

# A Bayesian Approach to Continuous Type Principal-Agent Problems

Mike G. Tsionas\*

Albert G. Assaf†

Ruijun Bu‡

## Abstract

Singham (2019) [Singham, D.I. (2019). Sample Average Approximation for the Continuous Type Principal-Agent Problem. *European Journal of Operational Research*, Volume 275, Issue 3, 16 June 2019, Pages 1050-1057] proposed an important advance in the numerical solution of continuous type principal-agent problems using Monte Carlo simulations from the distribution of agent “types” followed by bootstrapping. In this paper, we propose a Bayesian approach to the problem which produces nearly the same results without the need to rely on optimization or lower and upper bounds for the optimal value of the objective function. Specifically, we cast the problem in terms of maximizing the posterior expectation with respect to a suitable posterior measure. In turn, we use efficient Markov Chain Monte Carlo techniques to perform the computations.

**Key Words:** Pricing; Principal-Agent Models; Bayesian Analysis; Markov Chain Monte Carlo.

**Acknowledgments:** The authors are grateful to an anonymous referee for several useful remarks on earlier versions of the paper, as well as to Dashi Singham.

---

\*Lancaster University Management School, LA1 4YX, U.K, [m.tsionas@lancaster.ac.uk](mailto:m.tsionas@lancaster.ac.uk)

†Isenberg School of Management, University of Massachusetts-Amherst, 90 Campus Center Way, 209A Flint Lab, Amherst, MA 01003, USA, [assaf@isenberg.umass.edu](mailto:assaf@isenberg.umass.edu)

‡The University of Liverpool Management School, L69 7ZH, U.K, [Ruijunbu@liverpool.ac.uk](mailto:Ruijunbu@liverpool.ac.uk)

# 1 Introduction

In principal-agent problems (Laffont and Martimort, 2009), an agent has a “type” which is given by a random variable  $\theta$  with finite support, say  $[\underline{\theta}, \bar{\theta}]$ . The “type” of the agent is, in this paper, the demand of the agent for a given product offered by the principal. The probability density function (pdf) of  $\theta$  is  $f(\theta)$ . We assume that this density (across different “types” or agents) is known to both the agent and the principal. However, the precise value of  $\theta$  is unknown to the principal. The principal’s problem is to maximize expected profit by offering the agent  $q$  units of the product at total price  $\tau$ . In the continuous version of the problem, the principal offers the agent contract functions  $q(\theta)$  and  $\tau(\theta)$ . In turn, the agent chooses the quantity and price option depending on his / her private demand. Singham (2019) provides a method for bounding the optimal profit and finding solution estimates for the continuous version of the principal-agent problem. This method can also be used when the density  $f(\theta)$  is not simple enough to yield analytical solutions, or even when  $f(\theta)$  is not available but can be estimated using a sample of random numbers from the corresponding distribution. For an example in principal-agent problems, see Singham and Cai (2017) where a sample-average approximation is used (in the context of an example). More general models are considered in Singham (2019).

The principal’s objective is to offer the agent a contract  $\{q(\theta), \tau(\theta)\}$  which specifies a quantity  $q(\theta)$  to be delivered at total price  $\tau(\theta)$  for each given demand. This contract maximizes agent’s expected profit  $\mathbb{E}[\tau(\theta) - \sigma(q(\theta))]$  with respect to random variable  $\theta$ , where  $\sigma(q(\theta))$  is principal’s cost (increasing function of  $q$ ). **The agent has a value function  $v(q, \theta)$  for receiving  $q$  units of the product when his/her demand or “type” is  $\theta$ .** We denote the excess utility to the agent by  $\Delta(\theta) = v(q(\theta), \theta) - \tau(\theta)$ , **where  $\tau(\theta)$  is total cost for demanding  $\theta$ .** We can rewrite the decision variables for the principal using  $\{q(\theta), \Delta(\theta)\}$  instead of  $\{q(\theta), \tau(\theta)\}$ .<sup>1</sup> Singham (2019) proposed the following principal-agent problem originally formulated in Appendix 3.1 of Laffont and Martimort (2009):

$$\begin{aligned}
 & \max_{\{q(\theta), \Delta(\theta)\}, \theta \in [\underline{\theta}, \bar{\theta}]} \int_{\underline{\theta}}^{\bar{\theta}} \{v(q(\theta), \theta) - \Delta(\theta) - \sigma(q(\theta))\} f(\theta) d\theta \\
 & \text{s.t } \Delta(\underline{\theta}) = 0, \quad (IR) \\
 & \Delta(\theta) \geq \Delta(\theta') + v(q(\theta'), \theta) - v(q(\theta'), \theta'), \quad \theta \geq \theta', \quad \theta, \theta' \in [\underline{\theta}, \bar{\theta}], \quad (IC) \\
 & q(\theta) \geq q(\theta'), \quad \theta \geq \theta', \quad \theta, \theta' \in [\underline{\theta}, \bar{\theta}]. \quad (MON)
 \end{aligned} \tag{1}$$

In this problem, (IR) is the individual rationality constraint, (IC) is the incentive compatibility constraint, and (MON) is the monotonicity constraint. Singham (2019) observed that if  $\{\theta^{(m)}, m = 1, \dots, M\}$  is a random

---

<sup>1</sup>Standard assumptions are:  $v(q, \theta)$  is nonnegative, twice differentiable, concave in  $q$ , increasing in  $q$  and  $\theta$ ,  $v(0, \theta) = 0$ , the single-crossing property holds, and  $\partial v(q, \theta) / \partial q$  is increasing in  $\theta$ . Moreover, Singham (2019) assumes that  $v(q, \theta)$  and  $s(q)$  are bounded over  $q$  and  $\theta \in [0, \bar{\theta}]$ .

sample from the distribution whose pdf is  $f(\theta)$ , and  $\theta^{(1)} \leq \dots \leq \theta^{(M)}$ , an approximation to the problem in (1) is:

$$\begin{aligned} \max_{\{q_m, \Delta_m\}_{m=1, \dots, M}} \quad & M^{-1} \sum_{m=1}^M \{v(q_m, \theta^{(m)}) - \Delta_m - \sigma(q_m)\} \\ \text{s.t.} \quad & \Delta_1 = 0, \\ & \Delta_m \geq \Delta_{m-1} + v(q_{m-1}, \theta^{(m)}) - v(q_{m-1}, \theta^{(m-1)}), \quad m = 2, \dots, M, \\ & q_m \geq q_{m-1}, \quad m = 2, \dots, M. \end{aligned} \tag{2}$$

Here,  $q_m \equiv q(\theta^{(m)})$  and  $\Delta_m \equiv \Delta(q(\theta^{(m)}))$  are decision variables defined as discrete points in the contract function. Singham (2019) proposed a bootstrapping approximation based on lower and upper bounds for the objective function and showed that the algorithm works well in practice. In this paper, we cast the problem in terms of Bayesian inference and show that use of the bootstrap as in Singham (2019) can be avoided. Implementation of Bayesian inference is based on simple Markov Chain Monte Carlo (MCMC) methods using an appropriate likelihood function of the problem in (1). The Bayesian approach delivers posterior standard deviations and marginal posterior densities of highest posterior density (HPD) intervals for the parameters of interest which, in this instance, are  $q_m$  and  $\Delta_m$ . The HPD interval is defined, here, as the 95% Bayes probability interval which is, by definition, an interval that contains the parameter(s) with posterior probability 95%.

## 2 Proposed method

Assume  $F : X \rightarrow \mathbb{R}$ , where  $X \subseteq \mathbb{R}^n$ , and  $F(x)$  is a continuous function that has a global maximum at  $x^* \in X$ . Then by Pincus (1968) we know that

$$x^* = \lim_{h \rightarrow \infty} \frac{\int_X x \exp\{hF(x)\} dx}{\int_X \exp\{hF(x)\} dx}. \tag{3}$$

Therefore, if  $p(x; h) = \frac{\exp\{hF(x)\}}{\int_X \exp\{hF(x')\} dx'}$ , we have  $x^* = \lim_{h \rightarrow \infty} \mathbb{E}(X)$ , when the pdf of  $X$  is given by  $p(x; h)$ , and  $\mathbb{E}(\cdot)$  denotes expectation with respect to the distribution whose density is  $p(x; h)$ . In turn, to estimate the global optimum we can proceed as follows.

- *First*, fix a value for  $h$ .
- *Second*, draw a number of samples from  $p(x; h)$ , denoted by  $\{x_h^{(s)}, s = 1, \dots, S\}$ .
- *Third*, average these samples to estimate the global optimum, viz.  $x^* = \lim_{h \rightarrow \infty} \lim_{S \rightarrow \infty} S^{-1} \sum_{s=1}^S x_h^{(s)}$ .

Here, we propose a different approach relative to Singham (2019). Specifically, we provide a general Bayesian approach to the problem using Markov Chain Monte Carlo (MCMC) to provide statistical inferences. Let  $\{\tilde{\theta}_1, \dots, \tilde{\theta}_M\}$

be a partition of  $[\underline{\theta}, \bar{\theta}]$  with  $\tilde{\theta}_1 = \underline{\theta} \leq \dots \leq \tilde{\theta}_M = \bar{\theta}$ . Notice that this is not a random sample from the distribution whose density is  $p(x; h)$ . Using the Pincus (1968) approximation to (2), the objective function is the expectation corresponding to the density:

$$p(\Delta, q; h) \propto \prod_{m=1}^M \frac{\exp \left\{ h [v_m - \Delta_m - \sigma(q_m)] f(\tilde{\theta}_m) \right\}}{\sum_{m'=1}^M \exp \left\{ h [v_{m'} - \Delta_{m'} - \sigma(q_{m'})] f(\tilde{\theta}_{m'}) \right\}} \cdot p(\Delta, q), \quad (4)$$

where  $\Delta = [\Delta_m, m = 1, \dots, M]$ ,<sup>2</sup>  $q = [q_m, m = 1, \dots, M]$ ,  $v_m = v(q_m, \tilde{\theta}_m)$ ,  $m = 1, \dots, M$ . Here,  $\Delta$  and  $q$  are parameters of the problem, viz. they are the unknowns in (2) with respect to which we optimize. We adopt the following prior:

$$p(\Delta, q) \propto \mathbb{I}((\Delta, q) \in \mathcal{R}), \quad (5)$$

where  $\mathcal{R}$  is the space of restrictions defined by (2), and  $\mathbb{I}(\cdot)$  is the indicator function.

Two important points: i) We do *not* actually need approximations, say  $\hat{\Delta}(\cdot)$  and  $\hat{q}(\cdot)$  as  $\Delta$  and  $q$  are treated as parameters, ii) we use a partition  $\{\tilde{\theta}_1, \dots, \tilde{\theta}_M\}$  and, therefore, we do *not* need to draw a large number of samples from  $f(\theta)$  to solve (2). Since  $f(\tilde{\theta}_m)$  and  $\sigma(q_m)$  and the summation in the denominator of (4) are constant, we define  $f_m = f(\tilde{\theta}_m)$  and  $\sigma_m = \sigma(q_m)$ . Moreover, let us define the parameter vector  $\Theta = [\Delta', q']'$  (where  $\Delta_1 = 0$ ) so that the problem simplifies to:

$$p(\Theta; h) \propto \prod_{m=1}^M \left\{ \exp \{ h [v_m - \Delta_m - \sigma_m] \} f_m \right\} \cdot p(\Theta). \quad (6)$$

The problem then becomes to draw a number of (not necessarily iid) samples, say  $\{\Theta^{(s)}, s = 1, \dots, S\}$  from the distribution of parameters in  $\Theta$  whose non-normalized density is given by (6). To proceed, we use a Hit-and-Run Markov Chain Monte Carlo (MCMC) algorithm to draw a large sample  $\{\Theta^{(s)}, s = 1, \dots, S\}$ . We have to keep in mind that the first element of  $\Delta^{(s)}$ , say  $\Delta_1^{(s)} = 0 \forall s = 1, \dots, S$ .

The purpose of Algorithm 1 is to generate posterior samples for the parameters in  $\Theta$ . Specifically, we use a Hit-and-Run Markov Chain Monte Carlo (MCMC) algorithm to draw a large sample  $\{\Theta^{(s)}, s = 1, \dots, S\}$ . The Hit-and-Run algorithm relies on the idea that one can obtain candidates that are uniformly distributed in the set defined by the restrictions. For more details on the Hit-and-Run algorithm, we refer to Belisle et al. (2013) and Chen and Schmeiser (1996). Our MCMC implementation is described in Algorithm 1.

It is known that MCMC converges under quite mild conditions, namely that the posterior is continuous and positive everywhere (Mengersen and Tweedie, 1996). Therefore, the selection of initial values  $\Theta^{(0)}$  is not essential. In our application, we draw  $\Theta^{(0)}$  from the prior. We verify convergence by drawing 100 different initial conditions and running Algorithm 1 starting from these values.

---

<sup>2</sup>Notice that  $\Delta_1 = 0$  from the statement of (2).

---

**Algorithm 1** MCMC

---

1. Let  $\{\tilde{\theta}_1, \dots, \tilde{\theta}_M\}$  be a partition of  $[\underline{\theta}, \bar{\theta}]$  with  $\tilde{\theta}_1 = \underline{\theta} \leq \dots \leq \tilde{\theta}_M = \bar{\theta}$ .
  2. Let  $\Delta = [\Delta_m, m = 2, \dots, M]$ ,  $q = [q_m, m = 1, \dots, M]$ ,  $v_m = v(q_m, \tilde{\theta}_m)$ ,  $m = 1, \dots, M$ . Also define  $\Theta = [\Delta', q']'$ . For  $s = 1, 2, \dots, S$ , do:
  3. Suppose  $\Theta^{(0)}$  is given. Generate a candidate  $\Theta^c$  drawn uniformly in  $\mathcal{R}$  using the Hit-and-Run algorithm:
    - 3a. Generate a direction  $\delta$  from a uniform distribution on the surface of the unit sphere.
    - 3b. Draw  $\lambda$  uniformly in the set  $\{\lambda \in \mathbb{R} | \Theta^c + \lambda\delta \in \mathcal{R}\}$ .
    - 3c. Set  $\Theta^c = \Theta^{(s)} + \lambda\delta$ .
    - 3d. Accept the candidate with the Metropolis-Hastings probability:  $\min\left\{1, \frac{p(\Theta^c; h)}{p(\Theta^{(s)}; h)}\right\}$ , and set  $\Theta^{(s+1)} = \Theta^c$ . Otherwise, set  $\Theta^{(s+1)} = \Theta^{(s)}$ .
- 
- next  $s$ .
- 

The optimal contract is approximated using the MCMC draws, viz:

$$\bar{q}_m = S^{-1} \sum_{s=1}^S q_m^{(s)}, \quad \bar{\Delta}_m = S^{-1} \sum_{s=1}^S \Delta_m^{(s)} \quad \forall m = 1, \dots, M, \quad (7)$$

where  $\{q_m^{(s)}, \Delta_m^{(s)}, s = 1, \dots, S\}$  is the available MCMC sample for all  $m = 1, \dots, M$ . From (4) or (6) the objective function is computed, for each MCMC draw, as:

$$O_m^{(s)} \equiv v_m^{(s)} - \Delta_m^{(s)} - \sigma_m^{(s)} \quad \forall m = 1, \dots, M, s = 1, \dots, S. \quad (8)$$

Taking the average across MCMC samples, we obtain the final value of the objective function:

$$\bar{O}_m = S^{-1} \sum_{s=1}^S O_m^{(s)} \quad \forall m = 1, \dots, M. \quad (9)$$

Provided the problem has a global maximum, it is easy to show that the posterior is finitely integrable as the likelihood is bounded and the prior is proper. Therefore, (7) and (9) converge to the actual values by a law of large numbers. These approximations can be made arbitrarily accurate provided  $S$  increases.


### 3 Illustration

We consider the second example in Singham (2019). We have:

$$v(q, \theta) = \alpha\theta(1 - e^{-q}), \quad \sigma(q) = \beta q^2 - \gamma q, \quad (10)$$

where  $\alpha = 10, \beta = 5, \gamma = 3$  and  $f(\theta)$  is a triangular distribution with minimum  $\underline{\theta} = 0.3$ , maximum  $\bar{\theta} = 1$  and mode at 0.8. We partition the support  $[\underline{\theta}, \bar{\theta}]$  using  $M = 10$  points in total (the endpoints being 0.3 and 1). We use

Figure 1: Optimal contract ( $M = 10$ )



Notes: Dotted lines represent 95% Bayes probability intervals which are computed point-wise for each value.

$S = 150,000$  iterations of MCMC omitting the first 50,000 to mitigate the impact of possible start-up effects. After some initial experimentation, we set  $h = 10^3$  as increasing the value of  $h$  produces identical results for all quantities of interest (optimal contract and objective function) to at least three decimal places. Moreover, we examined the optimal contract and the objective function when starting MCMC from 100 different, randomly chosen, sets of initial conditions. Results were the same with those reported here, to at least three decimal places indicating that MCMC converges very fast and explores the posterior in a thorough way.

The results for the optimal contract ( $q(\theta), \Delta(\theta)$ ) are reported in Figure 1 and they are quite similar to those reported in Figure 2 of Singham (2019).<sup>3</sup>


The optimal value of the objective function is 2.113 which is well within the range [2.111, 2.115] of lower and upper bounds provided by Singham (2019) in Table 2 when 10,000 Monte Carlo simulations are used. Our posterior standard error is 0.0068 which is quite close to 0.007 reported by Singham (2019). The total execution time on a desktop PC (Intel® Core™ i9-7900X CPU @ 3.30 GHz, Windows 10) running `fortran77` was less than a minute, with the bulk of the computation falling on rejection sampling for step 3b of Algorithm 1. No parallelization was used in this case.

In Figure 2, we report the optimal contract when  $M = 150$ , along with 95% Bayes probability intervals (also known as **HPD** intervals in this instance). These intervals contain the parameter with posterior probability 95%. The results are, essentially, the same as those reported in Figure 1. Total execution time was slightly over 1.5 minutes when parallel computing was used after translating software programs to `fortran90`.<sup>4</sup>

<sup>3</sup>The choice of a small  $M$  would only work with linear interpolation in cases where the optimal contracts could be approximated in a piece-wise linear way and the value of  $M$  chosen at key points to match potential non-smooth points. As computations can be easily carried out using parallel programming, considering large values of  $M$  is not prohibitive.

<sup>4</sup>We have used `F77_T0_F90` which is a `fortran90` program which can convert `fortran77` codes to `fortran90` format. See [https://people.sc.fsu.edu/~jburkardt/f\\_src/f77\\_to\\_f90/f77\\_to\\_f90.html](https://people.sc.fsu.edu/~jburkardt/f_src/f77_to_f90/f77_to_f90.html)

Figure 2: Optimal contract ( $M = 150$ )



Notes: Dotted lines represent 95% Bayes probability intervals which are computed point-wise for each value.


## 4 Illustration: Carbon capture

We consider the first example in Singham (2019) which has an analytical solution and is, therefore, useful to examine the effect of the number of MCMC draws on the optimal contract. This is an implementation based on Singham and Cai (2017), see also Singham et al. (2015). In their description: “[T]he agent is a power plant who emits CO<sub>2</sub> into the atmosphere and demands the service of the principal, who is a CO<sub>2</sub> storage operator. The principal offers a menu of options to the agent, with each option corresponding to a different demand value from a continuous distribution. The agent demands units of CO<sub>2</sub> 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<sub>2</sub> to the atmosphere. If the agent participates and purchases  $q$  units of storage, he has a linear cost  $\gamma q$  to capture CO<sub>2</sub> emissions before sending them to storage”. The value function of the agent is given as:

$$v(q, \theta) = p(\theta) - p(\theta - q) - \gamma q, \quad (11)$$

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  $\sigma(q) = \beta q$  corresponding to  $q$  units of CO<sub>2</sub> storage. We set  $\alpha = 286$ ,  $\beta = 13$ ,  $\gamma = 45$  (after taking account of the units of measurement), and  $f(\theta)$  is a uniform distribution in the interval  $[\underline{\theta}, \bar{\theta}]$  where  $\underline{\theta} = 0.3$ , and  $\bar{\theta} = 1$ . We partition the support  $[\underline{\theta}, \bar{\theta}]$  using  $M = 150$  points in total (the endpoints being 0.3 and 1). After careful examination, we set  $h = 10^3$  as increasing the value of  $h$  produces identical results for all quantities of interest (optimal contract and objective function) to at least three decimal places. We used 100 different initial conditions to check convergence and robustness. The number of draws is set to 100 or 10,000 based on a burn-in phase of 100 MCMC passes.

Figure 3: Carbon capture optimal contract



Notes: In panels (a) and (b) we reported optimal values of  $q$  and  $\Delta$  using  $M = 150$  points in the support of  $\theta$ . In panels (c) and (d) reported are the results for  $q$  and  $\Delta$ , respectively, using 100 draws, along with 95% HPD interval.

The results for the optimal contract  $(q(\theta), \Delta(\theta))$  are reported in panels (a) and (b) of Figure 3 (for  $M = 150$  but different number of MCMC draws) and they are quite similar to those reported in Figure 1 of Singham (2019). The optimal posterior mean contract along with 95% HPD intervals, are reported in panels (c) and (d) of Figure 3.

The optimal value of the objective function is 17.2489 which is quite close to the exact solution (17.25). Our posterior standard error is 0.057 which is quite close to 0.059 reported by Singham (2019) in her Table 1 for  $N = 100,000$ . The total execution time on a desktop PC (Intel® Core™ i9-7900X CPU @ 3.30 GHz, Windows 10) running `fortran77` was less than two minutes, with the bulk of the computation falling on rejection sampling for step 3b of Algorithm 1.

## Concluding remarks

In this work, we follow Singham (2019) to cast the continuous principal-agent problem in terms of maximizing the posterior expectation with respect to a suitable posterior measure. We propose the use efficient Markov Chain Monte Carlo (MCMC) techniques to perform the computations, improving on the bootstrapping / Monte Carlo approach of Singham (2019) providing almost the same results at a fraction of time. Moreover, we do not need random draws from the distribution of “types” followed by optimization to deliver upper and lower values for the objective function. The required number of draws is close to 10,000 as Singham (2019) has shown. With the same number of draws from the posterior measure and a fixed number of points in the support of the distribution of “types”, MCMC provides almost the same results with those reported in Singham (2019, Example 2, Section 4.2). Part of the attraction of MCMC is that Bayes probability intervals are available for the objective function as well as the optimal contract.

## References

- Belisle, C.J.P. , H.E. Romeijn & R.L. Smith (1993). Hit-and-Run algorithms for generating multivariate distributions, *Mathematics of Operations Research* 18, 255–266.
- Chen, M.-H. & Schmeiser, B.W. (1996). General Hit-and-Run Monte Carlo Sampling for Evaluating Multi-dimensional Integrals, *Operations Research Letters* 19 (4), 161–169.
- Laffont, J.-J. & D. Martimort (2009). *The Theory of Incentives: The Principal-Agent Model*. Princeton University Press.
- Pinkus, M. (1968). A closed form solution for certain programming problems. *Operations Research* 16 (3), 690–694.
- Singham, D., & W. Cai (2017). *Sample Average Approximations for the Continuous Type Principal- Agent Problem: An Example*. In W. K. V. Chan, A. D’Ambrogio, G. Zacharewicz, N. Mustafee, G. Wainer, and E. Page, editors, *Proceedings of the 2017 Winter Simulation Conference.*, pages 2010–2020, Piscataway, New Jersey, 2017. Institute of Electrical and Electronics Engineers, Inc.
- Singham, D., W. Cai, & J. White (2015). *Optimal Carbon Capture and Storage Contracts using Historical CO2 Emissions Levels*. *Energy Systems*, 6 (3), 331–360.
- Singham, D. I. (2019). *Sample Average Approximation for the Continuous Type Principal-Agent Problem*. *European Journal of Operational Research*, 275 (3), 1050–1057.
- Singham, D. I., & W. Cai (2017). *Sample Average Approximations for the Continuous Type Principal- Agent Problem: An Example*. In W. K. V. Chan, A. D’Ambrogio, G. Zacharewicz, N. Mustafee, G. Wainer, & E. Page, Editors, *Proceedings of the 2017 Winter Simulation Conference*, pp. 2010–2020, Piscataway, New Jersey. Institute of Electrical and Electronics Engineers, Inc.