2017 IEEE 56th Annual Conference on Decision and Control (CDC)
December 12-15, 2017, Melbourne, Australia

Plant Gain Estimation in Online Advertising Processes

Niklas Karlsson[1]

***Abstract*— Application of control theory to auction networks is important since it offers computationally efficient solutions to otherwise intractable problems. It is also of great significance due to the widespread use of auctions for resource allocation, e.g. in online advertising. But controlling agents in many large auction networks is challenging because of an uncertain and discontinuous plant. In this paper we utilize bid randomization to make the plant effectively continuous and present novel theoretical results for how to estimate the plant gain given a model uncertainty of the competitive landscape.**

I. INTRODUCTION

Online advertising presents very interesting and extremely challenging problems in feedback control and computational systems. Good solutions have the prospect of making a huge positive business impact because of the large size and rapid growth of the industry.

An early publication on feedback control applied to online advertising is available in [1], wherein several important challenges are outlined but detailed solutions are omitted. A more comprehensive and up-to-date overview of the control problem is available in [2]. There the author proposes a bid randomization technique [3] to overcome some of the challenges. One of the challenges is the need to estimate extremely small event rates [4], and another is to control the event rate of a discontinuous plant [5]. The fact that the plant is unknown, dynamic, nonlinear, and in general discontinuous is a characteristic property of online advertising processes and is a fundamental challenge in the development of feedback control solutions. Different approaches at indirectly estimating and controlling the plant are proposed in [6], [7]. The former of these papers makes use of bid randomization, but both are based exclusively on local feedback information to estimate the plant gain. Adaptive estimation and control is always challenging since it requires persistent excitation [8], but is particularly difficult when the unknown plant is nonlinear and potentially discontinuous. A systematic methodology for off-line modelling of advertising plants is proposed in [9]. The methodology is conveniently used to simulate realistic plants in a test bed for control algorithms.

In this paper we shall estimate the gain of an advertising plant based on a global nominal plant model and a given model uncertainty. The plant model and uncertainty may be derived using techniques from [9], but may alternatively be computed by simply counting the number of available ad *impressions* (the opportunity of showing an ad creative) that were available at different price points for some sampled subset of historical data. Our main contribution is four theorems translating the nominal plant model and model uncertainty into statistical properties of the plant gain. To our best knowledge there is no previous work resembling what is presented in this paper.

The paper is organized as follows. First, we formulate the problem in Section II. Next, in Section III, the Heisenberg bid randomization technique is introduced. The basic rules for impression allocation in online advertising are explained in Section IV. The key results of the paper are presented in Section V in the form of four theorems, and are illustrated in a simulated example in Section VI. Finally, the paper is wrapped up with some concluding remarks in Section VII.

II. PROBLEM FORMULATION

The ultimate objective is to manage an advertisement budget to the satisfaction of the advertiser in the common setup where impressions are awarded via an open *Real Time Bidding* (RTB) exchange. In an RTB exchange any advertiser can submit bids for available impressions. An advertiser typically wants to maximize the total returned value of the available budget while satisfying one or more constraints. The most common constraint is on the smooth delivery of the budget, e.g. to deliver a certain number of \$US worth of impressions per day. Due to the massive scale and dynamic nature of the competitive landscape, we are limited to decentralized and feedback-based techniques [2].

Optimal feedback-based control of an ad campaign involves adjusting the bid price for different impression opportunities up or down until the total value is maximized without any constraint being violated [2]. The feedback data is noisy, but more importantly, the delivery response to adjustments in bid prices is discontinuous and varies dramatically depending on the competitive landscape. See Figure 1 for an example

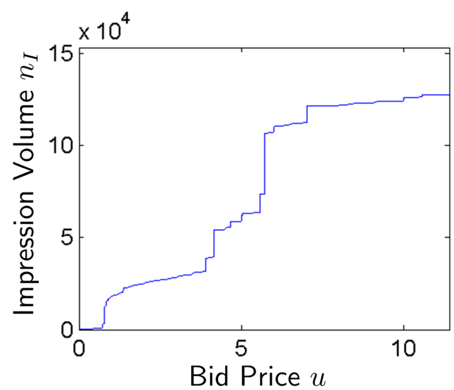

Fig. 1. An example of impression volume versus bid price relationship.

[1]N. Karlsson is Vice President of R&D at Aol Platforms, 395 Page Mill Road, Palo Alto, CA 94306, USA niklas.karlsson at teamaol.com

of the characteristic staircase relationship between bid price and awarded impression count. The discontinuous nature of the process makes an analytical treatment of the problem cumbersome, but a method to turn the system effectively continuous was proposed in [3], [2]. See Section III for more details on this method. The now continuous control problem is still challenging in both theory and practice since the process remains nonlinear and the competitive landscape is a priori unknown. Note, no matter how short the system delay is, closed loop stability demands that the loop gain is less than one around the optimal operating point [10]. But to design a controller that leads to a sufficiently small loop gain we must know the process sensitivity (or plant gain).

Assume we have access to a historical snapshot of the competitive landscape describing how many impression opportunities were available at different price points. The snapshot may be the result of temporal filtering and future projection, but most importantly, it is at best an estimate of what the competitive landscape will be in the future. The objective of this paper is to estimate the plant gain between bid price and impression volume at different operating points of the bid price based on an uncertain estimate of the competitive landscape. It is outside the scope of this paper to utilize the estimated plant gain for control design.

III. Heisenberg Bidding

Heisenberg bidding [3], [2] operates by randomly perturbing a *nominal bid price* u_p according to some statistical distribution. Heisenberg bidding can be implemented with other probability distributions, but in this paper it is defined by the gamma distribution (Appendix A) parameterized by u_p and a *bid uncertainty* u_u, to generate a *final bid price* u used in the market clearing. In particular, u is a realization of a random variable U defined by

$$U \sim \text{Gamma}\left(\frac{1}{u_u^2}, \frac{1}{u_p u_u^2}\right) \quad \text{if} \quad u_p, u_u > 0 \quad (1)$$

and $U = u_p$, otherwise. In terms of the shape parameter α and the inverse scale parameter β, Heisenberg bidding is defined by $\alpha = 1/u_u^2$ and $\beta = 1/(u_p u_u^2)$. Hence, $\text{E}U = u_p$ and $\text{Var}\,(U) = u_p^2 u_u^2$. In other words, $u_u = \text{Std}(U)/\text{E}U$, where $\text{Std}(U)$ is the standard deviation of U.

Figure 2 shows three examples of Heisenberg bid distributions. The bid uncertainty u_u defines the level of perturbation of the nominal bid. If $u_u = 0$ there is no perturbation at all. The smaller u_u is, the spikier is the *probability density function* (PDF).

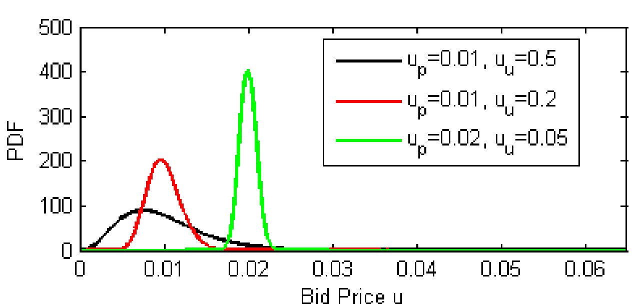

Fig. 2. Three examples of Gamma probability density functions parameterized by the nominal bid price u_p and the bid uncertainty u_u.

Depending on the value of the bid uncertainty, the primary plant input-output relationship (u_p to n_I) can be made arbitrarily smooth. This is illustrated in Figure 3, which shows the result of adding the dimension of bid uncertainty to the plant in Figure 1.

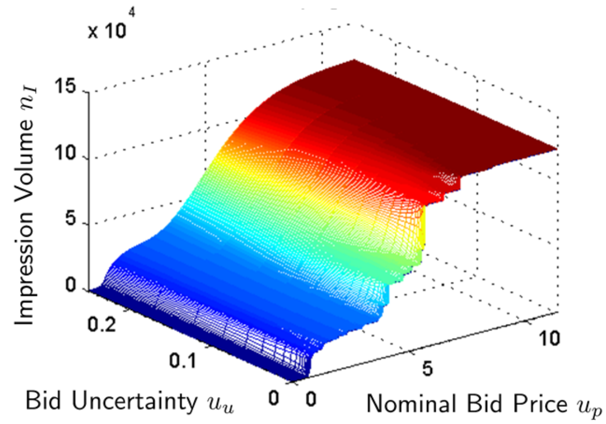

Fig. 3. The impact of using bid uncertainty on the plant in Figure 1. Note that n_I is discontinuous with respect to u_p only when u_u equals zero.

IV. Impression Allocation in Online Advertising

Consider an advertising campaign interested in impressions from an inventory pool subdivided into n segments. A segment may represent a type of users, e.g. "Female Users between 20 and 30 years of Age," "Tech Geeks," or "Internet Users that have not seen a specific ad before," but may as well be defined as "Any User" or "A Unique User".

The impression allocation for segment i is governed by a sealed standard auction [11], where u is the bid price used in the auction. The highest competing bid price is denoted u_i, and the available number of impressions $n_{I,i}^{tot}$. The number of awarded impressions from segment i is consequently $n_{I,i} = \mathbb{I}_{\{u \geq u_i\}} n_{I,i}^{tot}$, where for simplicity we assume the auction is always won if $u \geq u_i$ and where $\mathbb{I}_X$ is the indicator function satisfying $\mathbb{I}_X = 1$, if $X = \text{true}$, and $\mathbb{I}_X = 0$, if $X = \text{false}$.

The campaign-level impression volume n_I is obtained by summing across all segments targeted by the campaign:

$$n_I = \sum_i \mathbb{I}_{\{u \geq u_i\}} n_{I,i}^{tot}. \quad (2)$$

However, u, u_i, and $n_{I,i}^{tot}$ are often random or unknown and better described as realizations of random variables U, U_i, and $N_{I,i}^{tot}$. Hence, n_I is a realization of the random variable $N_I = \sum_i \mathbb{I}_{\{U \geq U_i\}} N_{I,i}^{tot}$. We choose to generate u via Heisenberg bidding (Section III); i.e., the bid price for each impression auction is an independent sample from a probability density function defined by (1), and parameterized by u_p and u_u. Assume that u_i and $n_{I,i}^{tot}$ are constant, but known only in Bayesian sense as samples from probabilistic belief functions of U_i and $N_{I,i}^{tot}$. Conditioned on U_i and $N_{I,i}^{tot}$, each auction for a segment i impression is a Bernoulli experiment with success rate equal to $\Pr(U \geq U_i|U_i)$ while $N_{I,i} = \mathbb{I}_{\{U \geq U_i\}} N_{I,i}^{tot}$ is a binomial random variable (Appendix B) defined by $N_{I,i} \sim \text{Binomial}\left(N_{I,i}^{tot}, \Pr(U \geq U_i|U_i)\right)$.

The success rate is used extensively in the following derivations and is referred to as the win rate. Define *win rate*, W_i, as the probability that an auction for a segment i impression is won by outbidding all other bidders, hence $W_i = \Pr(U \geq U_i|U_i)$. But according to the scaling property of the gamma distribution stating that whenever $X \sim \text{Gamma}(\alpha, \beta)$, then $\beta X \sim \text{Gamma}(\alpha, 1)$ [12], it follows that $W_i = \Pr(U/(u_u^2 u_p) \geq U_i/(u_u^2 u_p)|U_i) = \Pr(\acute{U} \geq U_i/(u_u^2 u_p)|U_i)$, where $\acute{U} \sim \text{Gamma}(1/u_u^2, 1)$. Since the cumulative density function $F_{\acute{U}}(\acute{u}_i) = \Pr(\acute{U} \leq \acute{u}_i)$ for any $\acute{u}_i$ it follows that

$$W_i = 1 - F_{\acute{U}}\left(\left.\frac{U_i}{u_u^2 u_p}\right| U_i\right). \quad (3)$$

Define *win rate sensitivity*, $W_{u_p,i}$, as the derivative of W_i with respect to u_p. It is straight-forward to show that

$$W_{u_p,i} = f_{\acute{U}}\left(\left.\frac{U_i}{u_u^2 u_p}\right| U_i\right) \frac{U_i}{u_u^2 u_p^2}. \quad (4)$$

To turn $W_{u_p,i}$ into a more useful format we make use of the following theorem:

Theorem 4.1:
If $X \sim \text{Gamma}(\alpha, \beta)$, then $f_X(x)x = \alpha f_{\grave{X}}(x)/\beta$, where $\grave{X} \sim \text{Gamma}(\alpha + 1, \beta)$.

Proof: Replace $f_X(x)$ with the expression defining the Gamma probability density function

$$\begin{aligned} f_X(x)x &= \frac{\beta^\alpha}{\Gamma(\alpha)} x^{\alpha-1} e^{-\beta x} x \\ &= \frac{\Gamma(\alpha+1)}{\Gamma(\alpha)\beta}\left[\frac{\beta^{\alpha+1}}{\Gamma(\alpha+1)} x^\alpha e^{-\beta x}\right]. \end{aligned}$$

The expression in $[\cdot]$ is the probability density function of a $\text{Gamma}(\alpha+1, \beta)$ random variable, and $\Gamma(\alpha+1) = \alpha\Gamma(\alpha)$ [13]. It follows that

$$f_X(x)x = \frac{\alpha}{\beta} f_{\grave{X}}(x),$$

where $\grave{X} \sim \text{Gamma}(\alpha+1, \beta)$, which completes the proof. ■

A direct application of Theorem 4.1 to (4) yields

$$W_{u_p,i} = \frac{1}{u_u^2 u_p} f_{\tilde{U}}\left(\left.\frac{U_i}{u_u^2 u_p}\right| U_i\right), \quad (5)$$

where $\tilde{U} \sim \text{Gamma}\left(1/u_u^2 + 1, 1\right)$.

For a particular realization of U_i the win rate and the win rate sensitivity are denoted w_i and $w_{u_p,i}$, respectively, and obtained by replacing U_i by u_i in (3) and (5).

The campaign-level impression volume is given by the sum of awarded impressions in individual segments; i.e., $N_I = \sum_i N_{I,i}$, where the conditional $N_{I,i}$s are independent $\text{Binomial}\left(N_{I,i}^{tot}, W_i\right)$ random variables. Its expected value is referred to as the *impression rate* and is denoted Y; i.e.,

$$Y = \text{E}\left(N_I|N_{I,i}^{tot}, W_i, \forall i\right) = \sum_i N_{I,i}^{tot} W_i. \quad (6)$$

The variance $\text{Var}\left(N_I|N_{I,i}^{tot}, W_i, \forall i\right) = \sum_i N_{I,i}^{tot} W_i(1-W_i)$ is not discussed beyond this paragraph, but note if the campaign is competitive in at least some sufficiently large segment of impressions ($N_{I,i}^{tot}$ and W_i are sufficiently large for some i), then $\text{Std}\left(N_I|N_{I,i}^{tot}, W_i, \forall i\right)/\text{E}\left(N_I|N_{I,i}^{tot}, W_i, \forall i\right)$ is small and $Y \in \mathbb{R}_{\geq 0}$ is an excellent approximation of $N_I \in \mathbb{Z}_{\geq 0}$ conditioned on $N_{I,i}^{tot}$ and W_i, for $i = 1, \ldots, n$.

The *impression rate sensitivity*, Y_{u_p}, is defined by $Y_{u_p} = dY/du_p$ and since $N_{I,i}^{tot}$ is independent of u_p it follows that

$$Y_{u_p} = \sum_i N_{I,i}^{tot} W_{u_p,i}. \quad (7)$$

In the context of feedback control Y_{u_p} represents the plant gain at the control signal value u_p. Knowledge of Y_{u_p} and its statistical properties can be used to select a controller gain that results in a stable and robust closed loop system according to the Nyquist criteria. Last but not least, note that Y and Y_{u_p} are random variables in Bayesian sense (they are unknown constant values).

V. Plant Gain Estimation

The remainder of the paper is focused on the statistical inference of impression rate sensitivity, or *plant gain*, which is the more commonly used term in the control system community. Keep in mind that Y_{u_p} is viewed as a random variable only in Bayesian sense since N_i^{tot} and U_i, for $i = 1, 2, \ldots$, are not known precisely. Our ultimate goal is to derive formulas for the Bayesian $\text{E}Y_{u_p}$ and $\text{Var}Y_{u_p}$, where the expectation and variance operators are with respect to the unknown N_i^{tot} and U_i based on some assumed statistical properties of their belief functions. With this goal we prove the following theorems.

Theorem 5.1: Expected Win Rate Sensitivity
Assume a randomized bidding strategy $U \sim \text{Gamma}\left(1/u_u^2, 1/(u_u^2 u_p)\right)$ with $u_p, u_u > 0$. If the highest competing bid price, u_i, is only known probabilistically as a realization of the random variable $U_i \sim \text{Gamma}\left(1/\sigma_u^2, 1/(\sigma_u^2 \bar{u}_i)\right)$, for known values of $\sigma_u, \bar{u}_i > 0$, then the *expected win rate sensitivity* is

$$\text{E}W_{u_p,i} = \frac{u_u^2 \sigma_u^2 f_{Z_1}\left(\frac{u_u^2 u_p}{\sigma_u^2 \bar{u}_i + u_u^2 u_p}\right)}{u_p\left(u_u^2 + \sigma_u^2 + u_u^2 \sigma_u^2\right)\left(u_u^2 + \sigma_u^2\right)}, \quad (8)$$

where $Z_1 \sim \text{Beta}\left(1/\sigma_u^2 + 1, 1/u_u^2 + 1\right)$.

Proof: The expected win rate sensitivity is given by

$$\text{E}W_{u_p,i} = \int_0^\infty w_{u_p,i}(u_i) f_{U_i}(u_i) du_i,$$

where $w_{u_p,i}(u_i)$ is a realization of (5) and $f_{U_i}(u_i)$ is the probability density function of a $\text{Gamma}\left(1/\sigma_u^2, 1/(\sigma_u^2 \bar{u}_i)\right)$ random variable. Plugging in the expressions for $w_{u_p,i}(u_i)$ and $f_{U_i}(u_i)$ yields

$$\begin{aligned} \text{E}W_{u_p,i} = &\int_0^\infty \frac{1}{u_u^2 u_p} \frac{1}{\Gamma(1/u_u^2+1)}\left(\frac{u_i}{u_u^2 u_p}\right)^{1/u_u^2} e^{-u_i/(u_u^2 u_p)} \\ &\cdot \frac{(1/(\sigma_u^2 \bar{u}_i))^{1/\sigma_u^2}}{\Gamma(1/\sigma_u^2)} u_i^{1/\sigma_u^2 - 1} e^{-u_i/(\sigma_u^2 \bar{u}_i)} du_i. \end{aligned}$$

Collect all factors independent of u_i and move them outside the integral sign.

$$\mathrm{E}W_{u_p,i} = \frac{\int_0^\infty u_i^{1/u_u^2+1/\sigma_u^2-1} e^{-(1/(u_u^2 u_p)+1/(\sigma_u^2 \bar{u}_i))u_i} du_i}{\Gamma(1/u_u^2+1)\Gamma(1/\sigma_u^2)(u_u^2 u_p)^{1/u_u^2+1}(\sigma_u^2\bar{u}_i)^{1/\sigma_u^2}}.$$

The integrand is recognized as the kernel of a Gamma$(1/u_u^2 + 1/\sigma_u^2, 1/(u_u^2 u_p) + 1/(\sigma_u^2\bar{u}_i))$ probability distribution, hence the integral over $(0,\infty)$ must equal $\Gamma(1/u_u^2 + 1/\sigma_u^2)/[1/(u_u^2 u_p) + 1/(\sigma_u^2\bar{u}_i)]^{1/u_u^2+1/\sigma_u^2}$. Replacement of the integral with this expression and straight-forward rearrangement of the right hand side yields

$$\mathrm{E}W_{u_p,i} = \frac{\Gamma\left(\frac{1}{u_u^2}+\frac{1}{\sigma_u^2}\right)}{u_u^2 u_p \Gamma\left(\frac{1}{u_u^2}+1\right)\Gamma\left(\frac{1}{\sigma_u^2}\right)} \cdot \left(\frac{u_u^2 u_p}{\sigma_u^2\bar{u}_i + u_u^2 u_p}\right)^{1/\sigma_u^2}\left(1-\frac{u_u^2 u_p}{\sigma_u^2\bar{u}_i + u_u^2 u_p}\right)^{1/u_u^2}.$$

But $0 < u_u^2 u_p/(\sigma_u^2\bar{u}_i + u_u^2 u_p) < 1$, hence the expression for $\mathrm{E}W_{u_p,i}$ contains the kernel of a Beta $\left(1/\sigma_u^2+1, 1/u_u^2+1\right)$ distribution (Appendix C) evaluated at $u_u^2 u_p/(\sigma_u^2\bar{u}_i + u_u^2 u_p)$. Replacing the kernel with an expression containing the corresponding beta PDF yields

$$\mathrm{E}W_{u_p,i} = \frac{\Gamma\left(\frac{1}{u_u^2}+\frac{1}{\sigma_u^2}\right)}{u_u^2 u_p \Gamma\left(\frac{1}{u_u^2}+1\right)\Gamma\left(\frac{1}{\sigma_u^2}\right)} \cdot B\left(\frac{1}{\sigma_u^2}+1, \frac{1}{u_u^2}+1\right) f_{Z_1}\left(\frac{u_u^2 u_p}{\sigma_u^2\bar{u}_i + u_u^2 u_p}\right),$$

where $Z_1 \sim$ Beta $\left(1/\sigma_u^2+1, 1/u_u^2+1\right)$. Now, use the known identities $\Gamma(x+1) = x\Gamma(x)$, for $x > 0$, and $B(x,y) = \Gamma(x)\Gamma(y)/\Gamma(x+y)$ [13] together with various cancelations to obtain

$$\mathrm{E}W_{u_p,i} = \frac{\frac{1}{u_u^2}\cdot\frac{1}{\sigma_u^2} f_{Z_1}\left(\frac{u_u^2 u_p}{\sigma_u^2\bar{u}_i + u_u^2 u_p}\right)}{u_p\left(\frac{1}{u_u^2}+\frac{1}{\sigma_u^2}+1\right)\left(\frac{1}{u_u^2}+\frac{1}{\sigma_u^2}\right)}.$$

This can be further simplified as

$$\mathrm{E}W_{u_p,i} = \frac{u_u^2\sigma_u^2 f_{Z_1}\left(\frac{u_u^2 u_p}{\sigma_u^2\bar{u}_i + u_u^2 u_p}\right)}{u_p\left(u_u^2+\sigma_u^2+u_u^2\sigma_u^2\right)\left(u_u^2+\sigma_u^2\right)},$$

where $Z_1 \sim$ Beta $\left(1/\sigma_u^2+1, 1/u_u^2+1\right)$, which completes the proof. ■

Theorem 5.2: Expected Impression Rate Sensitivity Assume a randomized bidding strategy $U \sim$ Gamma $\left(1/u_u^2, 1/(u_u^2 u_p)\right)$ with $u_p, u_u > 0$ and consider impressions in n segments denoted $i = 1,\ldots,n$. Suppose the highest competing bid price u_i and the total number of available impressions n_i^{tot}, $i = 1,\ldots,n$, are known only probabilistically as realizations of random variables U_i and N_i^{tot}, where $U_i \sim$ Gamma $\left(1/\sigma_u^2, 1/(\sigma_u^2\bar{u}_i)\right)$ and $\mathrm{E}N_i^{tot} = \bar{n}_i^{tot}$, for known values of $\sigma_u, \bar{u}_i$, and $\bar{n}_i^{tot}$. If U_i and N_i^{tot} are independent for all i, then the expected impression rate sensitivity is

$$\mathrm{E}Y_{u_p} = \frac{u_u^2\sigma_u^2\sum_{i=1}^n \bar{n}_i^{tot} f_{Z_1}\left(\frac{u_u^2 u_p}{\sigma_u^2\bar{u}_i + u_u^2 u_p}\right)}{u_p\left(u_u^2+\sigma_u^2+u_u^2\sigma_u^2\right)\left(u_u^2+\sigma_u^2\right)},$$

where $Z_1 \sim$ Beta $\left(1/\sigma_u^2+1, 1/u_u^2+1\right)$.

Proof: The expected impression rate sensitivity is given by the expected value of (7). Since N_i^{tot} and $W_{u_p,i}$ are independent it follows that

$$\mathrm{E}Y_{u_p} = \sum_{i=1}^n \mathrm{E}\left(N_i^{tot}\right)\mathrm{E}\left(W_{u_p,i}\right).$$

Replace $\mathrm{E}\left(N_i^{tot}\right)$ with $\bar{n}_i^{tot}$ and $\mathrm{E}\left(W_{u_p,i}\right)$ with (8) yields

$$\mathrm{E}Y_{u_p} = \frac{u_u^2\sigma_u^2\sum_{i=1}^n \bar{n}_i^{tot} f_{Z_1}\left(\frac{u_u^2 u_p}{\sigma_u^2\bar{u}_i + u_u^2 u_p}\right)}{u_p\left(u_u^2+\sigma_u^2+u_u^2\sigma_u^2\right)\left(u_u^2+\sigma_u^2\right)},$$

where $Z_1 \sim$ Beta $\left(1/\sigma_u^2+1, 1/u_u^2+1\right)$, which completes the proof. ■

Theorem 5.3: Second Moment of Win Rate Sensitivity Assume a randomized bidding strategy $U \sim$ Gamma $\left(1/u_u^2, 1/(u_u^2 u_p)\right)$ with $u_p, u_u > 0$. If the highest competing bid price u_i is known only probabilistically as a realization of U_i, where $U_i \sim$ Gamma $\left(1/\sigma_u^2, 1/(\sigma_u^2\bar{u}_i)\right)$ for known values of $\sigma_u, \bar{u}_i$, then the second moment of the win rate sensitivity is

$$\mathrm{E}W_{u_p,i}^2 = \frac{2^{1-2/u_u^2}}{B\left(\frac{1}{u_u^2},\frac{1}{u_u^2}\right)}\frac{u_u^2\sigma_u^2 f_{Z_2}\left(\frac{u_u^2 u_p}{2\sigma_u^2\bar{u}_i + u_u^2 u_p}\right)}{u_p^2\left(2\sigma_u^2+u_u^2+u_u^2\sigma_u^2\right)\left(2\sigma_u^2+u_u^2\right)},$$

where $Z_2 \sim$ Beta $\left(1/\sigma_u^2+1, 2/u_u^2+1\right)$.

Proof: The second moment of the win rate sensitivity is given by

$$\mathrm{E}W_{u_p,i}^2 = \int_0^\infty w_{u_p,i}^2(u_i) f_{U_i}(u_i) du_i,$$

where $w_{u_p,i}(u_i)$ is a realization of (5) and $f_{U_i}(u_i)$ is the probability density function of a Gamma $\left(1/\sigma_u^2, 1/(\sigma_u^2\bar{u}_i)\right)$ random variable. Plugging in the expressions for $w_{u_p,i}(u_i)$ and $f_{U_i}(u_i)$ yields

$$\mathrm{E}W_{u_p,i}^2 = \int_0^\infty \left(\frac{1}{u_u^2 u_p}\frac{1}{\Gamma(1/u_u^2+1)}\left(\frac{u_i}{u_u^2 u_p}\right)^{1/u_u^2} e^{-u_i/(u_u^2 u_p)}\right)^2 \cdot \frac{(1/(\sigma_u^2\bar{u}_i))^{1/\sigma_u^2}}{\Gamma(1/\sigma_u^2)} u_i^{1/\sigma_u^2-1} e^{-u_i/(\sigma_u^2\bar{u}_i)} du_i.$$

Collect all factors independent of u_i and move them outside the integral sign.

$$\mathrm{E}W_{u_p,i}^2 = \frac{\int_0^\infty u_i^{2/u_u^2+1/\sigma_u^2-1} e^{-(2/(u_u^2 u_p)+1/(\sigma_u^2\bar{u}_i))u_i} du_i}{\left(u_u^2 u_p\Gamma(1/u_u^2+1)(u_u^2 u_p)^{1/u_u^2}\right)^2\Gamma(1/\sigma_u^2)(\sigma_u^2\bar{u}_i)^{1/\sigma_u^2}}.$$

The integrand is recognized as the kernel of a Gamma$(2/u_u^2 + 1/\sigma_u^2, 2/(u_u^2 u_p) + 1/(\sigma_u^2\bar{u}_i))$ probability distribution, hence the integral over $(0,\infty)$ must equal

$\Gamma\left(2/u_u^2+1/\sigma_u^2\right)/\left(2/(u_u^2u_p)+1/(\sigma_u^2\bar{u}_i)\right)^{2/u_u^2+1/\sigma_u^2}$. Replacement of the integral with this expression and straightforward rearrangement of the right hand side yields

$$\mathrm{E}W_{u_p,i}^2 = \frac{\Gamma\left(\frac{2}{u_u^2}+\frac{1}{\sigma_u^2}\right)}{2^{2/u_u^2}u_p^2\Gamma\left(\frac{1}{u_u^2}\right)^2\Gamma\left(\frac{1}{\sigma_u^2}\right)} \cdot\left(1-\frac{u_u^2u_p}{2\sigma_u^2\bar{u}_i+u_u^2u_p}\right)^{2/u_u^2}\left(\frac{u_u^2u_p}{2\sigma_u^2\bar{u}_i+u_u^2u_p}\right)^{1/\sigma_u^2}.$$

But $0<u_u^2u_p/(2\sigma_u^2\bar{u}_i+u_u^2u_p)<1$, hence the expression for $\mathrm{E}W_{u_p,i}^2$ contains the kernel of a Beta $\left(1/\sigma_u^2+1,2/u_u^2+1\right)$ distribution evaluated at $u_u^2u_p/(2\sigma_u^2\bar{u}_i+u_u^2u_p)$. Replacing the kernel with an expression containing the corresponding beta PDF yields

$$\mathrm{E}W_{u_p,i}^2 = \frac{\Gamma\left(\frac{2}{u_u^2}+\frac{1}{\sigma_u^2}\right)}{2^{2/u_u^2}u_p^2\Gamma\left(\frac{1}{u_u^2}\right)^2\Gamma\left(\frac{1}{\sigma_u^2}\right)} \cdot B\left(\frac{1}{\sigma_u^2}+1,\frac{2}{u_u^2}+1\right)f_{Z_2}\left(\frac{u_u^2u_p}{2\sigma_u^2\bar{u}_i+u_u^2u_p}\right),$$

where $Z_2\sim$ Beta $\left(1/\sigma_u^2+1,2/u_u^2+1\right)$. Now, use the previously referenced properties of the $\Gamma(x)$ and $B(x,y)$ functions together with various cancelations to obtain

$$\mathrm{E}W_{u_p,i}^2 = \frac{2^{1-2/u_u^2}}{B\left(\frac{1}{u_u^2},\frac{1}{u_u^2}\right)}\frac{\sigma_u^2u_u^2f_{Z_2}\left(\frac{u_u^2u_p}{2\sigma_u^2\bar{u}_i+u_u^2u_p}\right)}{u_p^2\left(2\sigma_u^2+u_u^2+\sigma_u^2u_u^2\right)\left(2\sigma_u^2+u_u^2\right)},$$

where $Z_2\sim$ Beta $\left(1/\sigma_u^2+1,2/u_u^2+1\right)$, which completes the proof. ■

Theorem 5.4: Variance of Impression Rate Sensitivity Assume a randomized bidding strategy $U\sim$ Gamma $\left(1/u_u^2,1/(u_u^2u_p)\right)$ with $u_p,u_u>0$ and consider impressions in n segments denoted $i=1,\ldots,n$. Suppose the highest competing bid price u_i and the total number of available impressions n_i^{tot}, $i=1,\ldots,n$, are known only probabilistically as realizations of random variables U_i and N_i^{tot}, where $U_i\sim$ Gamma $\left(1/\sigma_u^2,1/(\sigma_u^2\bar{u}_i)\right)$, $\mathrm{E}N_i^{tot}=\bar{n}_i^{tot}$, and $\mathrm{Var}N_i^{tot}=\sigma_I^2(\bar{n}_i^{tot})^2$ for known values of $\sigma_u,\bar{u}_i,\sigma_I,\bar{n}_i^{tot}$. If $U_i\perp U_j, N_i^{tot}\perp N_j^{tot}$, and $U_i\perp N_i^{tot}$, for all $i\neq j$, ($\perp$ denotes independence), and $\sigma_u,\bar{u}_i>0$, then the variance of the impression volume sensitivity is

$$\mathrm{Var}Y_{u_p} = -\frac{u_u^4\sigma_u^4\sum_{i=1}^n\left(\bar{n}_i^{tot}\right)^2f_{Z_1}^2\left(\frac{u_u^2u_p}{\sigma_u^2\bar{u}_i+u_u^2u_p}\right)}{\left[u_p\left(u_u^2+\sigma_u^2+u_u^2\sigma_u^2\right)\left(u_u^2+\sigma_u^2\right)\right]^2} + \frac{2^{1-2/u_u^2}\left(1+\sigma_I^2\right)}{B\left(\frac{1}{u_u^2},\frac{1}{u_u^2}\right)}\frac{u_u^2\sigma_u^2\sum_{i=1}^n\left(\bar{n}_i^{tot}\right)^2f_{Z_2}\left(\frac{u_u^2u_p}{2\sigma_u^2\bar{u}_i+u_u^2u_p}\right)}{u_p^2\left(2\sigma_u^2+u_u^2+u_u^2\sigma_u^2\right)\left(2\sigma_u^2+u_u^2\right)},$$

where $Z_1\sim$ Beta $\left(1/\sigma_u^2+1,1/u_u^2+1\right)$ and $Z_2\sim$ Beta $\left(1/\sigma_u^2+1,2/u_u^2+1\right)$.

Proof: We are looking for the variance of (7) and because of the independence properties of U_i and N_i^{tot}, the variance equals the sum of the variance of all terms; i.e., $\mathrm{Var}Y_{u_p}=\sum_i\mathrm{Var}Y_{u_p,i}$, where $Y_{u_p,i}=N_{I,i}^{tot}W_{u_p,i}$. To evaluate the variance of each term $Y_{u_p,i}$ we make use of the conditional variance identity [12], which states that for random variables X and Y, $\mathrm{Var}X=\mathrm{E}\left(\mathrm{Var}\left(X|Y\right)\right)+\mathrm{Var}\left(\mathrm{E}\left(X|Y\right)\right)$. It follows that $\mathrm{Var}\left(Y_{u_p,i}\right)=\mathrm{E}\left(\mathrm{Var}\left(N_{I,i}^{tot}W_{u_p,i}|W_{u_p,i}\right)\right)+\mathrm{Var}\left(\mathrm{E}\left(N_{I,i}^{tot}W_{u_p,i}|W_{u_p,i}\right)\right)$. Use the independence of $N_{I,i}^{tot}$ and $W_{u_p,i}$ to obtain $\mathrm{Var}\left(Y_{u_p,i}\right)=\mathrm{Var}\left(N_{I,i}^{tot}\right)\mathrm{E}\left(W_{u_p,i}^2\right)+\left(\mathrm{E}N_{I,i}^{tot}\right)^2\mathrm{Var}\left(W_{u_p,i}\right)$. For any random variable X, we have that $\mathrm{Var}X=\mathrm{E}(X^2)-(\mathrm{E}X)^2$, hence

$$\mathrm{Var}\left(Y_{u_p,i}\right) = \mathrm{Var}\left(N_{I,i}^{tot}\right)\mathrm{E}\left(W_{u_p,i}^2\right) + \left(\mathrm{E}N_{I,i}^{tot}\right)^2\left(\mathrm{E}\left(W_{u_p,i}^2\right)-\left(\mathrm{E}W_{u_p,i}\right)^2\right).$$

By assumption $\mathrm{E}N_{I,i}^{tot}=\bar{n}_i^{tot}$ and $\mathrm{Var}N_{I,i}^{tot}=\sigma_I^2\left(\bar{n}_i^{tot}\right)^2$ while expressions for $\mathrm{E}W_{u_p,i}$ and $\mathrm{E}\left(W_{u_p,i}^2\right)$ were derived in Theorems 5.1 and 5.3. Combining these results and summing over i yields the expression for $\mathrm{Var}Y_{u_p}$ stated in this theorem, which completes the proof. ■

VI. Simulation Results

Consider a plant consisting of six segments and with nominal parameters defined in table I. Assume

TABLE I

Nominal Plant Parameters

i	1	2	3	4	5	6
$\bar{u}_i$	0.2	0.3	0.9	2.4	2.5	3
$\bar{n}_{I,i}^{tot}/10^3$	1	0.2	1.8	0.4	0.4	0.7

the true highest competing bid price u_i and available number of impressions $n_{I,i}^{tot}$ are random realizations from $U_i\sim$ Gamma $\left(1/\sigma_u^2,1/(\sigma_u^2\bar{u}_i)\right)$ and $N_{I,i}^{tot}\sim$ Gamma $\left(1/\sigma_I^2,1/(\sigma_I^2\bar{n}_{I,i}^{tot})\right)$, with the model uncertainty given by $\sigma_u=\sigma_I=0.1$. Consider a randomized bidding strategy $U\sim$ Gamma $\left(1/u_u^2,1/(u_u^2u_p)\right)$ with $u_u=0.05$.

The top subplot of Figure 4 shows 10 randomly generated "true" (red) and the nominal (black) impression rate curves Y. In the middle plot the actual (red), the nominal (black), and the expected (green: Theorem 5.2) impression rate sensitivity Y_{u_p} are displayed. The final plot graphs the standard deviation of the impression rate sensitivity (green: Theorem 5.4). Note how the model uncertainty is encoded in $\mathrm{E}Y_{u_p}$ and $StdY_{u_p}$ providing important information to a control system in situations where the nominal plant gain otherwise may underestimate the true plant gain. Hence, the results derived in this paper enable a higher performing robust control system.

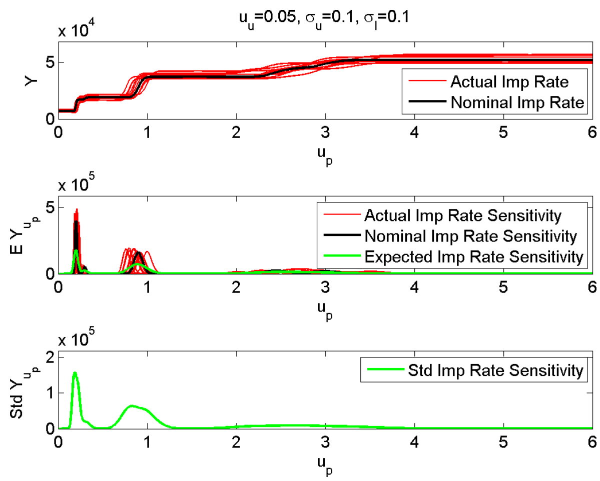

Fig. 4. Top: 10 randomly generated "true" response curves (red) and the nominal response curve (black). Middle: the true (red), the nominal (black), and the expected (green) impression rate sensitivity. Bottom: The standard deviation of the impression rate sensitivity (green).

VII. Conclusions

We have derived theoretical results for how to translate an uncertain competitive bidding landscape into a plant gain estimate (mean and variance). The results enable a robust control design, which is outside the scope of this paper. For software implementations of Theorems 5.2 and 5.4 we recommend making use of the following trivial identity to avoid numerical instability for small values of u_u:

$$\frac{2^{1-2/u_u^2}}{B\left(\frac{1}{u_u^2},\frac{1}{u_u^2}\right)} = e^{\left(1-\frac{2}{u_u^2}\right)\ln 2 - 2\ln\Gamma\left(\frac{1}{u_u^2}\right)+\ln\Gamma\left(\frac{2}{u_u^2}\right)}. \quad (9)$$

One future research direction is to use the results in this paper to compute a sound value of bid uncertainty u_u. The goal may be e.g. to reduce the necessary gain margin for a control system and thereby enhance the control performance.

Appendix

A. Gamma Distribution

The gamma distribution with parameters α and β is a continuous probability distribution. If the random variable X follows the gamma distribution we write $X \sim \text{Gamma}(\alpha, \beta)$. The probability density function of x is given by

$$f_X(x) = \frac{\beta^\alpha}{\Gamma(\alpha)} x^{\alpha-1} e^{-\beta x} \quad (10)$$

for $x > 0$, where $\Gamma(\alpha)$ is the gamma function defined by $\Gamma(\alpha) = \int_0^\infty e^{-t} t^{\alpha-1} dt$. Parameters $\alpha > 0$ and $\beta > 0$ are referred to as *shape* and *inverse scale*. The expected value and variance of X are $\text{E}(X) = \alpha/\beta$ while the variance is $\text{Var}(X) = \alpha/\beta^2$.

B. Binomial Distribution

The binomial distribution with parameters n and p is a discrete probability distribution. If the random variable X follows the binomial distribution we write $X \sim \text{Binomial}(n, p)$. The probability mass function of x is given by

$$f_X(x) = \binom{n}{x} p^x (1-p)^{n-x} \quad (11)$$

for $x = 0, 1, \ldots, n$. Parameters $n \in \{1, 2, \ldots\}$ and $p \in [0, 1]$ are referred to as *number of trials* and *success probability in each trial*, respectively. The expected value of X is $\text{E}(X) = np$ while the variance is $\text{Var}(X) = np(1-p)$.

C. Beta Distribution

The beta distribution with parameters α and β is a continuous probability distribution. If the random variable X follows the beta distribution we write $X \sim \text{Beta}(\alpha, \beta)$. The probability density function of x is given by

$$f_X(x) = \frac{x^{\alpha-1}(1-x)^{\beta-1}}{B(\alpha, \beta)}, \quad (12)$$

for $0 < x < 1$, where $B(\alpha, \beta)$ is the beta function (also called the Euler integral) defined by $B(\alpha, \beta) = \int_0^1 x^{\alpha-1}(1-x)^{\beta-1} dx$. Parameters $\alpha > 0$ and $\beta > 0$ are referred to as *shape* parameters. The expected value and variance of X are $\text{E}(X) = \alpha/(\alpha+\beta)$ and $\text{Var}(X) = \alpha\beta/[(\alpha+\beta)^2(\alpha+\beta+1)]$.

Acknowledgment

The author would like to thank Dr. Qixing Zheng for his careful verification of the proofs and for suggesting (9) to enhance the numerical stability of Theorems 5.3 and 5.4.

References

[1] Niklas Karlsson and Jianlong Zhang. Applications of feedback control in online advertising. *Proceedings of the 2013 American Control Conference, Washington, DC, June 17 - 19*, pages 6008 – 6013, 2013.

[2] Niklas Karlsson. Control problems in online advertising and benefits of randomized bidding strategies. *European Journal of Control*, 30:31–49, July 2016.

[3] Niklas Karlsson. *Systems and Methods For Controlling Bidding For Online Advertising Campaigns, United States Patent Application 20100262497*. USPTO, 2009.

[4] Niklas Karlsson. Adaptive estimation of small event rates. *Proceedings of the 55th IEEE Conference on Decision and Control, Las Vegas, NV, USA, December 12 - 14*, pages 3732–3737, 2016.

[5] Niklas Karlsson and Qian Sang. Event rate control in online advertising. *To appear in the Proceedings of the 2017 American Control Conference, Seattle, WA, USA*, May 24-26 2017.

[6] Niklas Karlsson. Adaptive control using Heisenberg bidding. *Proceedings of the 2014 American Control Conference, Portland, USA, June 4-6, 2014*, pages 1304–1309, 2014.

[7] Jiaxing Guo and Niklas Karlsson. Model reference adaptive control of advertising systems. *To appear in the Proceedings of the 2017 American Control Conference, Seattle, WA, USA*, May 24-26 2017.

[8] Karl Johan Åström and Björn Wittenmark. *Adaptive Control*. Prentice Hall, second edition, 1994.

[9] Vahid Mardanlou, Niklas Karlsson, and Jiaxing Guo. Statistical plant modeling and simulation in online advertising. *To appear in the Proceedings of the 2017 American Control Conference, Seattle, WA, USA*, May 24-26 2017.

[10] Panos J. Antsaklis and Anthony N. Michel. *Linear Systems*. Birkhaüser, 2006.

[11] Vijay Krishna. *Auction Theory*. Academic Press, 2002.

[12] George Casella and Roger L. Berger. *Statistical Inference*. Duxbury, 2 edition, 2001.

[13] W. W. Bell. *Special Functions for Scientists and Engineers*. Dover, 2004.