Proceedings of the 2017 Winter Simulation Conference
W. K. V. Chan, A. D'Ambrogio, G. Zacharewicz, N. Mustafee, G. Wainer, and E. Page, eds.

SAMPLE AVERAGE APPROXIMATIONS FOR THE CONTINUOUS TYPE PRINCIPAL-AGENT PROBLEM: AN EXAMPLE

Dashi I. Singham

Operations Research Department
Naval Postgraduate School
1411 Cunningham Road
Monterey, CA 93943, USA

Wenbo Cai

Department of Mechanical and Industrial Engineering
New Jersey Institute of Technology
200 Central Ave
Newark, NJ 07102, USA

ABSTRACT

Principal-agent problems study contracts for goods or services that a principal (seller) should offer an agent (buyer). The goal is for the principal to optimize the quantity and price in the contract offered to an agent with uncertain demand, where the principal has estimated a distribution for the agent's demand. The agent's demand distribution can be discrete or continuous. A deterministic optimization solution to the discrete distribution problem delivers a contract with price and quantity options targeted towards each possible demand realization. When the demand distribution is continuous, the optimal contract becomes a continuous function of the demand space. This paper introduces a sample average approximation to the continuous distribution problem using methods for solving the discrete distribution problem. We explore using numerical results an example motivated by carbon capture and storage systems.

1 INTRODUCTION

Principal-agent (PA) problems find the optimal contract (quantity and price for goods or services) for a principal (seller) to offer an agent (buyer). The contract consists of one or more options, and each option stipulates the quantity q of the product to be transferred, and a total price t that the agent pays to receive q. The agent may have hidden preferences that affect their demand for the product, or may have demand uncertainty such that their demand is not realized until after the contract has been set. While the principal may not know a particular agent's demand, she may know their demand distribution. The principal determines the optimal contract options by maximizing her expected profit over different potential realizations of the agent's demand. Once the agent realizes his actual demand level, he makes a utility-maximizing decision whether or not to participate and purchase q units of the product at price t.

The original PA problems in the literature are solved analytically using deterministic optimization (Maskin and Riley 1984). Much work has been done to derive the form of the optimal contracts analytically for specific problem settings, such as supply chain coordination between retailers and manufacturers (Tsay 1999, Iyer, Schwarz, and Zenios 2005), procurement strategy (Cachon and Zhang 2006, Chaturvedi and Martínez-de-Albéniz 2011), production planning (Yang et al. 2009, Norde et al. 2016), and performance-based services (Kim et al. 2010, Huber and Spinler 2013). The current literature on discrete-type problems is mostly dominated by the analysis of symmetric agents who have the same likelihood of being either of two types of agents. The analytical solutions provided by Laffont and Martimort (2009) are widely cited in the service contracting literature (Özer and Wei 2006, Hasija, Pinker, and Shumsky 2008, Akan, Ata, and Lariviere 2011).

Analytical solutions yield managerial insights that explain how different factors and parameters affect the optimal contract. However, simplifying assumptions and small problem sizes are often used to obtain tractability. This means that the demand distribution of agents is often assumed to be discrete with very few possible values. The paper by Cai and Singham (2017) presents one of the first attempts to consider heterogenous agent distributions (with multiple possible demand distributions). That paper derives analytical solutions for the case when there are two possible discrete demand distributions, each with two possible demand levels, and solves larger problems numerically using a nonlinear solver. One of our goals is to be able to solve larger problems numerically while understanding the structure behind the solutions. Lovejoy (2006) makes a significant advancement by considering a finite number of agent types and each type takes on a different set of discrete values. Analyzing the structural properties of quasi-linear utility functions, Lovejoy (2006) is able to reduce the dimension of the problem under certain conditions to obtain tractability.

Let demand θ have a discrete distribution with N possible levels so $\theta_n, n = 1, \ldots, N$. The optimal **contract** can consist of multiple **options** $(q_n, t_n), n = 1, \ldots, N$ where each option is a quantity and price combination geared towards a particular demand level θ_n. The principal chooses these options to maximize her expected profit given the uncertainty in demand. The agent may select the option (q_n, t_n) from the contract that maximizes his utility once his actual demand level is realized. Alternatively, let θ be a continuous random variable. Then, the optimal contract consists of the functions $(q(\theta), t(\theta))$, where the quantity and price are functions of θ. The continuous demand principal agent problem has been studied in Mirrlees (1999). It is often difficult to find an analytical or numerical solution for the continuous distribution case, and this paper proposes a sample-average approximation solution to help solve this problem. We refer to these two distributional settings as the “discrete problem” and the “continuous problem.”

The continuous demand problem presented here attempts to maximize a concave objective profit function which is an expectation over one continuous nonnegative random variable. We propose a sample-average approximation (SAA) approach. We sample N_0 values of θ from the continuous distribution, and formulate a deterministic optimization problem using a discrete equally weighted distribution over the sampled values of θ. The solution to the discrete problem yields a set of options $(q_n, t_n), n = 1, \ldots, N_0$ which “approximates” the continuous distribution solution $(q(\theta), t(\theta))$, and we will explore the nature of this approximation.

We study a simple implementation of the continuous problem that yields an analytical solution, allowing us to analyze the quality of the discrete problem approximation. Studying this example helps support ongoing work to justify use of sample average approximation to solve general instances of the continuous PA problem. The solution to the continuous problem is a function, and we describe the challenges associated with applying functional SAA results to this problem.

2 THE PRINCIPAL-AGENT (PA) PROBLEM

This section describes the principal-agent problem. First, suppose that the agent’s demand θ for a product is deterministic, and that the principal offers the agent an option (q, t) which sells q units of the product for total price t. The principal has a cost function $s(q)$ and will obtain a profit $t - s(q)$ if the agent participates. The value function to the agent receiving quantity q of the product is $v(q, \theta)$, and we assume this value function is concave and increasing in q, for $q \le \theta$. The utility to the agent from participating in the contract is $v(q, \theta) - t$. The principal then has the following optimization problem:

$$(\bar{q}, \bar{t}) = \begin{array}{cc} \arg\max_{(q,t)} & t - s(q) \\ \text{s.t.} & v(q, \theta) - t \ge 0, \end{array}$$

where the constraint is imposed to ensure the agent participates. Call $\bar{q}$ the “efficient quantity,” and $\bar{t}$ the “full price” that would be offered if there was no uncertainty for an agent with demand θ. Typically, the prices and quantities offered to agents under uncertainty in θ will be smaller than the full price and efficient quantity, respectively. We next describe the discrete and continuous problem when θ has a distribution.

2.1 The Discrete Problem

In the discrete problem, the hidden demand of the agent can take values $\theta_n, n = 1, \ldots N$, where θ_n is ordered by increasing value. The probability of each demand level θ_n is $\pi_n, n = 1, \ldots, N$, with $\sum_n \pi_n = 1$. As described in the introduction, the principal's goal is to design a contract of multiple options $(q_n,t_n), n = 1,\dots,N$, with each option geared towards a particular demand level θ_n. The principal's optimization problem can be written as

$$\begin{array}{llll} \max\limits_{\{q_n,t_n\}} & \Phi = \sum_{n=1}^{N} \pi_n (t_n - s(q_n)) & & \\ \text{s.t.} & v(q_n,\theta_n) - t_n \geq 0 & n \in 1,\dots,N & (IR_n) \\ & v(q_n,\theta_n) - t_n \geq v(q_{n'},\theta_n) - t_{n'} & n,n' \in 1,\dots,N, n \neq n' & (IC_{nn'}) \\ & q_n \geq 0, t_n \geq 0 & n \in 1,\dots,N & (NN_n). \end{array} \tag{1}$$

The objective is to choose the contract options offered towards the agent's demand levels such that the expected profit is maximized. In order to ensure that the objective function is valid in using demand probabilities π_n, the constraints ensure that agents select the contract that is designed for them. Each IR_n constraint is the *individual rationality* constraint that ensures the agent receives nonnegative utility from choosing the option that is designed for his realized demand level. The $IC_{nn'}$ *incentive compatibility* constraints ensure that if an agent has demand θ_n, he will not benefit by choosing the option designed for $\theta_{n'}$. Finally, the NN_n constraints ensure nonnegativity in the contract values.

We note that the solution to (1) can often be easily obtained using a nonlinear solver, for example when $v(q,\theta)$ is concave and $s(q)$ is linear in q. Formulation (1) can be re-written with a concave objective and linear constraints by using the information rents $\Delta_n = v(q_n,\theta_n) - t_n$ as decision variables instead of t_n. The information rent is the decrease in price relative to the full price $\bar{t}_n$ (as $\bar{t}_n = v(q_n,\theta_n)$ in the deterministic case). Moreover, let $\rho(\theta_n,\theta_{n'},q_{n'}) = v(q_{n'},\theta_n) - v(q_{n'},\theta_{n'})$ represent the difference in utility between agents with demand levels θ_n and $\theta_{n'}$ when selecting quantity $q_{n'}$. This reformulation is presented as:

$$\begin{array}{llll} \max\limits_{\{q_n,\Delta_n\}} \Phi = & \sum\limits_{n=1}^{N} \pi_n \left[v(q_n,\theta_n) - \Delta_n - s(q_n) \right] & & \\ \text{s.t.} & \Delta_n \geq 0 & n \in 1,\dots,N & (IR_n) \\ & \Delta_n \geq \rho(\theta_n,\theta_{n'},q_{n'}) + \Delta_{n'} & n,n' \in 1,\dots,N, n \neq n' & (IC_{nn'}) \\ & q_n \geq 0 & n \in 1,\dots,N & (NN_n). \end{array} \tag{2}$$

Formulation (2) has constraints that are linear in the decision variables q_n and Δ_n when $\rho(\theta_n,\theta_{n'},q_{n'})$ is linear in $q_{n'}$, and the objective function is concave when $v(q_n,\theta_n)$ is concave and increasing in q and $s(q_n)$ is linear and increasing in q_n.

2.2 The Continuous Problem

When θ is a bounded nonnegative continuous random variable with density $f(\theta)$, the optimal contract is a function of θ and takes the form $(q(\theta),t(\theta))$ over the range of θ, $[\underline{\theta},\overline{\theta}]$. Let the information rent be defined as $\Delta(\theta) \equiv v(q(\theta),\theta) - t(\theta)$, then the principal's problem can be written as

$$\begin{array}{llll} \max\limits_{\{q(\theta),\Delta(\theta)\}} \Phi = & \int_{\underline{\theta}}^{\overline{\theta}} \left[v(q(\theta),\theta) - \Delta(\theta) - s(q(\theta)) \right] f(\theta) d\theta & & \\ \text{s.t.} & \Delta(\theta) \geq 0 & \theta \in [\underline{\theta},\overline{\theta}] & (IR_\theta) \\ & \Delta(\theta) \geq \rho(\theta,\theta',q(\theta')) + \Delta(\theta') & \theta,\theta' \in [\underline{\theta},\overline{\theta}], \theta' \neq \theta & (IC_{\theta\theta'}) \\ & q(\theta) \geq 0 & \theta \in [\underline{\theta},\overline{\theta}] & (NN_\theta). \end{array} \tag{3}$$

It can be quite difficult to obtain analytical solutions for $q(\theta),\Delta(\theta)$. Appendix A derives a partial analytical solution for (3).

3 EXAMPLE: CARBON CAPTURE AND STORAGE

We present an example of a PA problem. Our context is motivated by carbon capture and storage (CCS) systems. CCS is the process of capturing CO_2 from an emissions source, such as a power plant, and transporting it to a facility where it can be liquified and injected underground for storage. The idea is to prevent, at relatively low cost, excess CO_2 emissions from entering the atmosphere. CCS, in conjunction with renewable energy and other alternatives, can form a portfolio for reducing emissions. Klokk et al. (2010) and Middleton et al. (2012) describe alternative models for structuring CCS systems between emissions sources and storage operators.

We study a pay-at-the-gate model, where power plants pay an external agency to transport and store CO_2, as opposed to conducting the storage themselves (Esposito, Monroe, and Friedman 2011). The principal is the storage operator, who wants to maximize her profits from charging the agents (power plants/emissions sources) to store excess CO_2 over the allowable emissions limit. The agent has a random demand θ for CCS services based on his CO_2 emissions. If the agent selects an option with quantity q and price t, he will pay a price t to transport at most q units (Megatonnes of CO_2).

A carbon tax penalty is required to provide an incentive for power plants to participate in CCS. We assume a convex and increasing penalty function $p(x), x \geq 0$ applied where x is the emissions level above the allowable limit, and $p(x) = 0$ for all $x < 0$. Additionally, the emitter faces a cost of capturing $c(q)$ in preparation for delivering q units to the principal. The emitter's value function for a contract with quantity q when his demand is θ is defined as

$$v(q,\theta) = p(\theta) - p(\theta - q) - c(q) \qquad \text{for} \quad q \leq \theta. \tag{4}$$

We can assume $q \leq \theta$ because $v(q,\theta) = p(\theta) - c(q)$ decreases in q when $q > \theta$, so the emitter will never choose a q that is greater than θ. The first term on the right hand side of (4) is the foregone penalty from emitting θ units. The second term is the tax penalty paid on the excess emissions above the amount to be transported, q. The last term is the cost of capturing q units. In PA problems, the value function is required to meet a single-crossing property, which means that the derivative of $v(q,\theta)$ with respect to q increases as θ increases so that an agent with higher θ will have a higher valuation for larger values of q. When $p(x)$ is increasing and convex in x and $c(q)$ is linear and increasing in q, the single-crossing property is met. Finally, PA problems assume monotonicity in q_n such that $q_n \leq q_{n'}$ for $n < n'$.

We employ generic functions to enable analytical solutions and numerical computations. The penalty function used is $p(x) = \frac{\alpha}{2}x^2$, for $x \geq 0$ with $\alpha > 0$. The cost functions are linear in q: $c(q) = \gamma q$, $\gamma > 0$, and $s(q) = \beta q$, $\beta > 0$, where γ is the per unit cost of capturing CO_2 and β is the per unit cost of transporting and storing CO_2. Using the monotonicity property for q_n, formulation (2) can be reduced to

$$\begin{aligned}
\max_{\{q_n,\Delta_n\}} \ \Phi = \quad & \sum_{n=1}^{N} \pi_n \left[v(q_n,\theta_n) - \Delta_n - \beta q_n\right] && \\
\text{s.t.} \quad & \Delta_1 \geq 0 && (IR_1) \\
& \Delta_n \geq \alpha(\theta_n - \theta_{n'})q_{n'} + \Delta_{n'} \qquad \forall n = 2,...,N, n' = n-1 && (IC_{nn'}) \\
& q_n \geq q_{n'} \qquad \forall n = 2,...,N, n' = n-1 && (MON_{nn'}).
\end{aligned} \tag{5}$$

We note that Δ_n is also monotonic and increasing. Because (5) has far fewer constraints than (2), we perform computations using (5) to find solutions to the discrete problem. Next, we can use our specific forms of $v(q,\theta), s(q), c(q)$ and $\rho(\theta_n,\theta_{n'},q_{n'})$ in the continuous case. Appendix B shows that the optimal quantity and information rent functions are:

$$q^*(\theta) = \theta + \frac{\gamma+\beta}{\alpha} - \frac{\overline{F}(\theta)}{f(\theta)}, \qquad \Delta^*(\theta) = \int_{\underline{\theta}}^{\theta} \alpha q^*(\vartheta) d\vartheta.$$

Additionally, if we assume that the density of θ is uniform on the range $[\underline{\theta},\overline{\theta}]$, then we can derive the threshold that determines whether the agent participates as:

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

where if $\theta \geq \theta_c$ the agent will purchase CCS services from the principal, otherwise he will not participate. Appendix B derives the optimal quantity and information rent functions as

$$q^*(\theta) = \begin{cases} 0 & \text{if } \theta < \theta_c \\ 2\theta - \overline{\theta} - \frac{\gamma+\beta}{\alpha} & \text{if } \theta \geq \theta_c, \end{cases} \tag{6}$$

$$\Delta^*(\theta) = \begin{cases} 0 & \text{if } \theta < \theta_c \\ \alpha(\theta^2 - \theta_c^2) - (\alpha\overline{\theta} + \gamma + \beta)(\theta - \theta_c) & \text{if } \theta \geq \theta_c. \end{cases} \tag{7}$$

4 SAMPLE AVERAGE APPROXIMATION

We propose using sample average approximation (SAA) to solve the PA problem when θ has a continuous distribution. As Section 2.2 suggests, it is difficult to derive the optimal contract function $(q(\theta),t(\theta))$ for a general form of the profit and cost functions, though it is possible for some specific situations as shown in (6) and (7). However, for a discrete distribution for θ, the optimal solution often can be readily calculated numerically using standard nonlinear optimization packages. We explore the ability of the discrete problem to approximate the continuous problem using sample average approximation.

We formulate an approximation to the continuous problem using the discrete problem with N_0 sampled values of θ_n from the continuous density $f(\theta)$. These sampled values of θ_n form a discrete empirical distribution with each sample having equal weight $1/N_0$. Then, we can solve the following problem with a sample average approximation in the objective:

$$\begin{array}{lll} \max\limits_{\{q_n,\Delta_n\}} \Phi = & \frac{1}{N_0}\sum\limits_{n=1}^{N_0}[v(q_n,\theta_n) - \Delta_n - \beta q_n] & \\ \text{s.t.} & \Delta_1 \geq 0 & (IR_1) \\ & \Delta_n \geq \alpha(\theta_n - \theta_{n'})q_{n'} + \Delta_{n'} \quad \forall n = 2,...,N_0, n' = n-1 & (IC_{nn'}) \\ & q_n \geq q_{n'} \quad \forall n = 2,...,N_0, n' = n-1 & (MON_{nn'}). \end{array} \tag{8}$$

We abuse notation and call the solutions (θ_n,q_n^*) and (θ_n,Δ_n^*) "function" approximations, and will observe in the numerical results below that as N_0 increases, these "functions" more closely approximate the functional solutions $(\theta,q^*(\theta))$ and $(\theta,\Delta^*(\theta))$ from the continuous case as derived in (6) and (7). Continuous functions can be approximated using piecewise linear interpolations between the discrete points in (θ_n,q_n^*) and (θ_n,Δ_n^*). However, the optimal quantity (θ_n,q_n^*) takes a step function form in the discrete case, with the steps increasing as θ increases. As shown by Maskin and Riley (1984), there exists a set of subintervals where the optimal quantity function is constant. This step function solution shows pooling, in which consecutive values of θ are assigned the same q, but have different information rents. To construct continuous functions based on the discrete solutions, we note that a semi-continuous feasible solution exists so that between demands θ_n and θ_{n+1}, the principal can offer either q_n^* or q_{n+1}^* (we omit the details for brevity).

We explore the behavior of the SAA estimator (8) when the demand distribution $f(\theta)$ is $Unif[\underline{\theta},\overline{\theta}]$ to compare the numerical results to the known solution. Emissions data is publicly available and we use data from eight plants in Illinois. Applying proposed 2015 CO_2 limits by the Environmental Protection Agency, we calculate the amount of emissions per month that would need to be stored using the historical data sets for the years 2000-2012. A number of distributions can be estimated for the storage demand, but for the purposes of this paper we specify a uniform distribution between 0.3 and 1 Megatonnes which encompasses the range of values across the eight power plants.

From past research on CCS systems (Cai et al. 2014, Singham et al. 2015) we estimate the cost and penalty parameters in the model. We calculate the cost to transport and store CO_2 as $\beta =$\$13/tonne, and the cost to the emitter of capturing CO_2 as $\gamma =$\$45/tonne. We can estimate a carbon tax of approximately \$100/tonne, and calibrate α by estimating the per unit tax to be \$100 when $\theta = 0.7$ Mt (Megatonnes). Solving using the total tax function $\frac{\alpha}{2}\theta^2$ yields $\alpha = 2.86 \times 10^{-4}$.

We model formulations (2), (5), and (8) using Pyomo (Hart et al. 2011, Hart et al. 2012) and apply the nonlinear solver IPOPT (Wächter and Biegler 2006) to obtain all numerical results. For various values of N_0, we take samples θ_n from $Unif(0.3, 1)$ and use those demand values with $\pi_n = 1/N_0$ to solve the discrete problem (8). This yields contract option solution "functions" (θ_n, q_n^*) and (θ_n, Δ_n^*). Figure 1 shows the discrete problem solution for $N_0 = 100$ and $N_0 =$1,000. The red solid line denotes the analytical solution to the continuous problem. Figure 2 shows the discrete solutions for increased sample sizes of $N_0 =$10,000 and $N_0 =$100,000. As we increase N_0, we observe (θ_n, q_n^*) and (θ_n, Δ_n^*) more closely approximating $(\theta, q^*(\theta))$ and $(\theta, \Delta^*(\theta))$, which are (6) and (7).

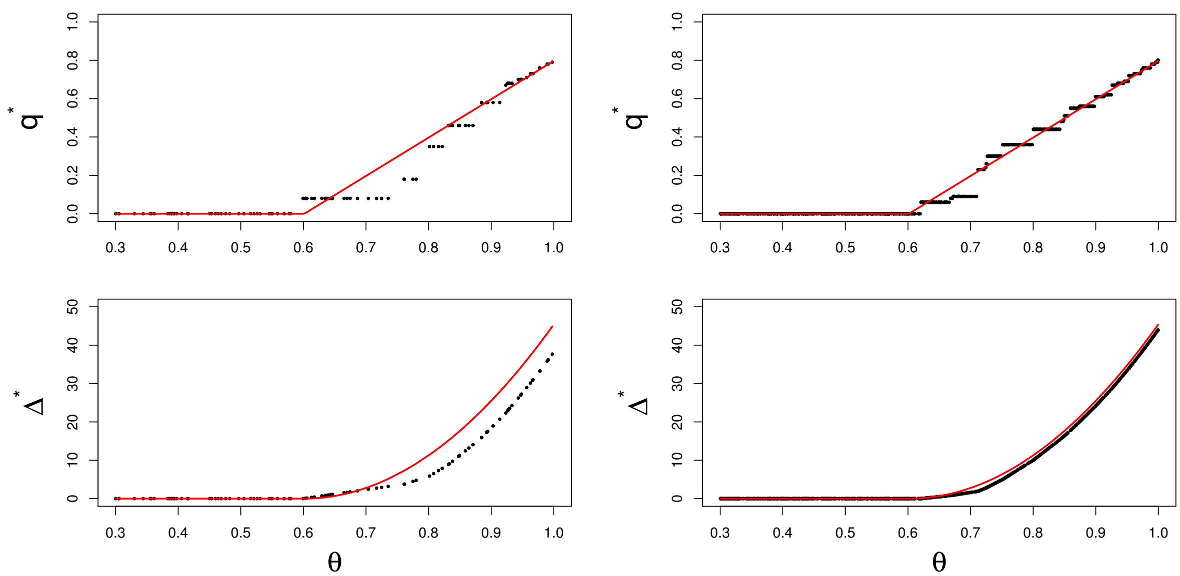

Figure 1: Discrete problem solution with $N_0 =$100 and $N_0 =$1,000. The red line shows the analytical results (6) and (7) for the continuous problem.

For large enough N_0, the discrete solution begins to resemble the piecewise linear solution from the continuous case. In the case of (θ_n, Δ_n^*), an approximate quadratic function is obtained for small sample sizes, but the estimate appears to be biased and converges from below to the optimal solution as N_0 increases. This means that the information rent will be smaller than optimal due to the perceived certainty in θ in the discrete case relative to the true density $f(\theta)$, so the prices offered will be higher than they should be in the continuous case.

There are many issues to consider before we can begin to establish the consistency of the SAA estimator (8). There exist some results showing the consistency of SAA estimators when the decision variables of the problem lie in a function space. One such result is Proposition 3.3 of Royset and Wets (2017), which shows the epi-convergence of sample average approximations for locally inf-integrable random lower-semicontinuous functions under some conditions. However, (8) has multiple complications that prevent

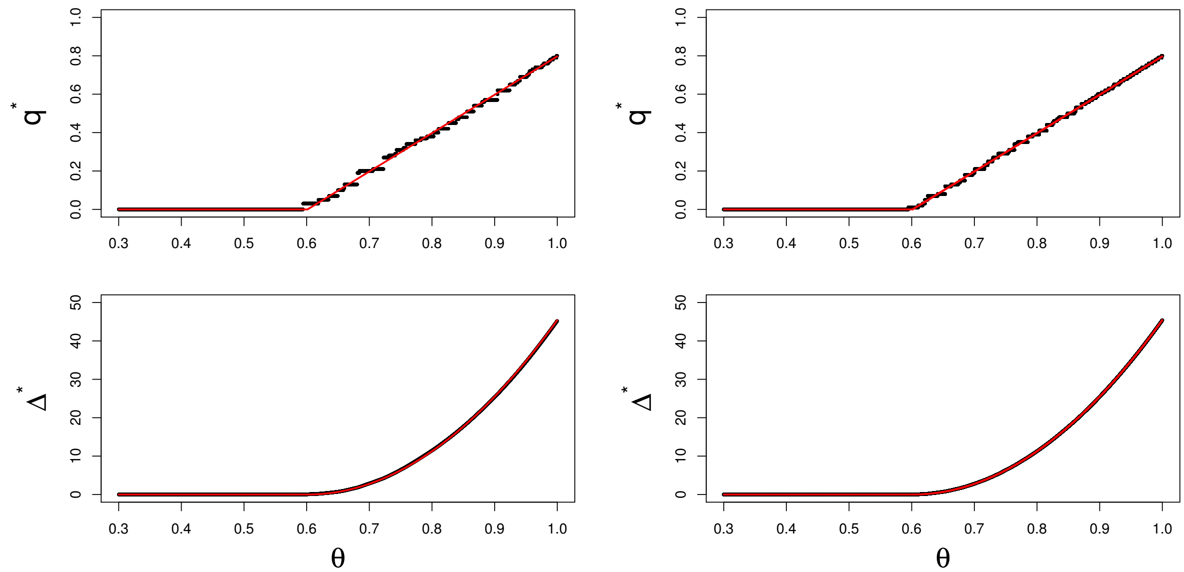

Figure 2: Discrete problem solution with N_0 =10,000 and N_0 =100,000. The red line shows the analytical results (6) and (7) for the continuous problem.

Table 1: Mean and half-width of the optimal profit (in millions of dollars) for the SAA estimator (8) with 100 independent replications used for each value of N_0.

N_0	Mean	95% Half-Width
100	18.28	0.36
1,000	17.41	0.13
10,000	17.29	0.04
100,000	17.27	0.01
Continuous	17.25	—

immediate application of prior results to show consistency. The first is that while the functions $q^*(\theta)$ and $\Delta^*(\theta)$ are continuous over the range of θ, the SAA estimates q_n^* and Δ_n^* are not continuous or defined for all θ. Additionally, the random generation of constraints associated with the samples θ_n cannot be immediately written as a sample average because they are functions of multiple particular values of θ (θ_n and $\theta_{n'}$). Future work will determine the conditions under which the SAA approximation can be quantified under a general PA setting.

To observe the behavior of SAA estimators over many replications, we replicate solving (8) for multiple sample paths and collect the maximum expected profit value. We report confidence intervals for the optimal profit to observe the bias in the estimate and the decreasing half-width as N_0 increases in Table 1. We observe the optimal profit approach the analytical optimal profit for the continuous problem. For well-defined SAA problems with finite solution spaces, the sample size N_0 can conservatively be chosen to guarantee with a given probability that the solution is close to optimal (Shapiro 2003), but due to the lack of a consistency result we are unable to offer direct guidance for our particular problem.

We note that solving the discrete problem for different functional forms of $v(q,\theta)$ and $s(q)$ is not difficult, though sometimes a mixed-integer nonlinear solver is needed. However, we cannot ensure the quality of the SAA solution will be similar to that given in the problem above. For example, when θ is assumed to have an exponential distribution, the unbounded range of the function space yields complications numerically because the range of the space in the discrete problem depends on the samples drawn, and the fit of the discrete solution can be poor at high demand levels where fewer values are sampled.

5 CONCLUSION

This paper suggests the use of sample average approximation to solve continuous principal-agent problems using sampled demand values in discrete principal-agent problems. The discrete problem, while non-linear, is often well-posed so that a solution can be obtained using a solver. We explore a specific implementation that yields analytical solutions for the continuous problem. We present an example that suggests consistency may exist under some simplified conditions. Future work needs to show what conditions are needed, and how the discrete solution can be transformed to a continuous function that approximates the true solution.

ACKNOWLEDGMENTS

This work was partially supported by National Science Foundation grants CMMI-1535762 and CMMI-1535831.

A CONTINUOUS PROBLEM PARTIAL SOLUTION

The derivation of the solution for the continuous problem follows the one provided in Tadelis and Segal (2005). We assume the distribution of θ is continuous and bounded over $[\underline{\theta},\overline{\theta}]$. Let $v_\theta(q(\theta),\theta)$ be the derivative of the value function with respect to θ. The problem in formulation (3) is equivalent to the following:

$$\begin{array}{llll}\max\limits_{\{q(\theta),\Delta(\theta)\}} \Phi = & \int_{\underline{\theta}}^{\overline{\theta}} [v(q(\theta),\theta)-\Delta(\theta)-s(q(\theta))]\, f(\theta)d\theta & & \\ \text{s.t.} & \Delta(\underline{\theta})=0 & & (IR_{\underline{\theta}}) \\ & \Delta'(\theta)=v_\theta(q(\theta),\theta) & \theta\in[\underline{\theta},\overline{\theta}] & (ICFOC_\theta) \\ & q'(\theta)\geq 0 & \theta\in[\underline{\theta},\overline{\theta}] & (MON_\theta).\end{array}$$

Constraints $(ICFOC_\theta)$ and $(IR_{\underline{\theta}})$ yield

$$\Delta^*(\theta)=\int_{\underline{\theta}}^{\theta} v_\vartheta(q(\vartheta),\vartheta)d\vartheta. \tag{9}$$

Thus, the objective function becomes

$$\begin{aligned}
& \max_{q(\theta)} \int_{\underline{\theta}}^{\overline{\theta}} \left[v(q(\theta),\theta) - \int_{\underline{\theta}}^{\theta} v_\vartheta(q(\vartheta),\vartheta)d\vartheta - s(q(\theta)) \right] f(\theta)d\theta \\
= \; & \max_{q(\theta)} \int_{\underline{\theta}}^{\overline{\theta}} [v(q(\theta),\theta)-s(q(\theta))]\, f(\theta)d\theta - \int_{\underline{\theta}}^{\overline{\theta}}\int_{\underline{\theta}}^{\theta} v_\vartheta(q(\vartheta),\vartheta)d\vartheta f(\theta)d\theta \\
= \; & \max_{q(\theta)} \int_{\underline{\theta}}^{\overline{\theta}} [v(q(\theta),\theta)-s(q(\theta))]\, f(\theta)d\theta - \left(\int_{\underline{\theta}}^{\theta} v_\vartheta(q(\vartheta),\vartheta)d\vartheta \cdot F(\theta)\Big|_{\underline{\theta}}^{\overline{\theta}} - \int_{\underline{\theta}}^{\overline{\theta}} v_\theta(q(\theta),\theta)F(\theta)d\theta \right) \\
= \; & \max_{q(\theta)} \int_{\underline{\theta}}^{\overline{\theta}} [v(q(\theta),\theta)-s(q(\theta))]\, f(\theta)d\theta - \left(\int_{\underline{\theta}}^{\overline{\theta}} v_\theta(q(\theta),\theta)d\theta - \int_{\underline{\theta}}^{\overline{\theta}} v_\theta(q(\theta),\theta)F(\theta)d\theta \right) \\
= \; & \max_{q(\theta)} \int_{\underline{\theta}}^{\overline{\theta}} \left[v(q(\theta),\theta) - s(q(\theta)) - v_\theta(q(\theta),\theta)\frac{\overline{F}(\theta)}{f(\theta)} \right] f(\theta)d\theta.
\end{aligned} \tag{10}$$

B A SPECIAL CASE

In the case where $p(x) = \frac{\alpha}{2}x^2$, $v(q,\theta) = p(\theta) - p(\theta - q) - c(q)$, $c(q) = \gamma q$, and $s(q) = \beta q$, we can compute $v(q,\theta) = \frac{\alpha}{2}\theta^2 - \frac{\alpha}{2}(\theta - q)^2 - \gamma q = \alpha\theta q - \frac{\alpha}{2}q^2 - \gamma q$, $v_\theta(q,\theta) = \alpha q$, $v_q(q,\theta) = \alpha(\theta - q) - \gamma$ and $\Delta^*(\theta) = \int_{\underline{\theta}}^{\theta} \alpha q(\vartheta)d\vartheta$.

The objective function in the last line of (10) achieves optimality when the integrand at each θ is maximized. Taking the derivative of the integrand with respect to q for a fixed θ and setting it to zero, we obtain the following:

$$\frac{\partial}{\partial q}\left[v(q,\theta) - s(q) - v_\theta(q,\theta)\frac{\overline{F}(\theta)}{f(\theta)}\right] = \alpha\theta - \alpha q - (\gamma+\beta) - \alpha\frac{\overline{F}(\theta)}{f(\theta)} = 0,$$

and thus the quantity selected becomes a function of θ, that is, $q(\theta) = \theta - \frac{\gamma+\beta}{\alpha} - \frac{\overline{F}(\theta)}{f(\theta)}$. In the special case where θ follows a distribution $Unif[\underline{\theta}, \overline{\theta}]$, $q(\theta) = \theta - \frac{\gamma+\beta}{\alpha} - (\overline{\theta} - \theta) = 2\theta - \overline{\theta} - \frac{\gamma+\beta}{\alpha}$.

Because the derivative of $q(\theta)$ with respect to θ is constant and equals to 2, the (MON_θ) constraint is satisfied. We can thus conclude that $q^*(\theta) = 2\theta - \overline{\theta} - \frac{\gamma+\beta}{\alpha}$. To ensure nonnegativity is satisfied, we set the optimal quantity as

$$q^*(\theta) = \begin{cases} 0 & \text{if } \theta < \theta_c, \text{ where } \theta_c = \max\left\{\underline{\theta}, \frac{1}{2}\left(\overline{\theta} + \frac{\gamma+\beta}{\alpha}\right)\right\} \\ 2\theta - \overline{\theta} - \frac{\gamma+\beta}{\alpha} & \text{if } \theta \geq \theta_c. \end{cases}$$

Using (9) we can derive

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

The principal's optimal profit is then

$$\Phi^* = -\frac{\alpha}{3}\left(\overline{\theta}^3 - \theta_c^3\right) + \alpha\left(\overline{\theta} - \theta_c\right)\left[\frac{1}{2}\overline{\theta}^2 + \frac{1}{2}\left(\frac{\gamma+\beta}{\alpha}\right)^2 + \theta_c\left(\theta_c - \frac{\gamma+\beta}{\alpha}\right)\right].$$

REFERENCES

Akan, M., B. Ata, and M. A. Lariviere. 2011. “Asymmetric Information and Economies-of-Scale in Service Contracting”. *Manufacturing & Service Operations Management* 13 (1): 58–72.

Cachon, G., and F. Zhang. 2006. “Procuring Fast Delivery: Sole Sourcing with Information Asymmetry”. *Management Science* 52 (6): 881–896.

Cai, W., and D. Singham. 2017. “A Principal-Agent Problem with Heterogeneous Demand Distributions for a Carbon Capture and Storage System”. Forthcoming in *European Journal of Operational Research*.

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

Chaturvedi, A., and V. Martínez-de-Albéniz. 2011. “Optimal Procurement Design in the Presence of Supply Risk”. *Manufacturing & Service Operations Management* 13 (2): 227–243.

Esposito, R., L. Monroe, and J. Friedman. 2011. “Deployment Models for Commercialized Carbon Capture and Storage”. *Environmental Science & Technology* 45 (1): 139–146.

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

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

Hasija, S., E. J. Pinker, and R. A. Shumsky. 2008. "Call Center Outsourcing Contracts Under Information Asymmetry". *Management Science* 54 (4): 793–807.

Huber, S., and S. Spinler. 2013. "Pricing of Full-Service Repair Contracts". *European Journal of Operational Research* 222:113–121.

Iyer, A., L. Schwarz, and S. Zenios. 2005. "A Principal-Agent Model for Product Specification and Production". *Management Science* 51:106–119.

Kim, S., M. Coen, S. Netessine, and S. Veeraraghavan. 2010. "Contracting for Infrequent Restoration and Recovery of Mission-Critical Systems". *Management Science* 56 (9): 1551–1567.

Klokk, Ø., P. Schreiner, A. Pages-Bernaus, and A. Tomasgard. 2010. "Optimizing a CO_2 Value Chain for the Norwegian Continental Shelf". *Energy Policy* 38:6604–6614.

Laffont, J.-J., and D. Martimort. 2009. *The Theory of Incentives: The Principal-Agent Model*. Princeton University Press.

Lovejoy, W. S. 2006. "Optimal Mechanisms with Finite Agent Types". *Management Science* 52 (5): 788–803.

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

Middleton, R., M. Kuby, R. Wei, G. Keating, and R. Pawar. 2012. "A Dynamic Model for Optimally Phasing in CO_2 Capture and Storage Infrastructure". *Environmental Modelling & Software* 37:193–205.

Mirrlees, J. 1999. "The Theory of Moral Hazard and Unobservable Behavior: Part I". *The Review of Economic Studies* 66 (1): 3–21.

Norde, H., U. Özen, and M. Slikker. 2016. "Setting the Right Incentives for Global Planning and Operations". *European Journal of Operational Research* 253:441–455.

Özer, O., and W. Wei. 2006. "Strategic Commitment for Optimal Capacity Decision Under Asymmetric Forecast Information". *Management Science* 52 (8): 1238–1257.

Royset, J.O. and Wets, R. J-B. 2017. "Constrained Maximum Likelihood Estimators for Densities". https://arxiv.org/abs/1702.08109.

Shapiro, A. 2003. "Monte Carlo Sampling Methods". *Handbooks in Operations Research and Management Science* 10:353–425.

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

Tadelis, S. and Segal, I. 2005. "Lectures in Contract Theory". http://faculty.haas.berkeley.edu/stadelis/econ_206_notes_2006.pdf.

Tsay, A. 1999, October. "The Quantity Flexibility Contract and Supplier-Customer Incentives". *Management Science* 45 (10): 1339–1358.

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

Yang, Z., G. Aydin, V. Babich, and D. Beil. 2009. "Supply Disruptions, Asymmetric Information, and a Backup Production Option". *Management Science* 55 (2): 192–209.

AUTHOR BIOGRAPHIES

DASHI I. SINGHAM is a Research Associate Professor at the Naval Postgraduate School in the Operations Research Department. Her main research area of interest is simulation analysis, with a focus on output analysis methods such as sequential confidence interval estimators and crossing probabilities for the cumulative sample mean. Recently, she has also worked in search theory and revenue management models. She obtained her Ph.D. from UC Berkeley in Industrial Engineering and Operations Research. Her e-mail address is dsingham@nps.edu.

WENBO CAI is an Assistant Professor in the department of Mechanical and Industrial Engineering at the New Jersey Institute of Technology. Her main research is in the area of operations management, with a focus on incorporating game theory frameworks in revenue management models to study the impact of consumer purchasing behaviors on supply chain management decisions. She obtained her Ph.D. in Industrial Engineering and Operations Research at UC Berkeley. Her email address is cai@njit.edu.