ is quite large.

4.4.2 Benchmark solution approaches

We compare the performance of the row-and-column generation algorithm to two benchmark approaches: a mixed-integer linear programming reformulation of (4.2) and a standard row generation algorithm for min-max problems.

Mixed-integer linear programming reformulation

By formulating the dual problem of the inner maximization problem in (4.11), we obtain the following mixed-integer linear program:

$$\begin{aligned}
\underset{\mathbf{y},\mathbf{z},\mathbf{w},\eta,\alpha,\mathbf{fl}}{\text{minimize}} \quad & \eta + \alpha + \sum_{j \in \mathcal{J}} \lambda_j w_j \\
\text{subject to} \quad & \alpha + \gamma_k \geq \sum_{i \in \mathcal{I}} \sum_{k \in \mathcal{K}} d_{ik} z_{ik}, \quad k \in \mathcal{K}, \\
& \eta - \frac{\gamma_k}{1-\beta} + \sum_{\{j | k \in \mathcal{K}_j\}} w_j \geq 0, \quad k \in \mathcal{K}, \\
& \mathbf{z} \in \mathbf{Z}(\mathbf{y}),\ \mathbf{y} \in \mathbf{Y}, \\
& \gamma_k \geq 0, \quad k \in \mathcal{K}, \\
& \alpha \geq 0.
\end{aligned} \tag{4.18}$$

If $\mathcal{K}$ is relatively small, then the MIP formulation (4.18) can be solved using standard integer programming solvers. However, this formulation can become computationally intractable if a large number of scenarios are used to discretize the uncertainty regions.

Row generation

A well known approach to solving min-max problems like (4.2) is to use a Benders-based row generation algorithm (Brown et al., 2006, 2009, Alderson et al., 2011). The row generation algorithm is conceptually similar to our row-and-column generation method outlined in Section 4.4.1, in that it also decomposes the problem into a master and subproblem which provide lower and upper bounds, respectively. However, a key difference with our approach is that the row generation algorithm carries the full set of columns (i.e., all scenarios in $\mathcal{K}$) in the master problem in each iteration. The number of variables and constraints in the master problem of the row generation algorithm in the s^{th} iteration is thus $O(|\mathcal{I}||\mathcal{K}|)$ instead of $O(|\mathcal{I}||\mathcal{R}|s)$. As a result, the row generation algorithm may have poor performance if $|\mathcal{K}|$ is large.

4.4.3 Comparison of solution approaches

To demonstrate the performance of the row-and-column generation algorithm, we generated and solved several random problem instances using rectangular uncertainty regions and a square lattice for the discretization. We use Euclidean distances for all parameters. We solve each of these instances using the three solution methods discussed above, and focus on the impact that the number of scenarios has on the total solution time. All problems were implemented in MATLAB R2011a using YALMIP as the modeling language and CPLEX 12.1 with default parameter settings as the solver, on a 2.4 GHz quad-core CPU.

In Table 4.1, we report numerical results on how the solution times of the MIP formulation (MIP), row generation algorithm (Row) and row-and-column generation algorithm (R+C) are impacted by the number of scenarios. We report results only for the cases where $\beta = 0$ and $\beta = 0.9$, and note that similar trends were observed for other values of β. We also compute the error bound ε and the associated upper bound on the continuous problem, $\bar{Z}_C$, using (4.6). Note that the error bounds for the $\beta = 0.9$ cases are larger than the $\beta = 0$ case, which is expected based on the bound expression given in (4.6).

Unlike the MIP and row generation algorithm, the performance of the row-and-column generation algorithm scales extremely well with the number of scenarios, and in many cases outperforms the benchmark approaches by one to two orders of magnitude with respect to the solution time. Note that the performance improvement is most pronounced at larger values of $|\mathcal{K}|$, due to the poor scalability of the MIP and row generation approaches. These results suggest that, when using the row-and-column generation algorithm, if a problem can be solved to optimality using a small number of scenarios, then the problem tends to remain tractable when using a fine discretization and an extremely large number of scenarios (e.g., just under 100,000 scenarios are used in the final row of Table 4.1).

Figure 4.2 depicts an example of the total number of scenarios generated in the master problem R-AED-MP and the associated optimality gap at each iteration of the row-and-column generation algorithm. The plot corresponds to the instance given in the second-last row of Table 1 ($\mathcal{I} = 100$, $\mathcal{J} = 20$, $P = 20$, $\mathcal{K} = 24,743$).

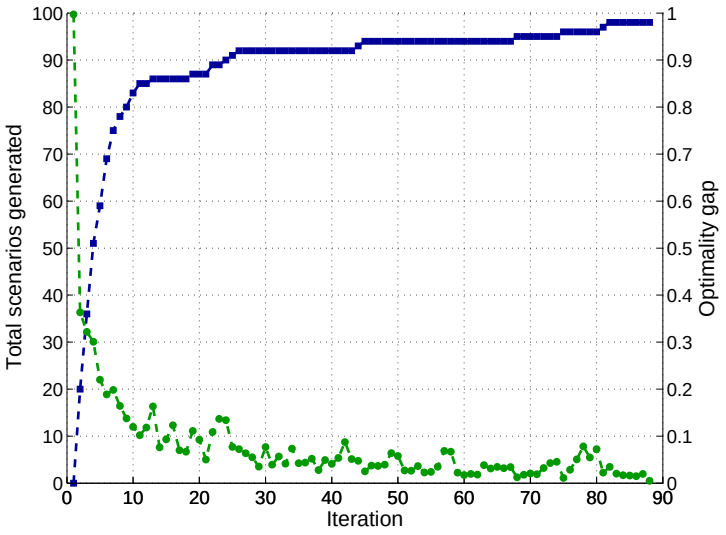

Figure 4.2: Example of total number of scenarios generated (blue squares) and optimality gap (green circles) for each iteration of the row-and-column generation algorithm ($|\mathcal{I}|= 100$, $|\mathcal{J}|= 20$, $P = 20$, $|\mathcal{K}|= 24,743$).

Table 4.1: Computational results showing effect of number of scenarios ($|\mathcal{K}|$) on error bound and solution times (in CPU seconds) of MIP, row generation and row-and-column generation algorithm. Asterisk (*) indicates instance did not solve in 10,000 CPU seconds.

					$\beta = 0$						$\beta = 0.9$					
$\|\mathcal{I}\|$	$\|\mathcal{J}\|$	P	$\|\mathcal{K}\|$	σ	Z_D	ε	$\bar{Z}_C$	MIP	Row	R+C	Z_D	ε	$\bar{Z}_C$	MIP	Row	R+C
25	5	2	456	5.0	124.6	5.7%	131.7	1	2	2	180.0	39.3%	250.7	12	35	44
25	5	2	1,864	2.5	126.0	2.8%	129.5	2	3	2	182.4	19.4%	217.8	56	45	42
25	5	2	7,446	1.3	126.4	1.4%	128.2	9	9	3	182.8	9.7%	200.4	107	135	49
25	5	2	29,621	0.6	126.9	0.7%	127.8	93	311	8	183.4	4.8%	192.2	*10,000	5,359	64
50	10	5	886	5.0	105.2	6.7%	112.3	6	12	4	178.0	39.7%	248.7	37	342	230
50	10	5	3,563	2.5	107.4	3.3%	110.9	94	104	13	181.8	19.4%	217.2	390	1,257	290
50	10	5	14,242	1.3	108.3	1.6%	110.0	*10,000	*10,000	33	183.1	9.7%	200.8	*10,000	*10,000	414
50	10	5	57,016	0.6	108.5	0.8%	109.4	*10,000	*10,000	49	181.3	4.9%	190.2	*10,000	*10,000	603
75	15	10	1,263	5.0	83.6	8.5%	90.7	23	1,875	174	139.3	50.8%	210.0	142	293	340
75	15	10	5,112	2.5	85.1	4.2%	88.6	2,415	*10,000	523	138.9	25.5%	174.2	2,088	9,799	614
75	15	10	20,470	1.3	85.4	2.1%	87.1	*10,000	*10,000	614	139.1	12.7%	156.8	*10,000	*10,000	1,534
75	15	10	81,940	0.6	85.8	1.0%	86.7	*10,000	*10,000	778	140.2	6.3%	149.0	*10,000	*10,000	1,610
100	20	20	1,511	5.0	60.0	11.8%	67.0	30	228	53	84.6	83.6%	155.3	9	174	87
100	20	20	6,158	2.5	61.4	5.8%	64.9	2,617	2,579	485	84.7	41.7%	120.1	129	454	172
100	20	20	24,743	1.3	62.1	2.8%	63.9	*10,000	*10,000	585	85.7	20.6%	103.4	*10,000	4,255	2,175
100	20	20	99,032	0.6	62.7	1.4%	63.5	*10,000	*10,000	1,739	73.0	12.1%	81.8	*10,000	*10,000	3,252

4.5 Case study setup

In this section, we discuss the setup of our numerical study on the placement of AEDs in public locations, including an overview of the data used, calibration of the uncertainty set, benchmark models, and a simulation-based validation method to evaluate the performance of the AED deployments produced by each model.

4.5.1 Data

Historical cardiac arrests. We consider the AED location problem in a densely populated region in the downtown core of Toronto, Canada, which roughly corresponds to the city's financial district. We obtained a cardiac arrest dataset through the Resuscitation Outcomes Consortium (ROC), which is a collaborative network of research institutions across the United States and Canada that collects cardiac arrest data for the purpose of improving clinical practice (Morrison et al., 2008). The ROC dataset contains the geographic coordinates of all cardiac arrests in the City of Toronto from January 2006 to April 2013. During this period, a total of 43 cardiac arrests occurred in public locations within the study region. These cardiac arrest locations are used to estimate the λ_j parameters in formulation R-AED.

Candidate sites. We constructed a set of candidate sites for AED placement by randomly selecting 120 public spaces (restaurants, cafes, shops, etc.) within the downtown region we consider. The addresses of these public spaces were obtained through the City of Toronto Employment Survey (City of Toronto, 2010), which collects data on all Toronto businesses on an annual basis. We use publicly accessible buildings such as restaurants and cafes since they represent the type of location where one might reasonably expect an AED to be accessible to the public through a public access defibrillation program. Figure 4.3 provides a visualization of the study region, historical cardiac arrests and the candidate sites for AED placement.

Distance parameters. To calculate the distance parameters, we converted the addresses of the candidate sites into Universal Transverse Mercator (UTM) coordinates. Coordinates in the UTM system are given in units of meters (unlike latitude and longitude coordinates), which allows for straightforward calculation of distances within the study region. We constructed the scenario set $\mathcal{K}$ by intersecting the study region with a square grid with a spacing of 20 meters, which resulted in approximately 2,600 scenarios. We generated the d_{ik} parameters by calculating the Euclidean distance between each candidate site and scenario pair. Although the Euclidean distance is only an approximation of the one-way distance a lay responder would have to travel to retrieve an AED, previous studies have shown Euclidean distance to be highly correlated with actual road distance, especially in urban areas where the road network is relatively dense (Bach, 1981, Jones et al., 2010, Boscoe et al., 2012). Moreover, the use of Euclidean distances is particularly justified in the context of AED placement, since the typical distance between a cardiac arrest and the nearest AED is low enough that lay responders can often take shortcuts that do not necessarily appear on the road network (e.g., through buildings, unmarked pedestrian paths, or parking lots).

As a robustness check, we repeat a subset of experiments using rectilinear distances and actual walking distances. The walking distances are obtained using the Google Maps Distance Matrix Application Programming Interface (API), implemented in Python. The Distance Matrix API is a query tool that accepts pairs of locations as input (as latitude-longitude coordinates or street addresses) and returns the distance between the two locations

for a given mode of transportation (e.g., walking or driving) (Google, 2016).

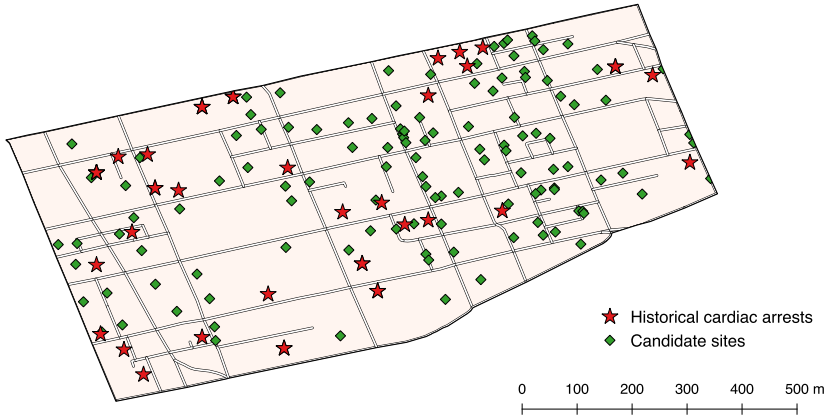

Figure 4.3: Locations of historical cardiac arrests and candidate sites for AED placement.

4.5.2 Estimation of uncertainty set parameters

One of the key modeling decisions in our approach is to construct a set of uncertainty regions from the full service area. One simple approach is to partition the service area into a set of identically shaped cells and allow each cell to represent an uncertainty region. We propose that in the context of public access defibrillation, it may be most intuitive from a policy making perspective to use pre-existing geographic divisions within the city, such as city blocks, postal/zip codes, or census tracts. Further, in settings where historical cardiac arrest data is unavailable, one may have to rely on alternate demographic information that are only recorded in aggregate as a proxy for cardiac arrest incidence. For example, the City of Toronto collects socio-economic data at the neighbourhood level for planning purposes (City of Toronto, 2016), which naturally implies a set of uncertainty regions.

Once a set of uncertainty regions is selected, the corresponding probabilities $\lambda_1, \ldots, \lambda_{|\mathcal{J}|}$ must be estimated from historical cardiac arrest data. We propose the following approach, which produces consistent (i.e., asymptotically optimal) estimates of the uncertainty region probabilities.

Proposition 4.3. *Let $a_1, \ldots, a_n \in \mathcal{A}$ be the locations of n cardiac arrest events, and assume that $a_1, \ldots, a_n$ are drawn i.i.d. from the distribution μ. For each $j \in \mathcal{J}$, let*

$$\hat{\lambda}_j^n = \frac{1}{n} \sum_{c=1}^{n} \mathbb{1}\{a_c \in \mathcal{A}_j\}. \tag{4.19}$$

Then $|\hat{\lambda}_j^n - \lambda_j| \longrightarrow 0$ as $n \longrightarrow \infty$ almost surely. Further, if the sets $\Xi'_1, \ldots, \Xi'_{|\mathcal{R}|}$ are all non-empty, then the uncertainty set $\mathbf{U}$ constructed from $\hat{\lambda}_j^n$, $j \in \mathcal{J}$ is guaranteed to be non-empty.

The estimator in Proposition 4.3 is simply the fraction of cardiac arrest events that have occurred in each uncertainty region. Since the uncertainty regions are fixed, the convergence

result in Proposition 4.3 follows directly from the strong law of large numbers (to be precise, note that the arrival of cardiac arrests in each region can, with some modification, be modeled as a multinomial trials process, which has a well-known maximum likelihood estimator (Bickel and Doksum, 2015)). The assumption that the cardiac arrest locations are drawn i.i.d. from a common distribution is a standard statistical assumption, and can be interpreted as requiring the probability of a cardiac arrest event in different parts of the city to be stable over time. Fortunately, cardiac arrests in Toronto have been observed to exhibit this temporal stability (Chan et al., 2014).

Note also that these properties hold for *any* chosen configuration of the uncertainty regions. In other words, once the uncertainty regions are selected, the estimate $\hat{\lambda}^n$ is guaranteed to converge to the true arrival probabilities. This convergence is generally known as statistical *consistency*, which is considered a fundamental requirement for an estimator (Bickel and Doksum, 2015).

In addition to estimating the arrival probabilities, we can also use the Wilson score method as a heuristic to estimate an associated 95% confidence interval for each arrival probability (Wilson, 1927). These confidence intervals can also be used to guide the design of the uncertainty regions. For example, we might require that the uncertainty regions be selected such that estimation error of each of the arrival probabilities $\lambda_1, \ldots, \lambda_{|\mathcal{J}|}$ is no greater than some threshold δ with 95% confidence. We emphasize, however, that the Wilson score method is only a heuristic for estimating the confidence intervals, meaning the estimation error is not always guaranteed to be less than δ with 95% confidence. As an alternative, we also note that the problem of designing uncertainty regions based on data is closely related to the well-studied problem in statistics of optimally selecting histogram bin widths (e.g., Wand (1997)), although we do not formally explore that connection in this work.

The service area in our numerical study spans only two census tracts and three postal codes zones. Since using too few uncertainty regions may fail to capture important aspects of the cardiac arrest distribution, we instead identified uncertainty regions by using the road network to divide the study area into 15 approximately equal-area uncertainty regions. The estimates of the arrival probabilities for the 15 uncertainty regions obtained using this approach are reported in Table 4.5.2, along with their approximate 95% confidence intervals, calculated using the Wilson score method (Wilson, 1927).

4.5.3 Benchmarks: Sample average approximation and ex-post models

We compare the performance of the AED deployment produced by the robust formulation R-AED with two benchmark models. The first benchmark is a *nominal* model, which optimizes only over historical cardiac arrest locations without taking uncertainty into account. The nominal model represents the the sample average approximation approach to the AED problem, since it optimizes directly with respect to a historical sample. The nominal model is given by formulation N-AED, where the scenarios $\{\xi_1, \ldots, \xi_{|\mathcal{K}|}\}$ are given by the historical cardiac arrest locations, and $u_k = 1/|\mathcal{K}|$ for all $k \in \mathcal{K}$. The second benchmark is an *ex-post* model, which optimizes directly over a set of simulated cardiac arrests, and thus has perfect foresight. The ex-post model is also given by formulation N-AED, except instead of optimizing over historical cardiac arrest locations, the ex-post model optimizes the AED deployment with respect to a simulated set of cardiac arrest locations, which we also evaluate the performace of the nominal and robust models against. The ex-post model can be thus interpreted

Table 4.2: Estimates of probabilities $\lambda_1, \ldots, \lambda_{|\mathcal{J}|}$ and approximate 95% confidence intervals for 15 uncertainty regions.

j	$\hat{\lambda}_j$	95% C.I.
1	0.12	(0.05, 0.24)
2	0.16	(0.07, 0.30)
3	0.09	(0.03, 0.22)
4	0.12	(0.04, 0.25)
5	0.07	(0.01, 0.19)
6	0.05	(0.01, 0.15)
7	0.05	(0.01, 0.15)
8	0.02	(0.00, 0.12)
9	0.02	(0.00, 0.12)
10	0.09	(0.03, 0.22)
11	0.02	(0.00, 0.12)
12	0.02	(0.00, 0.12)
13	0.07	(0.01, 0.18)
14	0.05	(0.01, 0.15)
15	0.05	(0.01, 0.15)

as the "best possible" model we could solve if the locations of all future cardiac arrests were known a priori. In all instances, we set $\beta = 0.9$ and the number of AEDs (P) to 30. All instances for all models were implemented using MATLAB R2011a via YALMIP, and solved using CPLEX 12.1 with default parameter settings on a single node of a computing cluster with a 2.9 GHz quad-core CPU.

4.5.4 Model validation

To assess the performance of the AED deployments produced by R-AED and N-AED, we used simulation to generate hypothetical cardiac arrest events, and computed various performance metrics associated with the distance distribution induced by each AED deployment and the set of simulated cardiac arrests. We then solve the ex-post model based on the simulated cardiac arrest locations, to gain insight into the best-possible deployment, for the given model parameterization and objective function.

To construct the validation sets, we first used kernel density estimation to estimate the underlying demand density, which is a well known semi-parametric technique for estimating a probability density function from a finite sample (Sheather and Jones, 1991, Terrell and Scott, 1992). Intuitively, kernel density involves centering a continuous density function (the kernel) at each data point, and then aggregating and normalizing all kernel functions to obtain a single probability density function. Kernel density estimation requires the specification of a kernel function as well as a parameter h, known as the *bandwidth*, which is typically proportional to the standard deviation of the kernel distribution. In short, the bandwidth represents the degree to which the data set is smoothed to obtain the estimated density,

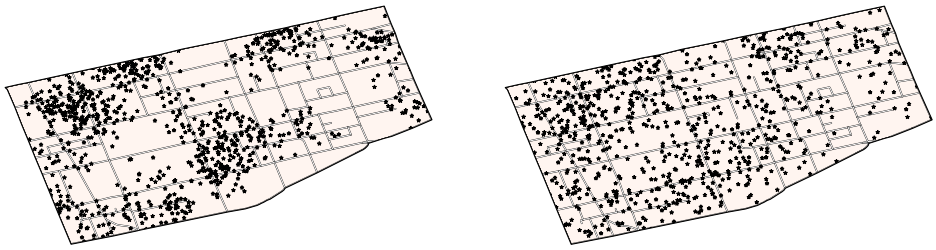

(a) h = 50m. (b) h = 100m.

Figure 4.4: Locations of 1,000 simulated cardiac arrests based on kernel density estimation using bandwidths of 50m and 100m.

with a higher bandwidth resulting in more aggressive smoothing. We note that using kernel density estimation to estimate the distribution of future cardiac arrests is supported by the fact that cardiac arrest locations in Toronto have been shown to exhibit temporal stability (Chan et al., 2014).

To assess model performance under different values of the bandwidth parameter h, we constructed four different validation sets. Specifically, we used Gaussian kernels with $h = 10$, 50, 100 and 150 meters to estimate four demand densities, each of which is used to simulate a separate validation set. The case where $h = 10$ can be interpretted as a "low uncertainty" environment, since the low bandwidth causes a majority of simulated events to fall near a historical cardiac arrest location. Conversely, the case where $h = 150$ can be interpretted as a "high uncertainty" environment, since the larger bandwidth leads to some cardiac arrests being simulated far from any historical location. By varying the bandwidth parameter from 10 to 150, we are able to evaluate the performance of the AED deployment under a variety of possible demand distributions. For each of the four estimated densities, we simulated a total of 50 validation sets containing 100 cardiac arrests each. We do not impose any constraints on where cardiac arrests may be simulated (i.e., a small fraction may occur beyond the service area). Figure 4.4 illustrates the effect of the bandwidth on the estimated demand density. Note that the distribution of the simulated cardiac arrests is notably more uniform for a bandwidth of 100 meters compared to the distribution for a bandwidth of 50 meters, due to increased smoothing.

4.6 Results

The numerical results are organized as follows. In Section 4.6.1, we compute the average value (over 50 validation sets) of four performance metrics: the mean, VaR, CVaR and maximum of the distance distribution induced by each AED deployment. We also report worst-case performance of each deployment, which is given by the maximum value of each performance metric over 50 validation sets. In Section 4.6.2, we evaluate the models using a different but related measure of AED deployment performance, known as cardiac arrest coverage. Lastly, in Section 4.6.3 we do a robustness check by presenting results under alternate distance measures, namely, actual walking distances obtained via Google Maps and the rectilinear

distance metric.

4.6.1 Performance metrics

Figure 4.5 shows the mean, VaR, CVaR and maximum value of the distance distribution for the nominal, robust and ex-post models (note that the y-axis differs between the plots). The performance metrics are computed for each of the four demand distributions (i.e., four bandwidth parameters), and are averaged over 50 validation sets. Since Figure 4.5 shows the average value of each performance metric, we can interpret it as depicting the performance of each model under "typical" cardiac arrest locations.

In the cases where $h = 10$, the AED deployments produced by all three models perform similarly. This is not a surprising observation, since in the $h = 10$ case, which represents a low uncertainty environment, the vast majority of simulated cardiac arrests occur within the vicinity of one of the 43 historical cardiac arrest locations. However, it is worth noting that the robust solution performs at least as well as the nominal solution in the $h = 10$ case. In other words, the performance of the robust model does not appear to suffer when the level of uncertainty is low, despite optimizing AED locations to protect against uncertainty. Moreover, in the cases where $h \geq 50$, the robust solution substantially outperforms the nominal solution on all four performance metrics. For example, in the case where $h = 100$, the robust solution outperforms the nominal solution on the mean distance metric by 9% (86m vs 95m) and on the CVaR metric by 15% (223m vs 189m). The performance difference between the robust and nominal models was found to be statistically significant ($p < 0.05$) in all comparisons where $h \geq 50$. Examples of AED deployments under each model for $\beta = 0$ and $\beta = 0.9$ are given in Section 4.6.4.

Note that the ex-post model achieves the best performance in all instances, which is expected since it optimizes the AED deployment directly over the simulated cardiac arrest locations (note that the ex-post model is solved independently for each of the 50 trials, since it takes the simulated cardiac arrest locations as input). The gap in performance between the nominal and ex-post models may therefore be interpreted as the degradation in performance due to the uncertainty in cardiac arrest locations. Figure 4.5 shows that the robust deployment substantially closes this performance gap in many cases. For example, for the CVaR metric with $h = 100$, using the robust deployment decreases the performance gap with the ex-post model from 44 meters (223m vs 179m) down to 10 meters (189m vs 179m), representing a gap improvement of 77%. In this setting, it may be convenient to think of the performance gap as being analogous to the usual notion of optimality gap. However, we emphasize that the ex-post solution represents a theoretical lower bound on performance only for the CVaR metric, since all models optimize for CVaR only. For the remaining three performance metrics, it is possible for the nominal or robust model to outperform the ex-post model by chance.

Table 4.3 shows the performance gap achieved by the nominal and robust deployment with respect to the ex-post deployment. For example, the first row and column of Table 4.3 indicates that the performance of the nominal solution on the mean distance metric in the $h = 10$ case is within 12% of the performance achieved by the ex-post solution, whereas the gap from the robust solution is within 7% (averaged over 50 trials). The difference in performance between the nominal and robust solutions is most pronounced at the tail of the distribution, which is unsurprising since the models optimize for CVaR. Note that for the

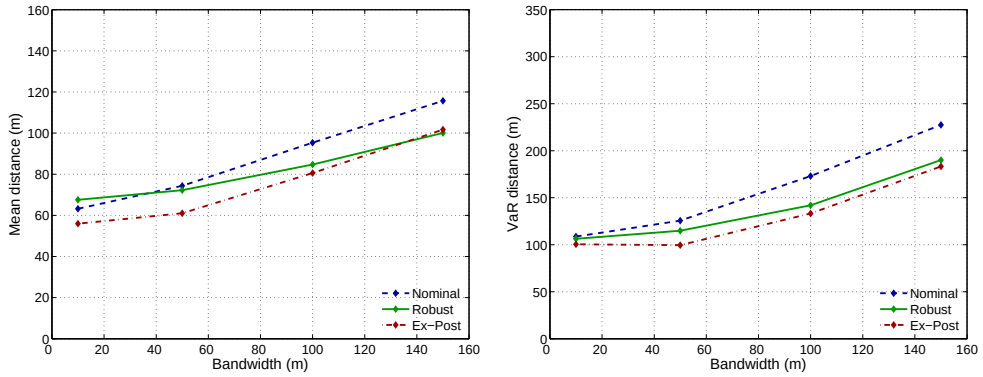

(a) Mean.

(b) VaR.

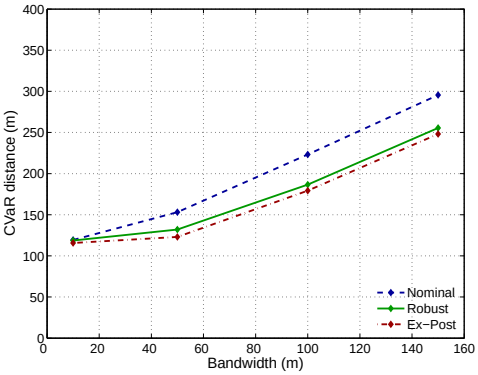

(c) CVaR.

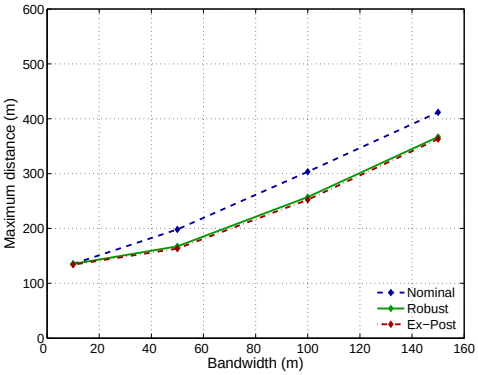

(d) Maximum.

Figure 4.5: Average performance over 50 validation sets of nominal, robust and ex-post solutions under increasing kernel bandwidth ($\beta = 0.9$).

Table 4.3: Average performance gap of nominal and robust solutions with respect to ex-post solution.

	Mean		VaR		CVaR		Max	
h	Nom.	Rob.	Nom.	Rob.	Nom.	Rob.	Nom.	Rob.
10	12%	7%	8%	5%	3%	3%	1%	1%
50	18%	11%	21%	13%	20%	8%	18%	2%
100	15%	6%	23%	9%	20%	5%	17%	3%
150	12%	1%	19%	4%	16%	4%	12%	3%

Table 4.4: Worst case performance gap of nominal and robust solutions with respect to ex-post solution.

	Mean		VaR		CVaR		Max	
h	Nom.	Rob.	Nom.	Rob.	Nom.	Rob.	Nom.	Rob.
10	30%	24%	28%	19%	10%	7%	6%	6%
50	29%	20%	34%	29%	31%	19%	41%	20%
100	29%	17%	40%	26%	27%	16%	31%	24%
150	20%	10%	31%	20%	24%	12%	26%	26%

CVaR and maximum distance metrics, the performance gap of the robust solutions is at most 8%, whereas the gap of the nominal solution can be as high as 20%.

We also consider the worst-case performance of each model by computing the maximum value of each performance metric over the 50 trials. As shown in Figure 4.6, the robust model generally outperforms the nominal model in the majority of high uncertainty instances ($h \geq 50$), and performs no worse than the nominal model in the low uncertainty instances ($h = 10$). Table 4.4 presents the worst-case performance gaps, which are obtained by computing the maximum relative gap in performance over 50 trials. As in Table 4.3, the robust solution substantially closes the worst-case performance gap in the vast majority of instances. For example, on the CVaR metric with $h = 100$, the worst-case performance gap is improved by 41% (27% down to 16%).

4.6.2 Cardiac arrest coverage

An alternate but related measure of performance in public access defibrillation is cardiac arrest *coverage*, which is the fraction of cardiac arrests that occur within a given distance of at least one AED (Chan et al., 2013). Figure 4.7 shows the cardiac arrest coverage under the nominal, robust and ex-post solutions at all distances between 0 and 250 meters, for each of the four bandwidths. The curves shown in Figure 4.7 are obtained by computing the average coverage at each distance over all 50 trials. Note that in the case where $h = 10$, all three models attain similar coverage at all distances, which is consistent with the results shown in Figure 4.5. Note that the robust solution dominates the nominal solution at all bandwidths, in the sense that it achieves a higher coverage level at all distances. Moreover, the coverage

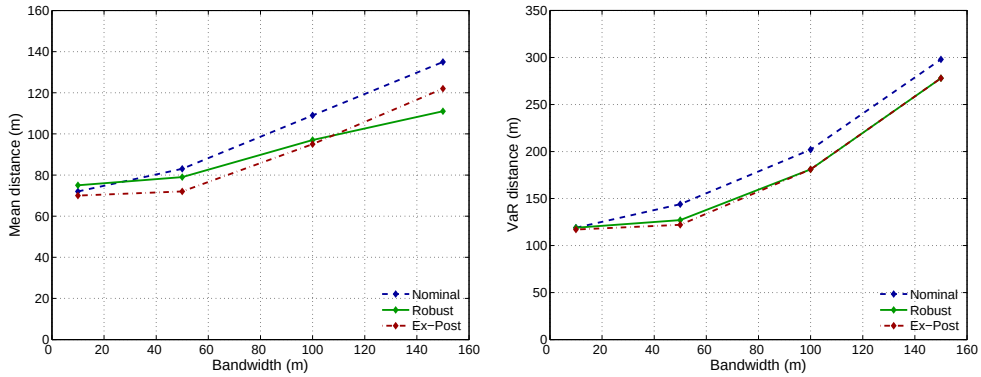

(a) Mean.

(b) VaR.

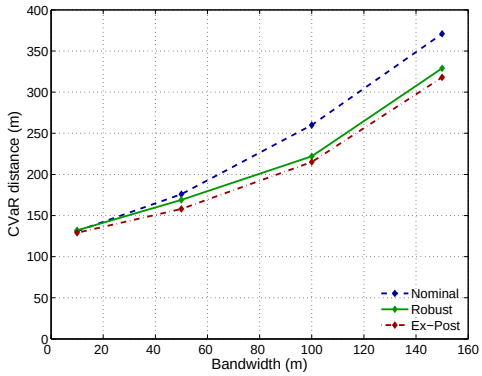

(c) CVaR.

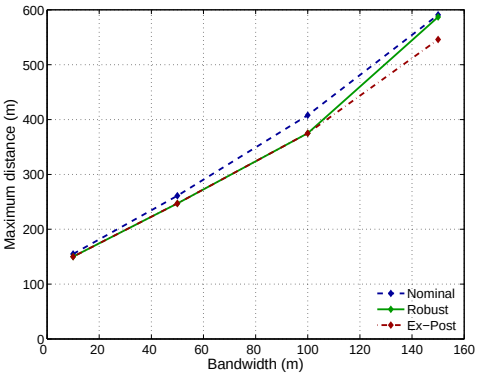

(d) Maximum.

Figure 4.6: Worst-case performance of nominal, robust and ex-post solutions on four distance metrics under increasing kernel bandwidth.

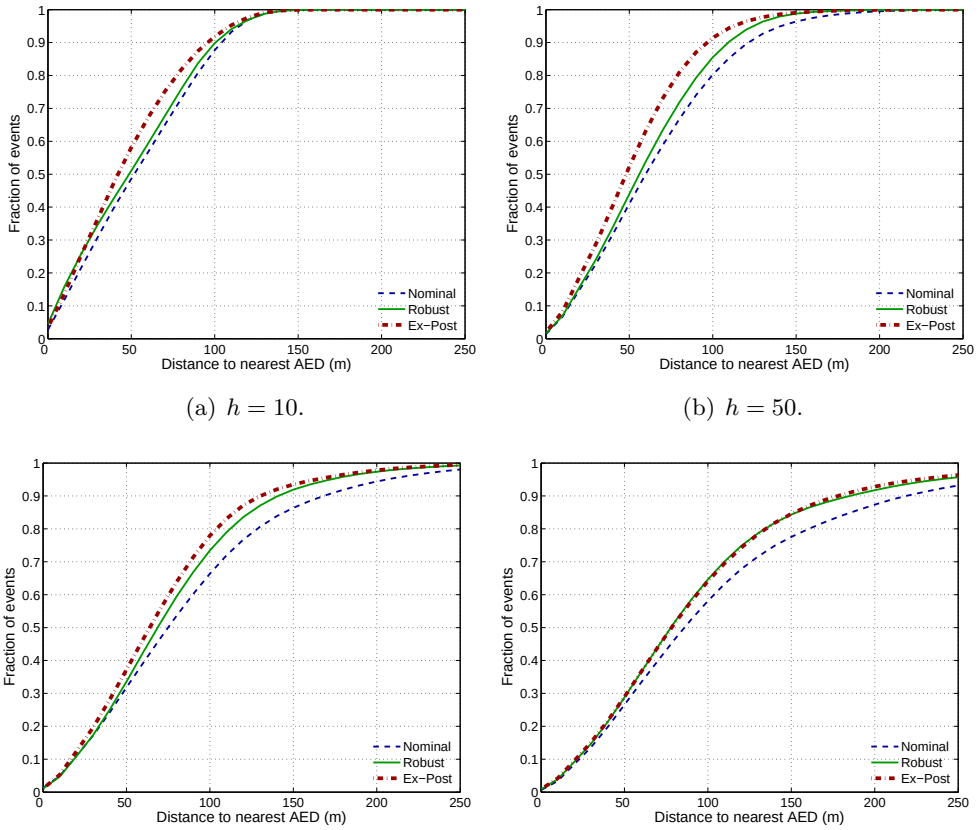

(a) $h = 10$. (b) $h = 50$.

(c) $h = 100$. (d) $h = 150$.

Figure 4.7: Cardiac arrest coverage for nominal, robust and ex-post solutions under four kernel bandwidths.

of the robust model is very close to the coverage obtained by the ex-post model, especially for high bandwidths. The improvement of the robust deployment over the nominal deployment appears most pronounced at higher distances ($\geq$ 100m), which is expected behavior since the AED locations are optimized for the tail of the distance distribution.

4.6.3 Alternate distance measures

We repeated a set of experiments using two alternate distance measures: the rectilinear distance, and actual walking distances based on city roads and walking paths. Table 4.5 shows the performance of each model when Euclidean, rectilinear and walking distances are used to compute the distances (for both the d_{ik} parameters and the performance metrics). When measured using actual walking distances, the robust and nominal deployments perform similarly on the mean distance metric. However, the robust model attains a lower VaR and CVaR, indicating improved performance by the robust deployment at the tail of the distance distribution. Table 4.6 provides the accompanying average and worst-case performance gaps of the nominal and robust solutions with respect to the ex-post solution. We found the results

to be consistent with our main finding in Section 4.6.1: the AED deployments produced by the robust model outperform the nominal model on all four performance metrics, both on average and in the worst-case realization of cardiac arrest locations. On the CVaR performance metric (which we emphasize is the most relevant point of comparison since all three models optimize for CVaR), the robust model cuts the performance gap in half, compared to the nominal model. It should be noted that the walking distances from the Google Maps API are likely conservative estimates of the actual distance a lay responder would have to travel, since it constrains the lay responder to walk on recognized pedestrian paths. In practice, a lay responder may be able to take shortcuts that are unmarked in Google Maps (e.g., parking lots, buildings, unmarked crosswalks). We therefore expect the actual lay responder distance to be somewhere between the Euclidean and Google Maps distances.

Table 4.5: Average and worst-case performance of nominal, robust and ex-post solutions over 50 simulated validation sets using different distance measures ($\beta = 0.9$, $h = 100$).

		Mean			VaR			CVaR			Max		
		Nom.	Rob.	Ex.	Nom.	Rob.	Ex.	Nom.	Rob.	Ex.	Nom.	Rob.	Ex.
Average	Euclidean	95	86	81	173	147	133	223	189	179	303	260	252
	Rectilinear	121	107	101	216	190	164	280	248	223	386	342	316
	Walking	134	133	119	233	221	202	281	269	256	367	365	355
Worst-case	Euclidean	109	96	95	202	181	181	260	222	215	408	375	375
	Rectilinear	136	120	119	263	241	215	345	311	283	553	553	498
	Walking	150	148	137	256	243	236	326	315	307	558	558	572

Table 4.6: Average and worst-case performance gap of nominal and robust solutions with respect to ex-post solution under alternate distance measures.

		Mean		VaR		CVaR		Max	
		Nom.	Rob.	Nom.	Rob.	Nom.	Rob.	Nom.	Rob.
Average	Euclidean	15%	6%	23%	9%	20%	5%	17%	3%
	Rectilinear	17%	6%	24%	13%	21%	10%	18%	8%
	Walking	12%	11%	13%	8%	9%	5%	7%	3%
Worst-case	Euclidean	29%	17%	40%	26%	27%	16%	31%	24%
	Rectilinear	28%	18%	41%	27%	32%	20%	41%	23%
	Walking	21%	20%	31%	24%	20%	12%	24%	22%

As an additional check to compare Euclidean and actual walking distances, we performed the following experiment. We randomly generated 1,000 pairs of locations within the downtown region of the City of Toronto, which includes the service area in our study. For each pair of locations, we computed the Euclidean distance (based on the UTM coordinates of the locations) and the walking distance (via the Google Maps Distance Matrix API). We then performed an ordinary least squares regression on the distance data, which yielded the following model:

$$Dist_{\text{Walk}} = 1.341 \times Dist_{\text{Euclid}} + 21.1.$$

The R^2 coefficient of this model was 0.904, suggesting a strong linear relationship between Euclidean and actual walking distances. Both coefficients were statistically significant ($p < 0.05$). This finding is consistent with previous studies that have also found very high correlation between Euclidean and road network distances (Bach, 1981, Jones et al., 2010, Boscoe et al., 2012). These results suggest that Euclidean distances serve as a reasonable (constant factor) approximation for actual walking distances.

4.6.4 Visualization of defibrillator deployments

Figure 4.8 depicts an instance of the AED deployments produced by the nominal (SAA), robust and ex-post models, along with a set of simulated cardiac arrests (corresponding to the ex-post deployment) and the 43 historical cardiac arrest locations. Note that in this instance, the nominal model produces small "clusters" of AEDs, whereas the robust and ex-post models appear to position AEDs more efficiently with respect to the simulated cardiac arrests. This behavior is expected, since the nominal model only optimizes the tail of the empirical distance distribution, whereas the robust model accounts for potential deviation of the simulated cardiac arrests from the historical locations. Note also that the ex-post model achieves the most effective deployment of AEDs with respect to the simulated cardiac arrests, since it optimizes directly with respect to the simulated data.

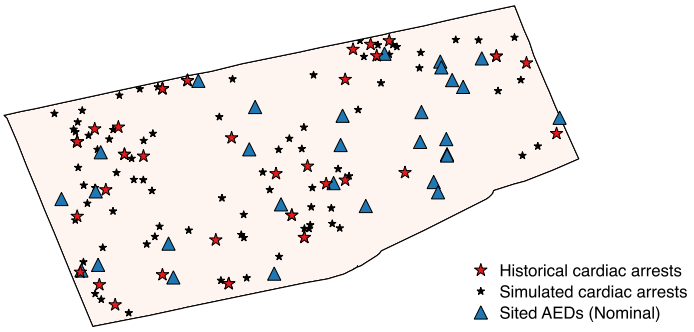

(a) Nominal.

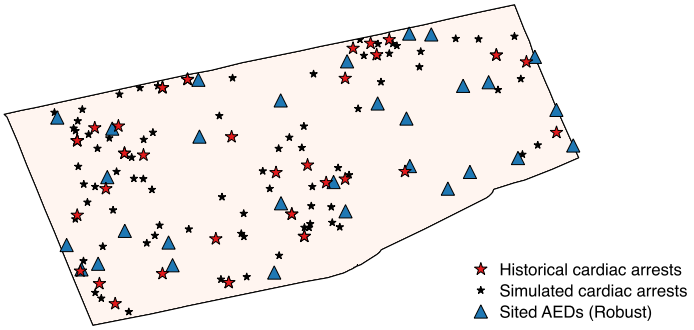

(b) Robust.

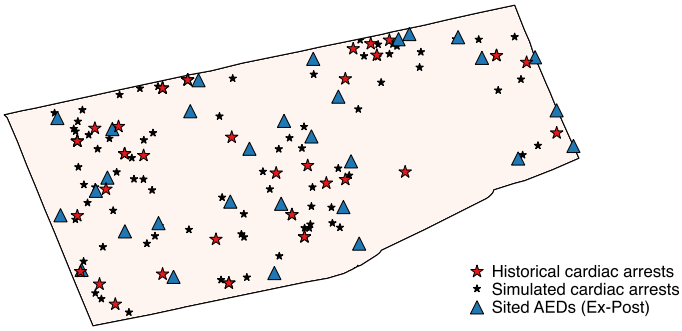

(c) Ex-post.

Figure 4.8: Example of nominal, robust and ex-post AED deployments ($P = 30$, $h = 100$, $\beta = 0.9$).

4.7 Discussion

Interestingly, our results indicate that the robust deployment performed better than the nominal deployment under *typical* demand realizations (Figure 4.5), in addition to the worst-case realizations (Figure 4.6). This might seem like a counterintuitive result, given that robust optimization models can perform poorly in typical realizations of the uncertainty due to their emphasis on optimizing a worst-case objective function. However, in the context of AED deployment (and facility location more generally), the observed performance gap between the robust and nominal models can likely be attributed to the nominal model overfitting to the set of historical cardiac arrest locations. In other words, the nominal model represents a naive sample average approximation in which AED placement is optimized only with respect to a small set of historical cardiac arrest locations. Our results suggest that if there is a nontrivial amount of uncertainty in the locations of future cardiac arrest, a sample average approximation based on limited historical data can lead to AED deployments that perform poorly when measured against out-of-sample cardiac arrests. On the other hand, our robust formulation produces an improved deployment by accounting for the possibility of cardiac arrests occuring in new locations not reflected in the historical data.

Survival outcomes. Our results suggest that a robust approach to AED deployment can improve cardiac arrest response by improving bystander accessibility to AEDs, especially for those patients who collapse far away from the nearest defibrillator. While a rigorous analysis of survival outcomes is beyond the scope of this work, the time sensitive nature of cardiac arrest treatment suggests that even modest improvements in AED accessibility (through shorter distances) can have a significant impact on survival outcomes. We also emphasize that all numerical results presented in Section 4.6 are with respect to the one-way distance to the nearest AED. However, since retrieval of an AED by a lay responder likely requires a round-trip, decreasing the distance between the cardiac arrest victim and the nearest AED by 30 meters implies that the trip made by the lay responder is shortened by 60 meters, which can have a material impact on the total time between collapse and treatment.

Comparison to best practices. Historically, best practices for AED deployment (e.g, as prescribed by the American Heart Association) have been purely retrospective, and do not account for the uncertainty in future cardiac arrest locations, which our results suggest can help improve the accessibility of AEDs. Existing guidelines also emphasize the importance of identifying building *types* that are at risk of cardiac arrest, and placing AEDs in those locations. Our proposed approach is markedly different from the current AHA recommendations. We instead focus on characterizing the distributional uncertainty in cardiac arrest locations, and optimally deploying AEDs in a manner that does not depend on building type, as long as the AED can be accessed by a lay responder. We also note that placing AEDs in a small set of high risk buildings ignores the shape of the distance distribution, and may lead to AED deployments where a large number of cardiac arrests occur far away from any AED. In contrast, our approach considers the entirety of the (worst-case) distance distribution, and allows the modeler to directly optimize various aspects of the distribution via the CVaR objective function.

Crowdsourcing lay responders. The existing focus on identifying high-risk buildings for AED placement is partially due to the current low rate of bystander intervention during cardiac arrests, and a prevailing belief that an AED must be located at the precise location of a cardiac arrest in order to be used. However, recent advances in mobile phone applications

show promise in recruiting lay responders (who may be up to 500m away) to perform CPR on victims of cardiac arrest (Ringh et al., 2015, Brooks et al., 2016, 2014). This "crowdsourcing" approach to cardiac arrest interventions can be extended to incorporate AEDs as well, by notifying lay responders of the locations of both the cardiac arrest victim and a nearby AED. We posit that in a setting where lay responders are recruited to intervene during a cardiac arrest, it is even more critical that AEDs are tactically located throughout an urban area, so that they can be quickly located and transported to a victim of cardiac arrest by a lay responder. Thus, our approach to strategically deploying AEDs can work synergistically with other innovations in the chain of survival to improve the overall response to cardiac arrest.

Risk measures in public access defibrillation. Our model also highlights the potential role that risk measures like CVaR can play in public access defibrillation. While VaR-based response time targets are common in ambulance dispatching (Pons and Markovchick, 2002), no similar guidelines exist in public access defibrillation. Given the potential impact on survival of protecting against large distances between cardiac arrest victims and nearby AEDs, it may be worthwhile to explicitly incorporate risk-based targets in the decision making process when determining AED locations. this work formally introduces the notion of risk to the AED deployment problem, by allowing decision makers to emphasize the tail of the distance (or response time) distribution when planning AED locations.

Future work. We highlight a few potential directions for future work. First, we have assumed throughout this paper that the set of uncertainty regions is given. It may be fruitful to investigate data-driven approaches to identifying an effective configuration of the uncertainty regions. Similarly, we note that the uncertainty set proposed in this paper can be generalized to permit the arrival probabilities to be uncertain as well. For example, we might allow each λ_j to reside within a 95% confidence interval that is also estimated from the data, which would retain the polyhedral structure of the uncertainty set. Our model can also be extended to incorporate additional considerations in public access defibrillation, such as lay responder behavior. However, we expect that our main finding in this paper – that accounting for cardiac arrest location uncertainty improves AED accessibility – would persist under such modifications to the model. Lastly, the main features of our modeling approach – a distributional uncertainty set that induces sparse worst-case distributions coupled with an efficient row-and-column generation algorithm – may find relevance in other applications with a similar problem structure as well.

Chapter 5

Conclusion

In this work we presented a new set of data-driven models and algorithms with practical applications to problems in healthcare operations and policy. We briefly summarize the key contributions of each chapter here.

In Chapter 2, we developed and analyzed a formulation for inverse optimization in the setting where noisy measurements of the optimal points of a convex optimization problem are available. Our approach requires solving a bilevel program, and we defined a new duality-based reformulation to convert this bilevel program into a single level program. We showed that our formulation as a bilevel program leads to statistical consistency, in contrast to existing heuristics. Although our formulation is NP-hard to solve, we provided two numerical algorithms that maintain the statistical consistency of our formulation. Finally, we demonstrated that our approach improves upon existing methods for inverse optimization through a series of numerical experiments using both synthetic and empirical data.

In Chapter 3, we proposed a new type of contract to address the challenges faced by the Medicare Shared Savings Program (MSSP). In addition to the shared savings payment, we proposed that the MSSP incorporate direct subsidies to partially reimburse ACO investments toward efficiency improvements. We showed that selecting the shared savings and subsidy rates appropriately yields a contract that dominates all possible contracts within the current, non-subsidized program, in the sense that it boosts both Medicare savings and ACO payments. We also quantified the improvement potential through structural estimation of the principal-agent model, and found that switching to the proposed contract can increase both Medicare savings and ACO payments. By contrast, optimizing the shared savings rate within the current, non-subsidized MSSP contract will decrease ACO payments substantially, with minimal benefits to Medicare. We also found that the proposed contract performs well in comparison to a non-parametric contract that represents the theoretical best performance that can be expected from the MSSP, under the estimated agent model. We also found that ACOs with low benchmarks face difficulties in generating savings. The average savings of an ACO with a benchmark greater than \$14,300 per beneficiary was estimated to be \$240 ($\approx$2%), whereas the average savings of an ACO with a benchmark of less than \$10,300 per beneficiary was estimated to be effectively zero. We believe these findings provide supporting evidence in favor of redesigning the benchmarking methodology. One possible approach may be to develop benchmarks partially based on regional averages, so that historically cost efficient ACOs are not inadvertently penalized through the assignment of a lower financial

benchmark.

In Chapter 4, we developed a data-driven optimization approach for deploying automated external defibrillators (AEDs) in public spaces while accounting for uncertainty in the locations of future cardiac arrests. Our approach involves constructing a distributional uncertainty set that consists of several uncertainty regions, where the parameters of the uncertainty set are calibrated based on historical cardiac arrest data. To approximate the arrival of cardiac arrests in continuous space, we discretize the service area into an extremely large set of scenarios, and develop a row-and-column generation algorithm that exploits the structure of the uncertainty set and scales gracefully in the number of scenarios. As an auxilliary result, we show that our formulation subsumes a large class of facility location problems, and unifies the classical p-median and p-center problems, as well as their robust analogues. Our numerical results suggest that hedging against cardiac arrest location uncertainty can lead to improved AED accessibility under both typical and worst-case realizations of the uncertainty. In particular, we found that our robust AED deployment outperformed a nominal (sample average approximation) deployment by between 9-20% with respect to AED retrieval distances, depending on the performance metric and underlying demand distribution. Further, we found that in many cases, the robust deployment performed nearly as well as an ex-post model with perfect foresight, and in many instances improved the performance gap with respect to the ex-post model by 40-70%. This finding suggests that our robust approach manages to avoid the performance loss due to uncertainty that is suffered by the sample average approach.

Bibliography

Aalami, H., Moghaddam, M. P., and Yousefi, G. (2010). Demand response modeling considering interruptible/curtailable loads and capacity market programs. *Applied Energy*, 87(1):243–250.

Adida, E., Mamani, H., and Nassiri, S. (2016). Bundled payment vs. fee-for-service: Impact of payment scheme on performance. *Management Science*.

Ahuja, R. K. and Orlin, J. B. (2001a). Inverse optimization. *Operations Research*, 49(5):771–783.

Ahuja, R. K. and Orlin, J. B. (2001b). Inverse optimization. *Operations Research*, 49(5):771–783.

Akira Health, Inc. (2017). Akira Health Accountable Care Organization.

Alderson, D. L., Brown, G. G., Carlyle, W. M., and Wood, R. K. (2011). Solving defender-attacker-defender models for infrastructure defense. *12th INFORMS Computing Society Conference*.

Anderson, G. F. and Steinberg, E. P. (1984). Hospital readmissions in the Medicare population. *New England Journal of Medicine*, 311(21):1349–1353.

Andritsos, D. and Tang, C. S. (2015). Incentive Programs for Reducing Readmissions When Patient Care is Co-Produced. *Available at SSRN 2666215*.

Arifoglu, K., Deo, S., and Iravani, S. M. (2012). Consumption externality and yield uncertainty in the influenza vaccine supply chain: Interventions in demand and supply sides. *Management Science*, 58(6):1072–1091.

ASHRAE (2013). *ANSI/ASHRAE Standard 55-2013: Thermal Environmental Conditions for Human Occupancy*. ASHRAE.

Aswani, A. (2015). Low-rank approximation and completion of positive tensors. arXiv:1412.0620.

Aswani, A., Gonzalez, H., Sastry, S., and Tomlin, C. (2013). Provably safe and robust learning–based model predictive control. *Automatica*, 49(5):1216–1226.

Aswani, A., Kaminsky, P., Mintz, Y., Flowers, E., and Fukuoka, Y. (2016). Predictive modeling of behavior in weight loss interventions. Submitted.

Aswani, A., Master, N., Taneja, J., Krioukov, A., Culler, D., and Tomlin, C. (2012a). Energy–efficient building HVAC control using hybrid system LBMPC. In *IFAC Conference on Nonlinear Model Predictive Control*.

Aswani, A., Master, N., Taneja, J., Krioukov, A., Culler, D., and Tomlin, C. (2012b). Quantitative methods for comparing different HVAC control schemes. In *International Conference on Performance Evaluation Methodologies and Tools*.

Aswani, A., Master, N., Taneja, J., Smith, V., Krioukov, A., Culler, D., and Tomlin, C. (2012c). Identifying models of HVAC systems using semi-parametric regression. In *American Control Conference*.

Aswani, A., Shen, Z.-J. M., and Siddiq, A. (2018a). Data-driven incentive design in the medicare shared savings program.

Aswani, A., Shen, Z.-J. M., and Siddiq, A. (2018b). Inverse optimization with noisy data. *Operations Research*.

Aswani, A. and Tomlin, C. (2012). Incentive design for efficient building quality of service. In *Allerton Conference on Communication, Control, and Computing*, pages 90–97.

Ata, B., Killaly, B. L., Olsen, T. L., and Parker, R. P. (2013). On hospice operations under medicare reimbursement policies. *Management Science*, 59(5):1027–1044.

Ata, B., Lee, D., and Tongarlak, M. H. (2012). Optimizing organic waste to energy operations. *Manufacturing & Service Operations Management*, 14(2):231–244.

Atamturk, A. and Zhang, M. (2007). Two-stage robust network flow and design under demand uncertainty. *Operations Research*, 55(4):662–673.

Audet, C., Hansen, P., Jaumard, B., and Savard, G. (1997). Links between linear bilevel and mixed 0–1 programming problems. *Journal of Optimization Theory and Applications*, 93(2):273–300.

Averbakh, I. (2003). Complexity of robust single facility location problems on networks with uncertain edge lengths. *Discrete Applied Mathematics*, 127(3):505–522.

Averbakh, I. and Bereg, S. (2005). Facility location problems with uncertainty on the plane. *Discrete Optimization*, 2(1):3–34.

Averbakh, I. and Berman, O. (2000). Algorithms for the robust 1-center problem on a tree. *European Journal of Operational Research*, 123(2):292–302.

Bach, L. (1981). The problem of aggregation and distance for analyses of accessibility and access opportunity in location-allocation models. *Environment and Planning A*, 13(8):955–978.

Bajari, P., Benkard, C. L., and Levin, J. (2007). Estimating dynamic models of imperfect competition. *Econometrica*, 75(5):1331–1370.

Balasubramanian, S. and Bhardwaj, P. (2004). When not all conflict is bad: Manufacturing-marketing conflict and strategic incentive design. *Management Science*, 50(4):489–502.

Bard, J. F. and Moore, J. T. (1990). A branch and bound algorithm for the bilevel programming problem. *SIAM Journal on Scientific and Statistical Computing*, 11(2):281–292.

Baron, O., Milner, J., and Naseraldin, H. (2011). Facility location: A robust optimization approach. *Production and Operations Management*, 20(5):772–785.

Bartlett, P. and Mendelson, S. (2002). Rademacher and gaussian complexities: Risk bounds and structural results. *J. Mach. Learn. Res.*

Bastani, H., Bayati, M., Braverman, M., Gummadi, R., and Johari, R. (2016). Analysis of medicare pay-for-performance contracts. *Available at SSRN 2839143*.

Becker, L., Eisenberg, M., Fahrenbruch, C., and Cobb, L. (1998). Public locations of cardiac arrest implications for public access defibrillation. *Circulation*, 97(21):2106–2109.

Beil, D. R. and Wein, L. M. (2003). An inverse-optimization-based auction mechanism to support a multiattribute rfq process. *Management Science*, 49(11):1529–1545.

Ben-Tal, A., Den Hertog, D., De Waegenaere, A., Melenberg, B., and Rennen, G. (2013). Robust solutions of optimization problems affected by uncertain probabilities. *Management Science*, 59(2):341–357.

Ben-Tal, A., Ghaoui, L. E., and Nemirovski, A. (2009). *Robust optimization*. Princeton University Press.

Berge, C. (1963). *Topological Spaces: including a treatment of multi-valued functions, vector spaces, and convexity*. Courier Dover Publications.

Bertsimas, D., Brown, D., and Caramanis, C. (2011). Theory and applications of robust optimization. *SIAM Review*, 53(3):464–501.

Bertsimas, D., Gupta, V., and Paschalidis, I. C. (2012). Inverse optimization: a new perspective on the black-litterman model. *Operations research*, 60(6):1389–1403.

Bertsimas, D., Gupta, V., and Paschalidis, I. C. (2015). Data-driven estimation in equilibrium using inverse optimization. *Mathematical Programming*, 153(2):595–633.

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

Bertsimas, D. and Thiele, A. (2006). Robust and data-driven optimization: Modern decision-making under uncertainty. *INFORMS Tutorials in Operations Research: Models, Methods, and Applications for Innovative Decision Making*.

Bertsimas, D. and Tsitsiklis, J. N. (1997). *Introduction to linear optimization*, volume 6. Athena Scientific Belmont, MA.

Berwick, D. M. (2011). Launching accountable care organizations: The proposed rule for the Medicare Shared Savings Program. *New England Journal of Medicine*, 364(16):e32.

Berwick, D. M. and Hackbarth, A. D. (2012). Eliminating waste in US health care. *Journal of the American Medical Association*, 307(14):1513–1516.

Besbes, O., Iancu, D. A., and Trichakis, N. (2016). Dynamic pricing under debt: Spiraling distortions and efficiency losses. Technical report, Working Paper.

Bickel, P. and Doksum, K. (2006). *Mathematical Statistics: Basic Ideas And Selected Topics*, volume 1. Pearson Prentice Hall, 2nd edition.

Bickel, P. J. and Doksum, K. A. (2015). *Mathematical statistics: basic ideas and selected topics*, volume 2. CRC Press.

Birge, J. R. and Louveaux, F. (2011). *Introduction to stochastic programming*. Springer Science & Business Media.

Birge, J. R. and Wets, R. J.-B. (1987). Computing bounds for stochastic programming problems by means of a generalized moment problem. *Mathematics of Operations Research*, 12(1):149–162.

Bonnans, J. and Shapiro, A. (2000). *Perturbation Analysis of Optimization Problems*. Springer.

Bonnans, J. F. and Ioffe, A. (1995a). Second-order sufficiency and quadratic growth for nonisolated minima. *Mathematics of Operations Research*, 20(4):801–817.

Bonnans, J. F. and Ioffe, A. D. (1995b). Quadratic growth and stability in convex programming problems with multiple solutions. *J. Convex Anal*, 2(1-2):41–57.

Borgers, T., Strausz, R., and Krahmer, D. (2015). *An introduction to the theory of mechanism design*. Oxford University Press, USA.

Boscoe, F. P., Henry, K. A., and Zdeb, M. S. (2012). A nationwide comparison of driving distance versus straight-line distance to hospitals. *The Professional Geographer*, 64(2):188–196.

Boyd, S. and Vandenberghe, L. (2009). *Convex optimization*. Cambridge university press.

Brooks, S., Worthington, H., Gonedalles, T., Bobrow, B., and Morrison, L. (2014). Implementation of the pulsepoint smartphone application for crowd-sourcing bystander resuscitation. *Critical Care*, 18(Suppl 1):P484.

Brooks, S. C., Hsu, J. H., Tang, S. K., Jeyakumar, R., and Chan, T. C. (2013). Determining risk for out-of-hospital cardiac arrest by location type in a canadian urban setting to guide future public access defibrillator placement. *Annals of emergency medicine*, 61(5):530–538.

Brooks, S. C., Simmons, G., Worthington, H., Bobrow, B. J., and Morrison, L. J. (2016). The pulsepoint respond mobile device application to crowdsource basic life support for patients with out-of-hospital cardiac arrest: Challenges for optimal implementation. *Resuscitation*, 98:20–26.

Brotcorne, L., Laporte, G., and Semet, F. (2003). Ambulance location and relocation models. *European journal of operational research*, 147(3):451–463.

Brown, G., Carlyle, M., Salmeron, J., and Wood, K. (2006). Defending critical infrastructure. *Interfaces*, 36(6):530–544.

Brown, G. G., Carlyle, W. M., Harney, R. C., Skroch, E. M., and Wood, R. K. (2009). Interdicting a nuclear-weapons project. *Operations Research*, 57(4):866–877.

Building Robotics (2016). Comfy.

Burton, D. and Toint, P. L. (1992). On an instance of the inverse shortest paths problem. *Mathematical Programming*, 53(1-3):45–61.

Caffrey, S. L., Willoughby, P. J., Pepe, P. E., and Becker, L. B. (2002). Public use of automated external defibrillators. *New England Journal of Medicine*, 347(16):1242–1247.

Carlsson, J. G. and Delage, E. (2013). Robust partitioning for stochastic multivehicle routing. *Operations Research*, 61(3):727–744.

Carr, S. and Lovejoy, W. (2000). The inverse newsvendor problem: Choosing an optimal demand portfolio for capacitated resources. *Management Science*, 46(7):912–927.

Chan, T. C., Craig, T., Lee, T., and Sharpe, M. B. (2014). Generalized inverse multiobjective optimization with application to cancer therapy. *Operations Research*.

Chan, T. C., Li, H., Lebovic, G., Tang, S. K., Chan, J. Y., Cheng, H. C., Morrison, L. J., and Brooks, S. C. (2013). Identifying locations for public access defibrillators using mathematical optimization. *Circulation*, 127(17):1801–1809.

Chan, T. C. and Mišić, V. V. (2013). Adaptive and robust radiation therapy optimization for lung cancer. *European Journal of Operational Research*, 231(3):745–756.

Chan, T. C., Shen, Z.-J. M., and Siddiq, A. (2018). Robust defibrillator deployment under cardiac arrest location uncertainty via row-and-column generation. *Operations Research*.

Chatterjee, S. (2014). A new perspective on least squares under convex constraint. *Ann. Statist.*, 42(6):2340–2381.

Chemama, J., Cohen, M. C., Lobel, R., and Perakis, G. (2014). Consumer Subsidies with a Strategic Supplier: Commitment vs. Flexibility.

Chen, F. (2000). Sales-force incentives and inventory management. *Manufacturing & Service Operations Management*, 2(2):186–202.

Chen, G., Daskin, M. S., Shen, Z.-J. M., and Uryasev, S. (2006). The -reliable mean-excess regret model for stochastic facility location modeling. *Naval Research Logistics*, 53.

Chen, S. and Lee, H. (2016). Incentive Alignment and Coordination of Project Supply Chains. *Management Science*.

Chen, T., Klastorin, T., and Wagner, M. R. (2015). Incentive Contracts in Serial Stochastic Projects. *Manufacturing & Service Operations Management*, 17(3):290–301.

Chernew, M., McGuire, T., and McWilliams, J. M. (2014). Refining the ACO Program: Issues and Options. *Department of Healthcare Policy, Harvard Medical School*.

Chick, S. E., Mamani, H., and Simchi-Levi, D. (2008). Supply chain coordination and influenza vaccination. *Operations Research*, 56(6):1493–1506.

Cho, S.-H., Jang, H., Lee, T., and Turner, J. (2014). Simultaneous location of trauma centers and helicopters for emergency medical service planning. *Operations Research*, 62(4):751–771.

Church, R. L. and Scaparra, M. P. (2007). Protecting critical assets: the r-interdiction median problem with fortification. *Geographical Analysis*, 39(2):129–146.

City of Toronto (2010). Toronto employment survey (for 2009).

City of Toronto (2016). Demographics - neighbourhood profiles.

CMS (2016a). 2016 Press release: Physicians and health care providers continue to improve quality of care, lower costs.

CMS (2016b). Centers for Medicare and Medicaid Services.

CMS (2016c). Centers for Medicare and Medicaid Services.

CMS (2016d). Improving Quality of Care for Medicare Patients: Accountable Care Organizations.

CMS (2016e). Medicare Shared Savings Program Accountable Care Organizations Performance Year 2015 Results.

CMS (2017). Bundled Payments for Care Improvement (BPCI) Initiative: General Information.

Cohen, M. C., Lobel, R., and Perakis, G. (2015a). The impact of demand uncertainty on consumer subsidies for green technology adoption. *Management Science*.

Cohen, M. C., Perakis, G., and Thraves, C. (2015b). Competition and externalities in green technology adoption. *Available at SSRN 2607688*.

Cooper, L. (1978). Bounds on the weber problem solution under conditions of uncertainty. *Journal of Regional Science*, 18(1):87–92.

Crama, P., Reyck, B. D., and Degraeve, Z. (2008). Milestone payments or royalties? contract design for r&d licensing. *Operations Research*, 56(6):1539–1552.

Crosson, F. J. (2011). The accountable care organization: whatever its growing pains, the concept is too vitally important to fail. *Health Affairs*, 30(7):1250–1255.

Cui, T., Ouyang, Y., and Shen, Z.-J. M. (2010). Reliable facility location design under the risk of disruptions. *Operations Research*, 58(4):998–1011.

DeHoratius, N. and Raman, A. (2007). Store manager incentive design and retail performance: An exploratory investigation. *Manufacturing & Service Operations Management*, 9(4):518–534.

Delage, E. and Ye, Y. (2010). Distributionally robust optimization under moment uncertainty with application to data-driven problems. *Operations Research*, 58(3):595–612.

Dempe, S. (2002). *Foundations of bilevel programming*. Springer Science & Business Media.

Dempe, S., Kalashnikov, V., Pérez-Valdés, G., and Kalashnikova, N. (2015). *Bilevel Programming Problems*. Springer.

Douven, R., McGuire, T. G., and McWilliams, J. M. (2015). Avoiding unintended incentives in ACO payment models. *Health Affairs*, 34(1):143–149.

Drezner, Z. (1989). Stochastic analysis of the weber problem on the sphere. *Journal of the Operational Research Society*, pages 1137–1144.

Dupačová, J. (1980). On minimax decision rule in stochastic linear programming. *Studies in Mathematical Programming*, pages 47–60.

Eddy, D. M. and Shah, R. (2012). A simulation shows limited savings from meeting quality targets under the Medicare shared savings program. *Health Affairs*, 31(11):2554–2562.

Efron, B. and Tibshirani, R. (1994a). *An Introduction to the Bootstrap*. Chapman & Hall/CRC Monographs on Statistics & Applied Probability. Taylor & Francis.

Efron, B. and Tibshirani, R. J. (1994b). *An introduction to the bootstrap*. CRC press.

Engdahl, J. and Herlitz, J. (2005). Localization of out-of-hospital cardiac arrest in göteborg 1994–2002 and implications for public access defibrillation. *Resuscitation*, 64(2):171–175.

Erkin, Z., Bailey, M. D., Maillart, L. M., Schaefer, A. J., and Roberts, M. S. (2010). Eliciting patients' revealed preferences: An inverse markov decision process approach. *Decision Analysis*, 7(4):358–365.

Erkut, E., Ingolfsson, A., and Erdoğan, G. (2008). Ambulance location for maximum survival. *Naval Research Logistics*, 55(1):42–58.

Faragó, A., Szentesi, Á., and Szviatovszki, B. (2003). Inverse optimization in high-speed networks. *Discrete Applied Mathematics*, 129(1):83–98.

Farias, V. F., Jagabathula, S., and Shah, D. (2013). A nonparametric approach to modeling choice with limited data. *Management Science*, 59(2):305–322.

Federal Register (2011). Medicare Shared Savings Program: Accountable Care Organizations. Final rule. *Federal Register*, 76(212):67802.

Fetter, R. B. (1991). Diagnosis related groups: understanding hospital performance. *Interfaces*, 21(1):6–26.

Fisher, E. S., Shortell, S. M., Kreindler, S. A., Van Citters, A. D., and Larson, B. K. (2012). A framework for evaluating the formation, implementation, and performance of accountable care organizations. *Health Affairs*, 31(11):2368–2378.

Folke, F., Lippert, F., Nielsen, S., Gislason, G., Hansen, M., Schramm, T., and Sorensen, R. (2009). Location of cardiac arrest in a city center strategic placement of automated external defibrillators in public locations. *Circulation*, 120(6):510–517.

Francis, R., Lowe, T. J., and Tamir, A. (2000). Aggregation error bounds for a class of location models. *Operations research*, 48(2):294–307.

Fuloria, P. C. and Zenios, S. A. (2001). Outcomes-adjusted reimbursement in a health-care delivery system. *Management Science*, 47(6):735–751.

Gabrel, V., Lacroix, M., Murat, C., and Remli, N. (2014). Robust location transportation problems under uncertain demands. *Discrete Applied Mathematics*, 164(1):100–111.

Garber, A. M. and Skinner, J. (2008). Is American health care uniquely inefficient? Technical report, National Bureau of Economic Research.

Ghaoui, L. E., Oks, M., and Oustry, F. (2003). Worst-case value-at-risk and robust portfolio optimization: A conic programming approach. *Operations research*, 51(4):543–556.

Gibbons, R. (1998). Incentives in organizations. Technical report, National bureau of economic research.

Goh, J. and Sim, M. (2010). Distributionally robust optimization and its tractable approximations. *Operations research*, 58(4-part-1):902–917.

Google (2016). Google maps distance matrix api.

Green, P. E. and Srinivasan, V. (1990). Conjoint analysis in marketing: new developments with implications for research and practice. *The Journal of Marketing*, pages 3–19.

Greenshtein, E. and Ritov, Y. (2004). Persistence in high-dimensional linear predictor selection and the virtue of overparametrization. *Bernoulli*, 10(6):971–988.

Grinblatt, M. and Titman, S. (1989). Adverse risk incentives and the design of performance-based contracts. *Management Science*, 35(7):807–822.

Grossman, S. J. and Hart, O. D. (1983). An analysis of the principal-agent problem. *Econometrica: Journal of the Econometric Society*, pages 7–45.

Guajardo, J. A., Cohen, M. A., Kim, S.-H., and Netessine, S. (2012). Impact of performance-based contracting on product reliability: An empirical analysis. *Management Science*, 58(5):961–979.

Gundry, J. W., Comess, K. A., DeRook, F. A., Jorgenson, D., and Bardy, G. H. (1999). Comparison of naive sixth-grade children with trained professionals in the use of an automated external defibrillator. *Circulation*, 100(16):1703–1707.

Guo, P., Tang, C. S., Wang, Y., and Zhao, M. (2016). The impact of reimbursement policy on patient welfare, readmission rate and waiting time in a public healthcare system: Fee-for-service vs bundled payment. *Working paper, London Business School*.

Gupta, D. and Mehrotra, M. (2015). Bundled Payments For Healthcare Services: Proposer Selection and Information Sharing. *Operations Research*, 63(4):772–788.

Harris, M. and Raviv, A. (1979). Optimal incentive contracts with imperfect information. *Journal of economic theory*, 20(2):231–259.

Hart, O. D. and Holmström, B. (1986). *The theory of contracts*. Department of Economics, Massachusetts Institute of Technology.

Hastie, T., Tibshirani, R., and Friedman, J. (2009). *The Elements of Statistical Learning*. Springer-Verlag, 2nd edition.

Hastie, T., Tibshirani, R., Friedman, J., and Franklin, J. (2005). The elements of statistical learning: data mining, inference and prediction. *The Mathematical Intelligencer*, 27(2):83–85.

Haviv, I. and Regev, O. (2012). Tensor-based hardness of the shortest vector problem to within almost polynomial factors. *Theory of Computing*, 8(1):513–531.

Haywood, T. T. and Kosel, K. C. (2011). The ACO model – a three-year financial loss? *New England Journal of Medicine*, 364(14):e27.

Hazinski, M. F., Idris, A. H., Kerber, R. E., Epstein, A., Atkins, D., Tang, W., and Lurie, K. (2005). Lay rescuer automated external defibrillator (public access defibrillation) programs lessons learned from an international multicenter trial: Advisory statement from the american heart association emergency cardiovascular committee; the council on cardiopulmonary, perioperative, and critical care; and the council on clinical cardiology. *Circulation*, 111(24):3336–3340.

Heart and Stroke Foundation (2013). Statistics.

Hendee, W. R., Becker, G. J., Borgstede, J. P., Bosma, J., Casarella, W. J., Erickson, B. A., Maynard, C. D., Thrall, J. H., and Wallner, P. E. (2010). Addressing Overutilization in Medical Imaging. *Radiology*, 257(1):240–245.

Heuberger, C. (2004). Inverse combinatorial optimization: A survey on problems, methods, and results. *Journal of Combinatorial Optimization*, 8(3):329–361.

Hillar, C. and Lim, L.-H. (2013). Most tensor problems are np-hard. *J. ACM*, 60(6):45:1–45:39.

Hochbaum, D. S. (2003). Efficient algorithms for the inverse spanning-tree problem. *Operations Research*, 51(5):785–797.

Hölmstrom, B. (1979). Moral hazard and observability. *The Bell journal of economics*, pages 74–91.

Iyengar, G. and Kang, W. (2005). Inverse conic programming with applications. *Operations Research Letters*, 33(3):319–330.

Jennrich, R. I. (1969). Asymptotic properties of non-linear least squares estimators. *The Annals of Mathematical Statistics*, pages 633–643.

Jiang, H., Pang, S., and Savin, S. (2016). Capacity management for outpatient medical services under competition and performance-based incentives. *Working paper, Wharton School, University of Pennsylvania*.

Jiang, H., Pang, Z., and Savin, S. (2012). Performance-based contracts for outpatient medical services. *Manufacturing & Service Operations Management*, 14(4):654–669.

Jones, S. G., Ashby, A. J., Momin, S. R., and Naidoo, A. (2010). Spatial implications associated with using euclidean distance measurements and geographic centroid imputation in health care research. *Health services research*, 45(1):316–327.

Kalcsics, J., Nickel, S., Puerto, J., and Tamir, A. (2002). Algorithmic results for ordered median problems. *Operations Research Letters*, 30(3):149–158.

Keshavarz, A., Wang, Y., and Boyd, S. (2011). Imputing a convex objective function. In *Intelligent Control (ISIC), 2011 IEEE International Symposium on*, pages 613–619. IEEE.

Khanjari, N. E., Iravani, S., and Shin, H. (2013). The impact of the manufacturer-hired sales agent on a supply chain with information asymmetry. *Manufacturing & Service Operations Management*, 16(1):76–88.

Khouja, M. and Zhou, J. (2010). The effect of delayed incentives on supply chain profits and consumer surplus. *Production and Operations Management*, 19(2):172–197.

Laffont, J.-J. and Martimort, D. (2009). *The theory of incentives: the principal-agent model*. Princeton university press.

Lariviere, M. A. (2015). OM ForumSupply Chain Contracting: Doughnuts to Bubbles. *Manufacturing & Service Operations Management*.

Larsen, M., Eisenberg, M., Cummins, R., and Hallstrom, A. (1993). Predicting survival from out-of-hospital 309 cardiac arrest: a graphic model. *Annals of Emergency Medicine*, (22):1652–1658.

Lee, D. K. and Zenios, S. A. (2012). An evidence-based incentive system for Medicare's End-Stage Renal Disease Program. *Management science*, 58(6):1092–1105.

Levi, R., Perakis, G., and Romero, G. (2016). On the Effectiveness of Uniform Subsidies in Increasing Market Consumption. *Management Science*.

Liberatore, F., Scaparra, M. P., and Daskin, M. S. (2011). Analysis of facility protection strategies against an uncertain number of attacks: the stochastic r-interdiction median problem with fortification. *Computers & Operations Research*, 38(1):357–366.

Lieberman, S. M. and Bertko, J. M. (2011). Building regulatory and operational flexibility into accountable care organizations and shared savings. *Health Affairs*, 30(1):23–31.

Lim, M., Daskin, M., Bassamboo, A., and Chopra, S. (2010). A facility reliability problem: Formulation, properties, and algorithm. *Naval Research Logistics*, 57(1):58–70.

Liu, P. and Wu, S. (2014). An agent-based simulation model to study accountable care organizations. *Health care management science*, pages 1–13.

Losada, C., Scaparra, M. P., Church, R. L., and Daskin, M. S. (2012). The stochastic interdiction median problem with disruption intensity levels. *Annals of Operations Research*, 201(1):345–365.

Mamani, H., Chick, S. E., and Simchi-Levi, D. (2013). A game-theoretic model of international influenza vaccination coordination. *Management Science*, 59(7):1650–1670.

McGlynn, E. A., Asch, S. M., Adams, J., Keesey, J., Hicks, J., DeCristofaro, A., and Kerr, E. A. (2003). The quality of health care delivered to adults in the United States. *New England journal of medicine*, 348(26):2635–2645.

McWilliams, J. M., Chernew, M. E., Landon, B. E., and Schwartz, A. L. (2015). Performance differences in year 1 of pioneer accountable care organizations. *New England Journal of Medicine*, 372(20):1927–1936.

McWilliams, J. M., Hatfield, L. A., Chernew, M. E., Landon, B. E., and Schwartz, A. L. (2016). Early Performance of Accountable Care Organizations in Medicare. *New England Journal of Medicine*.

McWilliams, J. M., Landon, B. E., and Chernew, M. E. (2013). Changes in health care spending and quality for Medicare beneficiaries associated with a commercial ACO contract. *JAMA*, 310(8):829–836.

Melo, M. T., Nickel, S., and Saldanha-Da-Gama, F. (2009). Facility location and supply chain management–a review. *European Journal of Operational Research*, 196(2):401–412.

Milgate, K. and Cheng, S. B. (2006). Pay-for-performance: the MedPAC perspective. *Health Affairs*, 25(2):413–419.

Miller, H. D. (2009). From volume to value: better ways to pay for health care. *Health Affairs*, 28(5):1418–1428.

Miller, M. (2013). *Medicare Payment Advisor Commission: Hospital Policy Issues*.

Mirrlees, J. (1974). Notes on welfare economics, information and uncertainty. *Essays on economic behavior under uncertainty*, pages 243–261.

Moore, K. D. et al. (2011). *The work ahead: activities and costs to develop an accountable care organization*. American Hospital Association.

Morrison, L. J., Nichol, G., Rea, T. D., Christenson, J., Callaway, C. W., Stephens, S., Pirrallo, R. G., Atkins, D. L., Davis, D. P., Idris, A. H., et al. (2008). Rationale, development and implementation of the Resuscitation Outcomes Consortium Epistry - Cardiac Arrest. *Resuscitation*, 78(2):161–169.

Mozaffarian, D., Benjamin, E. J., Go, A. S., Arnett, D. K., Blaha, M. J., Cushman, M., Das, S. R., de Ferranti, S., Després, J.-P., Fullerton, H. J., et al. (2016). Executive summary: Heart disease and stroke statistics-2016 update: A report from the american heart association. *Circulation*, 133(4):447.

Myerson, R. B. (1981). Optimal auction design. *Mathematics of operations research*, 6(1):58–73.

Nahmias, S. and Cheng, Y. (2009). *Production and operations analysis*, volume 5. McGraw-Hill New York.

National Association of ACOs (2014). National ACO Survey.

Nickel, S. and Puerto, J. (1999). A unified approach to network location problems. *Networks*, 34(4):283–290.

OECD (2015). Organization for Economic Cooperation and Development Health Statistics.

Ordonez, F. and Zhao, J. (2007). Robust capacity expansion of network flows. *Networks*, 50(2):136–145.

Owen, S. H. and Daskin, M. S. (1998). Strategic facility location: A review. *European Journal of Operational Research*, 111(3).

PAD Trial Investigators (2004). Public-access defibrillation and survival after out-of-hospital cardiac arrest. *N Engl J Med*, 2004(351):637–646.

Page, R. L., Joglar, J. A., Kowal, R. C., Zagrodzky, J. D., Nelson, L. L., Ramaswamy, K., Barbera, S. J., Hamdan, M. H., and McKenas, D. K. (2000). Use of automated external defibrillators by a us airline. *New England Journal of Medicine*, 343(17):1210–1216.

Pell, J. P., Sirel, J. M., Marsden, A. K., Ford, I., and Cobbe, S. M. (2001). Effect of reducing ambulance response times on deaths from out of hospital cardiac arrest: cohort study. *BMJ*, 322(7299):1385–1388.

Plambeck, E. L. and Zenios, S. A. (2000). Performance-based incentives in a dynamic principal-agent model. *Manufacturing & Service Operations Management*, 2(3):240–263.

Pons, P. T. and Markovchick, V. J. (2002). Eight minutes or less: does the ambulance response time guideline impact trauma patient outcome? *The journal of emergency medicine*, 23(1):43–48.

Portner, M. E., Pollack, M. L., Schirk, S. K., and Schlenker, M. K. (2004). Out-of-hospital cardiac arrest locations in a rural community: where should we place aeds? *Prehospital and disaster medicine*, 19(04):352–355.

PulsePoint (2015). PulsePoint AED.

Raghu, T., Sen, P., and Rao, H. R. (2003). Relative performance of incentive mechanisms: Computational modeling and simulation of delegated investment decisions. *Management Science*, 49(2):160–178.

Ratliff, L. J., Dong, R., Ohlsson, H., and Sastry, S. S. (2014a). Incentive design and utility learning via energy disaggregation. In *19th World Congress of the International Federation of Automatic Control*.

Ratliff, L. J., Dong, R., Ohlsson, H., and Sastry, S. S. (2014b). Incentive design and utility learning via energy disaggregation. In *Proceedings of the 19th IFAC World Congress*, pages 3158–3163.

Ringh, M., Rosenqvist, M., Hollenberg, J., Jonsson, M., Fredman, D., Nordberg, P., Järnbert-Pettersson, H., Hasselqvist-Ax, I., Riva, G., and Svensson, L. (2015). Mobile-phone dispatch of laypersons for cpr in out-of-hospital cardiac arrest. *New England Journal of Medicine*, 372(24):2316–2325.

Rockafellar, R. T. and Uryasev, S. (2000). Optimization of conditional value-at-risk. *Journal of Risk*, 2:21–42.

Rockafellar, R. T. and Uryasev, S. (2002). Conditional value-at-risk for general loss distributions. *Journal of Banking and Finance*, 26(7):1443–1471.

Rockafellar, R. T. and Wets, R. J.-B. (1998). *Variational analysis*, volume 317. Springer.

Rosenthal, M. B., Cutler, D. M., and Feder, J. (2011). The ACO rulesstriking the balance between participation and transformative potential. *New England Journal of Medicine*, 365(4):e6.

Rosenthal, M. B., Fernandopulle, R., Song, H. R., and Landon, B. (2004). Paying for quality: providers incentives for quality improvement. *Health Affairs*, 23(2):127–141.

Saez-Gallego, J., Morales, J. M., Zugno, M., and Madsen, H. (2016). A data-driven bidding model for a cluster of price-responsive consumers of electricity. *IEEE Transactions on Power Systems*, PP(99):1–11.

Savva, N., Tezcan, T., and Yildiz, O. (2016). Yardstick competition for service systems. *Working paper, London Business School*.

Scaparra, M. P. and Church, R. L. (2008). A bilevel mixed-integer program for critical infrastructure protection planning. *Computers & Operations Research*, 35(6):1905–1923.

Scarf, H., Arrow, K., and Karlin, S. (1958). A min-max solution of an inventory problem. *Studies in the mathematical theory of inventory and production*, 10:201–209.

Schaefer, A. J. (2009). Inverse integer programming. *Optimization Letters*, 3(4):483–489.

Shapiro, A. and Kleywegt, A. (2002). Minimax analysis of stochastic problems. *Optimization Methods and Software*, 17(3):523–542.

Shavell, S. (1979a). On moral hazard and insurance. In *Foundations of Insurance Economics*, pages 280–301. Springer.

Shavell, S. (1979b). Risk sharing and incentives in the principal and agent relationship. *The Bell Journal of Economics*, pages 55–73.

Sheather, S. and Jones, M. (1991). A reliable data-based bandwidth selection method for kernel density estimation. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 683–690.

Sheffi, Y. (2005). The resilient enterprise: overcoming vulnerability for competitive advantage. *MIT Press Books*, 1.

Shen, Z.-J. M., Zhan, R. L., and Zhang, J. (2011). The reliable facility location problem: Formulations, heuristics, and approximation algorithms. *INFORMS Journal on Computing*, 23(3):470–482.

Shleifer, A. (1985). A theory of yardstick competition. *The RAND Journal of Economics*, pages 319–327.

Siddiq, A. (2013). *Robust facility location under demand location uncertainty*. PhD thesis.

Siddiq, A. A., Brooks, S. C., and Chan, T. C. (2013). Modeling the impact of public access defibrillator range on public location cardiac arrest coverage. *Resuscitation*, 84(7):904–909.

Simchi-Levi, D., Kaminsky, P., and Simchi-Levi, E. (2004). *Managing the supply chain: the definitive guide for the business professional*. McGraw-Hill Companies.

Snyder, L. V. (2006). Facility location under uncertainty: A review. *IIE Transactions*, 38(7):547–564.

So, K. C. and Tang, C. S. (2000). Modeling the impact of an outcome-oriented reimbursement policy on clinic, patients, and pharmaceutical firms. *Management Science*, 46(7):875–892.

Sohoni, M. G., Bassamboo, A., Chopra, S., Mohan, U., and Sendil, N. (2010). Threshold incentives over multiple periods and the sales hockey stick phenomenon. *Naval Research Logistics (NRL)*, 57(6):503–518.

Starfield, B. (2000). Is US health really the best in the world? *Journal of the American Medical Association*, 284(4):483–485.

Sun, C. L., Demirtas, D., Brooks, S. C., Morrison, L. J., and Chan, T. C. (2016). Overcoming spatial and temporal barriers to public access defibrillators via optimization. *Journal of the American College of Cardiology*, 68(8):836–845.

Tao, T. (2012). *Topics in Random Matrix Theory*. Graduate studies in mathematics. American Mathematical Society.

Taylor, T. A. and Xiao, W. (2014). Subsidizing the distribution channel: Donor funding to improve the availability of malaria drugs. *Management Science*, 60(10):2461–2477.

Terrell, G. R. and Scott, D. W. (1992). Variable kernel density estimation. *The Annals of Statistics*, pages 1236–1265.

Troutt, M. D., Pang, W.-K., and Hou, S.-H. (2006). Behavioral estimation of mathematical programming objective function coefficients. *Management science*, 52(3):422–434.

Tversky, A. and Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science*, 211(4481):453–458.

Valenzuela, T., Roe, D., Nichol, G., Clark, L., Spaite, D., and Hardman, R. (2000). Outcomes of rapid defibrillation by security officers after cardiac arrest in casinos. *New England Journal of Medicine*, (343):1206–1209.

van der Vaart, A. (2000). *Asymptotic Statistics*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press.

Vershynin, R. (2012). *Compressed Sensing*, chapter Introduction to the non-asymptotic analysis of random matrices, pages 210–268. Cambridge University Press.

Wald, A. (1949). Note on the consistency of the maximum likelihood estimate. *Ann. Math. Statist.*, 20(4):595–601.

Wand, M. (1997). Data-based choice of histogram bin width. *The American Statistician*, 51(1):59–64.

Wang, L. (2009). Cutting plane algorithms for the inverse mixed integer linear programming problem. *Operations Research Letters*, 37(2):114–116.

Weisfeldt, M. L., Sitlani, C. M., Ornato, J. P., Rea, T., Aufderheide, T. P., Davis, D., Dreyer, J., Hess, E. P., Jui, J., Maloney, J., et al. (2010). Survival after application of automatic external defibrillators before arrival of the emergency medical systemevaluation in the resuscitation outcomes consortium population of 21 million. *Journal of the American College of Cardiology*, 55(16):1713–1720.

Wennberg, J. E., Fisher, E. S., Skinner, J. S., et al. (2002). Geography and the debate over Medicare reform. *Health Affairs*, 21(2):10–10.

Whang, S. (1992). Contracting for software development. *Management science*, 38(3):307–324.

Wiesemann, W., Kuhn, D., and Sim, M. (2014). Distributionally robust convex optimization. *Operations Research*, 62(6):1358–1376.

Wilensky, G. R. (2013). Developing a viable alternative to Medicares physician payment strategy. *Health Affairs*, pages 10–1377.

Wilson, E. B. (1927). Probable inference, the law of succession, and statistical inference. *Journal of the American Statistical Association*, 22(158):209–212.

Xu, H., Caramanis, C., and Mannor, S. (2012). A distributional interpretation of robust optimization. *Mathematics of Operations Research*, 37(1):95–110.

Yamin, D. and Gavious, A. (2013). Incentives' effect in influenza vaccination policy. *Management Science*, 59(12):2667–2686.

Zeng, B. and Zhao, L. (2013). Solving two-stage robust optimization problems using a column-and-constraint generation method. *Operations Research Letters*, 41(5):457–461.

Zhang, D., Gurvich, I., Mieghem, J. V., Park, E., Young, R., and Williams, M. (2016a). Hospital Readmissions Reduction Program: An Economic and Operational Analysis. *Management Science*.

Zhang, H., Wernz, C., and Hughes, D. R. (2016b). Modeling and designing health care payment innovations for medical imaging. *Health Care Management Science*, pages 1–15.

Zhang, H. and Zenios, S. (2008). A dynamic principal-agent model with hidden information: Sequential optimality through truthful state revelation. *Operations Research*, 56(3):681–696.

Zhang, J. and Liu, Z. (1996). Calculating some inverse linear programming problems. *Journal of Computational and Applied Mathematics*, 72(2):261–273.

Zhang, J. and Xu, C. (2010). Inverse optimization for linearly constrained convex separable programming problems. *European Journal of Operational Research*, 200(3):671–679.

Zhao, L. and Zeng, B. (2012). Robust unit commitment problem with demand response and wind energy. In *Power and Energy Society General Meeting, 2012 IEEE*, pages 1–8. IEEE.

Appendices

Appendix A

Proofs

A.1 Proofs from Chapter 2

Proof of Theorem 2.1. We prove this by showing a reduction from the problem of computing the best rank-1 approximation of an order 3 tensor (which is NP-hard (Hillar and Lim, 2013)) to IOP. Consider any $\psi \in \mathbb{R}^{r_1 \times r_2 \times r_3}$, where $r_1, r_2, r_3 \in \mathbb{Z}_+$. This defines ψ to be an order 3 tensor. We define $\rho = r_1 + r_2 + r_3$, and suppose the parameter vector is given by $\theta = (a, b, c) \in \Theta = \mathbb{R}^\rho$, where $a \in \mathbb{R}^{r_1}$, $b \in \mathbb{R}^{r_2}$, and $c \in \mathbb{R}^{r_3}$. Also define the discrete set $\mathcal{U} = [r_1] \times [r_2] \times [r_3]$, and suppose that $u = (\alpha, \beta, \gamma)$ is uniformly distributed over $\mathcal{U}$. Furthermore, suppose y is a random variable given by $\psi_{\alpha,\beta,\gamma}$, which means that y is dependent on u since $u = (\alpha, \beta, \gamma)$. Then we define the following forward optimization problem

$$\mathcal{S}(u, \theta) = \arg\min_x \Big(x - a_\alpha \cdot b_\beta \cdot c_\gamma\Big)^2. \tag{A.1}$$

This forward optimization problem is a quadratic program (QP) when (u, θ) is fixed, and so the solution set is $\mathcal{S}(u, \theta) = a_\alpha b_\beta c_\gamma$. Note that the solution set consists of a single point. Next, observe that

$$\min_{\theta \in \Theta} \; Q(\theta) = \min_{\theta \in \mathbb{R}^\rho} \frac{1}{\rho} \sum_{\alpha=1}^{r_1} \sum_{\beta=1}^{r_2} \sum_{\gamma=1}^{r_3} \Big(\psi_{\alpha,\beta,\gamma} - a_\alpha \cdot b_\beta \cdot c_\gamma\Big)^2, \tag{A.2}$$

where we have converted the expectation into a weighted sum using the fact that u is uniformly distributed over $\mathcal{U}$. Observe that (A.2) is the problem of computing the best rank-1 approximation to an order 3 tensor (Hillar and Lim, 2013). □

Proof of Corollary 2.1. We show this result using the same construction used to prove Theorem 2.1. In particular, observe that if $\{u_1, \ldots, u_n\} = \mathcal{U}$, then IOP–SAA is equivalent to IOP , which is NP-hard by Theorem 2.1. Finally, note that the condition $\{u_1, \ldots, u_n\} = \mathcal{U}$ occurs with nonzero probability since the set $\mathcal{U}$ is finite and since the u_i are sampled uniformly from $\mathcal{U}$. □

Proof of Proposition 2.1. We show this using a counterexample. Suppose FOP is $\min\{x^2 - (\theta + u) \cdot x \mid x \in [0, 10]\}$, and note its solution set $\mathcal{S}(u, \theta) = \min\{\frac{u+\theta}{2}, 10\}$ is single-valued.

Assume the distribution of u is

$$u = \begin{cases} 0, & \text{with probability (w.p.) } \frac{1}{2} \\ 20, & \text{w.p. } \frac{1}{2} \end{cases} \tag{A.3}$$

and that the distribution of w is

$$w = \begin{cases} -1, & \text{w.p. } \frac{1}{2} \\ +1, & \text{w.p. } \frac{1}{2} \end{cases} \tag{A.4}$$

Finally, suppose $y = \mathcal{S}(u, \theta) + w$, $\Theta = \{\theta \in \mathbb{R} : 0 \leq \theta \leq 10\}$, and $\theta_0 = 10$. By construction, this problem satisfies **A1**,**A2**,**IC**. Also, observe that the joint distribution of (u, y) is

$$(u, y) = \begin{cases} (\ 0,\ 4), & \text{w.p. } \frac{1}{4} \\ (\ 0,\ 6), & \text{w.p. } \frac{1}{4} \\ (20,\ 9), & \text{w.p. } \frac{1}{4} \\ (20, 11), & \text{w.p. } \frac{1}{4} \end{cases} \tag{A.5}$$

We show that both VIA and KKA are not estimation consistent for this problem. We begin with VIA. This approach solves

$$\begin{array}{l} \min\limits_{\theta \in \Theta} \ \frac{1}{n} \sum_{i=1}^{n} \epsilon_i^2 \\ \text{s.t. } \nabla f(y_i, u_i, \theta) \cdot (x_i - y_i) \geq -\epsilon_i, \forall x_i \in [0, 10], \qquad \forall i \in [n] \end{array} \tag{A.6}$$

The constraint

$$\nabla f(y_i, u_i, \theta) \cdot (x_i - y_i) \geq -\epsilon_i, \forall x_i \in [0, 10] \tag{A.7}$$

is a variational inequality, and VIA exactly reformulates this using linear duality. We operate with the original variational inequality since the reformulation in VIA is exact and does not change the solution. If $y_i = 4$, then a straightforward calculation gives that (A.7) is equivalent to the constraint: $\epsilon_i \geq 4 \cdot (8 - \theta)$ if $\theta \leq 8$, and $\epsilon_i \geq -6 \cdot (8 - \theta)$ if $\theta > 8$. If $y_i = 6$, then (A.7) is equivalent to the constraint $\epsilon_i \geq 6 \cdot (12 - \theta)$. If $y_i = 9$, then (A.7) is equivalent to the constraint $\epsilon_i \geq 2 + k$. Finally, if $y_i = 11$, then (A.7) is equivalent to the constraint: $\epsilon_i \geq 11 \cdot (2 - \theta)$ if $\theta \leq 2$, and $\epsilon_i \geq 2 - \theta$ if $\theta > 2$. Next, we solve the problem $\min\{\epsilon_i^2 \mid (A.7)\}$ for each possible value of y_i and θ. If $y_i = 4$, then the minimum is $16 \cdot (8 - \theta)^2$ if $\theta \leq 8$, and $36 \cdot (8 - \theta)^2$ if $\theta > 8$. If $y_i = 6$, then the minimum is $36 \cdot (12 - \theta)^2$. If $y_i = 9$, then minimum is $(2 + \theta)^2$. If $y_i = 11$, then the minimum is $121 \cdot (2 - \theta)^2$ if $\theta \leq 2$, and 0 if $\theta > 2$. Thus, we have

$$4 \cdot \mathbb{E}(\epsilon_i^2) = \begin{cases} 36 \cdot (12 - \theta)^2 + (2 + \theta)^2 + 121 \cdot (2 - \theta)^2 + 16 \cdot (8 - \theta)^2, & \text{if } \theta \leq 2 \\ 36 \cdot (12 - \theta)^2 + (2 + \theta)^2 + 16 \cdot (8 - \theta)^2, & \text{if } \theta \in (2, 8] \\ 36 \cdot (12 - \theta)^2 + (2 + \theta)^2 + 36 \cdot (8 - \theta)^2, & \text{if } \theta > 8 \end{cases} \tag{A.8}$$

Finally, we solve the optimization problem $\min\{\mathbb{E}(\epsilon_i^2) \mid \theta \in [0, 10]\}$. A simple calculation gives that the minimum occurs at $\theta^* = \frac{718}{73} \approx 9.8356$. However, the minimizer of (A.6) will

converge in probability to θ^*, because (i) we can exactly reformulate (A.6) as

$$\begin{aligned}
&\min_{\theta\in\Theta} \tfrac{1}{n}\textstyle\sum_{i=1}^n \epsilon_i^2 \\
&\text{s.t. } \epsilon_i^2 = \begin{cases} 16\cdot(8-\theta)^2\cdot\mathbf{1}(\theta\le 8) + 36\cdot(8-\theta)^2\cdot\mathbf{1}(\theta>8), & \text{if } y_i = 4 \\ 36\cdot(12-\theta)^2, & \text{if } y_i = 6 \\ (2+\theta)^2, & \text{if } y_i = 9 \\ 121\cdot(2-\theta)^2\cdot\mathbf{1}(\theta\le 2), & \text{if } y_i = 11 \end{cases} \qquad \forall i\in[n]
\end{aligned} \tag{A.9}$$

which (ii) implies we can apply the uniform law of large numbers (Jennrich, 1969) since ϵ_i^2 as defined in (A.9) is a continuous function, and thus (iii) we get convergence of the minimizer from a standard consistency result in statistics (see for instance Theorem 5.7 in (van der Vaart, 2000) or Theorem 5.2.3 in (Bickel and Doksum, 2006)). This shows VIA is not estimation consistent, since $\theta_0 = 10$. Next, we consider KKA. This approach solves

$$\begin{aligned}
&\min_{\theta\in\Theta} \tfrac{1}{n}\textstyle\sum_{i=1}^n \|\epsilon_i\|^2 \\
&\text{s.t. } \nabla f(y_i,u_i,\theta) - \lambda_{i1} + \lambda_{i2} = \epsilon_{i1} \\
&\qquad -\lambda_{i1}\cdot y_i = \epsilon_{i2} \\
&\qquad \lambda_{i2}\cdot(y_i-10) = \epsilon_{i3} \\
&\qquad \lambda_i \ge 0
\end{aligned} \tag{A.10}$$

We first solve the problem (A.10), with $n = 1$, for each possible value of y_i and θ. If $y_i = 4$, then the minimum is $\frac{16}{17}\cdot(8-\theta)^2$ if $\theta \le 8$, and $\frac{36}{37}\cdot(8-\theta)^2$ if $\theta > 8$. If $y_i = 6$, then the minimum is $\frac{36}{37}\cdot(12-\theta)^2$. If $y_i = 9$, then the minimum is $\frac{1}{2}\cdot(2+\theta)^2$. If $y_i = 11$, then the minimum is $\frac{121}{122}\cdot(2-\theta)^2$ if $\theta \le 2$, and $\frac{1}{2}\cdot(2-\theta)^2$ if $\theta > 2$. Thus, we have

$$4\cdot\mathbb{E}(\|\epsilon_i\|^2) = \begin{cases} \frac{36}{37}\cdot(12-\theta)^2 + \frac{1}{2}\cdot(2+\theta)^2 + \frac{121}{122}\cdot(2-\theta)^2 + \frac{16}{17}\cdot(8-\theta)^2, & \text{if } \theta\le 2 \\ \frac{36}{37}\cdot(12-\theta)^2 + \frac{1}{2}\cdot(2+\theta)^2 + \frac{1}{1}\cdot(2-\theta)^2 + \frac{16}{17}\cdot(8-\theta)^2, & \text{if } \theta\in(2,8] \\ \frac{36}{37}\cdot(12-\theta)^2 + \frac{1}{2}\cdot(2+\theta)^2 + \frac{1}{2}\cdot(2-\theta)^2 + \frac{36}{37}\cdot(8-\theta)^2, & \text{if } \theta > 8 \end{cases} \tag{A.11}$$

Finally, we solve the optimization problem $\min\{\mathbb{E}(\|\epsilon_i\|^2) \mid \theta\in[0,10]\}$. A simple calculation gives that the minimum occurs at $\theta^* = \frac{12080}{1833} \approx 6.5903$. However, the minimizer of (A.10) will converge in probability to θ^*, because (i) we can exactly reformulate (A.10) as

$$\begin{aligned}
&\min_{\theta\in\Theta} \tfrac{1}{n}\textstyle\sum_{i=1}^n \|\epsilon_i\|^2 \\
&\text{s.t. } \|\epsilon_i\|^2 = \begin{cases} \frac{16}{17}\cdot(8-\theta)^2\cdot\mathbf{1}(\theta\le 8) + \frac{36}{37}\cdot(8-\theta)^2\cdot\mathbf{1}(\theta>8), & \text{if } y_i = 4 \\ \frac{36}{37}\cdot(12-\theta)^2, & \text{if } y_i = 6 \\ \frac{1}{2}\cdot(2+\theta)^2, & \text{if } y_i = 9 \\ \frac{121}{122}\cdot(2-\theta)^2\cdot\mathbf{1}(\theta\le 2) + \frac{1}{2}\cdot(2-\theta)^2, & \text{if } y_i = 11 \end{cases} \qquad \forall i\in[n]
\end{aligned} \tag{A.12}$$

which (ii) implies we can apply the uniform law of large numbers (Jennrich, 1969) since

$\|\epsilon_i\|^2$ as defined in (A.12) is a continuous function, and thus (iii) we get convergence of the minimizer from a standard consistency result in statistics (see for instance Theorem 5.7 in (van der Vaart, 2000) or Theorem 5.2.3 in (Bickel and Doksum, 2006)). This shows that KKA is not estimation consistent, since $\theta_0 = 10$. □

Proof of Corollary 2.2. It suffices to show that risk consistency is necessary for estimation consistency in the counterexample given in the proof of Proposition 2.1. First note that the risk function

$$Q(\theta) = \mathbb{E}\Big(\|y - \min\{\tfrac{u+\theta}{2}, 10\}\|^2\Big) = \tfrac{1}{4} \cdot \Big((4 - \tfrac{\theta}{2})^2 + (6 - \tfrac{\theta}{2})^2 + (9-10)^2 + (11-10)^2\Big) \tag{A.13}$$

is continuous since $\Theta = \{\theta \in \mathbb{R} : 0 \leq \theta \leq 10\}$. Now suppose a sequence $\hat{\theta}_n$ is estimation consistent. Since $\hat{\theta}_n \xrightarrow{p} \theta_0$, by continuity of $Q(\theta)$ and the continuous mapping theorem (van der Vaart, 2000), we have $Q(\hat{\theta}_n) \xrightarrow{p} Q(\theta_0)$. Since $\arg\min\{Q(\theta) \mid \theta \in \Theta\} = 10 = \theta_0$, and $\hat{\theta}_n \longrightarrow \theta_0$ we have that $\hat{\theta}_n$ converges to a minimizer of $Q(\theta)$. Hence $\hat{\theta}_n$ is risk consistent. □

Proof of Proposition 2.2. Because the feasible set $\Phi(u, \theta)$ is convex for fixed u, θ by **A1** and has a nonempty interior by **R1**, this means $\Phi(u, \theta)$ is continuous in θ by Example 5.10 from (Rockafellar and Wets, 1998). Thus, we can apply the Berge Maximum Theorem (Berge, 1963) to FOP. This implies $\mathcal{S}(u, \theta)$ is upper hemicontinuous in θ for fixed $u \in \mathcal{U}$. However, $\mathcal{S}(u, \theta)$ consists of a single point for fixed $u \in \mathcal{U}$ and $\theta \in \Theta$, because the objective function is strictly convex and since **R1** holds. Consequently, $\mathcal{S}(u, \theta)$ is a continuous single-valued function for fixed $u \in \Theta$ (see for instance Theorem 2.6 in (Rockafellar and Wets, 1998)). Thus, we can apply the Berge Maximum Theorem to RISK–SAA, and this implies that $Q_n(\theta)$ as defined in RISK–SAA is continuous. □

Proof of Proposition 2.7. The solution set $\mathcal{S}(u, \theta)$ is nonempty under **A1**,**R1** (see for instance Theorem 1.9 of (Rockafellar and Wets, 1998)). Pick any $x_i \in \mathcal{S}(u_i, \theta)$, and let λ_i be such that x_i, λ_i satisfy (2.5) – this λ_i exists by Proposition 2.3. Next, consider the sets

$$\begin{aligned} \overline{\mathcal{S}}(u_i, \theta; \epsilon) &= \{x : f(x, u_i, \theta) - h(\lambda_i, u_i, \theta) \leq \epsilon,\ g(x, u_i, \theta) \leq \epsilon\} \\ \mathcal{S}(u_i, \theta; \epsilon) &= \{x : f(x, u_i, \theta) - h(\lambda, u_i, \theta) \leq \epsilon,\ g(x, u_i, \theta) \leq \epsilon,\ \lambda \geq 0\} \end{aligned} \tag{A.14}$$

and note that $\overline{\mathcal{S}}(u_i, \theta; \epsilon) = \mathcal{S}(u_i, \theta; \epsilon)$, since by optimality of λ_i with respect to the dual problem we have $h(\lambda, u_i, \theta) \leq h(\lambda_i, u_i, \theta)$ for all $\lambda \geq 0$. Observe that the functions $f(x_i, u_i, \theta), g(x_i, u_i, \theta)$ are continuous and convex from **A1**, and the point x_i belongs to the interior of $\overline{\mathcal{S}}(u_i, \theta; \epsilon)$ since it satisfies (2.5). Thus, we can apply Example 5.10 from Rockafellar and Wets (1998): This yields that $\overline{\mathcal{S}}(u_i, \theta; \epsilon)$ is continuous in θ, ϵ for any $\epsilon > 0$, and so we also get continuity of $\mathcal{S}(u_i, \theta; \epsilon)$ by its equality to $\overline{\mathcal{S}}(u_i, \theta; \epsilon)$. Since R–DB–RISK–SAA can be written as $Q_n(\theta; \epsilon) = \min_{x_i}\{\frac{1}{n}\sum_{i=1}^n \|y_i - x_i\|^2 \mid x_i \in \mathcal{S}(u_i, \theta; \epsilon),\ \forall i \in [n]\}$, we are able to apply the Berge Maximum Theorem (Berge, 1963). This implies continuity of $Q_n(\theta; \epsilon)$ in θ, ϵ for any $\epsilon > 0$. □

Proof of Proposition 2.8. Let $\mathcal{C}_n(\theta, \epsilon)$ be the feasible set of R–DB–RISK–SAA, and define $(X, \Lambda) = \{x_i, \lambda_i, \forall i \in [n]\}$. Suppose $(X, \Lambda) \in \mathcal{C}_n(\theta, \alpha)$, where $\alpha \geq 0$. Then for any $\beta \geq \alpha$ we must have $(X, \Lambda) \in \mathcal{C}_n(\theta, \beta)$ by the definition of the constraints in R–DB–RISK–SAA. This means that

$$\mathcal{C}_n(\theta, \epsilon_1) \supseteq \mathcal{C}_n(\theta, \epsilon_2) \supseteq \cdots \tag{A.15}$$

As a result, the set $\mathcal{D}_n(\theta, \epsilon_\nu) = \{\theta, X, \Lambda : \theta \in \Theta \text{ and } (X, \Lambda) \in \mathcal{C}_n(\theta, \epsilon_\nu)\}$ is also monotone nonincreasing:

$$\mathcal{D}_n(\theta, \epsilon_1) \supseteq \mathcal{D}_n(\theta, \epsilon_2) \supseteq \cdots \tag{A.16}$$

Also, the feasible set $\Phi(u, \theta)$ is convex for fixed u, θ by **A1** and has a nonempty interior by **R1**. This means $\Phi(u, \theta)$ is continuous in θ by Example 5.10 from (Rockafellar and Wets, 1998), and so we can apply the Berge Maximum Theorem (Berge, 1963) to FOP. This implies $\mathcal{S}(u, \theta)$ is upper hemicontinuous in θ for fixed $u \in \mathcal{U}$. By Remark 3.2 of (Dempe et al., 2015), this means $Q_n(\theta)$ is lower semicontinuous. Thus, by Proposition 7.4.d of (Rockafellar and Wets, 1998) we have that the extended real-valued function $\{Q_n(\theta; \epsilon_\nu) \mid \theta \in \Theta\}$ epiconverges to the extended real-valued function $\{Q_n(\theta) \mid \theta \in \Theta\}$. The result then follows from Exercise 7.32.d and Theorem 7.33 of (Rockafellar and Wets, 1998). □

Lemma A.1. *Suppose* **R4** *holds. Then for* $t > c_1 \cdot \gamma$ *we have*

$$\mathbb{P}\Big(|\gamma^{-m} \cdot \tfrac{1}{n} \textstyle\sum_{j=1}^{n} K(\frac{u_j - u_i}{\gamma}) - \mu(u_i)| > t\Big) \leq 2 \exp\Big(-2c_2 n \gamma^{2m} \cdot (t - c_1 \cdot \gamma)^2\Big), \tag{A.17}$$

where $c_1, c_2 > 0$ *are constants.*

Proof. Recall $\mu(u)$ is the probability density function of u, and note that

$$\begin{aligned}
|\mu(u_i) - \mathbb{E}[\gamma^{-m} K(\tfrac{u-u_i}{\gamma})|u_i]| &= |\mu(u_i) - \gamma^{-m} \textstyle\int_{\mathbb{R}^m} K(\frac{u-u_i}{\gamma}) \mu(u) du| \\
&= |\mu(u_i) - \gamma^{-m} \textstyle\int_{\mathbb{R}^m} K(s) \mu(u_i + \gamma s) \gamma^m ds| \\
&= |\mu(u_i) - \textstyle\int_{\mathbb{R}^m} K(s)(\mu(u_i) + \gamma \nabla \mu(u_i + \beta\gamma s)^T s) ds| \\
&= |\textstyle\int_{\mathbb{R}^m} K(s) \nabla \mu(u_i + \beta\gamma s)^T s ds| \cdot \gamma \\
&\leq c_1 \cdot \gamma,
\end{aligned} \tag{A.18}$$

where the second line follows from a change of variables $s = (u - u_i)/\gamma$, the third line follows from the multivariate form of Taylor's Theorem with some $\beta \in [0, 1]$, the fourth line follows because a Kernel function has the property $\int K(u) du = 1$, and the fifth line follows by setting $c_1 = \max_{u \in \mathcal{U}} |\int_{\mathbb{R}^m} K(s) \nabla \mu(u)^T s ds|$. Note this c_1 term is finite because (i) a kernel function has the property that its support is finite (i.e., $K(u) = 0$ for $\|u\| > 1$), and (ii) $\mu(u)$ is a continuously differentiable probability density function by **R4**. Next, note that by Hoeffding's inequality (Vershynin, 2012) we have for $t > 0$ that

$$\mathbb{P}\Big(|\gamma^{-m} \cdot \tfrac{1}{n} \textstyle\sum_{j=1}^{n} K(\frac{u_j - u_i}{\gamma}) - \mathbb{E}[\gamma^{-m} K(\frac{u-u_i}{\gamma})|u_i]| > t\Big) \leq 2 \exp\Big(-2c_2 n \gamma^{2m} t^2\Big), \tag{A.19}$$

where $c_2 = (\max_u K(u))^2$. Combining (A.18) and (A.19) gives the desired result. □

Lemma A.2. *Suppose* **A1** *and* **R1**–**R4** *hold. Then for* $t > c_3 \cdot \gamma^{1/2} + c_4 \cdot \gamma$ *we have*

$$\mathbb{P}\Big(\|\gamma^{-m} \cdot \tfrac{1}{n}\sum_{j=1}^{n} y_j \cdot K(\tfrac{u_j - u_i}{\gamma}) - \mu(u_i)\mathcal{S}(u_i,\theta_0)\| > t\Big) \leq 2\exp\Big(-2c_5 n\gamma^{2m}\cdot(t - c_3\cdot\gamma^{1/2} - c_4\cdot\gamma)\Big). \tag{A.20}$$

where $c_3, c_4, c_5 > 0$ *are constants.*

Proof. First, note that $\mathcal{S}(u,\theta)$ consists of a single point from the strict convexity assumption in **R3**. Next, note that having **A1** and **R1**–**R4** means that Proposition 4.41 of (Bonnans and Shapiro, 2000) holds: This means for $\gamma > 0$ sufficiently small we have

$$\|\mathcal{S}(u,\theta_0) - \mathcal{S}(u_i,\theta_0)\| \leq \alpha\cdot\gamma^{1/2}, \tag{A.21}$$

where $\alpha > 0$ is a constant, whenever $\|u - u_i\| \leq \gamma$. Next, recall that y_i conditioned on u_i has distribution $\mathcal{S}(u_i,\theta_0) + w_i$ under **IC**. Moreover, we have

$$\mathbb{E}[\gamma^{-m} y K(\tfrac{u-u_i}{\gamma})|u_i] = \mathbb{E}[\gamma^{-m}\mathcal{S}(u,\theta_0)K(\tfrac{u-u_i}{\gamma})|u_i], \tag{A.22}$$

since $\mathbb{E}(w_i) = 0$ and w_i is independent of u_i. Thus, we have

$$\begin{aligned}
&\|\mu(u_i)\mathcal{S}(u_i,\theta_0) - \mathbb{E}[\gamma^{-m} y K(\tfrac{u-u_i}{\gamma})|u_i]\| \\
&\quad= \|\mu(u_i)\mathcal{S}(u_i,\theta_0) - \gamma^{-m}\textstyle\int_{\mathbb{R}^m} K(\tfrac{u-u_i}{\gamma})\mu(u)\mathcal{S}(u,\theta_0)du\| \\
&\quad= \|\mu(u_i)\mathcal{S}(u_i,\theta_0) - \gamma^{-m}\textstyle\int_{\mathbb{R}^m} K(s)\mu(u_i+\gamma s)\mathcal{S}(u_i+\gamma s,\theta_0)\gamma^m ds\| \\
&\quad= \|\mu(u_i)\mathcal{S}(u_i,\theta_0) - \textstyle\int_{\mathbb{R}^m} K(s)(\mu(u_i) + \gamma\nabla\mu(u_i+\beta\gamma s)^T s)(\mathcal{S}(u_i,\theta_0)+ \\
&\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad \mathcal{S}(u_i+\gamma s,\theta_0) - \mathcal{S}(u_i,\theta_0))ds\| \\
&\quad= \|\textstyle\int_{\mathbb{R}^m} K(s)\mu(u_i)(\mathcal{S}(u_i+\gamma s,\theta_0) - \mathcal{S}(u_i,\theta_0))ds + \int_{\mathbb{R}^m} K(s)\gamma\nabla\mu(u_i+\beta\gamma s)^T s\mathcal{S}(u,\theta_0)ds\| \\
&\quad\leq c_3\cdot\gamma^{1/2} + c_4\cdot\gamma,
\end{aligned} \tag{A.23}$$

where the second line follows from a change of variables $s = (u - u_i)/\gamma$, the third line follows from the multivariate form of Taylor's Theorem with some $\beta \in [0,1]$, the fourth line follows because a Kernel function has the property $\int K(u)du = 1$, and the fifth line follows from (A.21) and by setting $c_3 = \alpha\cdot\max_{u\in\mathcal{U}}|\int_{\mathbb{R}^m} K(s)\mu(u)ds|$ and $c_4 = \max_{u\in\mathcal{U}}(|\int_{\mathbb{R}^m} K(s)\nabla\mu(u)^T s ds|\cdot\|\mathcal{S}(u,\theta_0)\|)$. Note the c_3, c_4 terms are finite because (i) a kernel function has the property that its support is finite (i.e., $K(u) = 0$ for $\|u\| > 1$), (ii) $\mu(u)$ is a continuously differentiable probability density function by **R4**, and (iii) $\mathcal{S}(u,\theta_0)$ is bounded by **R1**. Next, note that y is a sub-exponential random variable (Vershynin, 2012) since (i) $\mathcal{S}(u,\theta_0)$ is a bounded random variable by **R1**, and (ii) w is sub-exponential by **R4**. Hence, by Hoeffding's inequality for sub-exponential random variables (Vershynin, 2012) we have for $t > 0$ that

$$\mathbb{P}\Big(\|\gamma^{-m}\cdot\tfrac{1}{n}\sum_{j=1}^{n} y_j\cdot K(\tfrac{u_j-u_i}{\gamma}) - \mathbb{E}[\gamma^{-m} y K(\tfrac{u-u_i}{\gamma})|u_i]\| > t\Big) \leq 2\exp\Big(-2c_5 n\gamma^{2m}t\Big), \tag{A.24}$$

for some $c_5 > 0$. Combining (A.23) and (A.24) gives the desired result. □

Proof of Proposition 2.9. The first part follows from the strict convexity assumption in

R3, and the third part follows directly from the second part. And so we focus on proving the second part. We will prove this using a truncation argument (see for instance (Tao, 2012)).

First, note that the function $\psi(x,y) = x/y$ over the domain $(x,y) \in [-M,M] \times [\sigma, \sigma+1]$ is Lipschitz continuous with constant $L_1 = \sqrt{(M^2+(\sigma+1)^2)}/\sigma^2$. Suppose we choose $M = \max_{u\in\mathcal{U}} \|\mu(u)\mathcal{S}(u,\theta_0)\|+1$. As a result, using Lemma A.1 and Lemma A.2 we have

$$\begin{aligned}
&\mathbb{P}\Big(\|\overline{x}_i - \tfrac{\mu(u_i)\mathcal{S}(u_i,\theta_0)}{\sigma+\mu(u_i)}\| > t\Big) \\
&\quad\leq \mathbb{P}\Big(|\gamma^{-m}\cdot \tfrac{1}{n}\textstyle\sum_{j=1}^n K(\frac{u_j-u_i}{\gamma}) - \mu(u_i)| > t/L_1\Big) + \mathbb{P}\Big(\|\gamma^{-m}\cdot \tfrac{1}{n}\textstyle\sum_{j=1}^n y_j\cdot K(\frac{u_j-u_i}{\gamma})\| > M\Big) + \\
&\qquad\quad \mathbb{P}\Big(\|\gamma^{-m}\cdot \tfrac{1}{n}\textstyle\sum_{j=1}^n y_j\cdot K(\frac{u_j-u_i}{\gamma}) - \mu(u_i)\mathcal{S}(u_i,\theta_0)\| > t/L_1\Big) \\
&\quad\leq 2\exp\Big(-2c_2 n\gamma^{2m}\cdot(t/L_1 - c_1\cdot\gamma)^2\Big) + 2\exp\Big(-2c_2 n\gamma^{2m}\cdot(1-c_1\cdot\gamma)^2\Big) + \\
&\qquad\quad 2\exp\Big(-2c_5 n\gamma^{2m}\cdot(t/L_1 - c_3\cdot\gamma^{1/2} - c_4\cdot\gamma)\Big),
\end{aligned} \tag{A.25}$$

for $t > \max\{c_1\cdot\gamma, c_3\cdot\gamma^{1/2} + c_4\cdot\gamma\}$. Next, observe that the function $\psi(x,y)$ over the domain

$$(x,y) \in [\min_{u\in\mathcal{U}} \mu(u)\mathcal{S}(u,\theta), \max_{u\in\mathcal{U}} \mu(u)\mathcal{S}(u,\theta)] \times [\min_{u\in\mathcal{U}} \mu(u), \max_{u\in\mathcal{U}} \mu(u)], \tag{A.26}$$

is Lipschitz continuous with some constant $L_2 > 0$ since (i) the denominator of ψ is bounded away from zero because of **R4**, and (ii) the numerator of ψ is bounded by **R1**,**R4**. Thus, we have

$$\mathbb{P}\Big(\|\overline{x}_i - \mathcal{S}(u_i,\theta_0)\| > t\Big) \leq \mathbb{P}\Big(\|\overline{x}_i - \tfrac{\mu(u_i)\mathcal{S}(u_i,\theta_0)}{\sigma+\mu(u_i)}\| > t - \sigma/L_2\Big), \tag{A.27}$$

for $t > \sigma/L_2$. Suppose we choose $\gamma = O(n^{-2/(8m+1)})$, $\sigma = O(\gamma)$, and $t = n^{-1/(16m+2)}$. Then combining (A.25) and (A.27) gives that for sufficiently large n we have

$$\mathbb{P}\Big(\|\overline{x}_i - \mathcal{S}(u_i,\theta_0)\| > n^{-1/(16m+2)}\Big) \leq c_6\exp\Big(-c_7 n^{1/2}\Big), \tag{A.28}$$

where $c_6, c_7 > 0$ are constants. And so combining the union bound with (A.28) gives

$$\begin{aligned}
\mathbb{P}\Big(\max_{i\in[n]} \|\overline{x}_i - \mathcal{S}(u_i,\theta_0)\| > n^{-1/(16m+2)}\Big) &\leq n\mathbb{P}\Big(\|\overline{x}_i - \mathcal{S}(u_i,\theta_0)\| > n^{-1/(16m+2)}\Big) \\
&\leq c_6\exp\Big(-c_7 n^{1/2} + \log n\Big).
\end{aligned} \tag{A.29}$$

The final implication of the result follows by noting that $n^{-2/(8m+1)} \to 0$ and $c_6\exp(-c_7 n^{1/2} + \log n) \to 0$ as $n\to\infty$. □

Proof of Theorem 2.2. Proposition 2.7 gives continuity of $Q_n(\theta;\epsilon)$. Thus, we can apply the uniform law of large numbers (Jennrich, 1969), which gives

$$\sup_{\theta\in\Theta} |Q_n(\theta;\epsilon) - Q(\theta;\epsilon)| \xrightarrow{p} 0. \tag{A.30}$$

Consider any $\theta_0 \in \arg\min\{Q(\theta;\epsilon) \mid \theta \in \Theta\}$ and any $\theta_1 \in \arg\min\{Q_n(\theta;\epsilon) \mid \theta \in \Theta\}$. By assumption $Q_n(\theta_1;\epsilon) \leq Q_n(\theta_0;\epsilon)$, and so we have

$$Q(\hat{\theta}_n;\epsilon) + Q_n(\hat{\theta}_n;\epsilon) - Q(\hat{\theta}_n;\epsilon) + Q_n(\theta_1;\epsilon) - Q_n(\hat{\theta}_n;\epsilon) \leq Q(\theta_0;\epsilon) + Q_n(\theta_0;\epsilon) - Q(\theta_0;\epsilon). \tag{A.31}$$

Rearranging terms gives

$$Q(\hat{\theta}_n;\epsilon) - Q(\theta_0) \leq |Q_n(\hat{\theta}_n;\epsilon) - Q(\hat{\theta}_n;\epsilon)| + |Q_n(\theta_1;\epsilon) - Q_n(\hat{\theta}_n;\epsilon)| + |Q_n(\theta_0;\epsilon) - Q(\theta_0;\epsilon)|. \tag{A.32}$$

Recall (i) $Q(\theta_0;\epsilon) \leq Q(\hat{\theta}_n;\epsilon)$ by definition of θ_0, (ii) $Q_n(\theta;\epsilon)$ is continuous, and (iii) $\hat{\theta}_n$ is nearly-optimal for R–IOP–SAA in probability. Thus, combining these facts with (A.30) and (A.32) gives that $Q(\hat{\theta}_n;\epsilon) - Q(\theta_0;\epsilon) \xrightarrow{p} 0$. This is the desired result. □

Proof of Theorem 2.3. Proposition 2.2 gives continuity of $Q_n(\theta)$. The remainder of the proof is identical to Theorem 2.2. □

Proof of Theorem 2.4. Because the feasible set $\Phi(u,\theta)$ is convex for fixed u,θ by **A1** and has a nonempty interior by **R1**, this means $\Phi(u,\theta)$ is continuous in θ by Example 5.10 from (Rockafellar and Wets, 1998). Thus, we can apply the Berge Maximum Theorem (Berge, 1963) to FOP. This implies $\mathcal{S}(u,\theta)$ is upper hemicontinuous in θ for fixed $u \in \mathcal{U}$. By Remark 3.2 of Dempe et al. (2015), this means $Q_n(\theta)$ is lower semicontinuous. Thus, we can apply Theorem 5.14 of (van der Vaart, 2000). The result follows from the conclusion of Theorem 5.14 of (van der Vaart, 2000) if we can show (i) $\theta_0 \in \arg\min\{Q(\theta) \mid \theta \in \Theta\}$, and that (ii) θ_0 is the unique solution. First, note $Q(\theta) = \mathbb{E}(\min_{x\in\mathcal{S}(u,\theta)}\|\xi - x\|^2) + \mathbb{E}(w^2)$, since ξ, x is almost surely independent of w because by **IC** we have that (i) ξ, u are independent of w, and (ii) $\mathcal{S}(u,\theta)$ is almost surely single-valued. Since by **IC** we have $\xi \in \mathcal{S}(u,\theta_0)$, this means that $Q(\theta_0) = \mathbb{E}(w^2)$ and that $\theta_0 \in \arg\min\{Q(\theta) \mid \theta \in \Theta\}$. Next, consider any $\theta \in \Theta \setminus \theta_0$. Then by **IC** we have $\mathbb{E}[\min_{x\in\mathcal{S}(u,\theta)}\|\xi - x\|^2 \mid u \in \mathcal{U}(\theta)] > 0$ since $\xi \in \mathcal{S}(u,\theta_0)$ and $\text{dist}(\mathcal{S}(u,\theta),\mathcal{S}(u,\theta_0)) > 0$ for each $u \in \mathcal{U}(\theta)$. Because $\mathbb{P}(u \in \mathcal{U}(\theta)) > 0$ from **IC**, this means $\mathbb{E}(\min_{x\in\mathcal{S}(u,\theta)}\|\xi - x\|^2) > 0$ for any $\theta \in \Theta \setminus \theta_0$. Consequently, we have $Q(\theta) > Q(\theta_0)$ for any $\theta \in \Theta \setminus \theta_0$. □

Proof of Theorem 2.5. By Corollary 2.3, there exists $E, Z > 0$ such that if

$$\hat{\theta}_n \in z\text{-}\arg\min\{Q_n(\theta;\epsilon) \mid \theta \in \Theta\}$$

for any $0 \leq z \leq Z$ and $0 \leq \epsilon \leq E$, then $\text{dist}(\hat{\theta}_n, \arg\min\{Q_n(\theta) \mid \theta \in \Theta\}) < d$. Suppose we choose $z = Z$. Because $Q_n(\theta;\epsilon)$ is continuous in θ by Proposition 2.7, there exists $\Delta > 0$ such that for any $0 < \delta \leq \Delta$ we have

$$\min\{Q_n(\theta;\epsilon) - Q_n(\theta_0;\epsilon) \mid \theta \in \mathcal{T}(\delta)\} < z, \tag{A.33}$$

[3]Technically, this theorem applies to maximizing upper semicontinuous functions, but the results and proof trivially extend to the case of minimizing lower semicontinuous functions.

where $\theta_0 \in \arg\min\{Q_n(\theta;\epsilon) \mid \theta \in \Theta\}$. By construction, we have

$$\arg\min\{Q_n(\theta;\epsilon) \mid \theta \in \mathcal{T}(\delta)\} \subseteq z\text{-}\arg\min\{Q_n(\theta;\epsilon) \mid \theta \in \Theta\}. \tag{A.34}$$

Next, note the enumeration algorithm returns a solution $\hat{\theta}_n \in \arg\min\{Q_n(\theta;\epsilon) \mid \theta \in \mathcal{T}(\delta)\}$, which also satisfies $\hat{\theta}_n \in z\text{-}\arg\min\{Q_n(\theta;\epsilon) \mid \theta \in \Theta\}$. The result follows from applying the first line of the proof. □

Proof of Theorem 2.6. By Corollary 2.3, there exists $E, Z > 0$ such that if

$$\hat{\theta}_n \in z\text{-}\arg\min\{Q_n(\theta;\epsilon) \mid \theta \in \Theta\}$$

for any $0 \leq z \leq Z$ and $0 \leq \epsilon \leq E$, then $\text{dist}(\hat{\theta}_n, \arg\min\{Q_n(\theta) \mid \theta \in \Theta\}) < d$. Suppose we choose $z = Z$ and $\epsilon = 0$, and note that $Q_n(\theta;0) = Q_n(\theta)$ by their definitions. Because $Q_n(\theta)$ is continuous in θ by Proposition 2.2, there exists $\Delta > 0$ such that for any $0 < \delta \leq \Delta$ we have

$$\min\{Q_n(\theta) - Q_n(\theta_0) \mid \theta \in \mathcal{T}(\delta)\} < z, \tag{A.35}$$

where $\theta_0 \in \arg\min\{Q_n(\theta) \mid \theta \in \Theta\}$. By construction, we have

$$\arg\min\{Q_n(\theta) \mid \theta \in \mathcal{T}(\delta)\} \subseteq z\text{-}\arg\min\{Q_n(\theta) \mid \theta \in \Theta\}. \tag{A.36}$$

Next, note the enumeration algorithm returns a solution $\hat{\theta}_n \in \arg\min\{Q_n(\theta) \mid \theta \in \mathcal{T}(\delta)\}$, which also satisfies $\hat{\theta}_n \in z\text{-}\arg\min\{Q_n(\theta) \mid \theta \in \Theta\}$. The result follows from the first line of the proof. □

Proof of Theorem 2.7. Note that $\min\{-h(\lambda,u,\theta) \mid \lambda \geq 0\} = -f(\mathcal{S}(u,\theta),u,\theta)$ by strong duality (which holds because of **A1**,**R1** (Bonnans and Shapiro, 2000)). Next, consider the function

$$R(\theta) = \mathbb{E}\Big(\min_{\lambda\geq 0} f(\mathcal{S}(u,\theta_0),u,\theta) - h(\lambda,u,\theta)\Big) = \mathbb{E}\Big(f(\mathcal{S}(u,\theta_0),u,\theta) - f(\mathcal{S}(u,\theta),u,\theta)\Big), \tag{A.37}$$

its sample average approximation

$$R_n(\theta) = \frac{1}{n}\sum_{i=1}^{n}\Big(\min_{\lambda_i\geq 0} f(\mathcal{S}(u_i,\theta_0),u_i,\theta) - h(\lambda_i,u_i,\theta)\Big) = \frac{1}{n}\sum_{i=1}^{n}\Big(f(\mathcal{S}(u_i,\theta_0),u_i,\theta) - f(\mathcal{S}(u_i,\theta),u_i,\theta)\Big), \tag{A.38}$$

and its semiparametric approximation

$$\overline{R}_n(\theta) = \frac{1}{n}\sum_{i=1}^{n}\Big(\min_{\lambda_i\geq 0} f(\overline{x}_i,u_i,\theta) - h(\lambda_i,u_i,\theta)\Big) = \frac{1}{n}\sum_{i=1}^{n}\Big(f(\overline{x}_i,u_i,\theta) - f(\mathcal{S}(u_i,\theta),u_i,\theta)\Big). \tag{A.39}$$

Note that $\min\{\overline{R}_n(\theta) \mid \theta \in \Theta\}$ is simply a reformulation of SP–IOP–RISK–SAA. Next, observe

that $\mathbb{E}[f(\mathcal{S}(u,\theta_0),u,\theta)-f(\mathcal{S}(u,\theta),u,\theta)|u\in\mathcal{U}(\theta)]>0$ since (i) $f(x,u,\theta)$ is twice continuously differentiable in x by **R3**, and (ii) $\text{dist}(\mathcal{S}(u,\theta),\mathcal{S}(u,\theta_0))>0$ for each $u\in\mathcal{U}(\theta)$ by **IC**. Consequently, we have $R(\theta)>0$ for $\theta\in\Theta\setminus\theta_0$. As shown in the proof for Proposition 2.2, $\mathcal{S}(u,\theta)$ is continuous in θ. And so $R_n(\theta)$ and $\overline{R}_n(\theta)$ are continuous because (i) $f(x,u,\theta)$ is twice continuously differentiable in x,θ by **R3**. Next, recall that $\mathcal{U}$ is bounded by **R4**, Θ is bounded by **R2**, $f(x,u,\theta)$ is twice continuously differentiable in x,θ by **R3**, and the feasible set of FOP is absolutely bounded by **R1**. This means there exists $L>0$ such that for all $\theta\in\Theta$ we have $\max_{i\in[n]}|f(\overline{x},u_i,\theta)-f(\mathcal{S}(u_i,\theta_0),u_i,\theta)|\leq Ln^{-1/(18m)}$ whenever $\max_{i\in[n]}\|\overline{x}_i-\mathcal{S}(u_i,\theta_0)\|\leq n^{-1/(18m)}$ (which occurs with probability at least $1-k_1\exp(-k_2n^{1/4})$ by Proposition 2.9). Thus, we have that $\sup_{\theta\in\Theta}|R_n(\theta)-\overline{R}_n(\theta)|\xrightarrow{p}0$. Now consider any $\hat{\theta}_n\in\arg\min\{\overline{R}_n(\theta)\mid\theta\in\Theta\}$, and note that the estimate $\hat{\theta}_n$ returned by the semiparametric algorithm satisfies this property by construction. By definition we have $\overline{R}_n(\hat{\theta}_n)\leq\overline{R}_n(\theta_0)$, which can be rewritten as

$$R_n(\hat{\theta}_n)+\overline{R}_n(\hat{\theta}_n)-R_n(\hat{\theta}_n)\leq R_n(\theta_0)+\overline{R}_n(\theta_0)-R_n(\theta_0). \tag{A.40}$$

Thus, we have

$$R_n(\hat{\theta}_n)\leq R_n(\theta_0)+|\overline{R}_n(\hat{\theta}_n)-R_n(\hat{\theta}_n)|+|\overline{R}_n(\theta_0)-R_n(\theta_0)|. \tag{A.41}$$

We have thus shown all the conditions required to apply Theorem 5.14 of (van der Vaart, 2000), which gives $\hat{\theta}_n\xrightarrow{p}\theta_0$. Now let $\overline{\theta}_n\in\arg\min\{Q_n(\theta)\mid\theta\in\Theta\}$. By Theorem 2.4, we have $\overline{\theta}_n\xrightarrow{p}\theta_0$. This means that $|\overline{\theta}_n-\hat{\theta}_n|\xrightarrow{p}0$. □

A.2 Proofs from Chapter 3

For the purposes of this section, let the expected shared savings payment be

$$R(x,\alpha)=\mathbb{E}[r(y,\alpha)]=\int_{-\infty}^{\infty}r(y,\alpha)h(y|x)dy$$

and let the expected subsidy payment be

$$S(x,\beta)=\mathbb{E}[s(y,\beta)]=\int_{-\infty}^{\infty}s(y,\beta)h(y|x)dy$$

We first present a lemma which is useful in the proofs to follow.

Lemma A.3. *The derivative of the expected shared savings payment with respect to x is*

$$\begin{aligned}(\partial/\partial x)R(x,\alpha)=(1-\alpha)\left(hg(-h-x)+G(-h-x)-G(C_\ell-x)\right)\\+\alpha\left(hg(h-x)+G(C_u-x)-G(h-x)\right).\end{aligned}$$

The derivative of the expected subsidy payment with respect to x is

$$(\partial/\partial x)S(x,\beta)=\beta\left((\partial/\partial x)c(x,\theta)(1-G(h-x))+c(x,\theta)g(h-x)\right).$$

Proof. Consider the shared savings term first. Integrating over the piecewise expression given in (3.2), we have

$$\begin{aligned} R(x,\alpha) &= \int_{-\infty}^{\infty} r(y,\alpha)h(y|x)dy \\ &= (1-\alpha)C_\ell \int_{-\infty}^{C_\ell} h(y|x)dy + (1-\alpha)\int_{C_\ell}^{-h} yh(y|x)dy \\ &\quad + \alpha \int_{h}^{C_u} yh(y|x)dy + \alpha C_u \int_{C_u}^{\infty} h(y|x)dy. \end{aligned}$$

Since $y = x + \xi$, we may apply a change of variables and write the expression above equivalently as

$$\begin{aligned} R(x,\alpha) = (1-\alpha)C_\ell \int_{-\infty}^{C_\ell - x} g(\xi)d\xi + (1-\alpha)\int_{C_\ell - x}^{-h-x} (x+\xi)g(\xi)d\xi \\ + \alpha \int_{h-x}^{C_u - x} (x+\xi)g(\xi)d\xi + \alpha C_u \int_{C_u - x}^{\infty} g(\xi)d\xi. \end{aligned}$$

Applying Leibniz's rule to each of the four terms above yields

$$\begin{aligned} (\partial/\partial x)R(x,\alpha) =& -(1-\alpha)C_\ell g(C_\ell - x) + (1-\alpha)(hg(-h-x) + C_\ell g(C_\ell - x) \\ &+ G(-h-x) - G(C_\ell - x)) + \alpha(-C_u g(C_u - x) + hg(h-x) \\ &+ G(C_u - x) - G(h-x)) + \alpha C_u g(C_u - x) \\ =& (1-\alpha)\left(hg(-h-x) + G(-h-x) - G(C_\ell - x)\right) \\ &+ \alpha\left(hg(h-x) + G(C_u - x) - G(h-x)\right), \end{aligned}$$

where the second line is obtained by canceling out terms. For the subsidy term, we integrate over the expression given in (3.3) to obtain

$$S(x,\beta) = \beta c(x,\theta)\int_h^{\infty} h(y|x)dy = \beta c(x,\theta)\int_{h-x}^{\infty} g(\xi)d\xi.$$

Applying the chain rule to this expression yields

$$\begin{aligned} (\partial/\partial x)S(x,\beta) &= \beta\left((\partial/\partial x)c(x,\theta)(1-G(h-x)) + c(x,\theta)(\partial/\partial x)\left(\int_{h-x}^{\infty} g(\xi)d\xi\right)\right) \\ &= \beta\left((\partial/\partial x)c(x,\theta)(1-G(h-x)) + c(x,\theta)g(h-x)\right). \quad \square \end{aligned}$$

Proof of Lemma 3.1. To show that $u(x,\alpha,\beta,\theta)$ is strictly concave over $[0,\bar{x}]$ for any θ, it suffices to show that $(\partial/\partial x)u(x,\alpha,\beta,\theta) < 0$ for all $x \in [0,\bar{x}]$ and $\theta \in \Theta$. Pick some x and θ so that $(\partial^2/\partial x^2)u(x)$ exists (note that $(\partial^2/\partial x^2)u(x)$ may be undefined for certain values of x due to the possible non-differentiability of the density $g(\xi)$ at some points. In these cases we may consider the left and right derivatives of $u(x,\alpha,\beta,\theta)$ and apply a similar argument). By twice differentiating the expression for $u(x,\alpha,\beta,\theta)$ given in (3.4), we have

$$(\partial^2/\partial x^2)u(x) = (\partial^2/\partial x^2)R(x,\alpha) + (\partial^2/\partial x^2)S(x,\beta) - (\partial^2/\partial x^2)c(x,\theta).$$

We consider each of the three terms in this equation in sequence. For the shared savings term, $(\partial^2/\partial x^2)R(x,\alpha)$, differentiating the expression for $(\partial/\partial x)R(x,\alpha)$ given in Lemma A.3 yields

$$\begin{aligned}(\partial^2/\partial x^2)R(x,\alpha) = (1-\alpha)\left(-hg'(-h-x) - g(-h-x) + g(C_\ell - x)\right) \\ + \alpha\left(-hg'(h-x) - g(C_u - x) + g(h-x)\right).\end{aligned}$$

Since g is increasing over $(-\infty, 0)$ and $C_\ell \leq -h$, we have $g(C_\ell - x) \leq g(-h-x)$ for all x. Further, we have $g'(-h-x) \geq 0$ and thus $-hg'(-h-x) \leq 0$ for all x. It follows that $(1-\alpha)(-hg'(-h-x) - g(-h-x) + g(C_\ell - x)) \leq 0$. Hence we may drop the first term in the expression for $(\partial^2/\partial x^2)R(x,\alpha)$ above to obtain

$$(\partial^2/\partial x^2)R(x,\alpha) \leq \alpha\left(-hg'(h-x) - g(C_u - x) + g(h-x)\right).$$

Now by dropping the negative term $-g(C_u - x)$ and noting that $-hg'(h-x) \leq h\bar{g}'$ and $g(h-x) \leq \bar{g}$, we have $(\partial^2/\partial x^2)R(x,\alpha) \leq \bar{\alpha}(h\bar{g}' + \bar{g})$. For the subsidy term $(\partial^2/\partial x^2)S(x,\beta)$, we differentiate the expression for $(\partial/\partial x)S(x,\beta)$ given in Lemma A.3 to obtain

$$\begin{aligned}(\partial^2/\partial x^2)S(x,\beta) &= \beta\left((1-G(h-x))(\partial^2/\partial x^2)c(x,\theta) + 2g(h-x)(\partial/\partial x)c(x,\theta) - c(x,\theta)g'(h-x)\right), \\ &\leq \bar{\beta}\left((\partial^2/\partial x^2)c(x,\theta) + 2\bar{g}(\partial/\partial x)c(x,\theta)\right),\end{aligned}$$

where the inequality follows by noting that $1 - G(h-x) \leq 1$ and $g(h-x) \leq \bar{g}$. Combining the bounds for $(\partial^2/\partial x^2)R(x,\alpha)$ and $(\partial^2/\partial x^2)S(x,\beta)$, we can now write

$$\begin{aligned}(\partial^2/\partial x^2)u(x) &\leq \bar{\alpha}\left(h\bar{g}' + \bar{g}\right) + \bar{\beta}\left((\partial^2/\partial x^2)c(x,\theta) + 2\bar{g}(\partial/\partial x)c(x,\theta)\right) - (\partial^2/\partial x^2)c(x,\theta) \\ &= \bar{\beta}\left(h\bar{g}' + \bar{g}\right) + 2\beta\bar{g}(\partial/\partial x)c(x,\theta) - (1-\bar{\beta})(\partial^2/\partial x^2)c(x,\theta), \\ &< 0,\end{aligned}$$

where the final strict inequality follows from Assumption 3.3. □

Proof of Lemma 3.2. Let α and β be fixed. We first prove that $x(\theta)$ is strictly increasing over $[\theta_{\alpha,\beta}, \bar{\theta}]$. We wish to show that $(d/d\theta)x(\theta) > 0$ for any $\theta \in [\theta_{\alpha,\beta}, \bar{\theta}]$. Pick any such θ. Since $x(\theta) > 0$, it follows that $x(\theta)$ is a solution to the first order condition

$$(\partial/\partial x)u(x,\alpha,\beta,\theta) = (\partial/\partial x)R(x,\alpha) + (\partial/\partial x)S(x,\beta) - \gamma - (\partial/\partial x)c(x,\theta) = 0.$$

By the Implicit Function Theorem, we have

$$(d/d\theta)x(\theta) = -(\partial^2/\partial x\partial\theta)u(x,\alpha,\beta,\theta)/(\partial^2/\partial x^2)u(x,\alpha,\beta,\theta).$$

Since $R(x,\alpha)$ and $S(x,\beta)$ do not depend on θ, we have

$$-(\partial^2/\partial x\partial\theta)u(x) = (\partial^2/\partial x\partial\theta)c(x,\theta) < 0,$$

where the inequality follows from part (iii) of Assumption 3.1. Further, by Proposition 3.1, we also have $(\partial^2/\partial x^2)u(x,\alpha,\beta,\theta) < 0$, and thus $(d/d\theta)x(\theta) > 0$ over $[\theta_{\alpha,\beta}, \bar{\theta}]$. It remains to show that x is non-decreasing over $[\underline{\theta}, \theta_{\alpha,\beta}]$, which follows immediately from the definition of

$\theta_{\alpha,\beta}$ and the observation that x is bounded below by 0. □

Proof of Lemma 3.3. The proof proceeds in a manner similar to Lemma 3.2. Pick some $\alpha \in \mathcal{A}$ and $\theta \in [\theta_{\alpha,0}, \bar{\theta}]$. Now pick any $\beta \in [\underline{\beta}, \bar{\beta}]$. Since $\theta \geq \theta_{\alpha,0}$, we have $x(\beta) > 0$. It follows that $x(\beta)$ is a solution to the first-order condition

$$(\partial/\partial x)R(x,\alpha) + (\partial/\partial x)S(x,\beta) - \gamma - (\partial/\partial x)c(x,\theta) = 0.$$

By the Implicit Function Theorem, we have

$$(d/d\beta)x(\theta) = -(\partial^2/\partial x \partial \beta)u(x,\alpha,\beta,\theta)/(\partial^2/\partial x^2)u(x).$$

Since $R(x,\alpha)$ and $c(x,\theta)$ do not depend on β, we have

$$-(\partial^2/\partial x\partial\beta)u(x) = -(\partial^2/\partial x\partial\beta)S(x,\beta) = -(\partial/\partial x)c(x,\theta)(1-G(h-x)) + c(x,\theta)g(h-x) < 0.$$

Further, by Lemma 3.1, we have $(\partial^2/\partial x^2)u(x,\alpha,\beta,\theta) < 0$. It follows that $(d/d\beta)x(\theta) > 0$. To show that $u(x(\theta))$ is increasing in β, note that $S(x,\beta)$ is strictly increasing in β, and therefore so is the ACO's payoff $u(x)$ for any x. It follows that the ACO's optimal payoff $u(x(\theta))$ must be strictly increasing in β as well. □

Proof of Proposition 3.1. We first show that $\lim_{\theta\to 0} x(\theta) = 0$ for any α. Pick some α, and suppose that for $\epsilon > 0$ we have $x(\theta) > \epsilon$ for all $\theta > 0$. Clearly, the shared savings term $R(x(\theta),\alpha)$ is bounded above by $x(\theta)$. The ACO's optimal payoff can then be bounded above by

$$u(x(\theta)) \leq x(\theta) - \gamma x(\theta) - (1-\beta)c(x(\theta),\theta).$$

Since $x(\theta) > \epsilon$ for all $\theta > 0$, and since by Assumption 3.1 $c(x,\theta)$ is increasing in x, we have $c(\epsilon,\theta) \leq c(x(\theta),\theta)$ for all $\theta > 0$. Also by Assumption 3.1, we have $c(\epsilon,\theta) \longrightarrow \infty$ and thus $\lim_{\theta\to 0} c(x(\theta),\theta) = \infty$. It follows that $\lim_{\theta\to 0} u(x(\theta)) = -\infty$. However, it is straightforward to show that $u(0,\theta)$ is bounded from below, which yields a contradiction. We conclude that $\lim_{\theta\to 0} x(\theta) = 0$. We now show that there exists $\theta_0' > 0$ such that $\beta(\theta) = 0$ is an optimal subsidy rate for $\theta \leq \theta_0'$. This follows immediately from the fact that

$$\lim_{\theta\to 0} v(\alpha,0,\theta) = -R(0,\alpha) \geq -R(0,\alpha) - \lim_{\theta\to 0}\beta c(x(\theta),\theta) = \lim_{\theta\to 0} v(\alpha,\beta,\theta)$$

for any $\beta > 0$ and $\alpha > 0$. We now show that there exists θ_0 such that $\alpha(\theta) \geq 1/2$ for all $\theta \leq \theta_0$, for sufficiently large C_u and $|C_\ell|$. Using the expression for $R(x,\alpha)$ given in Lemma A.3, we have

$$\begin{aligned}\lim_{C_u\to\infty}\lim_{C_\ell\to-\infty}\lim_{\theta\to 0} R(x(\theta),\alpha) &= (1-\alpha)\int_{-\infty}^{-h}\xi g(\xi)d\xi + \alpha\int_h^{\infty}\xi g(\xi)d\xi \\ &= (1-\alpha)\mathbb{E}[\xi|\xi\leq -h] + \alpha\mathbb{E}[\xi|\xi\geq h] \\ &= (2\alpha-1)\mathbb{E}[\xi|\xi\geq h].\end{aligned}$$

Note that $u(x) \leq R(x,\alpha)$ for all $x \geq 0$. Therefore, for C_u and $|C_\ell|$ sufficiently large, we have

$\lim_{\theta\to 0} u(x(\theta)) \leq \lim_{\theta\to 0} R(x(\theta),\alpha) < 0$ if $(2\alpha-1)<0$. Since the participation constraint requires $u(x(\theta)) \geq \bar{u}(\theta) \geq 0$, it follows that $\alpha(\theta) \geq 1/2$ for sufficiently small θ. □

Proof of Proposition 3.2. Let $\theta_L > 0$ be fixed. We first show that $\beta^*(\theta_H) \geq \beta^*(\theta_L)$ for sufficiently large θ_H. It suffices to show that for any $\alpha \in \mathcal{A}$, $\lim_{\theta\to\infty} (d/d\beta)v(x(\theta),\beta,\theta)|_{\beta=\beta^*(\theta_L)} > 0$. Taking the total derivative of v with respect to β yields

$$(d/d\beta)v(x(\theta),\beta,\theta) = (\partial/\partial x)v(x,\beta,\theta)(d/d\beta)x(\theta) + (\partial/\partial x)\partial v(x,\beta,\theta).$$

First, expanding the term $(\partial/\partial x)v(x,\beta,\theta)$, for all $x \in [0,\bar{x}]$ we have

$$\begin{aligned}(\partial/\partial x)v(x,\beta,\theta) &= 1-(\partial/\partial x)R(x,\alpha)-(\partial/\partial x)S(x,\beta)\\ &= 1-(\partial/\partial x)R(x,\alpha)-\beta\left((\partial/\partial x)c(x,\theta)(1-G(h-x))+c(x,\theta)g(h-x)\right)\\ &\geq 1-(\partial/\partial x)R(x,\alpha)-\beta\left((\partial/\partial x)c(x,\theta)+c(x,\theta)\bar{g}\right)\\ &\geq \delta-\beta\left((\partial/\partial x)c(x,\theta)+c(x,\theta)\bar{g}\right)\end{aligned}$$

where in the final inequality we have $\delta = 1-(\partial/\partial x)R(x,\alpha) > 0$ for all $x \in [0,\bar{x}]$ and $\alpha \in \mathcal{A}$ by Assumption 3.4. Since by Assumption 3.1 the functions $c(x,\theta)$ and $(\partial/\partial x)c(x,\theta)$ are both decreasing in θ, it follows that

$$\lim_{\theta\to\infty}\left((\partial/\partial x)c(x,\theta)+c(x,\theta)\bar{g}\right) = 0$$

for all $x \in [0,\bar{x}]$. Hence

$$\lim_{\theta\to\infty} (\partial/\partial x)v(x,\beta,\theta)|_{\beta=\beta^*(\theta_L)} > 0$$

for all $x \in [0,\bar{x}]$. Further, since by Proposition 3.3, $(d/d\beta)x(\theta) > 0$ for sufficiently large θ, it follows that

$$\lim_{\theta\to\infty} (\partial/\partial x)v(x,\beta,\theta)(d/d\beta)x(\theta)|_{\beta=\beta^*(\theta_L)} > 0.$$

Now note that for the second term, we have

$$\lim_{\theta\to\infty} (\partial/\partial\beta)v(x,\beta,\theta) = \lim_{\theta\to\infty} c(x,\theta)(1-G(h-x)) \leq \lim_{\theta\to\infty} c(\bar{x},\theta) = 0.$$

It follows that

$$\lim_{\theta\to\infty} (\partial/\partial\beta)v(x(\beta),\beta,\theta)|_{\beta=\beta^*(\theta_L)} > 0.$$

Thus $\beta^*(\theta_H) \geq \beta^*(\theta_L)$. We now show that $\alpha^*(\theta_H) \leq \alpha^*(\theta_L)$. Suppose $\alpha^*(\theta_H) > \alpha^*(\theta_L)$. Since the ACO's payoff $u(x)$ is strictly increasing in α and β, if $\beta^*(\theta_H) \geq \beta^*(\theta_L)$ and $\alpha^*(\theta_H) > \alpha^*(\theta_L)$, then a type θ_L ACO could earn a higher payoff by reporting its type to be θ_H instead of θ_L, which violates incentive compatibility. Therefore, $\alpha^*(\theta_H) \leq \alpha^*(\theta_L)$. □

Proof of Proposition 3.3. We first establish two supporting results. For conciseness, we suppress dependence of $x(\cdot)$ on α and β and dependence of $v(\cdot)$ on α. First, we show that Medicare's savings function $v(x(\theta),\beta,\theta)$ is continuous in β. Note that $v(x,\beta,\theta)$ is continuous in x and β. By the Berge Maximum Theorem, $x(\theta)$ is upperhemicontinuous in β. Since by Lemma 3.1 the optimal savings $x(\theta)$ is unique, $x(\theta)$ is also continuous in β. Hence $v(x(\theta),\beta,\theta)$

is continuous in β.

Next, we show that for all $\alpha \in \mathcal{A}$ and $\theta \in \Theta$, there exists $\beta_s > 0$ such that $(\partial/\partial x)\partial v(x, \beta, \theta) > 0$. Pick some α and θ. Now define $\delta = 1 - (\partial/\partial x)R(x, \alpha) > 0$, where the inequality follows from Assumption 3.4. Following the proof of Proposition 3.2, we have

$$(\partial/\partial x)v(x, \beta, \theta) \geq \delta - \beta\left((\partial/\partial x)c(\bar{x}, \theta) + c(\bar{x}, \theta)\bar{g}\right).$$

Since $\delta > 0$, $(\partial/\partial x)v(x, \beta, \theta) > 0$ for sufficiently small β.

We now prove the main result. We wish to show that there exists some $\theta_s > 0$ such that for any $\theta \geq \theta_s$ and $\alpha \in \mathcal{A}$, there exists $\tilde{\beta} > 0$ such that $v(x(\theta), 0, \theta) < v(x(\theta), \tilde{\beta}, \theta)$. Since $v(x(\theta), \beta, \theta)$ is continuous in β, for any $\tilde{\beta}$ we can apply the Mean Value Theorem to write

$$v(x(\theta), \tilde{\beta}, \theta) = v(x(\theta), 0, \theta) + \tilde{\beta} \cdot (d/d\beta)v(x(\theta), \beta, \theta)|_{\beta=\eta\tilde{\beta}}$$

for some $\eta \in [0, 1]$. Therefore, to show $v(x(\theta), 0, \theta) < v(x(\theta), \tilde{\beta}, \theta)$ for each $\theta \geq \theta_s$ and $\tilde{\beta}$, it suffices to show that

$$\tilde{\beta}(d/d\beta)v(x(\theta), \beta, \theta)|_{\beta=\eta\tilde{\beta}} > 0$$

for sufficiently large θ and an associated $\tilde{\beta}$ and $\eta \in [0, 1]$. We do so by showing that there exists θ_s such that for all $\theta \geq \theta_s$, there exists a constant $\tilde{\beta}$ such that $(d/d\beta)v(x(\theta), \beta, \theta) > 0$ for all $\beta \in [0, \tilde{\beta}]$ and $\alpha \in \mathcal{A}$. Pick $\tilde{\beta} = \beta_s$. Taking the total derivative yields

$$(d/d\beta)v(x(\theta), \beta, \theta) = (\partial/\partial x)v(x, \beta, \theta)(d/d\beta)x(\theta) + (\partial/\partial\beta)v(x, \beta, \theta).$$

Note that $(\partial/\partial x)v(x, \beta, \theta) > 0$ for all $\alpha \in \mathcal{A}$ if $\beta \leq \beta_s$ from the earlier argument, and by Lemma 3.3, $(d/d\beta)x(\theta) > 0$ for all β if $\theta \geq \theta_{\alpha,0}$. It follows that

$$(\partial/\partial x)v(x, \beta, \theta)(d/d\beta)x(\theta) > 0$$

for any $\alpha \in \mathcal{A}$, $\theta \geq \theta_{\alpha,0}$ and $\beta \in [0, \beta_s]$. We also have $(\partial/\partial\beta)v(x, \beta, \theta) = -c(x, \theta)(1-G(h-x))$, which, by Assumption 3.1, can be made arbitrarily small by picking θ to be large. It follows that for all $\alpha \in \mathcal{A}$, $\beta \in [0, \beta_s]$, $\lim_{\theta\to\infty} (d/d\beta)v(x(\theta), \beta, \theta) > 0$. Thus there exists θ_s such that for all $\theta \geq \theta_s$ and $\alpha \in \mathcal{A}$, there exists $\tilde{\beta}$ such that $v(x(\theta), 0, \theta) < v(x(\theta), \tilde{\beta}, \theta)$. It follows that $V_s^*(\theta) > V_0^*(\theta)$. Since for any $x > 0$, the ACO's payoff $u(x)$ is strictly increasing in β, it follows that $U_s^*(\theta) > U_0^*(\theta)$ as well. □

Proof of Proposition 3.4. We prove that the model is identifiable for a single benchmark group, since the proof extends in a straightforward manner to multiple benchmark groups. Let μ be fixed. We wish to show that if $h(y|\mu, \lambda, \sigma) = h(y|\mu, \tilde{\lambda}, \sigma')$ for all y, then we must have $\lambda = \tilde{\lambda}$ and $\sigma = \sigma'$ (Bickel and Doksum, 2015). Suppose there exists parameters (λ, σ) and $(\tilde{\lambda}, \sigma')$ such that $(\lambda, \sigma) \neq (\tilde{\lambda}, \sigma')$ and $h(y|\mu, \lambda, \sigma) = h(y|\mu, \tilde{\lambda}, \sigma')$ for all y. Thus $h(y|\mu, \lambda, \sigma) = h(y|\mu, \tilde{\lambda}, \sigma')$ for all $y \geq x(\bar{\theta})$ as well. Writing $h(y|\mu, \lambda, \sigma)$ in terms of the shock and type densities, we have $h(y|\mu, \lambda, \sigma) = \int_\Theta g(y - x(\theta)|\sigma, \theta)f(\theta|\lambda)d\theta$. Since g is the Laplace density,

we have

$$\int_\Theta (1/(2\sigma))e^{\sqrt{2}(y-x(\theta))/\sigma} f(\theta|\lambda)d\theta = \int_\Theta (1/(2\sigma'))e^{\sqrt{2}(y-x(\theta))/\sigma'} f(\theta|\tilde{\lambda})d\theta \quad \text{for all } y \geq x(\bar{\theta}). \tag{A.42}$$

We consider two cases: $\sigma \neq \sigma'$ and $\sigma = \sigma'$. First suppose $\sigma \neq \sigma'$. For conciseness, define

$$C = \int_\Theta (1/(2\sigma))e^{-\sqrt{2}x(\theta)/\sigma} f(\theta|\lambda)d\theta$$

and

$$C' = \int_\Theta (1/(2\sigma'))e^{-\sqrt{2}x(\theta)/\sigma'} f(\theta|\tilde{\lambda})d\theta.$$

Note that C and C' are constant with respect to y. Equation (A.42) then implies

$$Ce^{\sqrt{2}y/\sigma} = C'e^{\sqrt{2}y/\sigma'}$$

for all $y \geq x(\bar{\theta})$. Taking the natural logarithm of both sides and rearranging yields

$$\left(\sqrt{2}/\sigma - \sqrt{2}/\sigma'\right) y + \ln C - \ln C' = 0$$

for all $y \geq x(\bar{\theta})$. This linear equation is zero over all $y \geq x(\bar{\theta})$ only if $\left(\sqrt{2}/\sigma - \sqrt{2}/\sigma'\right) = 0$, which yields a contradiction. Now consider the case where $\sigma = \sigma'$. Then we must have $\lambda \neq \tilde{\lambda}$. Cancelling out the common $(1/(2\sigma))e^{\sqrt{2}y/\sigma}$ terms on both sides of (A.42) and rearranging, we have

$$\int_\Theta e^{-\sqrt{2}x(\theta)/\sigma} \left(f(\theta|\lambda) - f(\theta|\tilde{\lambda})\right) d\theta = 0,$$

which violates the assumption in the proposition statement. □

Proof of Corollary 3.1. Note $\Theta = [\underline{\theta}, \bar{\theta}]$. Let $\underline{\theta}$ be fixed. Pick some $(\lambda, \tilde{\lambda}, \sigma) \in \Lambda \times \Lambda \times \Sigma$ such that $\lambda \neq \tilde{\lambda}$, and suppose the following equality holds:

$$\int_\Theta e^{-\sqrt{2}x(\theta)/\sigma} \left(f(\theta|\lambda) - f(\theta|\tilde{\lambda})\right) = 0 \tag{A.43}$$

We first show that if (A.43) holds, then $\bar{\theta}$ is unique. Since $e^{-\sqrt{2}x(\theta)/\sigma}$ and $f(\theta|\lambda)$ are both strictly positive, the equality

$$\int_{[\theta_1,\theta_2]} e^{-\sqrt{2}x(\theta)/\sigma} \left(f(\theta|\lambda) - f(\theta|\tilde{\lambda})\right) d\theta = 0$$

can only hold for an interval $[\theta_1, \theta_2] \subset \mathbb{R}_+$ if $f(\theta|\lambda)$ and $f(\theta|\tilde{\lambda})$ intersect at some point in (θ_1, θ_2). Further, since $f(\theta|\lambda)$ and $f(\theta|\tilde{\lambda})$ are exponential densities, they can intersect at most once over any interval (θ_1, θ_2). Now suppose there exists multiple values of $\bar{\theta}$ such that (A.43)

holds. Let $\bar{\theta}^1$ and $\bar{\theta}^2$ be two such values, where $\bar{\theta}^1 < \bar{\theta}^2$. Then we have

$$\int_{[\underline{\theta},\bar{\theta}^1]} e^{-\sqrt{2}x(\theta)/\sigma} \left(f(\theta|\lambda) - f(\theta|\tilde{\lambda}) \right) d\theta = 0$$

and

$$\int_{[\underline{\theta},\bar{\theta}^2]} e^{-\sqrt{2}x(\theta)/\sigma} \left(f(\theta|\lambda) - f(\theta|\tilde{\lambda}) \right) d\theta = 0$$

which implies

$$\int_{[\bar{\theta}^1,\bar{\theta}^2]} e^{-\sqrt{2}x(\theta)/\sigma} \left(f(\theta|\lambda) - f(\theta|\tilde{\lambda}) \right) d\theta = 0.$$

It follows that $f(\theta|\lambda)$ and $f(\theta|\tilde{\lambda})$ intersect in both $(\underline{\theta}, \bar{\theta}^1)$ and $(\bar{\theta}^1, \bar{\theta}^2)$, a contradiction. Thus, there is at most one value $\bar{\theta}$ such that (A.43) holds. Let this parameter be $\bar{\theta}(\lambda, \tilde{\lambda}, \sigma)$. Since Λ and Σ are discrete, there are a finite number of such $\bar{\theta}(\lambda, \tilde{\lambda}, \sigma)$. Selecting $\bar{\theta} > \sup_{\lambda,\tilde{\lambda}\in\Lambda,\sigma\in\Sigma} \bar{\theta}(\lambda, \tilde{\lambda}, \sigma)$ implies (A.43) cannot hold for any $(\lambda, \tilde{\lambda}, \sigma) \in \Lambda \times \Lambda \times \Sigma$, which yields the result. □

Proof of Proposition 3.5. For conciseness let $\lambda = (\lambda_1, \ldots, \lambda_m)$ in what follows. Let $h(y|\mu, \lambda, \sigma)$ be the savings density for an ACO with benchmark μ, given λ and σ. Letting $\mathcal{L}(\lambda, \sigma|\mu, y)$ be the likelihood function, we can write

$$\begin{aligned}
\mathcal{L}(\lambda, \sigma|\mu, y) &= \prod_{i=1}^{n} h(y_i|\mu_i, \lambda, \sigma) \\
&= \prod_{i=1}^{n} \int_{\Theta} h(y_i|\mu_i, \lambda, \sigma, \theta) f(\theta|\mu_i, \lambda, \sigma) d\theta \\
&= \prod_{i=1}^{n} \int_{\Theta} g(y_i - x(\theta)|\mu_i, \lambda, \sigma, \theta) f(\theta|\mu_i, \lambda, \sigma) d\theta \\
&= \prod_{i=1}^{n} \int_{\Theta} g(y_i - x(\theta)|\mu_i, \lambda(\mu_i), \sigma, \theta) f(\theta|\lambda(\mu_i)) d\theta \\
&= \prod_{i=1}^{n} \int_{\Theta} g(y_i - x(\theta)|\sigma, \theta) f(\theta|\lambda(\mu_i)) d\theta
\end{aligned}$$

The first line follows by definition of the likelihood function and the independence assumption given in Assumption 3.5. The second line follows from conditioning on θ. The third line follows from noting that $y = x(\theta) + \xi$ and re-writing the savings distribution in terms of the shock density g. The fourth line follows from noting that $f(\theta|\mu_i, \lambda, \sigma) = f(\theta|\lambda(\mu_i))$, since $\lambda(\mu_i)$ fully defines the type density of an ACO. The final line follows by observing that $y_i - x(\theta)$ depends only on σ and θ (in addition to α and β, which are fixed throughout). Taking the logarithm of both sides, we obtain

$$\log \mathcal{L}(\lambda, \sigma|\mu, y) = \sum_{i=1}^{n} \log \left(\int_{\Theta} g\left(y_i - x(\theta)|\sigma, t\right) f\left(t|\lambda(\mu_i)\right) d\theta \right),$$

as desired. □

Proof of Proposition 3.6. It suffices to show that for any $\epsilon > 0$, there exists $\bar{T} > 0$ such that $|\tilde{x}_t - x_t^*| < \epsilon$ for all $t \geq \bar{T}$. First we show that for any ϵ, there exists $\tilde{T}$ such that $\tilde{x}_t \leq \epsilon$ for all $t \geq \tilde{T}$. Suppose there exists $\epsilon > 0$ and a subsequence $\tilde{x}_{t_k}$ such that $\tilde{x}_{t_k} > \epsilon$ for all $k \geq 0$. Let $k(t) = \sup\{k \geq 0 | t_k \leq t\}$. Since $\tilde{x}_t \geq 0$ for all $t \geq 0$, for any t we have $\sum_{i=0}^{t} \tilde{x}_i \geq \sum_{k=0}^{k(t)} \tilde{x}_{t_k}$. Then we have

$$\lim_{t\to\infty} s_t = \lim_{t\to\infty} \sum_{i=0}^{t} \tilde{x}_i \geq \lim_{t\to\infty} \sum_{k=0}^{k(t)} \tilde{x}_{t_k} \geq \lim_{t\to\infty} \sum_{k=0}^{k(t)} \epsilon = \infty.$$

Since $c(s_t)$ is strictly increasing by Assumption 3.1, it follows that $\lim_{t\to\infty} u_t(\tilde{x}_t) = -\infty$, which cannot be optimal. Therefore, there must exist $\tilde{T} > 0$ such that $\tilde{x}_t \leq \epsilon$ for all $t \geq \tilde{T}$. By a parallel argument, for any $\epsilon > 0$ there exists T^* such that $x_t^* \leq \epsilon$ for all $t \geq T^*$. The result follows by taking $\bar{T} = \max\{T^*, \tilde{T}\}$. □

A.3 Proofs from Chapter 4

We first present two lemmas that are helpful in the proof of Theorem 4.1.

Lemma A.4. *Let* $\mathbf{y}$ *be fixed. Define* $\mathcal{U}^* = \operatorname{argmax}_{\mu\in\mathcal{U}} \text{CVaR}_\mu[d(\xi, \mathbf{y}(\xi))]$. *Then for any set of facility locations* $\mathbf{y}$*, there exists a worst-case distribution* $\tilde{\mu} \in \mathcal{U}^*$ *supported on* $a_1, \ldots, a_{|\mathcal{R}|}$*, where* $a_r = \operatorname{argmax}_{a\in\mathcal{A}'_r} d(a, \mathbf{y}(a))$*, for all* $r = 1, \ldots, |\mathcal{R}|$.

Proof. For any $\bar{\mathcal{A}} \subset \mathcal{A}$, define

$$\lambda(\mu, \bar{\mathcal{A}}) = P_\mu(\xi \in \bar{\mathcal{A}}) = \int_{\bar{\mathcal{A}}} f_\mu(\xi) d\xi, \tag{A.44}$$

where f_μ is the density function associated with distribution μ. Pick a distribution $\mu^* \in \mathcal{U}^*$. Now construct $\tilde{\mu}$ to be the distribution with mass $\lambda(\mu^*, \mathcal{A}'_1), \ldots, \lambda(\mu^*, \mathcal{A}'_{|\mathcal{R}|})$ placed at the locations $a_1, \ldots, a_{|\mathcal{R}|}$. We now show that $\tilde{\mu}$ is also in $\mathcal{U}^*$. We do so by first showing that $\tilde{\mu}$ is in $\mathcal{U}$, and then by showing that the CVaR of the distance distribution under the constructed $\tilde{\mu}$ is at least as large as the CVaR under μ^*. For the first step, let $\mathcal{R}_j$ index the subregions that comprise $\mathcal{A}_j$. Now note that for each $j \in \mathcal{J}$, we have

$$P_{\tilde{\mu}}(\xi \in \mathcal{A}_j) = \sum_{r\in\mathcal{R}_j} P_{\tilde{\mu}}(\xi \in \mathcal{A}'_r) = \sum_{r\in\mathcal{R}_j} \lambda(\tilde{\mu}, \mathcal{A}'_r) = \sum_{r\in\mathcal{R}_j} P_{\mu^*}(\xi \in \mathcal{A}'_r) = P_{\mu^*}(\xi \in \mathcal{A}_j) = \lambda_j.$$

The first equality is due to the fact that $\bigcup_{r\in\mathcal{R}_j} \mathcal{A}'_r = \mathcal{A}_j$. The second equality follows from the definition of $\lambda(\mu, \bar{\mathcal{A}})$ in (A.44). The third equality follows by the construction of $\tilde{\mu}$. The fourth equality is again due to $\bigcup_{r\in\mathcal{R}_j} \mathcal{A}'_r = \mathcal{A}_j$. The final equality holds since $\mu^* \in \mathcal{U}$. Since from the above equations, $P_{\tilde{\mu}}(\xi \in \mathcal{A}_j) = \lambda_j$, we have $\tilde{\mu} \in \mathcal{U}$. It remains to show that

$\text{CVaR}_{\mu^*} \leq \text{CVaR}_{\tilde{\mu}}$. Writing CVaR explicitly, we have

$$\begin{aligned}
\text{CVaR}_{\mu^*}[d(\xi, \mathbf{y}(\xi))] &= \min_{\alpha}\ \alpha + \frac{1}{1-\beta}\int_{\mathcal{A}} [d(\xi, \mathbf{y}(\xi)) - \alpha]^+ f_{\mu^*}(\xi) d\xi \\
&= \min_{\alpha}\ \alpha + \frac{1}{1-\beta}\sum_{r\in\mathcal{R}}\int_{\mathcal{A}'_r} [d(\xi, \mathbf{y}(\xi)) - \alpha]^+ f_{\mu^*}(\xi) d\xi \\
&\leq \min_{\alpha}\ \alpha + \frac{1}{1-\beta}\sum_{r\in\mathcal{R}}\int_{\mathcal{A}'_r} [d(a_r, \mathbf{y}(a_r)) - \alpha]^+ f_{\mu^*}(\xi) d\xi \\
&= \min_{\alpha}\ \alpha + \frac{1}{1-\beta}\sum_{r\in\mathcal{R}} [d(a_r, \mathbf{y}(a_r)) - \alpha]^+ \int_{\mathcal{A}'_r} f_{\mu^*}(\xi) d\xi \\
&= \min_{\alpha}\ \alpha + \frac{1}{1-\beta}\sum_{r\in\mathcal{R}} [d(a_r, \mathbf{y}(a_r)) - \alpha]^+ \lambda(\tilde{\mu}, \mathcal{A}'_r) \\
&= \text{CVaR}_{\tilde{\mu}}[d(\xi, \mathbf{y}(\xi))].
\end{aligned}$$

The first line follows from the definition of CVaR. The second line is due to the fact that the subregions $\mathcal{A}'_r$, $r \in \mathcal{R}$ form a partition of $\mathcal{A}$. The third line follows from the definition of a_r as a point in $\mathcal{A}'_r$ that maximizes distance to the nearest AED. The fourth line follows since $[d(a_r, \mathbf{y}(a_r)) - \alpha]^+$ is constant with respect to the variable ξ. The fifth line follows from the definition of $\lambda(\tilde{\mu}, \mathcal{A}'_r)$ as the total probability mass in $\mathcal{A}'_r$. The final line follows by construction of $\tilde{\mu}$ and the definition of CVaR. □

For the second lemma, define $z_C(\mathbf{y}) = \max_{\mu\in\mathcal{U}} \text{CVaR}_{\mu}[d(\xi, \mathbf{y}(\xi))]$ to be the worst-case CVaR under the solution $\mathbf{y}$ in the continuous case. Similarly, let $z_D(\mathbf{y})$ be the worst-case CVaR in the discrete case (i.e., the optimal value of (4.4) given $\mathbf{y}$). The next lemma bounds the difference between the the worst-case continuous and discrete CVaR, for a given deployment $\mathbf{y}$. Note that Lemma 1 is used in the proof.

Lemma A.5. *For any* $\mathbf{y}$, $z_C(\mathbf{y}) - z_D(\mathbf{y}) \leq \ell(\sigma)/(1-\beta)$.

Proof. Let μ^* be a worst-case continuous distribution with mass $\lambda'_1, \ldots, \lambda'_{|\mathcal{R}|}$ on the locations $p_1, \ldots, p_{|\mathcal{R}|}$ (cf Lemma A.4). Construct a discrete distribution by placing mass $\lambda'_1, \ldots, \lambda'_{|\mathcal{R}|}$ on the locations $\xi_1, \ldots, \xi_{|\mathcal{R}|}$, where $\xi_r = \underset{\xi\in\Xi'_r}{\text{argmin}}\ d(\xi, p_r)$ for all r. We can now

write

$$\begin{aligned}
z_C(\mathbf{y}) - z_D(\mathbf{y}) &= \left(\min_{\alpha} \ \alpha + \frac{1}{1-\beta} \sum_{r \in \mathcal{R}} \lambda'_r [d(a_r, \mathbf{y}(a_r)) - \alpha]^+ \right) - z_D(\mathbf{y}) \\
&\leq \left(\min_{\alpha} \ \alpha + \frac{1}{1-\beta} \sum_{r \in \mathcal{R}} \lambda'_r [d(\xi_r, \mathbf{y}(\xi_r)) + \ell(\sigma) - \alpha]^+ \right) - z_D(\mathbf{y}) \\
&\leq \left(\min_{\alpha} \ \alpha + \frac{1}{1-\beta} \sum_{r \in \mathcal{R}} \lambda'_r [d(\xi_r, \mathbf{y}(\xi_r)) - \alpha]^+ \right) + \frac{1}{1-\beta} \sum_{r \in \mathcal{R}} \lambda'_r \ell(\sigma) - z_D(\mathbf{y}) \\
&= z_D(\mathbf{y}) + \frac{1}{1-\beta} \sum_{r \in \mathcal{R}} \lambda'_r \ell(\sigma) - z_D(\mathbf{y}) \\
&= \frac{\ell(\sigma)}{1-\beta}
\end{aligned}$$

The first line follows from the definition of $z_C(\mathbf{y})$. The second line follows from the triangle inequality and since $d(\xi_r, a_r) \leq \ell(\sigma)$ by definition of $\ell(\sigma)$. The third line follows from pulling $\ell(\sigma)$ out of the CVaR expression. The fourth line follows by definition of $z_D(\mathbf{y})$. The fifth line follows from cancelling out $z_D(\mathbf{y})$ terms and noting that $\sum_{r \in \mathcal{R}} \lambda'_r = 1$. □

Proof of Theorem 4.1. Note that $z_D(\mathbf{y}) \leq z_C(\mathbf{y})$ for any $\mathbf{y}$, since $\Xi'_r \subset \mathcal{A}'_r$ for all $r \in \mathcal{R}$. Now define $\mathbf{y}_C \in \underset{\mathbf{y}}{\operatorname{argmin}}\ z_C(\mathbf{y})$ and $\mathbf{y}_D \in \underset{\mathbf{y}}{\operatorname{argmin}}\ z_D(\mathbf{y})$. Using Lemma A.5, we can now write

$$z_D(\mathbf{y}_D) \leq z_D(\mathbf{y}_C) \leq z_C(\mathbf{y}_C) = Z_C \leq z_C(\mathbf{y}_D) \leq z_D(\mathbf{y}_D) + \ell(\sigma)/(1-\beta) = Z_D + \ell(\sigma)/(1-\beta),$$

which implies $Z_C - Z_D \leq \ell(\sigma)/(1-\beta)$. Since $Z_D \leq Z_C$, we have

$$\Delta = \frac{Z_C - Z_D}{Z_C} \leq \frac{Z_C - Z_D}{Z_D} \leq \frac{\ell(\sigma)}{Z_D(1-\beta)}.$$

It follows that if $\ell(\sigma) \leq \varepsilon(1-\beta)Z_D$, then $\Delta \leq \varepsilon$. □

Proof of Proposition 4.1. Since each scenario k belongs to exactly one set $\mathcal{K}_j$, in the context of formulation R-AED this implies that the uncertainty regions are disjoint. The set $\mathcal{U}$ is then separable in j, meaning the uncertainty in the discrete distribution $\mathbf{u}$ can be interpretted equivalently as uncertainty in $|\mathcal{J}|$ *conditional* distributions of $\mathbf{u}$ – one for each uncertainty region. We may now interpret x_j^k as the unknown probability that the next cardiac arrest arrives at location k, conditioned on the event $\{\xi \in \Xi_j\}$. Note that due to our choice of the uncertainty set $\mathcal{U}$, each vector $\mathbf{x}_j$ is only known to reside in the $|\mathcal{K}_j - 1|$-dimensional probability simplex, which is given precisely by the sets $\mathbf{X}_1, \ldots, \mathbf{X}_{|\mathcal{J}|}$. Also, since each scenario k only belongs to one of the sets $\mathcal{K}_1, \ldots, \mathcal{K}_{|\mathcal{J}|}$, it follows that $u_k = \sum_{j \in \mathcal{J}} \lambda_j x_{jk}$, for each $k \in \mathcal{K}$. By substituting $\sum_{j \in \mathcal{J}} \lambda_j x_{jk}$ for u_k, we can write formulation

R-AED equivalently as

$$\begin{aligned} \min_{\mathbf{y},\mathbf{z}} \max_{\mathbf{x}} \min_{\alpha} \ & \alpha + \frac{1}{(1-\beta)} \sum_{j\in\mathcal{J}} \sum_{k\in\mathcal{K}_j} \lambda_j x_{jk} \max\left\{ \sum_{i\in\mathcal{I}} d_{ik} z_{ik} - \alpha, 0 \right\} \\ \text{subject to } & \mathbf{x} \in \mathbf{X}, \mathbf{z} \in \mathbf{Z}(\mathbf{y}), \mathbf{y} \in \mathbf{Y}, \\ & \alpha \geq 0. \end{aligned} \tag{A.45}$$

Due to the separability of the scenario sets in j, we may take the x_{jk} terms inside the max operator in the objective to obtain another reformulation of R-AED:

$$\begin{aligned} \min_{\mathbf{y},\mathbf{z}} \max_{\mathbf{x}} \min_{\alpha} \ & \alpha + \frac{1}{(1-\beta)} \sum_{j\in\mathcal{J}} \lambda_j \max\left\{ \sum_{i\in\mathcal{I}} \sum_{k\in\mathcal{K}_j} d_{ik} z_{ik} x_{jk} - \alpha, 0 \right\} \\ \text{subject to } & \mathbf{x} \in \mathbf{X}, \mathbf{z} \in \mathbf{Z}(\mathbf{y}), \mathbf{y} \in \mathbf{Y}, \\ & \alpha \geq 0. \end{aligned} \tag{A.46}$$

We now show equivalence of (A.46) and the robust p-median problem. Let $(\mathbf{y}^*, \mathbf{z}^*, \mathbf{x}^*)$ be an optimal solution for formulation (A.46). Let $Z(\alpha)$ be the objective value of (A.46) for the solution $(\mathbf{y}^*, \mathbf{z}^*, \mathbf{x}^*, \alpha)$. Since the feasible sets of (A.46) and (4.8) are identical, all that remains is to show that their optimal values are equal. For conciseness define $g(\mathbf{x}_j, \mathbf{z}_j) := \sum_{i\in\mathcal{I}} \sum_{k\in\mathcal{K}_j} d_{ik} z_{ik} x_{jk}$. We first observe that for $\beta = 0$ and any $\alpha \leq \min_{j\in\mathcal{J}} g(\mathbf{x}_j^*, \mathbf{z}_j^*)$, the objective value of (A.46),

$$\begin{aligned} Z(\alpha) &= \alpha + \sum_{j\in\mathcal{J}} \lambda_j \left(g(\mathbf{x}_j^*, \mathbf{z}_j^*) - \alpha \right) \\ &= \alpha + \sum_{j\in\mathcal{J}} \lambda_j g(\mathbf{x}_j^*, \mathbf{z}_j^*) - \frac{1}{n}(n\alpha) \\ &= \sum_{j\in\mathcal{J}} \lambda_j g(\mathbf{x}_j^*, \mathbf{z}_j^*) \\ &= \sum_{i\in\mathcal{I}} \lambda_j \sum_{j\in\mathcal{J}} \sum_{k\in\mathcal{K}_j} z_{ik}^* d_{ik} x_{jk}^*, \end{aligned}$$

is equivalent to the objective value of (4.8). Note that in the second equality above we use the fact that $\lambda_j = 1/n$ for all $j \in \mathcal{J}$. Thus to show equivalence in the optimal values of formulations R-AED and (4.8) when $\beta = 0$, it suffices to show that an optimal α^* satisfies

$\alpha^* \leq \min_{j \in \mathcal{J}} g(\mathbf{x}_j^*, \mathbf{z}_j^*)$. Suppose that $\alpha^* > \min_{j \in \mathcal{J}} g(\mathbf{x}_j^*, \mathbf{z}_j^*)$. Then,

$$\begin{aligned} Z(\alpha^*) &= \alpha^* + \sum_{j \in \mathcal{J}} \lambda_j \max\left\{g(\mathbf{x}_j^*, \mathbf{z}_j^*) - \alpha^*, 0\right\} \\ &> \alpha^* + \sum_{j \in \mathcal{J}} \lambda_j \left(g(\mathbf{x}_j^*, \mathbf{z}_j^*) - \alpha^*\right) \\ &= \sum_{j \in \mathcal{J}} \lambda_j g(\mathbf{x}_j^*, \mathbf{z}_j^*) \\ &= \sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{J}} \sum_{k \in \mathcal{K}_j} \lambda_j z_{ik}^* d_{ik} x_{jk}^*, \end{aligned}$$

which contradicts the optimality of α^*. For the robust p-center problem, let $(\mathbf{y}^*, \mathbf{z}^*, \mathbf{x}^*)$ again be an optimal solution to formulation (A.46). Let $Z(\alpha)$ be the objective value of (A.46) for the solution $(\mathbf{y}^*, \mathbf{z}^*, \mathbf{x}^*, \alpha)$. If $\alpha = \max_{j \in \mathcal{J}} \{g(\mathbf{x}_j^*, \mathbf{z}_j^*)\}$, then the objective of formulation (A.46),

$$\begin{aligned} Z(\alpha) &= \alpha + \frac{1}{(1-\beta)} \sum_{j \in \mathcal{J}} \lambda_j \max\ \left\{g(\mathbf{x}_j^*, \mathbf{z}_j^*) - \alpha, 0\right\} \\ &= \max_{j \in \mathcal{J}} \{g(\mathbf{x}_j^*, \mathbf{z}_j^*)\} + \frac{1}{(1-\beta)} \sum_{j \in \mathcal{J}} \lambda_j \max\ \left\{g(\mathbf{x}_j^*, \mathbf{z}_j^*) - \max_{j \in \mathcal{J}} \{g(\mathbf{x}_j^*, \mathbf{z}_j^*)\}, 0\right\} \\ &= \max_{j \in \mathcal{J}} \{g(\mathbf{x}_j^*, \mathbf{z}_j^*)\} \\ &= \max_{j \in \mathcal{J}} \left\{\sum_{i \in \mathcal{I}} \sum_{k \in \mathcal{K}_j} z_{ik}^* d_{ik} x_{jk}^*\right\}, \end{aligned}$$

is equivalent to the objective of (4.9). We now wish to show that an optimal α^* satisfies $\alpha^* = \max_{j \in \mathcal{J}} \{g(\mathbf{x}_j^*, \mathbf{z}_j^*)\}$ for $\beta > (n-1)/n$. We do so by showing that neither $\alpha^* > \max_{j \in \mathcal{J}} \{g(\mathbf{x}_j^*, \mathbf{z}_j^*)\}$ nor $\alpha^* < \max_{j \in \mathcal{J}} \{g(\mathbf{x}_j^*, \mathbf{z}_j^*)\}$ can hold. First, suppose $\alpha^* > \max_{j \in \mathcal{J}} \{g(\mathbf{x}_j^*, \mathbf{z}_j^*)\}$. Then,

$$\begin{aligned} Z(\alpha^*) &= \alpha^* + \frac{1}{(1-\beta)} \sum_{j \in \mathcal{J}} \lambda_j \max\ \left\{g(\mathbf{x}_j^*, \mathbf{z}_j^*) - \alpha^*, 0\right\} \\ &= \alpha^* \\ &> \max_{j \in \mathcal{J}} \{g(\mathbf{x}_j^*, \mathbf{z}_j^*)\}, \end{aligned}$$

which contradicts the optimality of α^*. Now suppose $\alpha^* < \max_{j \in \mathcal{J}} \{g(\mathbf{x}_j^*, \mathbf{z}_j^*)\}$, and let $\delta = \max_{j \in \mathcal{J}} \{g(\mathbf{x}_j^*, \mathbf{z}_j^*)\} - \alpha^*$. It follows that $\sum_{j \in \mathcal{J}} \max\{g(\mathbf{x}_j^*, \mathbf{z}_j^*) - \alpha^*, 0\} \geq \delta$. Since $\lambda_j = 1/n$ for

all $j \in \mathcal{J}$, we can write

$$\begin{aligned} Z(\alpha^*) &= \alpha^* + \frac{1}{(1-\beta)n} \sum_{j \in \mathcal{J}} \max \left\{ g(\mathbf{x}_j^*, \mathbf{z}_j^*) - \alpha^*, 0 \right\} \\ &\geq \alpha^* + \frac{1}{(1-\beta)n} \delta \\ &= \alpha^* + \delta + \frac{1}{(1-\beta)n} \delta - \delta \\ &= \max_{j \in \mathcal{J}} \left\{ g(\mathbf{x}_j^*, \mathbf{z}_j^*) \right\} + \left(\frac{1}{(1-\beta)n} - 1 \right) \delta \\ &> \max_{j \in \mathcal{J}} \left\{ g(\mathbf{x}_j^*, \mathbf{z}_j^*) \right\}, \end{aligned}$$

where the final strict inequality follows from $\beta > (n-1)/n$. Again, this contradicts the optimality of α^*. □

Proof of Proposition 4.2. Note that the number of variables in the master problem formulation R-AED-MP is $|\mathcal{I}|(|\mathcal{K}_+|+1)+1$, and the number of constraints is $(2|\mathcal{I}|+1)|\mathcal{K}_+|+|\mathcal{I}|+|\mathcal{S}|+2$. Hence the number of variables and constraints in iteration s is $O(|\mathcal{I}||\mathcal{K}_+|)$. It remains to show that $|\mathcal{K}_+| \leq (|\mathcal{J}|+1)s$ in the s^{th} iteration of the row-and-column generation algorithm. Since the new variables and constraints that are added to R-AED-MP in iteration s correspond directly to non-zero elements of the vector $\mathbf{p}_s$ obtained at a solution to R-AED-SP (cf Algorithm 1), it suffices to show that any $\mathbf{p}$ obtained at a solution to R-AED-SP contains at most $|\mathcal{J}|+1$ non-zero elements. We first prove that any $\mathbf{u}$ obtained at an optimal solution to R-AED-SP is also an optimal solution to the following problem:

$$\begin{aligned} \underset{\mathbf{u}}{\text{maximize}} \quad & \sum_{k \in \mathcal{K}} d_{ik} \bar{z}_{ik} u_k \\ \text{subject to} \quad & \sum_{k \in \mathcal{K}} u_k = 1, \\ & \sum_{k \in \mathcal{K}_j} u_k = \lambda_j, \quad j \in \mathcal{J}, \\ & u_k \geq 0. \end{aligned} \tag{A.47}$$

By way of contradiction, suppose a solution $\mathbf{p}, \mathbf{u}$ is optimal to R-AED-SP but $\mathbf{u}$ is suboptimal to (A.47). Then there exists k_1 and k_2 such that $\sum_{i \in \mathcal{I}} d_{ik_2} \bar{z}_{ik_2} > \sum_{i \in \mathcal{I}} d_{ik_1} \bar{z}_{ik_1}$. Further, there must exist $\epsilon > 0$ and a vector $\tilde{\mathbf{u}} \in \mathcal{U}$ such that $\tilde{u}_{k_1} = u_{k_1} - \epsilon$, $\tilde{u}_{k_2} = u_{k_2} + \epsilon$, and $\tilde{u}_k = u_k$ for all $k \in \mathcal{K} \setminus \{k_1, k_2\}$. Let $\tilde{\mathbf{p}}$ be the solution obtained by solving R-AED-SP with $\tilde{\mathbf{u}}$ as a parameter. Now note in R-AED-SP, since $\sum_{i \in \mathcal{I}} d_{ik_2} \bar{z}_{ik_2} > \sum_{i \in \mathcal{I}} d_{ik_1} \bar{z}_{ik_1}$, the constraints $p_{k_2} \leq u_{k_2}/(1-\beta)$ and $\tilde{p}_{k_2} \leq \tilde{u}_{k_2}/(1-\beta)$ must be binding (otherwise we could slightly increase p_{k_2} while slightly decreasing p_{k_1} by the same amount to improve the objective function, with an identical argument applying to $\tilde{p}_{k_2}$). Since $p_{k_2} = u_{k_2}/(1-\beta)$ and $\tilde{p}_{k_2} = \tilde{u}_{k_2}/(1-\beta)$, by the construction of $\tilde{\mathbf{u}}$ it follows that $\tilde{p}_2 - p_2 \geq \epsilon/(1-\beta)$ and $p_1 - \tilde{p}_1 \leq \epsilon/(1-\beta)$. Since

$\sum_{i\in\mathcal{I}} d_{ik_2}\bar{z}_{ik_2} > \sum_{i\in\mathcal{I}} d_{ik_1}\bar{z}_{ik_1}$, the solution $\tilde{\mathbf{p}}, \tilde{\mathbf{u}}$ improves the objective by

$$\left(\sum_{i\in\mathcal{I}} d_{ik_2}\bar{z}_{ik_2} - \sum_{i\in\mathcal{I}} d_{ik_1}\bar{z}_{ik_1}\right)\frac{\epsilon}{1-\beta} > 0$$

over the solution $\mathbf{p}, \mathbf{u}$, which contradicts the optimality of $\mathbf{p}, \mathbf{u}$ with respect to R-AED-SP. Hence $\mathbf{u}$ must be optimal to (A.47) as well.

Since (A.47) is a standard form linear program with $|\mathcal{J}|+1$ equality constraints and $|\mathcal{K}|$ variables, at most $|\mathcal{J}|+1$ elements of a basic feasible solution to (A.47) can be non-zero (Bertsimas and Tsitsiklis, 1997). Further, since by assumption there are no ties in the d_{ik} parameters, every optimal solution to (A.47) must be a basic feasible solution. It follows that every $\mathbf{u}$ obtained at a solution of R-AED-SP has no more than $|\mathcal{J}|+1$ non-zero elements. Since $p_k > 0$ implies $u_k > 0$ for any feasible $\mathbf{p}, \mathbf{u}$, it follows that any $\mathbf{p}$ obtained at a solution to R-AED-SP cannot have more than $|\mathcal{J}|+1$ non-zero elements. □

Proof of Proposition 4.3. Let the subregions $\mathcal{A}'_1, \ldots, \mathcal{A}'_{|\mathcal{R}|}$ be the partitioning of $\mathcal{A}$ implied by the uncertainty regions $\mathcal{A}_1, \ldots, \mathcal{A}_{|\mathcal{J}|}$. Let $\lambda'_r = P(\xi \in \mathcal{A}'_r)$ for all $r \in \mathcal{R}$ under the true distribution. Since the subregions are disjoint, the number of historical cardiac arrests that fall in each of the regions $\mathcal{A}'_1, \ldots, \mathcal{A}'_{|\mathcal{R}|}$ is simply the outcome of n independent trials from a multinomial distribution with parameters $\lambda'_1, \ldots, \lambda'_{|\mathcal{R}|}$. Now define

$$\hat{\lambda}'_r = \frac{1}{n}\sum_{c=1}^{n} \mathbb{1}\{a_c \in \mathcal{A}'_r\}, \quad r \in \mathcal{R},$$

is a maximum likelihood estimator of $\lambda'_1, \ldots, \lambda'_r$, and that $|\hat{\lambda}'_r - \lambda'_r| \longrightarrow 0$ with probability 1 as $n \longrightarrow \infty$ (e.g., Example 1.6.7. in Bickel and Doksum (2015)). Now let $\mathcal{R}_j$ be the set of subregions that comprise region $\mathcal{A}_j$. We now have $\hat{\lambda}_j = \sum_{r\in\mathcal{R}_j} \hat{\lambda}'_r$, and thus $|\hat{\lambda}_j - \lambda_j| = |\sum_{r\in\mathcal{R}_j}(\hat{\lambda}'_r - \lambda'_r)| \leq \sum_{r\in\mathcal{R}_j}|\hat{\lambda}'_r - \lambda'_r|$ for each $j \in \mathcal{J}$, which implies $|\hat{\lambda}_j - \lambda_j| \longrightarrow 0$ with probability 1 for each $j \in \mathcal{J}$ as $n \longrightarrow \infty$. To see that $\hat{\lambda}$ ensures that $\mathbf{U}$ is non-empty, construct a distribution $\mathbf{u} \in \mathbf{U}$ as follows. First, assign each of the cardiac arrest locations $a_1, \ldots, a_n$ to the nearest scenario location within the same subregion. Then for each $k \in \mathcal{K}$, set u_k equal to the proportion of cardiac arrests that are assigned to location k. Thus $\mathbf{u}$ is non-negative and $\sum_{k\in\mathcal{K}} u_k = 1$ by construction. Lastly, for each $r \in \mathcal{R}$, $\sum_{k\in\mathcal{K}'_r} u_k$ is equal to the proportion of cardiac arrests within subregion $\mathcal{A}'_r$, which is also equal to $\hat{\lambda}'_r$ by definition. Thus

$$\sum_{k\in\mathcal{K}_j} u_k = \sum_{r\in\mathcal{R}_j}\sum_{k\in\mathcal{K}'_r} u_k = \sum_{r\in\mathcal{R}_j} \hat{\lambda}'_r = \hat{\lambda}_j.$$

□