Sample Average Approximation for the Continuous Type Principal-Agent Problem

Preprint · April 2018

1 author:

Dashi Singham
Naval Postgraduate School

Sample Average Approximation for the Continuous Type Principal-Agent Problem

D. I. Singham*[1]

[1]Operations Research Department

Naval Postgraduate School

Monterey, CA 93943

dsingham@nps.edu

April 4, 2018

Abstract

We develop a method for finding approximate solutions to the continuous agent type principal-agent problem when analytical methods are not available. The solution is calculated by solving a discrete agent type version of the problem using sample average approximation and bootstrapping. We show how a solution to the approximate problem can be used to derive a lower bound and expected upper bound for the optimal objective function, and evaluate the error associated with the approximation. Numerical examples illustrate convergence in the approximate solution to the true solution as the number of samples increases. This works yields a method for obtaining some tractability in continuous type principal-agent problems where solutions were previously unavailable.

Keywords: pricing; mechanism design; principal-agent models; sample average approximation

1 Introduction

In the principal-agent problem, the principal optimizes the terms of an exchange with an agent who may have some private characteristic θ that is unknown to the principal. For example, the agent

*Corresponding author. Email to: dsingham@nps.edu

may have private demand θ for a product produced by the principal. While the exact value of θ may be unknown to the principal, both parties know the distribution of θ across different agents. The principal seeks to maximize her expected profit given uncertainty in θ by offering the agent quantity q units of the product at total price t.

This paper studies the principal-agent problem when the distribution of θ is continuous over a bounded range $[\underline{\theta}, \overline{\theta}]$ with density $f(\theta)$. We refer to this setting as the "continuous problem" and the setting where θ is a discrete random variable as the "discrete problem." In the continuous problem, the principal offers the agent contract functions $q(\theta)$ and $t(\theta)$, so the agent chooses the quantity and price option depending on his private demand θ. In the discrete problem, the principal offers an option (q_m, t_m) for each possible realization of the random variable $\theta_m, m = 1, \ldots, M$ according to the revelation principle.

Many principal-agent results rely on the ability to derive analytical solutions for the contract options. This paper explores the case where such analytical solutions are intractable either because the formulation is too complex, or the space of possible values for θ is too large. For the continuous problem, the solution can be found analytically for some specific functions $f(\theta)$. See Laffont and Martimort (2009) for the foundations behind the continuous problem, and Singham and Cai (2017) for a specific solution example. Performing optimization over a function space for $q(\theta)$ and $t(\theta)$ is generally a difficult problem. This paper provides a method for bounding the optimal profit and finding solution estimates for the continuous problem. This method can be used when the density $f(\theta)$ is too complex to yield analytical solutions, or when $f(\theta)$ may not be available but data samples of θ are present to estimate the distribution.

We refer to both the continuous problem formulation and its optimal objective value as Φ. We approximate the continuous problem using an empirical distribution with M discrete samples from $f(\theta)$ when M is very large. We call this discrete formulation and its optimal objective value Φ^M. Φ^M can be solved numerically, but for arbitrarily large M becomes computationally intractable due to the large number of decision variables and constraints. We show how a sample average approximation (SAA) to Φ^M, based on a smaller sample size N bootstrapped from the M samples, yields an upper bound in expectation on Φ^M. Call this SAA problem and its optimal objective value $\tilde{\Phi}^N$. The solution to $\tilde{\Phi}^N$ can be interpolated to provide a feasible solution which is a lower bound to the continuous problem Φ. Without needing to solve Φ^M, we show how a series of smaller

$\tilde{\Phi}^N$ problems can be used to bound the solution to Φ^M which closely approximates the continuous problem Φ.

Given recent advances in computing, optimization methods can be applied in new ways to solve traditional contracting problems. Bertsimas and Thiele (2005) show how historical data can be used to formulate linear programs to solve newsvendor formulations. Of particular interest is the work by Cecchini et al. (2013) which formulates and solves bilevel nonlinear programs to solve principal-agent problems numerically. The bilevel structure is apparent from the principal's problem incorporating the agent's optimization over his set of possible actions. The authors motivate the work by citing the limitations of the assumptions of the traditional principal-agent model, quoting Lambert (2006) who argues that limitations of relying on closed-form results limit the type and complexity of models that can be solved. Dempe (1995) originally suggested modeling the principal-agent problem as a bilevel program, and Cecchini et al. (2013) construct a version of the ellipsoid method for obtaining numerical solutions. This paper focuses on solving continuous principal-agent problems, but there has also been recent work to expand the types of discrete problems that can be solved using numerical methods. Cai and Singham (2018) developed a nonlinear programming formulation to solve principal-agent problems when agents were subject to one of multiple discrete demand distributions. The principal faces the adverse selection problem with regards to the possible demand distribution and both the principal and the agent face stochastic uncertainty within the distribution. Finally, Singham and Cai (2017) present the idea of using sample average approximation for principal-agent models and show an example without formalizing the reasoning. This paper aims to fully demonstrate and formalize this idea.

Straightforward sample average approximation methods cannot be directly applied to the continuous problem Φ in the usual way because the solution space of the discrete approximation Φ^M is fundamentally different from that of the continuous problem. The solution to Φ lies in a continuous function space, while any discrete problem has a finite-sized solution space which depends on the number of sampled values used. As N and M increase, the number of decision variables and constraints increases in the principal-agent problem, and so the feasible space also changes and is different from that of Φ. We construct Φ^M as a way of compiling all discrete problems on the same space, allowing us to invoke known SAA convergence properties. The main result is that we can use an easily solvable discrete problem to obtain useful information about an intractable continuous problem.

Section 2 details the formulations for the continuous and discrete problems. Section 3 presents the approximation problem and bounds. Section 4 presents numerical examples that demonstrate convergence of the optimal objective value as the number of bootstrap samples increases. Section 5 concludes.

2 Formulation

Maskin and Riley (1984) establish the conditions for adverse selection to be studied in the principal-agent setting, and show the nature of nonlinear pricing schemes for the principal using quantity discounts. Burnetas et al. (2007) study asymmetric information between a supplier and a retailer where there is uncertainty in the demand distribution and quantity discounts can be used to improve the supplier's profits. Babich et al. (2012) study contracting options with buyback between a supplier and a retailer where the retailer has private knowledge of the demand for the product, which can follow a discrete or continuous distribution. Our formulation models a similar setting and focuses on the effect of agent demand uncertainty in the adverse selection problem without considering a moral hazard in terms of agent effort, though we note the principal can can balance investing in monitoring both demand and effort to improve profits (Kung and Chen 2012).

First, we describe the continuous problem and the standard assumptions made. See Appendix 3.1 of Laffont and Martimort (2009) for a complete reference on continuous principal-agent problems. Let θ be a bounded continuous random variable for the demand of the agent on $[\underline{\theta}, \overline{\theta}]$ with bounded density $f(\theta)$. The principal faces a nonnegative, differentiable, increasing cost function $s(q)$ for producing q units of the product. The agent has a value function $v(q, \theta)$ for receiving q units of the product when his demand is θ. Standard assumptions include $v(0, \theta) = 0$, and $v(q, \theta)$ is nonnegative, twice-differentiable and convex in q, and increasing in q and θ. To ensure good behavior of our sampling approximation, we also assume that the functions $v(q, \theta)$ and $s(q)$ are bounded over the space of $q, \theta \in [0, \overline{\theta}]$.

The principal's objective is to offer the agent contract functions $(q(\theta), t(\theta))$ that maximize her expected profit $E[t(\theta) - s(q(\theta))]$ with respect to the random variable θ. Let the information rent (excess utility to the agent) be defined as $\Delta(\theta) \equiv v(q(\theta), \theta) - t(\theta)$. We adopt the convention to rewrite the decision variables for the principal using $(q(\theta), \Delta(\theta))$ instead of $(q(\theta), t(\theta))$. Then, the principal's problem for a continuous demand distribution is

$$\Phi = \max_{\{q(\theta),\Delta(\theta)\}_{\theta\in[\underline{\theta},\overline{\theta}]}} \int_{\underline{\theta}}^{\overline{\theta}} \left[v(q(\theta),\theta) - \Delta(\theta) - s(q(\theta))\right] f(\theta)d\theta \tag{2.1}$$

$$\begin{aligned}
\text{subject to} \quad & \Delta(\underline{\theta}) = 0 && && (IR_{\underline{\theta}})\\
& \Delta(\theta) \geq \Delta(\theta') + v(q(\theta'),\theta) - v(q(\theta'),\theta') && \theta,\theta' \in [\underline{\theta},\overline{\theta}], \theta' \leq \theta && (IC_{\theta\theta'})\\
& q(\theta) \geq q(\theta') && \theta,\theta' \in [\underline{\theta},\overline{\theta}], \theta' \leq \theta && (MON_{\theta\theta'}).
\end{aligned}$$

We call the optimal profit value to the continuous demand problem Φ, and abuse notation and also refer to this formulation as Φ. The optimal profit is a probability-weighted integral over the possible demand values, where the profit to the principal $t(\theta) - s(q(\theta))$ is rewritten using the utility function $v(q(\theta), \theta)$ and information rent $\Delta(\theta)$. The optimization occurs over the space of almost everywhere (a.e.) differentiable functions $q(\theta)$ and $\Delta(\theta)$ on the domain $[\underline{\theta}, \overline{\theta}]$.

It is known that the principal can restrict herself to contracts that are *incentive feasible*, which refers to two types of constraints needed to ensure optimality given asymmetric information. First *individual rationality* constraints ensure the agents achieve nonnegative utility from participating, i.e., $\Delta(\theta) \geq 0$. The $IR_{\underline{\theta}}$ constraint ensures individual rationality is met for the smallest demand agent $\underline{\theta}$ (the information rent will always be zero for this agent), and $IC_{\theta\theta'}$ ensures information rent is nonnegative for other values of θ. Second, the *incentive compatibility* constraints $IC_{\theta\theta'}$ ensure the agent of type θ prefers the contract designed for him more than those designed for other types $\theta' \leq \theta$, i.e., $v(q(\theta), \theta) - t(\theta) \geq v(q(\theta'), \theta) - t(\theta')$. Finally, it can be shown that $q(\theta)$ should be increasing according to a monotonicity constraint $MON_{\theta\theta'}$, and this combined with the $IC_{\theta\theta'}$ constraint implies that a local incentive compatibility check implies global incentive compatibility.

Suppose problem Φ does not yield an analytical solution, but M samples drawn from $f(\theta)$ are available, either from a simulator or through observations of real demand. Without loss of generality, assume samples $\theta_1, \ldots, \theta_M$ are arranged in increasing order. In this case, we formulate the discrete problem with these M samples using a sample average objective:

$$
\begin{aligned}
\Phi^M = \max_{\{q_m, \Delta_m\}_{m=1,\dots,M}} \quad & \frac{1}{M}\sum_{m=1}^{M}\left[v(q_m,\theta_m) - \Delta_m - s(q_m)\right] && && (2.2) \\
\text{subject to} \quad & \Delta_1 = 0 && && (IR_1) \\
& \Delta_m \geq \Delta_{m-1} + v(q_{m-1},\theta_m) - v(q_{m-1},\theta_{m-1}) && m \in 2,...,M && (IC_{m,m-1}) \\
& q_m \geq q_{m-1} && m \in 2,...,M && (MON_{m,m-1}).
\end{aligned}
$$

Each sample θ_m, $m = 1,\dots,M$, has empirical weight $1/M$. The constraint set is formed from these sampled θ_m values. The size of the decision variable space is now finite ($q_m, \Delta_m, m = 1,\dots,M$), and needs to be no larger than the space of possible demand values according to the revelation principle. We call the discrete formulation (2.2) and its optimal profit solution Φ^M. The discrete problem is often solvable using a nonlinear or mixed-integer nonlinear solver, for example, when $v(q_m,\theta_m) - \Delta_m - s(q_m)$ is concave. In Cai and Singham (2018), a more complex version of this formulation is solved with multiple possible discrete demand distributions.

While taking a large sample M and solving Φ^M may yield a promising approximation to Φ, there are some technical challenges to showing directly that the optimal solution of the sample average problem converges to that of the continuous problem. Unlike in standard sample average approximation methods, the decision variable space increases as M increases. The number and exact form of the constraints also increases and changes with increasing M. Finally, the discrete problem solution yields a sequence of contract points (q_m, Δ_m), and this is not a solution to the continuous problem. The next section presents an approximation method for addressing these challenges.

3 Approximate Problem Formulation

3.1 Large M-problem and bootstrap

Recalling the assumptions on boundedness for the terms in the objective functions of Φ and Φ^M, we note the objective function of Φ^M is a Monte Carlo approximation of the integral objective function in Φ. The error, for a fixed $q(\theta), \Delta(\theta)$, between the objective in Φ and that of

$$\frac{1}{M}\sum_{m=1}^{M} v(q(\theta_m),\theta_m) - \Delta(\theta_m) - s(q(\theta_m)) \tag{3.1}$$

is $\mathcal{O}(1/\sqrt{m})$. The objective profit function (3.1) converges uniformly to the objective function in Φ because of a.e. differentiability and boundedness of the included terms. Establishing concretely the convergence of the optimal objective function value in Φ^M to the optimal value of Φ is outside the scope of this paper and relies on functional law of large numbers results, but we know that a Monte Carlo approximation for Φ using Φ^M for a given solution $q(\theta), \Delta(\theta)$ has error $\mathcal{O}(1/\sqrt{m})$. We fix a large sample size M for which we believe the error in the continuous objective function using a discrete approximation for the distribution is sufficiently small.

While M may be large enough to provide a close approximation to Φ, the discrete problem may not be computable directly if there are too many samples. The numerical results at the end of the paper will give an idea of how large the sample size can be while still having reasonable run times. It is usually much easier to generate the M samples than to solve the corresponding optimization problem Φ^M. We next derive bounds for Φ^M using a bootstrap formulation. Take N empirical samples $\tilde{\theta}_1, \ldots, \tilde{\theta}_N$ with replacement from the M samples to solve problem $\tilde{\Phi}^N$ in (3.2), Again, assume that the samples $\tilde{\theta}_1, \ldots, \tilde{\theta}_N$ are sorted in increasing order. We assume $N < M$ and $\tilde{\Phi}^N$ is solvable numerically. The bootstrap formulation is

$$\begin{aligned}
\tilde{\Phi}^N = \max_{\{\tilde{q}_n,\tilde{\Delta}_n\}_{n=1,\ldots,N}} \quad & \frac{1}{N}\sum_{n=1}^{N}\Big[v(\tilde{q}_n,\tilde{\theta}_n) - \tilde{\Delta}_n - s(\tilde{q}_n)\Big] && && (3.2)\\
\text{subject to} \quad & \tilde{\Delta}_1 = 0 && && (IR_1)\\
& \tilde{\Delta}_n \geq \tilde{\Delta}_{n-1} + v(\tilde{q}_{n-1},\tilde{\theta}_n) - v(\tilde{q}_{n-1},\tilde{\theta}_{n-1}) && \forall n = 2,...,N && (IC_{n,n-1})\\
& \tilde{q}_n \geq \tilde{q}_{n-1} && \forall n = 2,...,N && (MON_{n,n-1}).
\end{aligned}$$

The decision variables $\tilde{q}_n, \tilde{\Delta}_n$ correspond to the bootstrap samples $\tilde{\theta}_n$. We next show how solving $\tilde{\Phi}^N$ yields bounds on Φ and Φ^M.

3.2 Lower bound for Φ

A feasible lower bound for Φ can be constructed using a solution to $\tilde{\Phi}^N$. This lower bound solution, $(\hat{q}(\theta), \hat{\Delta}(\theta))$, can be constructed by using the points of the discrete solution to $\tilde{\Phi}^N$ as a skeleton:

$$\hat{q}(\theta) = \begin{cases} 0, & \underline{\theta} \leq \theta < \tilde{\theta}_1 \\ \tilde{q}_{n-1}, & \tilde{\theta}_{n-1} \leq \theta < \tilde{\theta}_n, \quad n = 2, \ldots, N \\ \tilde{q}_N, & \tilde{\theta}_N \leq \theta < \overline{\theta} \end{cases}$$

$$\hat{\Delta}(\theta) = \begin{cases} 0, & \underline{\theta} \leq \theta < \tilde{\theta}_1 \\ \tilde{\Delta}_{n-1} + v(\tilde{q}_{n-1}, \theta) - v(\tilde{q}_{n-1}, \tilde{\theta}_{n-1}), & \tilde{\theta}_{n-1} \leq \theta < \tilde{\theta}_n, \quad n = 2, \ldots, N \\ \tilde{\Delta}_N + v(\tilde{q}_N, \theta) - v(\tilde{q}_N, \tilde{\theta}_N), & \tilde{\theta}_N \leq \theta < \overline{\theta}. \end{cases}$$

Proposition 3.1. The solution $(\hat{q}(\theta), \hat{\Delta}(\theta))$ is feasible for Φ.

Proof. The solution $(\tilde{q}_n, \tilde{\Delta}_n)_{n=1,\ldots,N}$ is feasible for $\tilde{\Phi}^N$, and $\hat{q}(\theta), \hat{\Delta}(\theta)$ takes the same values on the discrete skeleton $\theta = \tilde{\theta}_n$ as $\tilde{\Phi}^N$. By inspection, we see that $\hat{q}(\theta)$ is nonnegative and monotonically increasing.

It remains to show that $\hat{\Delta}(\theta)$ is increasing and meets the $IC_{\theta\theta'}$ constraint for all $\theta' < \theta$. Recall the assumption that $v(q, \theta)$ is increasing in θ and q. Thus, $\hat{\Delta}(\theta)$ is increasing on the points of the discrete skeleton. $\hat{\Delta}(\theta)$ is also increasing within each interval $[\tilde{\theta}_{n-1}, \tilde{\theta}_n]$, and does not exceed $\tilde{\Delta}_n = \hat{\Delta}(\tilde{\theta}_n)$ because of the $IC_{n,n-1}$ constraint. Thus, $\hat{\Delta}(\theta)$ is increasing.

To show global incentive compatibility, we need to show local incentive compatibility. We show that for any $\theta_a, \theta_b \in [\tilde{\theta}_{n-1}, \tilde{\theta}_n]$ with $\theta_a < \theta_b$, the $IC_{\theta_a\theta_b}$ constraint is met. Write the following:

$$\hat{\Delta}(\theta_a) = \tilde{\Delta}_{n-1} + v(\tilde{q}_{n-1}, \theta_a) - v(\tilde{q}_{n-1}, \tilde{\theta}_{n-1})$$
$$\hat{\Delta}(\theta_b) = \tilde{\Delta}_{n-1} + v(\tilde{q}_{n-1}, \theta_b) - v(\tilde{q}_{n-1}, \tilde{\theta}_{n-1}).$$

To prove incentive compatibility, we need to show

$$\hat{\Delta}(\theta_b) \geq \hat{\Delta}(\theta_a) + v(q(\theta_a), \theta_b) - v(q(\theta_a), \theta_a).$$

Substituting for $\hat{\Delta}(\theta_a)$ and $\hat{\Delta}(\theta_b)$ yields:

$$\tilde{\Delta}_{n-1} + v(\tilde{q}_{n-1}, \theta_b) - v(\tilde{q}_{n-1}, \tilde{\theta}_{n-1}) \geq \tilde{\Delta}_{n-1} + v(\tilde{q}_{n-1}, \theta_a) - v(\tilde{q}_{n-1}, \tilde{\theta}_{n-1}) + v(q(\theta_a), \theta_b) - v(q(\theta_a), \theta_a)$$

$$v(\tilde{q}_{n-1}, \theta_b) \geq v(\tilde{q}_{n-1}, \theta_a) + v(q(\theta_a), \theta_b) - v(q(\theta_a), \theta_a)$$

$$v(\tilde{q}_{n-1}, \theta_b) \geq v(\tilde{q}_{n-1}, \theta_a) + v(\tilde{q}_{n-1}, \theta_b) - v(\tilde{q}_{n-1}, \theta_a)$$

All terms canceling yields the result. □

Thus, the solution to any problem $\tilde{\Phi}^N$ yields a feasible lower bound solution $(\hat{q}(\theta), \hat{\Delta}(\theta))$ for Φ.

3.3 Upper bound in expectation for Φ^M

We use sample average approximation to find an upper bound in expectation for Φ^M. First, we show how the solution to $\tilde{\Phi}^N$ can be used to find a feasible solution to Φ^M. Consider the following solution to $(\hat{q}_m, \hat{\Delta}_m)$ to Φ^M based on the solution to $\tilde{\Phi}^N$:

$$\hat{q}_m = \begin{cases} 0, & \underline{\theta} \leq \theta_m < \tilde{\theta}_1 \\ \tilde{q}_{n-1}, & \tilde{\theta}_{n-1} \leq \theta_m < \tilde{\theta}_n, \quad n = 2, \ldots, N \\ \tilde{q}_N, & \tilde{\theta}_N \leq \theta_m < \overline{\theta}, \end{cases} \tag{3.3}$$

$$\hat{\Delta}_m = \begin{cases} 0, & \underline{\theta} \leq \theta_m < \tilde{\theta}_1 \\ \tilde{\Delta}_{n-1} + v(\tilde{q}_{n-1}, \theta_m) - v(\tilde{q}_{n-1}, \tilde{\theta}_{n-1}), & \tilde{\theta}_{n-1} \leq \theta_m < \tilde{\theta}_n, \quad n = 2, \ldots, N \\ \tilde{\Delta}_N + v(\tilde{q}_N, \theta_m) - v(\tilde{q}_N, \tilde{\theta}_N), & \tilde{\theta}_N \leq \theta_m < \overline{\theta}. \end{cases} \tag{3.4}$$

The solution $\hat{q}_m, \hat{\Delta}_m$ is feasible for Φ^M following similar reasoning to that of Proposition 3.1. Let $m(\tilde{\theta}_n)$ be the index m of the ordered sample $(\theta_1, \ldots, \theta_M)$ that corresponds to the sample $\tilde{\theta}_n$, so $\theta_{m(\tilde{\theta}_n)} = \tilde{\theta}_n$. SAA often refers to optimization using a function that is the sample average approximation of an expectation. The sample average approximation to Φ^M using samples $\tilde{\theta}_n$ can be written as:

$$\tilde{\Phi}^M = \max_{\{q_m, \Delta_m\}_{m=1,\ldots,M}} \quad \frac{1}{N}\sum_{n=1}^{N}\left[v(q_{m(\tilde{\theta}_n)}, \tilde{\theta}_n) - \Delta_{m(\tilde{\theta}_n)} - s(q_{m(\tilde{\theta}_n)})\right] \tag{3.5}$$

$$\text{subject to} \quad \Delta_1 = 0 \qquad (IR_1)$$

$$\Delta_m \geq \Delta_{m-1} + v(q_{m-1}, \theta_m) - v(q_{m-1}, \theta_{m-1}) \quad \forall m = 2, ..., M \quad (IC_{m,m-1})$$

$$q_m \geq q_{m-1} \quad \forall m = 2, ..., M \quad (MON_{m,m-1}).$$

The goal is to show that as the bootstrap samples N increases, the solution to $\tilde{\Phi}^N$ approaches the solution for Φ^M. Formulation $\tilde{\Phi}^M$ will be used to demonstrate how the solution to $\tilde{\Phi}^N$ can be represented as a solution to Φ^M in the next proposition. Thus, the solution to the solvable problem $\tilde{\Phi}^N$ can be used to approximate Φ^M using $\tilde{\Phi}^M$ to place all instances of $\tilde{\Phi}^N$ on the same feasible space as Φ^M.

Proposition 3.2. The solution $(\hat{q}_m, \hat{\Delta}_m)$ based on the solution to $\tilde{\Phi}^N$ is optimal for $\tilde{\Phi}^M$.

Proof. Formulation $\tilde{\Phi}^M$ has the same constraint set as Φ^M, but the objective function only depends on the sampled $\tilde{\theta}_n$ values and the corresponding decision variables. Formulation $\tilde{\Phi}^M$ using $(\hat{q}_m, \hat{\Delta}_m)$ and $\tilde{\Phi}^N$ using $(\tilde{q}_n, \tilde{\Delta}_n)$ have the same objective function value, but the feasible region of $\tilde{\Phi}^M$ is a subset of the feasible region of $\tilde{\Phi}^N$ because the IC and MON constraints must hold for all M original samples rather than just the N bootstrap samples. Because the optimal solution to $\tilde{\Phi}^N$ can be used to construct a feasible solution $(\hat{q}_m, \hat{\Delta}_m)$ to $\tilde{\Phi}^M$ with the same objective function value, $(\hat{q}_m, \hat{\Delta}_m)$ is optimal for $\tilde{\Phi}^M$. □

Proposition 3.2 shows how we can use the solution to $\tilde{\Phi}^N$ to find the solution to $\tilde{\Phi}^M$. This is important because we can use properties of SAA to show that the solution to $\tilde{\Phi}^M$ converges to that of Φ^M as N increases. The next results formalizes this and presents the key idea needed to establish our bounds for Φ^M using $\tilde{\Phi}^N$.

Theorem 3.3. $\Phi^M \leq E[\tilde{\Phi}^N]$ and $\tilde{\Phi}^N \to \Phi^M$ as $N \to \infty$.

Proof. The result follows from standard SAA analysis (see 5.1 of Shapiro et al. (2009) for an overview). First, the discrete problem $\tilde{\Phi}^M$ is a SAA for Φ^M using the same decision variables and constraints, and the objective function of $\tilde{\Phi}^M$ uniformly converges to that of Φ^M, so $\tilde{\Phi}^M \to \Phi^M$

as $N \to \infty$. By Proposition 3.2, solving $\tilde{\Phi}^N$ yields the optimal value of $\tilde{\Phi}^M$. So as $N \to \infty$, $\tilde{\Phi}^N = \tilde{\Phi}^M \to \Phi^M$.

Next, we establish $E[\tilde{\Phi}^N]$ as an upper bound for Φ^M. Let $p(q, \Delta, \theta) = v(q, \theta) - \Delta - s(q)$ be the profit function in the objective summation. Let the feasible space of $\tilde{\Phi}^M$ and Φ^M be Ω^M. To simplify notation in what follows, we omit the full expression of indices in the decision variables so rather than $(q_m, \Delta_m)_{m=1,\ldots,M} \in \Omega^M$ we simply use $q_m, \Delta_m \in \Omega^M$. Define

$$\Phi^M = \max_{q_m, \Delta_m \in \Omega^M} \frac{1}{M} \sum_{m=1}^{M} p(q_m, \Delta_m, \theta_m) = \max_{q_m, \Delta_m \in \Omega^M} E_{\tilde{\theta}_n} \left[\frac{1}{N} \sum_{n=1}^{N} p(q_{m(\tilde{\theta}_n)}, \Delta_{m(\tilde{\theta}_n)}, \tilde{\theta}_n) \right]$$

because $\tilde{\theta}_n$ is sampled directly from θ_m. Note that for any candidate solution $(\bar{q}_m, \bar{\Delta}_m) \in \Omega^M$, we have

$$\frac{1}{N} \sum_{n=1}^{N} p(\bar{q}_{m(\tilde{\theta}_n)}, \bar{\Delta}_{m(\tilde{\theta}_n)}, \tilde{\theta}_n) \leq \max_{q_m, \Delta_m \in \Omega^M} \frac{1}{N} \sum_{n=1}^{N} p(q_{m(\tilde{\theta}_n)}, \Delta_{m(\tilde{\theta}_n)}, \tilde{\theta}_n) = \tilde{\Phi}^M. \tag{3.6}$$

Taking expectations of (3.6) with respect to $\tilde{\theta}_n$ and choosing the candidate solution to maximize the left hand side yields

$$\max_{\bar{q}_m, \bar{\Delta}_m \in \Omega^M} E_{\tilde{\theta}_n} \left[\frac{1}{N} \sum_{n=1}^{N} p(\bar{q}_{m(\tilde{\theta}_n)}, \bar{\Delta}_{m(\tilde{\theta}_n)}, \tilde{\theta}_n) \right] \leq E_{\tilde{\theta}_n}[\tilde{\Phi}^M]. \tag{3.7}$$

Note that the left hand side of (3.7) is Φ^M after changing notation. Also, we have $\tilde{\Phi}^M = \tilde{\Phi}^N$ from Proposition 3.2. Substituting into (3.7) yields $\Phi^M \leq E[\tilde{\Phi}^N]$ which gives the result. □

3.4 Algorithm

The reason for choosing a set of M fixed samples is to establish a fixed space for mapping problem $\tilde{\Phi}^N$ to $\tilde{\Phi}^M$ and Φ^M. Operating over the same space is critical in establishing a fixed constraint region as N increases and showing the convergence of a sequence of smaller SAA $\tilde{\Phi}^N$ problems to Φ^M. By solving the smaller problem $\tilde{\Phi}^N$, we can generate a lower bound for Φ, and an upper bound in expectation for Φ^M.

The error of estimation using a sample average approximation with N samples is $\mathcal{O}(1/\sqrt{n})$. As N increases, $E[\tilde{\Phi}^N] \to \Phi^M$. Because the difference in Φ and Φ^M is $\mathcal{O}(1/\sqrt{m})$ and the error in the sample average estimator is $\mathcal{O}(1/\sqrt{n})$, it is likely that the error in $E[\tilde{\Phi}^N]$ overwhelms the error between Φ^M and Φ when N is relatively small compared to M. Thus, the solution to $\tilde{\Phi}^N$ yields a lower bound and an approximate upper bound in expectation for Φ. We suggest the following algorithm leveraging these results to estimate Φ:

1. Sample M values of θ from $f(\theta)$, sort in increasing order as $\theta_m, m = 1, \ldots, M$.
2. Sample N values of θ_m with replacement ordered increasingly as $\tilde{\theta}_n, n = 1, \ldots N$.
3. Solve $\tilde{\Phi}^N$ using the samples from Step 2.
4. Generate the lower bound solution $(\hat{q}(\theta), \hat{\Delta}(\theta))$ from $\tilde{\Phi}^N$ with objective value $\underline{\Phi}$.
5. Repeat Steps 2-4 multiple times to estimate $E[\tilde{\Phi}^N]$ and collect multiple lower bounds.
6. Calculate the gap between the estimate of $E[\tilde{\Phi}^N]$ and the largest lower bound solution $\underline{\Phi}$. Increase N and repeat Steps 2–5 until the estimated optimality gap is small.

We note that there is much literature on estimating optimality gaps using the sample average approximation method to choose sample sizes N and the number of replications needed to estimate $E[\tilde{\Phi}^N]$ (see, for example, Bayraksan and Morton (2006), Bayraksan and Pierre-Louis (2012)).

4 Numerical Results

We demonstrate the performance of the algorithm with two examples. The first uses a past implementation of the principal-agent problem where an analytical solution to the continuous demand problem has been derived. The second example is a different implementation where the true solution is not known. We construct the formulation using Pyomo (Hart et al. 2011, 2012) and employ the nonlinear solver IPOPT (Wächter and Biegler 2006) for generating solutions. The computing time for solving the optimization problems using a single processor was minimal, ranging from near instantaneous for small N, to less than 10 minutes to solve $\tilde{\Phi}^N$ for N=100,000, including the time to numerically integrate to calculate the lower bound.

4.1 Example 1: A Carbon Capture Example with Known Solution

In the first example, we use an implementation based on past work solving principal-agent models from Singham and Cai (2017). In this context, the agent is a power plant who emits CO_2 into the atmosphere and demands the service of the principal, who is a CO_2 storage operator. The principal offers a menu of options to the agent, with each option corresponding a different demand value θ from a continuous distribution. The agent demands θ units of CO_2 storage based on his anticipated power usage. If the agent does not choose to store any carbon, he will pay a penalty $p(x) = \frac{\alpha}{2}x^2, x \geq 0$ and $p(x) = 0, x < 0$ for emitting x units of CO_2 to the atmosphere. If the agent participates and purchases q units of storage, he has a linear cost γq to capture CO_2 emissions before sending them to storage. We define the value function for the agent as

$$v(q, \theta) = p(\theta) - p(\theta - q) - \gamma q,$$

which is the penalty avoided by participating, minus the penalty actually paid on excess emissions over the amount stored, and minus the capture cost. The principal has a linear cost $s(q) = \beta q$ associated with delivering q units of CO_2 storage. Let the density of θ, $f(\theta)$, be uniform over $[\underline{\theta}, \overline{\theta}]$. The optimal contract offering for the principal was derived in Singham and Cai (2017) and employs a threshold policy. If agent demand is less than θ_c, the principal offers the agent a shutdown/non-participation option, while for demand values larger than θ_c the principal will offer a positive quantity and price to the agent. This threshold is derived as

$$\theta_c = \max\left\{\underline{\theta}, \frac{1}{2}\left(\overline{\theta} + \frac{\gamma + \beta}{\alpha}\right)\right\},$$

and the optimal contract to offer the agent is

$$q^*(\theta) = \begin{cases} 0 & \text{if } \theta < \theta_c \\ 2\theta - \overline{\theta} - \frac{\gamma+\beta}{\alpha} & \text{if } \theta \geq \theta_c, \end{cases} \tag{4.1}$$

$$\Delta^*(\theta) = \begin{cases} 0 & \text{if } \theta < \theta_c \\ \alpha(\theta^2 - \theta_c^2) - (\alpha\overline{\theta} + \gamma + \beta)(\theta - \theta_c) & \text{if } \theta \geq \theta_c. \end{cases} \tag{4.2}$$

Thus when $\theta > \theta_c$, the quantity offered the agent is linear with respect to θ and the information rent is quadratic with respect to θ. For $\theta < \theta_c$, the principal chooses not to service the agent because the demand is too low. Relying on our past research on pricing carbon capture and storage systems, (Cai et al. 2014, Singham et al. 2015), the model parameters are set as follows. The range of values for θ is $[\underline{\theta}, \overline{\theta}] = [0.3, 1]$ Megatonnes of CO_2 based on a sample of publicly available power plant emissions data. The capture cost for the agent is γ =\$45/tonne, and the storage cost for the principal is β =\$13/tonne. Assuming a carbon emissions tax of approximately \$100/tonne when $\theta = 0.7$ Megatonnes, the penalty function is calibrated so that $\alpha = 2.86 \times 10^{-4}$. The end result is that the optimal expected profit to the principal using (4.1) and (4.2) is \$17.25 million.

We can use this known optimal result to assess the quality of the algorithm proposed in Section 3.4. Table 1 shows the results for this example. Let the total discrete sample size be M=1,000,000. The problem $\tilde{\Phi}^N$ is solved using 100 replications for each tested value of the bootstrap sample size N. The major improvement over the results in Singham and Cai (2017) is that the greatest lower bound from the 100 replications is now reported so the optimal profit is now clearly bounded on both sides. As N increases, we see the bounds converging around the true optimal solution. The intuition for the SAA solution being an upper bound is that the principal is assuming less variability in demand using the sampled discrete distribution rather than the real continuous distribution, and hence overestimates her expected profit.

Table 1: Uniform demand distribution results for the CO_2 example using M=1 million. For each sample size N we solve and collect 100 replications from different bootstrap samples. The best lower bound, the expected upper bound, and the standard error of the upper bound estimates are presented.

N	Lower Bound	Upper Bound ($E[\tilde{\Phi}_N]$)	Std Err of $\tilde{\Phi}_N$
10	16.277	24.911	6.776
100	16.997	18.750	1.976
1,000	17.211	17.538	0.610
10,000	17.237	17.291	0.161
100,000	17.248	17.260	0.059
∞	17.250		—

The bounds are relatively tight for N=100,000. Recall that for small enough N the error in the bounds will outweigh the error in Φ^M with respect to Φ. If N becomes large, the bounds will likely converge around Φ^M which will differ from Φ depending on the size of M. Given good bounds on the optimal expected profit, we next turn to the quality of the contract solutions delivered to estimate q and Δ.

Figure 1 plots the optimal contract values computed for a single replication of the experiment

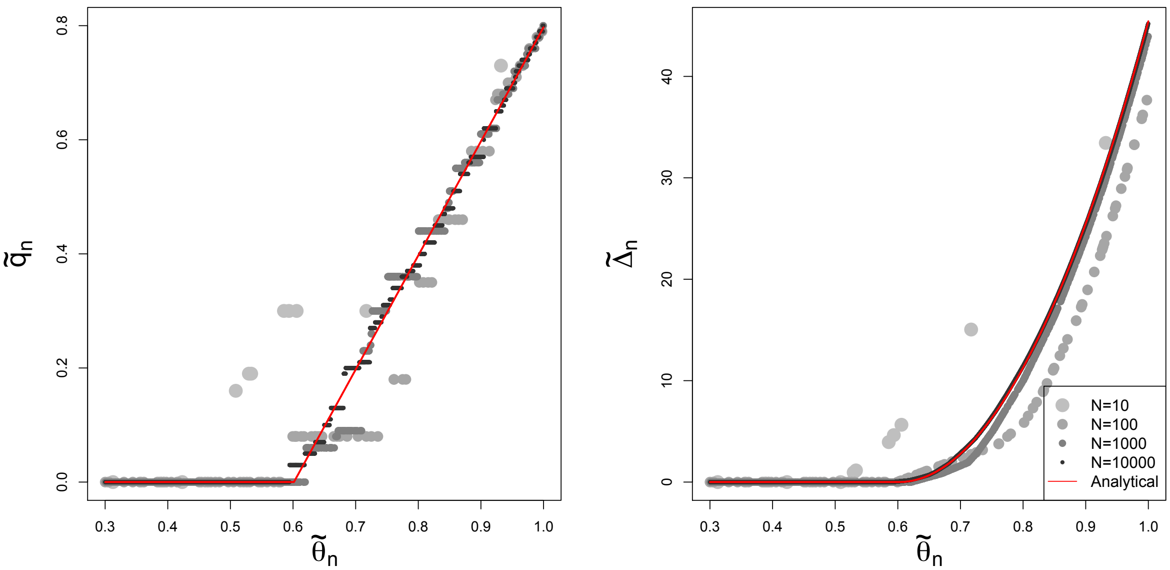

Figure 1: Optimal contract values $(\tilde{q}_N, \tilde{\Delta}_N)$ calculated for Example 1 (uniform distribution for θ) using single replications of different values of N. The red line is the analytical solution to the continuous problem.

for each value of N. The code produces optimal values of $\tilde{q}_N$ and $\tilde{\Delta}_N$ given a bootstrap sample $\tilde{\theta}_n, n = 1, \ldots, N$. We observe that as N increases, the contract values appear to "converge" to the known optimal solution for the continuous formulation, though as mentioned in Section 1 establishing functional convergence is difficult due to the discrete problem operating on a different space from the continuous problem. However, the discrete solution can be interpolated to see the similarity in the solution structure to the continuous problem.

The left plot of Figure 1 displays the optimal quantity functions for the discrete problem for different values of N, and the solution approaches the piecewise linear solution (red line). The right plot shows the optimal information rents, the solutions similarly approach the piecewise quadratic solution. In addition to the optimal expected profit bounds converging as N increases, this example demonstrates that the optimal contracts offered from the bootstrap may resemble the true optimal in structure. Thus, the solution of a single problem $\tilde{\Phi}^N$ can provide insight into the nature of the solution to the continuous problem.

4.2 Example 2: Triangular distribution

We can demonstrate algorithm performance for a more complicated formulation where analytical solutions may not be readily obtainable. The goal is to demonstrate that the solution can converge for more complicated functions. Consider the following arbitrary value and cost functions

$$v(q,\theta) = \alpha q(1 - e^{-\eta\theta} - q) \qquad s(q) = \beta q^2 - \gamma q.$$

Furthermore, let the demand density of θ be $f(\theta) = Tri(\underline{\theta}, \hat{\theta}, \overline{\theta})$, which is the triangular density function with minimum $\underline{\theta}$, mode $\hat{\theta}$, and maximum $\overline{\theta}$. We arbitrarily set the coefficients to be $\alpha = 200, \eta = 2, \beta = 10, \gamma = 3$. For the triangular distribution, we use $\underline{\theta} = 0.3, \hat{\theta} = 0.8, \overline{\theta} = 1$. Table 2 displays the experimental results varying the value of N. We again see as N increases the lower bound increases and the upper bound estimate decreases, with the variation in the upper bound estimate decreasing.

Table 2: Alternative example using a triangular density function. The best lower bound, the expected upper bound, and the standard error of the upper bound estimates are presented.

N	Lower Bound	Upper Bound ($E[\tilde{\Phi}_N]$)	Std Err of $\tilde{\Phi}_N$
10	18.720	23.240	2.251
100	19.038	19.780	0.884
1,000	19.087	19.269	0.243
10,000	19.107	19.136	0.084
100,000	19.112	19.127	0.026

Figure 2 displays the solutions for the optimal quantity and information rent menus for random realizations of $\tilde{\Phi}^N$. While we do not know the exact structure of the optimal solution, we observe as N increases the solutions may approach a smooth function of θ, implying that the underlying optimal contract has some structure. This example again lends hope that a single solution of the bootstrap problem for large enough N and M may allow us to closely bound the optimal profit and deliver an optimal solution that can be interpolated to be similar in shape to the true optimum.

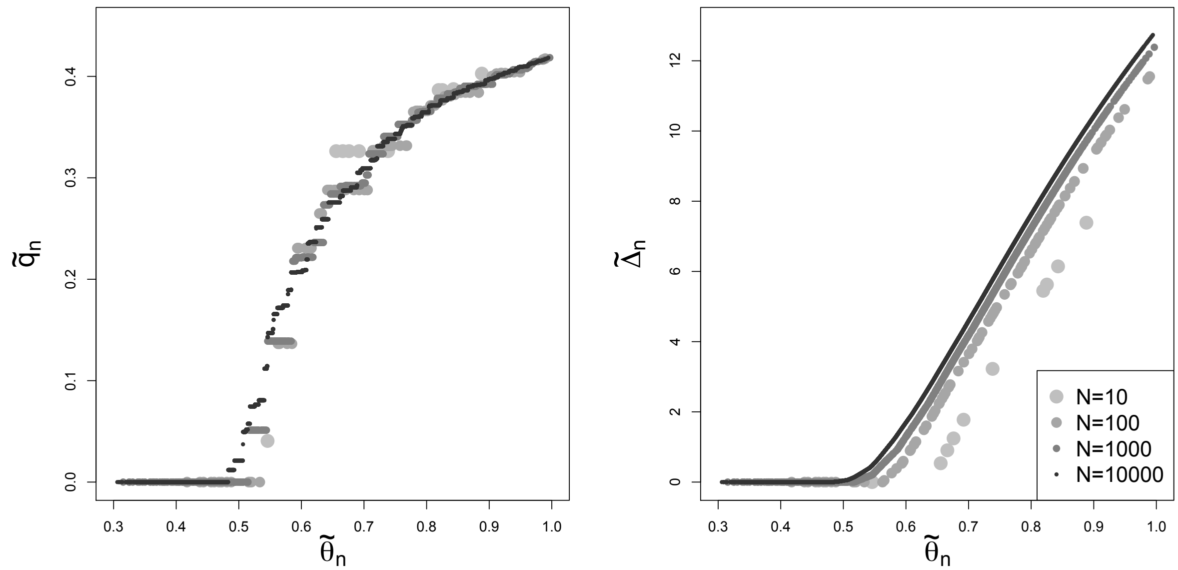

Figure 2: Optimal contract values $(\tilde{q}_N, \tilde{\Delta}_N)$ calculated for the triangular distribution example.

5 Conclusion

We present a method for computing approximate solutions to the continuous principal-agent problem when analytical solutions are not available. We rely on solutions to smaller discrete problems that are easily solvable computationally to provide bounds for the solution to the continuous problem. These discrete problems provide lower bound solutions to the continuous problem, and upper bounds in expectation to a large discrete problem that closely approximates the continuous problem. In addition to bounding the objective function value, the solutions from the discrete problem can be used to provide intuition about the structure of the continuous solution.

This method can also be used for solving the discrete agent-type problem when the number of agent types is too large to be solved computationally. We quantify the error associated with the approximations, and present numerical results that illustrate how the bounds improve as the sample size used in the approximate problem is increased. In the case where $f(\theta)$ is not available but data samples are present, then the approximate problem becomes the actual problem to be solved. Thus, this method can be used for analysis of continuous type principal-agent problems that previously were intractable.

Acknowledgements

We acknowledge the support of National Science Foundation grants CMMI-1535762 and CMMI-1535831 for this research, and many useful discussions with Wenbo Cai.

References

Babich, Volodymyr, Hantao Li, Peter Ritchken, Yunzeng Wang. 2012. Contracting with Asymmetric Demand Information in Supply Chains. *European Journal of Operational Research* **217**(2) 333–341.

Bayraksan, Güzin, David P Morton. 2006. Assessing Solution Quality in Stochastic Programs. *Mathematical Programming* **108**(2-3) 495–514.

Bayraksan, Güzin, Péguy Pierre-Louis. 2012. Fixed-Width Sequential Stopping Rules for a Class of Stochastic Programs. *SIAM Journal on Optimization* **22**(4) 1518–1548.

Bertsimas, Dimitris, Aurélie Thiele. 2005. A Data-Driven Approach to Newsvendor Problems. Tech. rep., Massachusetts Institute of Technology, Cambridge, MA.

Burnetas, Apostolos, Stephen M Gilbert, Craig E Smith. 2007. Quantity Discounts in Single-Period Supply Contracts with Asymmetric Demand Information. *IIE Transactions* **39**(5) 465–479.

Cai, W., D.I. Singham. 2018. Principal-Agent Problems with Heterogeneous Agent Demand Distributions. *European Journal of Operational Research* **264**(1) 239–256.

Cai, W., D.I. Singham, E.M. Craparo, J.A. White. 2014. Pricing Contracts Under Uncertainty in a Carbon Capture and Storage Framework. *Energy Economics* **34**(1) 56–62.

Cecchini, Mark, Joseph Ecker, Michael Kupferschmid, Robert Leitch. 2013. Solving Nonlinear Principal-Agent Problems using Bilevel Programming. *European Journal of Operational Research* **230**(2) 364–373.

Dempe, S. 1995. Computing Optimal Incentives via Bilevel Programming. *Optimization* **33**(1) 29–42.

Hart, W.E, C. Laird, J-P. Watson, D.L Woodruff. 2012. *Pyomo - Optimization Modeling in Python*, vol. 67. Springer Science & Business Media.

Hart, W.E, J-P. Watson, D.L. Woodruff. 2011. Pyomo: Modeling and Solving Mathematical Programs in Python. *Mathematical Programming Computation* **3**(3) 219–260.

Kung, Ling-Chieh, Ying-Ju Chen. 2012. The Resource Allocation Problem for Demand Forecasting and Performance Measurement. *Operations Research Letters* **40**(1) 61–68.

Laffont, Jean-Jacques, David Martimort. 2009. *The Theory of Incentives: The Principal-Agent Model*. Princeton University Press.

Lambert, Richard A. 2006. Agency Theory and Management Accounting. *Handbooks of Management Accounting Research* **1** 247–268.

Maskin, E., J. Riley. 1984. Monopoly with Incomplete Information. *RAND Journal of Economics* **15**(2) 171–196.

Shapiro, Alexander, Darinka Dentcheva, Andrzej Ruszczyński. 2009. *Lectures on Stochastic Programming: Modeling and Theory*. SIAM.

Singham, D.I., W. Cai. 2017. Sample Average Approximations for the Continuous Type Principal-Agent Problem: An Example. W. K. V. Chan, A. D'Ambrogio, G. Zacharewicz, N. Mustafee, G. Wainer, E. Page, eds., *Proceedings of the 2017 Winter Simulation Conference.*. Institute of Electrical and Electronics Engineers, Inc., Piscataway, New Jersey.

Singham, D.I., W. Cai, J.A. White. 2015. Optimal Carbon Capture and Storage Contracts using Historical CO_2 Emissions Levels. *Energy Systems* **6**(3) 331–360.

Wächter, A., L.T. Biegler. 2006. On the Implementation of an Interior-Point Filter Line-Search Algorithm for Large-Scale Nonlinear Programming. *Mathematical Programming* **106**(1) 25–57.