On the pricing of performance-based programmatic ad-buying contracts*

Joaquin Fernandez-Tapia[†] Olivier Guéant[‡] Jean-Michel Lasry[§]

Abstract

In this paper, we provide a mathematical framework for the rigorous pricing and risk management of performance-based programmatic ad-buying contracts. We mainly focus on the case of Real-Time Bidding (RTB) audience strategies, where ad inventory is purchased algorithmically through the participation to a large number of Vickrey auctions. Our approach is based on stochastic optimal control techniques. It is a general approach in that it makes it possible to consider a broad range of practical situations. In addition to the pricing of ad-buying contracts, we obtain results on both the optimal bidding strategy and the risk associated with each contract, the latter being obtained thanks to Monte Carlo simulations. Besides the mathematical framework itself, our goal is to show that mathematical and numerical tools exist for giving a fair price to performance-based ad-buying contracts – that are too rare in the industry, as of today – and to assess and manage the associated risk.

Keywords: Real-Time Bidding, Stochastic optimal control, Pricing, Risk management, Performance-based contracts, Monte Carlo simulations.

1 Introduction

Recent evolutions in technology have led to considerable changes in the advertising industry. An important example of these changes is in the way ad impressions (in short, the right to communicate a message to a potential customer) are purchased. Nowadays, a large part of

*The authors would like to thank Dominique Delport (Havas Media), Julien Laugel (MFG Labs), Pierre-Louis Lions (Collège de France), and Arnaud Parent (Havas Media) for the discussions they had on the topic.

†Université Pierre et Marie Curie, Laboratoire de Probabilités et Modèles Aléatoires (LPMA). 4 Place Jussieu, 75005 Paris. Joaquin acknowledges support in 2015 from the Chair "Économie et gestion des nouvelles données".

‡Ensae-Crest, 3 avenue Pierre Larousse, 92245 Malakoff Cedex, France. Member of the Scientific Advisory Board of Havas Media. The content of this article does not reflect the official opinion or the practices of Havas Media. Corresponding author: olivier.gueant@ensae.fr.

§Institut Louis Bachelier. Member of the Scientific Advisory Board of Havas Media. The content of this article does not reflect the official opinion or the practices of Havas Media.

the ad inventory on different platforms (display, mobile, TV, etc.) is purchased by companies – often through intermediaries, such as media-buying agencies – in a programmatic way. An example of this is *Real-Time Bidding* (RTB): advertisers can buy impressions in real-time by participating to auctions organized on market platforms called *ad exchanges*. Ad exchanges allow companies to target specific audiences at the right time and in the right context, thus expecting a better return-on-investment than in the case of *direct buys*, i.e. when impressions are purchased in bulk and over-the-counter.

In practice, the companies willing to advertise their products seldom set up by themselves algorithms for RTB. Instead, they rely on intermediaries providing the expertise, the quantitative methodologies, and the technological layers in order to access ad exchanges and implement optimal strategies and tactics. These intermediaries are often media agencies – whose scope and ability to optimize campaigns go far beyond digital advertising and RTB –, media trading desks, demand-side platforms (DSPs), or a combination of them.

One of the central questions at the interface between the technological evolutions, such as RTB, and the business models of the advertising industry, is the following:

"How media buying services should be priced?".

Nowadays, the pricing of ad-buying services is quite rudimentary. Budgets are globally chosen and split among different media, devices, audience targets, etc. Algorithms are then used to maximize the performance of advertising campaigns – measured by various key performance indicators (KPI) –, but the price eventually paid by the company advertising its products is almost always independent of the actual performance. In other words, the pricing and the execution processes are seldom interrelated. In particular, there is no measure or control of the risk of not meeting the expected benchmarks.

Some companies willing to communicate on their products, along with many participants in the advertising industry, have recently shown an interest for alternative pricing models: *performance-based pricing models*. Instead of spending a given budget on a strategy designed to maximize/minimize one or several KPIs, the company willing to advertise its products could pay the intermediary (the media agency or the trading desk) a price defined as a function of the performance of the campaign – for instance, as a function of the number of impressions, the number of clicks – linked to the cost-per-click (CPC) –, the number of acquisitions – linked to the cost-per-acquisition (CPA) –, etc.

Going from the current business environment to a new one where performance-based contracts account for a large part of the turnover will be a challenge for the whole industry: quantitative models will be more and more required, smarter and smarter risk management procedures will be needed, algorithms will need to be constantly improved, and not only to optimize average criteria. It is clear therefore that the media-buying industry has a lot to learn from the methods used in the financial industry (especially in the pricing of contingent claims and in risk management).

The goal of this article is to put forward a first mathematical framework for tackling the pricing problem associated with programmatic performance-based ad-buying services (in particular RTB), and for computing the optimal strategy to reach a desired level of performance. Our approach is interesting both for companies (it allows to have a transparent framework to quantify the value of ad-buying services) and for the ad trading desks of media agencies, as it allows to define the optimal strategy and control the associated risk. It will definitely serve as a stepping stone for more sophisticated mathematical models.

The academic literature on pricing issues in the context of programmatic advertising is quite scarce (e.g. [1, 3, 5]). The current literature fails in fact to provide a simple mathematical framework which is consistent from the auction level to the pricing of the whole strategy. Early attempts on mathematical optimization models for RTB optimization are [7, 8]. One of the first works to give an in-depth mathematical modeling of RTB can be found in [9] (to which this article and [10] are companion papers), but the focus is on strategies and not on pricing. In this paper, we consider a model close to the one presented in [9] (statistical modeling of RTB auctions and optimization through stochastic optimal control techniques), but we focus on the pricing of ad services and on the business consequences of the use of performance-based contracts.

In Section 2, we present the model and introduce the notations used throughout the paper. We also provide a general definition of performance-based contracts, and we present the indifference pricing approach considered in the paper. In Section 3, we show how the problem can be solved in the case of a risk-neutral intermediary, and we present two numerical examples. Thanks to Monte Carlo simulations, we also show the risk associated with performance-based contracts. In Section 4, we solve the problem in the case of a risk-averse intermediary with a constant absolute risk aversion coefficient. Two numerical examples are then discussed in the risk-averse case.

2 Notations, general definitions and model setup

2.1 Performance-based contracts

We consider an ad trader[1] connected to one or several ad exchanges. He receives requests to participate in auctions in order to purchase inventory and display some banners to the specific audiences he wants to target. The ad trader can controls its bid level for each auction he participates in. We assume that there are $J \geq 1$ different types (or sources) of inventories; either different segments of the audience, different ad formats, or different ad exchanges. From a practical perspective, the idea is to assume that each source of inventory represents an homogeneous audience (in terms of the information available to value the ad inventory).

[1] In practice, the ad trader should not be a human but rather an algorithm.

In this paper, we consider performance-based contracts between the advertiser and the intermediary. With such contracts, the price paid by the advertiser depends on the value of one or several KPIs at the end of the trading period or at the end of the whole campaign. These KPIs can be the number of impressions purchased,[2] the number of clicks, or the number of acquisitions of a product (following a click on a banner).[3] In this paper, we regard each click and/or each acquisition as a "conversion": when a banner is displayed, there may or may not be a conversion into something (a click or an acquisition) showing some form of interest of the (web) user for the product or for the brand.

Throughout the article, contracts are defined for a period of time $[0,T]$. We denote by I_T^j, for $j \in \{1,\dots,J\}$, the number of impressions at the end of the period for the source j. In the same way, we denote by C_T^j, for $j \in \{1,\dots,J\}$, the number of conversions at the end of the period for the source j. In this article, the payoff of performance-based contracts is a function of the tuple $(I_T^1,\dots,I_T^J,C_T^1,\dots,C_T^J)$.

From a financial perspective, the contracts work as follows:

- At time $t=0$, the company willing to advertise its products (hereafter, the advertiser), pays a price P (upfront payment) to the intermediary (media agency, trading desk, ad trader, etc).

- Between time $t=0$ and time $t=T$, the intermediary uses algorithms to buy ad inventory.

- At time $t=T$, the campaign associated with the contract is over and there is a final payment/payoff (positive or nonpositive) from the advertiser to the intermediary of the form

  $$g(I_T^1,\dots,I_T^J,C_T^1,\dots,C_T^J),$$

  where the function $g:\mathbb{N}^{2J}\to\mathbb{R}$ is a priori nondecreasing with respect to each coordinate.

This definition is very general. Several specific cases are particularly relevant, and will be considered as special cases throughout this paper:

1. Contracts with a target number of impressions $\overline{I}$. In that case, a natural payoff function is:

   $$g(i^1,\dots,i^J,c^1,\dots,c^J) = -\pi_I(i^1+\ldots+i^J-\overline{I})_-, \quad \pi_I>0.$$

   If the ad trader reaches the target number of impressions, then nothing happens. Otherwise, he has to compensate the advertiser for the impressions that have not been purchased: he pays π_I for each impression missing to reach the target.

[2] This is eventually related to the average cost per (thousand of) impressions, or cost-per-mille (CPM).
[3] The KPIs used in practice are the CPC (cost per click) and the CPA (cost per acquisition).

More generally, one can consider a payoff function of the form

$$g(i^1, \dots, i^J, c^1, \dots, c^J) = g_I(i^1 + \dots + i^J).$$

2. Contracts with a target number of conversions $\overline{C}$. In that case, a natural payoff function is:

$$g(i^1, \dots, i^J, c^1, \dots, c^J) = -\pi_C(c^1 + \dots + c^J - \overline{C})_-, \quad \pi_C > 0.$$

If the ad trader reaches the target number of conversions, then nothing happens. Otherwise, he has to compensate the advertiser for the conversions that have not been obtained: he pays π_C for each conversion missing to reach the target.

More generally, one can consider a payoff function of the form

$$g(i^1, \dots, i^J, c^1, \dots, c^J) = g_C(c^1 + \dots + c^J).$$

3. Contracts with both a target number of impressions and a target number of conversions. In that case, the payoff function is of the form:

$$g(i^1, \dots, i^J, c^1, \dots, c^J) = g_I(i^1 + \dots + i^J) + g_C(c^1 + \dots + c^J).$$

4. Contracts with payoffs depending on the different types of inventories and conversions. In that case, the payoff function is of the form:

$$g(i^1, \dots, i^J, c^1, \dots, c^J) = g_I^1(i^1) + \dots + g_I^J(i^J) + g_C^1(c^1) + \dots + g_C^J(c^J).$$

For technical reasons, we hereafter assume that g has at most linear growth at infinity, i.e. there exists κ such that

$$|g(i^1, \dots, i^J, c^1, \dots, c^J)| \leq \kappa \left(1 + i^1 + \dots + i^J + c^1 + \dots + c^J\right).$$

2.2 Setup of the model

In order to quantitatively define the optimal strategy for the ad trader, and then give a price to a given performance-based contract, we need to model the underlying auction process and its outcomes. For that purpose, we consider a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ equipped with a filtration $(\mathcal{F}_t)_{t\in\mathbb{R}_+}$ satisfying the usual conditions. We assume that all stochastic processes are defined on $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\in\mathbb{R}_+}, \mathbb{P})$.

2.2.1 Auctions

Auction requests are modeled with J marked Poisson processes: the arrival of requests from the source $j \in \{1, \dots, J\}$ is triggered by the jumps of the Poisson process $(N_t^j)_t$ with intensity $\lambda^j > 0$, and the marks $(p_n^j, \xi_n^j)_{n\in\mathbb{N}^*}$ correspond, for each auction request sent by the source j, to the highest bid p^j sent by another participant (i.e. the price to beat in order to win the auction), and to the occurrence $\xi^j \in \{0,1\}$ of a conversion related to the underlying unit of inventory purchased.

Every time the ad trader receives from the source j a request to participate in an auction (the source is known), the ad trader can bid a price: at time t, if he receives a request from the source j, we denote his bid by b^j_t. We assume that for each $j \in \{1,\dots,J\}$, the process $(b^j_t)_t$ is a predictable process with values in $\mathbb{R}_+ \cup \{+\infty\}$. In particular, we assume that the ad trader stands ready to bid (possibly a bid equal to 0 or $+\infty$) at all times.

If at time t the n^{th} auction associated with the source j occurs, the outcome of this auction is the following (corresponding to a Vickrey or second-price auction):

- If $b^j_t > p^j_n$, then the ad trader wins the auction: he pays the price p^j_n and his banner is displayed. In that case, the variable ξ^j_n is equal to 1 if there is a conversion, and is equal to 0 otherwise.[4]
- If $b^j_t \le p^j_n$, then another trader wins the auction. In that case, the variable ξ^j_n is not relevant.

We assume that, for each $j \in \{1,\dots,J\}$, $(p^j_n)_{n\in\mathbb{N}^*}$ are *i.i.d.* random variables distributed according to an absolutely continuous distribution. We denote by F^j the cumulative distribution function and by f^j the probability density function associated with the source j. We assume, for each $j \in \{1,\dots,J\}$, that:

- (H1) $\forall n \in \mathbb{N}^*$, p^j_n is almost surely positive. In particular, $F^j(0) = 0$.
- (H2) $\forall n \in \mathbb{N}^*$, $p^j_n \in L^1(\Omega)$.
- (H3) $\forall p > 0, f^j(p) > 0$.

We also assume that the random variables $(p^j_n)_{j\in\{1,\dots,J\},n\in\mathbb{N}^*}$ are all independent.

2.2.2 State variables

Each time the ad trader wins an auction, he pays the second best price and gets an impression. The cash spent by the ad trader to buy inventory is modeled by a process $(X_t)_t$. Its dynamics is:

$$dX_t = \sum_{j=1}^{J} p^j_{N^j_t} \mathbf{1}_{\left\{b^j_t > p^j_{N^j_t}\right\}} dN^j_t, \quad X_0 = 0.$$

For each $j \in \{1,\dots,J\}$, the number of impressions associated with the auction requests coming for the source j is modeled by an inventory process $(I^j_t)_t$. For each $j \in \{1,\dots,J\}$, the dynamics of $(I^j_t)_t$ is:

$$dI^j_t = \mathbf{1}_{\left\{b^j_t > p^j_{N^j_t}\right\}} dN^j_t, \quad I^j_0 = 0.$$

[4] In practice, there are sometimes issues with the attribution of acquisitions to a specific impression.

To simplify exposition, we write $I_t = (I_t^1, \dots, I_t^J) \in \mathbb{N}^J$, and we write the dynamics of the process $(I_t)_t$:

$$dI_t = \sum_{j=1}^{J} \mathbf{1}_{\left\{b_t^j > p^j_{N_t^j}\right\}} dN_t^j e^j,$$

where $(e^1, \dots, e^J)$ is the canonical basis of $\mathbb{R}^J$.

Conversions are modeled by the variables $(\xi_n^j)_{j\in\{1,\dots,J\},n\in\mathbb{N}^*}$. We assume that they are all independent and independent from the variables $(p_n^j)_{j\in\{1,\dots,J\},n\in\mathbb{N}^*}$. Moreover, we assume that for each $j \in \{1, \dots, J\}$, $(\xi_n^j)_{n\in\mathbb{N}^*}$ are *i.i.d.* random variables distributed according to a Bernoulli distribution with parameter $\nu^j \in [0, 1]$ (assumed to be known[5]). The parameters $\nu^j \in [0, 1]$ are known in practice as *conversion rates* and represent the probability to turn an impression into a conversion.

For each $j \in \{1, \dots, J\}$, the number of conversions associated with the auction requests coming for the source j is modeled by a process $(C_t^j)_t$. For each $j \in \{1, \dots, J\}$, the dynamics of $(C_t^j)_t$ is:

$$dC_t^j = \xi^j_{N_t^j} \mathbf{1}_{\left\{b_t^j > p^j_{N_t^j}\right\}} dN_t^j, \quad C_0^j = 0.$$

To simplify exposition, we write $C_t = (C_t^1, \dots, C_t^J) \in \mathbb{N}^J$, and we write the dynamics of the process $(C_t)_t$:

$$dC_t = \sum_{j=1}^{J} \xi^j_{N_t^j} \mathbf{1}_{\left\{b_t^j > p^j_{N_t^j}\right\}} dN_t^j e^j.$$

2.3 Indifference pricing of performance-based contracts

One of the main goals in this paper is to price performance-based contracts such as those discussed above. For that purpose, we are going to use an approach called indifference pricing. In this approach, the price of a contract is the price that makes the intermediary indifferent between (i) signing the contract with the advertiser and using its best algorithms to maximize (in utility terms) its final payoff and (ii) not signing the contract (and stay idle). Indifference pricing is used in finance for pricing contingent claims in incomplete markets (see [6, 11]), and it is perfectly suited for performance-based contracts in the case of the ad industry. In our case, pricing a contract with this approach requires first to compute the optimal strategy of the intermediary, given the payoff of the contract, and the problem is therefore a problem of stochastic optimal control.

In the risk-neutral case – considered in Section 3 – the problem faced by the intermediary is that of finding an optimal bidding strategy to maximize over $(b_t^1, \dots, b_t^J)_t \in \mathcal{A}^J$ the expected

[5] See [10] for a model in which conversion rates are unknown.

payoff

$$\mathbb{E}\left[P - X_T + g(I_T^1, \dots, I_T^J, C_T^1, \dots, C_T^J)\right],$$

where $\mathcal{A}$ is the set of predictable processes with values in $\mathbb{R}_+ \cup \{+\infty\}$.

In particular, the indifference price P^* of the contract is defined as:

$$P^* = \inf_{(b_t^1, \dots, b_t^J)_t \in \mathcal{A}^J} \mathbb{E}\left[X_T - g(I_T^1, \dots, I_T^J, C_T^1, \dots, C_T^J)\right].$$

Theoretically, this price can be positive or nonpositive, but the design of the contracts, i.e. the choice of g, should be such that it is indeed positive.

In addition to the case of a risk-neutral intermediary, we consider in Section 4 the case of a risk-averse one. More precisely, we consider the case of an intermediary with constant absolute risk aversion $\gamma > 0$. In that case, the problem of the intermediary is to maximize over $(b_t^1, \dots, b_t^J)_t \in \mathcal{A}^J$ the expected utility

$$\mathbb{E}\left[-\exp\left(-\gamma\left(P - X_T + g\left(I_T^1, \dots, I_T^J, C_T^1, \dots, C_T^J\right)\right)\right)\right],$$

and the indifference pricing approach leads to the following definition for the price of the contract:[6]

$$P^* = \frac{1}{\gamma}\log\left(\inf_{(b_t^1, \dots, b_t^J)_t \in \mathcal{A}^J} \mathbb{E}\left[\exp\left(\gamma\left(X_T - g\left(I_T^1, \dots, I_T^J, C_T^1, \dots, C_T^J\right)\right)\right)\right]\right). \tag{2.1}$$

3 Solution in the risk-neutral case

In this section, we consider the special case of a risk-neutral intermediary with the following control problem:[7]

$$\inf_{(b_t^1, \dots, b_t^J)_t \in \mathcal{A}^J} \mathbb{E}\left[X_T - g(I_T^1, \dots, I_T^J, C_T^1, \dots, C_T^J)\right]. \tag{3.1}$$

3.1 Hamilton-Jacobi equation and reduction to a system of ODEs

The value function associated with this problem is the function

$$U : (t, x, I, C) \in [0, T] \times \mathbb{R}_+ \times \mathbb{N}^J \times \mathbb{N}^J \mapsto \inf_{(b_s^1, \dots, b_s^J)_{s \geq t} \in \mathcal{A}_t^J} \mathbb{E}\left[X_T^{b,t,x} - g(I_T^{b,t,I}, C_T^{b,t,C})\right], \tag{3.2}$$

where $\mathcal{A}_t$ is the set of predictable processes on $[t, T]$ with values in $\mathbb{R}_+ \cup \{+\infty\}$, and where

$$dX_s^{b,t,x} = \sum_{j=1}^J p^j_{N_s^j} \mathbf{1}_{\left\{b_s^j > p^j_{N_s^j}\right\}} dN_s^j, \quad X_t^{b,t,x} = x,$$

[6] By definition, P^* is such that:

$$\sup_{(b_t^1, \dots, b_t^J)_t \in \mathcal{A}^J} \mathbb{E}\left[-\exp\left(-\gamma\left(P^* - X_T + g\left(I_T^1, \dots, I_T^J, C_T^1, \dots, C_T^J\right)\right)\right)\right] = -\exp(0) = -1.$$

[7] We consider hereafter a minimization problem instead of a maximization problem.

$$dI^{j\,b,t,I^j}_s = \mathbf{1}_{\left\{b^j_s > p^j_{N^j_s}\right\}} dN^j_s, \quad I^{j\,b,t,I^j}_t = I^j, \quad \forall j \in \{1, \ldots, J\},$$

and

$$dC^{j\,b,t,C^j}_s = \xi^j_{N^j_s} \mathbf{1}_{\left\{b^j_s > p^j_{N^j_s}\right\}} dN^j_s, \quad C^{j\,b,t,C^j}_t = C^j, \quad \forall j \in \{1, \ldots, J\}.$$

The associated Hamilton-Jacobi-Bellman equation is

$$-\partial_t u(t,x,I,C) - \sum_{j=1}^{J} \lambda^j \inf_{b^j \in \mathbb{R}_+} \int_0^{b^j} f^j(p) \left[(1-\nu^j)(u(t,x+p,I+e^j,C) - u(t,x,I,C))\right.$$
$$\left. +\nu^j(u(t,x+p,I+e^j,C+e^j) - u(t,x,I,C))\right] dp = 0, \tag{3.3}$$

with terminal condition

$$u(T,x,I^1,\ldots,I^J,C^1,\ldots,C^J) = x - g(I^1,\ldots,I^J,C^1,\ldots,C^J).$$

Eq. (3.3) is a non-standard integro-differential HJB equation in dimension $2J+2$. In order to find a solution to this equation, we consider the following ansatz:

$$u(t,x,I^1,\ldots,I^J,C^1,\ldots,C^J) = x + \theta_{I,C}(t),$$

where $(\theta_{I,C})_{(I,C)\in\mathbb{N}^J\times\mathbb{N}^J}$ is a family of functions defined on $[0,T]$.

With this ansatz, it is straightforward to see that Eq. (3.3) becomes

$$-\theta'_{I,C}(t) - \sum_{j=1}^{J} \lambda^j \inf_{b^j \in \mathbb{R}_+} \int_0^{b^j} f^j(p) \left[p + (1-\nu^j)(\theta_{I+e^j,C}(t) - \theta_{I,C}(t))\right.$$
$$\left. +\nu^j(\theta_{I+e^j,C+e^j}(t) - \theta_{I,C}(t))\right] dp = 0, \tag{3.4}$$

with terminal condition $\theta_{I,C}(T) = -g(I,C)$.

For each $j \in \{1,\ldots,J\}$, it is straightforward to see that the minimum of

$$b^j \mapsto \int_0^{b^j} f^j(p) \left[p + (1-\nu^j)(\theta_{I+e^j,C}(t) - \theta_{I,C}(t)) + \nu^j(\theta_{I+e^j,C+e^j}(t) - \theta_{I,C}(t))\right] dp$$

is reached at

$$b^{j*}_{I,C}(t) = \left(\theta_{I,C}(t) - \left((1-\nu^j)\theta_{I+e^j,C}(t) + \nu^j\theta_{I+e^j,C+e^j}(t)\right)\right)_+.$$

Therefore, after an integration by parts, Eq. (3.4) becomes

$$-\theta'_{I,C}(t) + \sum_{j=1}^{J} \lambda^j \int_0^{b^{j*}_{I,C}(t)} F^j(p)dp = 0, \tag{3.5}$$

with terminal condition $\theta_{I,C}(T) = -g(I,C)$.

3.2 Solution and special cases

3.2.1 A general result

We now state a theorem in which we prove the existence of a solution to Eq. (3.5) and solve both the pricing problem and the strategic problem faced by the ad trader.

Theorem 3.1. *Let us introduce the set*

$$E = \left\{ (a_{I,C})_{(I,C)\in\mathbb{N}^J\times\mathbb{N}^J}, \|a\| = \sup_{(I,C)\in\mathbb{N}^J\times\mathbb{N}^J} \frac{1}{1+\|I\|_1+\|C\|_1}|a_{I,C}| < +\infty \right\}$$

There exists a unique function $\theta \in C^1([0,T],E)$, such that the family of functions $(\theta_{I,J})_{(I,C)\in\mathbb{N}^J\times\mathbb{N}^J}$ is solution of Eq. (3.5).

The function $u : (t,x,I,C) \mapsto x + \theta_{I,C}(t)$ is the value function U of Eq. (3.2) associated with the control problem (3.1).

The optimal strategy of the ad trader is given in closed-loop by

$$b_t^{j*} = b_{I_{t-},C_{t-}}^{j*}(t) = \left(\theta_{I_{t-},C_{t-}}(t) - \left((1-\nu^j)\theta_{I_{t-}+e^j,C_{t-}}(t) + \nu^j\theta_{I_{t-}+e^j,C_{t-}+e^j}(t)\right)\right)_+.$$

The price of the contract is:

$$P^* = \theta_{0,0}(0).$$

Proof. Let us notice first that $(E,\|\|)$ is a Banach space and that, by assumption, $(-g(I,C))_{I,C} \in E$.

Let us introduce

$$G : (a_{I,C})_{I,C} \in E \mapsto \left(\sum_{j=1}^J \lambda^j \int_0^{b_{I,C}^{j*}} F^j(p)dp\right)_{I,C},$$

where

$$b_{I,C}^{j*} = \left(a_{I,C} - \left((1-\nu^j)a_{I+e^j,C} + \nu^j a_{I+e^j,C+e^j}\right)\right)_+.$$

The first step of the proof is to show that $\forall a \in E, G(a) \in E$.

For that purpose, let us notice that:

$$\begin{aligned} |G(a)_{I,C}| &= \left|\sum_{j=1}^J \lambda^j \int_0^{b_{I,C}^{j*}} F^j(p)dp\right| \\ &\le \sum_{j=1}^J \lambda^j |b_{I,C}^{j*}| \\ &\le \sum_{j=1}^J \lambda^j \left(|a_{I,C}| + (1-\nu^j)|a_{I+e^j,C}| + \nu^j|a_{I+e^j,C+e^j}|\right) \end{aligned}$$

$$\leq \sum_{j=1}^{J} \lambda^j \big(\|a\|(1+\|I\|_1+\|C\|_1)+(1-\nu^j)\|a\|(2+\|I\|_1+\|C\|_1)$$
$$+\nu^j\|a\|(3+\|I\|_1+\|C\|_1)\big)$$

Therefore,

$$\|G(a)\| \leq 4\sum_{j=1}^{J}\lambda^j\|a\| < +\infty,$$

and we have indeed $G : E \to E$.

The second step is to show that $G : E \to E$ is a Lipschitz function.

For that purpose, let us consider $a, \alpha \in E$. We have

$$G(a)_{I,C} - G(\alpha)_{I,C} = \sum_{j=1}^{J}\lambda^j \int_{\beta^{j*}_{I,C}}^{b^{j*}_{I,C}} F^j(p)dp,$$

where

$$b^{j*}_{I,C} = \left(a_{I,C} - \left((1-\nu^j)a_{I+e^j,C} + \nu^j a_{I+e^j,C+e^j}\right)\right)_+$$

and

$$\beta^{j*}_{I,C} = \left(\alpha_{I,C} - \left((1-\nu^j)\alpha_{I+e^j,C} + \nu^j \alpha_{I+e^j,C+e^j}\right)\right)_+ .$$

Therefore,

$$\begin{aligned}
|G(a)_{I,C} - G(\alpha)_{I,C}| &\leq \sum_{j=1}^{J}\lambda^j|\beta^{j*}_{I,C} - b^{j*}_{I,C}| \\
&\leq \sum_{j=1}^{J}\lambda^j\big(|a_{I,C}-\alpha_{I,C}| + (1-\nu^j)|a_{I+e^j,C}-\alpha_{I+e^j,C}| \\
&\qquad +\nu^j|a_{I+e^j,C+e^j}-\alpha_{I+e^j,C+e^j}|\big) \\
&\leq \sum_{j=1}^{J}\lambda^j\|a-\alpha\|\big((1+\|I\|_1+\|C\|_1)+(1-\nu^j)(2+\|I\|_1+\|C\|_1) \\
&\qquad +\nu^j(3+\|I\|_1+\|C\|_1)\big).
\end{aligned}$$

Consequently,

$$\|G(a)-G(\alpha)\| \leq 4\sum_{j=1}^{J}\lambda^j\|a-\alpha\|.$$

Eq. (3.5) is therefore a (backward) Cauchy problem for the function $\theta : [0,T] \to E$, which writes

$$\theta'(t) = G(\theta(t)), \qquad \theta(T) = (-g(I,C))_{I,C},$$

where $G : E \to E$ is a Lipschitz function. By Cauchy-Lipschitz theorem, there is a unique global solution $\theta \in C^1([0,T], E)$ to Eq. (3.5).

Now, we use a verification argument. We consider a given $t \in [0,T)$, and a given bidding strategy $(b^1_s, \ldots, b^J_s)_{s\ge t} \in \mathcal{A}^J_t$. We have:

$$u\left(T, X^{b,t,x}_{T-}, I^{b,t,I}_{T-}, C^{b,t,C}_{T-}\right) = u(t,x,I,C) + \int_t^T \partial_t u\left(s, X^{b,t,x}_{s-}, I^{b,t,I}_{s-}, C^{b,t,C}_{s-}\right) ds$$

$$+ \int_t^T \sum_{j=1}^J \left(u\left(s, X^{b,t,x}_{s-} + p^j_{N^j_s} \mathbf{1}_{\left\{b^j_s > p^j_{N^j_s}\right\}}, I^{b,t,I}_{s-} + \mathbf{1}_{\left\{b^j_s > p^j_{N^j_s}\right\}} e^j, C^{b,t,C}_{s-} + \xi^j_{N^j_s} \mathbf{1}_{\left\{b^j_s > p^j_{N^j_s}\right\}} e^j \right)\right.$$

$$\left. - u\left(s, X^{b,t,x}_{s-}, I^{b,t,I}_{s-}, C^{b,t,C}_{s-}\right)\right) dN^j_s.$$

Therefore,

$$u\left(T, X^{b,t,x}_{T-}, I^{b,t,I}_{T-}, C^{b,t,C}_{T-}\right) = u(t,x,I,C) + \int_t^T \theta'_{I^{b,t,I}_{s-}, C^{b,t,C}_{s-}}(s) ds$$

$$+ \int_t^T \sum_{j=1}^J \left(p^j_{N^j_s} \mathbf{1}_{\left\{b^j_s > p^j_{N^j_s}\right\}} + \theta_{I^{b,t,I}_{s-} + \mathbf{1}_{\left\{b^j_s > p^j_{N^j_s}\right\}} e^j, C^{b,t,C}_{s-} + \xi^j_{N^j_s} \mathbf{1}_{\left\{b^j_s > p^j_{N^j_s}\right\}} e^j}(s) \right.$$

$$\left. - \theta_{I^{b,t,I}_{s-}, C^{b,t,C}_{s-}}(s)\right) dN^j_s.$$

Because – by (H2) – $\forall j \in \{1, \ldots, J\}$, $\forall n \in \mathbb{N}^*$, $p^j_n \in L^1(\Omega)$, and because $\forall s \in [t,T], \theta(s) \in E$, we have:

$$\mathbb{E}\left[X^{b,t,x}_{T-} + \theta_{I^{b,t,I}_{T-}, C^{b,t,C}_{T-}}(T-)\right] = u(t,x,I,C) + \mathbb{E}\left[\int_t^T \theta'_{I^{b,t,I}_{s-}, C^{b,t,C}_{s-}}(s) ds \right.$$

$$+ \int_t^T \sum_{j=1}^J \int_0^{b^j} f^j(p) \left(p^j + (1-\nu^j)\theta_{I^{b,t,I}_{s-} + e^j, C^{b,t,C}_{s-}}(s) + \nu^j \theta_{I^{b,t,I}_{s-} + e^j, C^{b,t,C}_{s-} + e^j}(s)\right.$$

$$\left.\left. - \theta_{I^{b,t,I}_{s-}, C^{b,t,C}_{s-}}(s)\right) dp \lambda^j ds \right].$$

By definition of θ, we have therefore

$$\begin{aligned} \mathbb{E}\left[X^{b,t,x}_T - g\left(I^{b,t,I}_T, C^{b,t,C}_T\right)\right] &= \mathbb{E}\left[X^{b,t,x}_{T-} + \theta_{I^{b,t,I}_{T-}, C^{b,t,C}_{T-}}\right] \\ &\ge u(t,x,I,C), \end{aligned}$$

with equality when $\forall j \in \{1, \ldots, J\}, \forall s \in [t,T), b^j_s = b^{j*}_{I_{s-}, C_{s-}}(s)$.

Therefore, the value function of the control problem is indeed given by

$$u(t,x,I,C) = x + \theta_{I,C}(t),$$

and the optimal bidding strategy is given in closed-loop by

$$b_t^{j*} = b_{I_{t-},C_{t-}}^{j*}(t) = \left(\theta_{I_{t-},C_{t-}}(t) - \left((1-\nu^j)\theta_{I_{t-}+e^j,C_{t-}}(t) + \nu^j\theta_{I_{t-}+e^j,C_{t-}+e^j}(t)\right)\right)_+.$$

Finally, the indifference price of the contract is given by:

$$P^* = u(0,0,0,0) = \theta_{0,0}(0).$$

□

Theorem 3.1 is a general result which deserves a few comments. First, the optimal bidding strategy and the price of the contract depend on a single function θ, which is the solution of an ordinary differential equation (a priori in infinite dimension). Second, in many cases (see below), this ordinary differential equation boils down in fact to a finite system of (1-dimensional) ordinary differential equations that can be solved easily, at least numerically on a grid.

3.2.2 Special cases

Simple inventory-based performance contracts

Let us consider first the special case where

$$g(i^1,\ldots,i^J,c^1,\ldots,c^J) = g_I(i^1+\ldots+i^J).$$

In that case, by using the ansatz $\theta_{I,C}(t) = \tilde{\theta}_{I^1+\ldots+I^J}(t)$, Eq. (3.5) boils down to

$$-\tilde{\theta}'_\iota(t) + \int_0^{\tilde{b}^*_\iota(t)} \sum_{j=1}^J \lambda^j F^j(p)dp = 0, \tag{3.6}$$

with terminal condition $\tilde{\theta}_\iota(T) = -g_I(\iota)$, where

$$\tilde{b}^*_\iota(t) = \left(\tilde{\theta}_\iota(t) - \tilde{\theta}_{\iota+1}(t)\right)_+.$$

In particular, the optimal bidding strategy is the same across all sources of inventory, and it only depends on the total number of impressions already purchased, and on the time to horizon:

$$\forall j \in \{1,\ldots,J\}, b^{j*}_{I,C}(t) = \tilde{b}^*_{I^1+\ldots+I^J}(t).$$

It is also noteworthy that when $g_I(\iota) = -\pi_I(\iota - \overline{I})_-$, Eq. (3.6) boils down to a triangular system of ordinary differential equations indexed by $\iota \in \left\{0,\ldots,\overline{I}\right\}$. This system is nonlinear, but the solution can be approximated very easily on a grid.

Simple conversion-based performance contracts

Let us consider now the special case where

$$g(i^1, \ldots, i^J, c^1, \ldots, c^J) = g_C(c^1 + \ldots + c^J).$$

In that case, by using the ansatz $\theta_{I,C}(t) = \tilde{\theta}_{C^1+\ldots+C^J}(t)$, Eq. (3.5) boils down to

$$-\tilde{\theta}_c'(t) + \sum_{j=1}^{J} \lambda^j \int_0^{\tilde{b}_c^{j*}(t)} F^j(p)dp = 0, \tag{3.7}$$

with terminal condition $\tilde{\theta}_c(T) = -g_C(c)$, where

$$\tilde{b}_c^{j*}(t) = \nu^j \left(\tilde{\theta}_c(t) - \tilde{\theta}_{c+1}(t)\right)_+ .$$

In particular, the optimal bidding strategy is not the same across all sources of inventory. It depends on the total number of conversions already obtained, on the time to horizon, and on the probability of conversion associated with each source:

$$\forall j \in \{1, \ldots, J\}, b_{I,C}^{j*}(t) = \tilde{b}_{C^1+\ldots+C^J}^{j*}(t).$$

What is interesting is that the bidding strategy associated with each source of inventory is proportional to the probability of conversion associated with that source of inventory (see [9] for a similar result in a dual problem).

As above, it is noteworthy that when $g_C(c) = -\pi_C(c - \overline{C})_-$, Eq. (3.7) boils down to a triangular system of ordinary differential equations indexed by $c \in \{0, \ldots, \overline{C}\}$. This system is nonlinear, but the solution can be approximated very easily on a grid.

Inventory-and-conversion-based performance contracts

The third example we consider is related to payoff functions of the form:

$$g(i^1, \ldots, i^J, c^1, \ldots, c^J) = g_I(I^1 + \ldots + I^J) + g_C(c^1 + \ldots + c^J).$$

In that case, by using the ansatz $\theta_{I,C}(t) = \tilde{\theta}_{I^1+\ldots+I^J,C^1+\ldots+C^J}(t)$, Eq. (3.5) boils down to

$$-\tilde{\theta}_{\iota,c}'(t) + \sum_{j=1}^{J} \lambda^j \int_0^{\tilde{b}_{\iota,c}^{j*}(t)} F^j(p)dp = 0, \tag{3.8}$$

with terminal condition $\tilde{\theta}_{\iota,c}(T) = -g_I(\iota) - g_C(c)$, where

$$\tilde{b}_{\iota,c}^{j*}(t) = \left(\tilde{\theta}_{\iota,c}(t) - \left((1-\nu^j)\tilde{\theta}_{\iota+1,c}(t) + \nu^j \tilde{\theta}_{\iota+1,c+1}(t)\right)\right)_+ .$$

The optimal bidding strategy and the price of the contract can be obtained by solving a system of ODEs. The main points here are that the dimension of the indices is reduced from $2J$ to 2, and that there is (a priori) a mixed effect between ι and c.

It is also noteworthy that, when $g_I(\iota) = -\pi_I(\iota - \overline{I})_-$ and $g_C(c) = -\pi_C(c - \overline{C})_-$, Eq. (3.8) boils down to a system of ordinary differential equations indexed by $(\iota, c) \in \{0, \ldots, \overline{I}\} \times \{0, \ldots, \overline{C}\}$. This system is nonlinear, but the solution can be approximated very easily on a grid.

Performance contracts with separable payoffs

The last example we consider is related to payoff functions of the form:

$$g(i^1, \ldots, i^J, c^1, \ldots, c^J) = g_I^1(i^1) + \ldots + g_I^J(i^J) + g_C^1(c^1) + \ldots + g_C^J(c^J).$$

In that case, by using the ansatz $\theta_{I,C}(t) = \tilde{\theta}^1_{I^1,C^1}(t) + \ldots + \tilde{\theta}^J_{I^J,C^J}(t)$, Eq. (3.5) boils down to

$$\forall j \in \{1, \ldots, J\}, \quad -\tilde{\theta}^{j\prime}_{I^j,C^j}(t) + \lambda^j \int_0^{\tilde{b}^{j*}_{I^j,C^j}(t)} F^j(p)dp = 0, \tag{3.9}$$

with terminal condition $\forall j \in \{1, \ldots, J\}, \quad \tilde{\theta}^j_{I^j,C^j}(T) = -g_I^j(I^j) - g_C^j(C^j)$, where

$$\tilde{b}^{j*}_{I^j,C^j}(t) = \left(\tilde{\theta}^j_{I^j,C^j}(t) - \left((1-\nu^j)\tilde{\theta}^j_{I^j+1,C^j}(t) + \nu^j \tilde{\theta}^j_{I^j+1,C^j+1}(t)\right)\right)_+.$$

The optimal bidding strategy for the source j only depends on the number of impressions and on the number of conversions associated with the source j. In other words, the bidding strategy of the different sources are independent. Therefore, everything works as if there were J contracts, one for each source, and the price of the contract is the sum of the prices of these J contracts:

$$P^* = \tilde{\theta}^1_{0,0}(0) + \ldots + \tilde{\theta}^J_{0,0}(0).$$

3.3 Numerical examples

In this section we consider two numerical examples in order to illustrate the use of our model in the case of a risk-neutral intermediary. The two chosen examples represent baseline realistic situations: (i) a contract where the performance is measured in terms of the number of impressions but where the number of auctions is small (for example a campaign targeting a reduced segment of users or a campaign where the ads have an uncommon format), and (ii) a contract where the performance is measured only in terms of the number of conversions. In both cases, in order to simplify the exposition, we consider the simple case where there is only one source of inventory (i.e. $J = 1$).[8]

For each example, we provide the optimal bidding strategy as a function of the remaining time and the relevant state variables (the number of impressions or the number of conversions). Then, we consider Monte Carlo simulations for the auctions and we plot, by using

[8] Given the form of the solutions obtained in the previous paragraphs, this restriction is not a real issue.

the outcomes of the simulations, the empirical distribution of the money spent for the campaign (including the final payoff g) when one uses the optimal strategy. These Monte Carlo simulations enable to assess the risk associated with the contracts: sometimes the ad trader ends up paying more than expected (sometimes less also), because the realizations of the price to beat were higher than expected and/or because the number of conversions turned out to be lower than expected.

3.3.1 Inventory-based performance contracts

The first example we consider is that of a contract "guaranteeing" 5000 impressions over a 10-hour period on a market with low liquidity and expensive inventory. As mentioned above, this first case is relevant when the audience or the inventory is reduced, for instance when one targets a special segment of users, premium inventory or a particular format for the ad creative.

We consider that the average number of auctions per second is equal to 1 (i.e. $\lambda = 1\ \mathrm{s}^{-1}$). We consider that the prices to beat are distributed according to an exponential distribution

$$f(p) = \mu \exp(-\mu p) 1_{p \geq 0},$$

where $\mu = 200\ \$^{-1}$. This corresponds to an average price to beat of \$0.005, i.e. a high average price to beat of \$5 in terms of CPM.

The contract we consider is based on the promise to deliver 5000 impressions over a 10-hour period. In line with the first special case discussed above (inventory-based performance contracts) we consider

$$g_I(\iota) = -\pi_I(\iota - \bar{I})_-,$$

with $\pi_I = \$0.001$, and $\bar{I} = 5000$. In other words, for every impression missing to reach the 5000 target, the ad trader will pay \$0.001 to the advertiser (this corresponds to a maximum CPM of \$1).

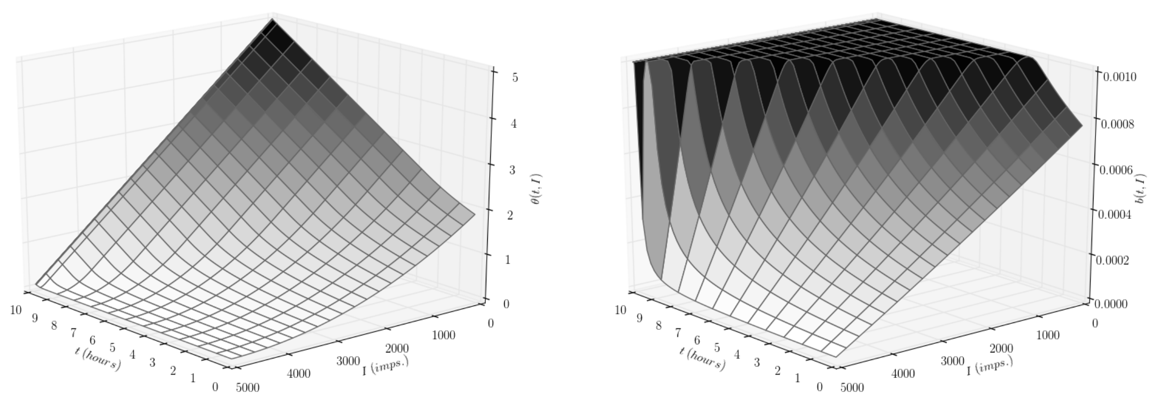

Figure 1: Left: the function $(\iota, t) \mapsto \tilde{\theta}_\iota(t)$. Right: the optimal bidding function $(\iota, t) \mapsto \tilde{b}^*_\iota(t)$.

The function $(\iota, t) \mapsto \tilde{\theta}_\iota(t)$ and the optimal bidding function $(\iota, t) \mapsto \tilde{b}^*_\iota(t)$ are plotted on Figure 1. We see on Figure 1 that there is a cap on the optimal bidding strategy. This cap, equal to \$0.001, corresponds to π_I, the penalty to be paid for each missing impression.

Figure 2 shows the evolution of the number of impressions, the cash spent and the bid level for a sample execution when one uses the optimal strategy.

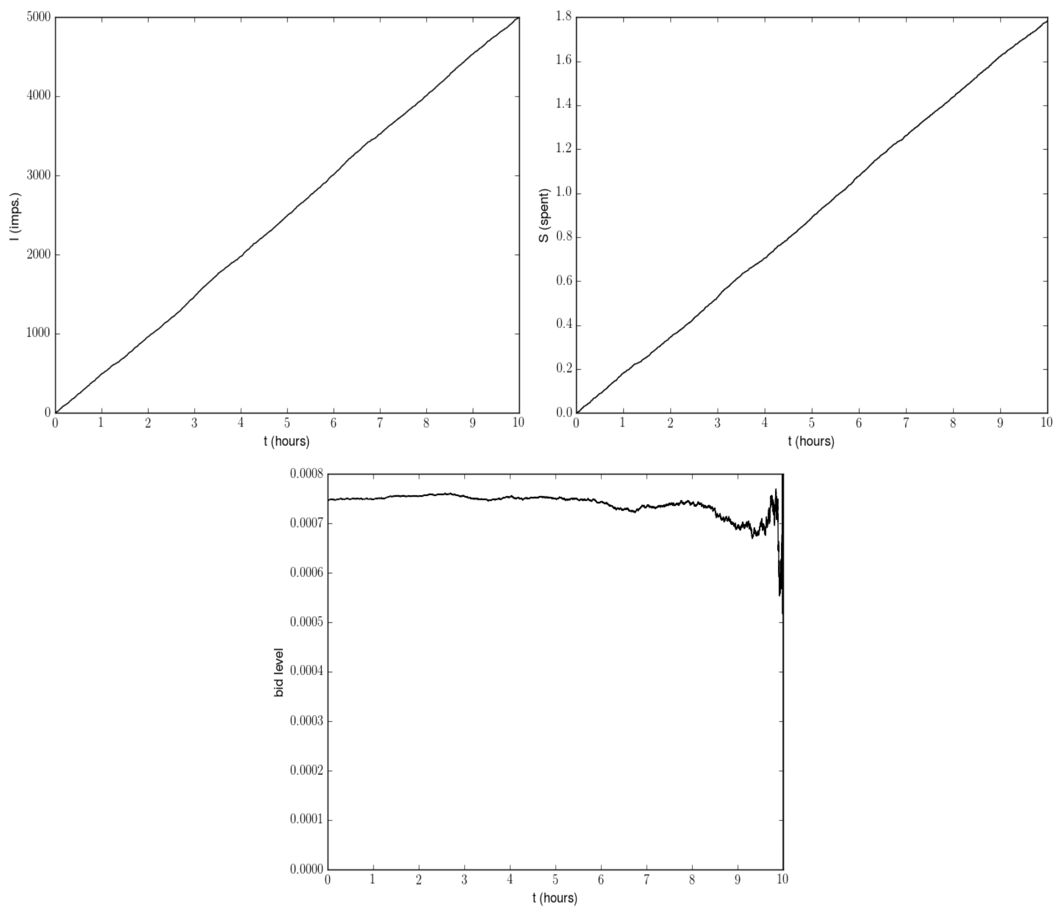

Figure 2: Evolution of the number of impressions, the cash spent and the bid level for an inventory-based performance contract.

The indifference price is here equal to $\tilde{\theta}_0(0) = \$1.8251$. By using Monte Carlo simulations (with 10000 draws) with the optimal bidding strategy found earlier, we recover an average spent of \$1.8249, and the standard deviation of the money spent is \$0.03122. The distribution of the money spent is plotted on Figure 3. On average the algorithm is able to buy 4997.78 impressions, with a standard deviation of 3.27 impressions.

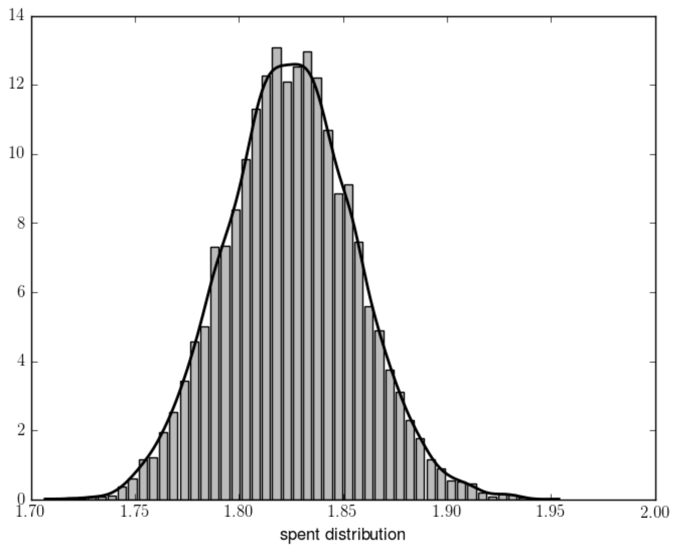

Figure 3: Distribution of the money spent over $[0,T]$ when the optimal bidding strategy is used – Monte Carlo simulations with 10000 draws.

3.3.2 Conversion-based performance contracts

The second example we consider is that of a contract "guaranteeing" 5000 conversions over a 10-hour period on a market with high liquidity and cheap inventory.

We consider that the average number of auctions per second is equal to 500 (i.e. $\lambda = 500\ \mathrm{s}^{-1}$). We consider that the prices to beat are distributed according to an exponential distribution

$$f(p) = \mu \exp(-\mu p) 1_{p\geq 0},$$

where $\mu = 1000\ \$^{-1}$. This corresponds to an average price to beat of \$0.001, i.e. an average price to beat of \$1 in terms of CPM.

The contract we consider is based on the promise to obtain 5000 conversions (here clicks) over a 10-hour period. The probability of conversion is $\nu = 0.001$. In line with the second special case discussed above (conversion-based performance contracts) we consider

$$g_C(c) = -\pi_C(c - \bar{C})_-,$$

with $\pi_C = \$0.4$, and $\bar{C} = 5000$. In other words, for each conversion missing to reach the 5000 target, the ad trader will pay \$0.4 to the advertiser (this corresponds to a maximum CPC or CPA of \$0.4).

The function $(c,t) \mapsto \tilde{\theta}_c(t)$ and the optimal bidding function $(c,t) \mapsto \tilde{b}_c^*(t)$ are plotted on Figure 4. We see on Figure 4 that there is a cap on the optimal bidding strategy. This cap is related to the probability of conversion and to the penalty to be paid for each missing conversion.

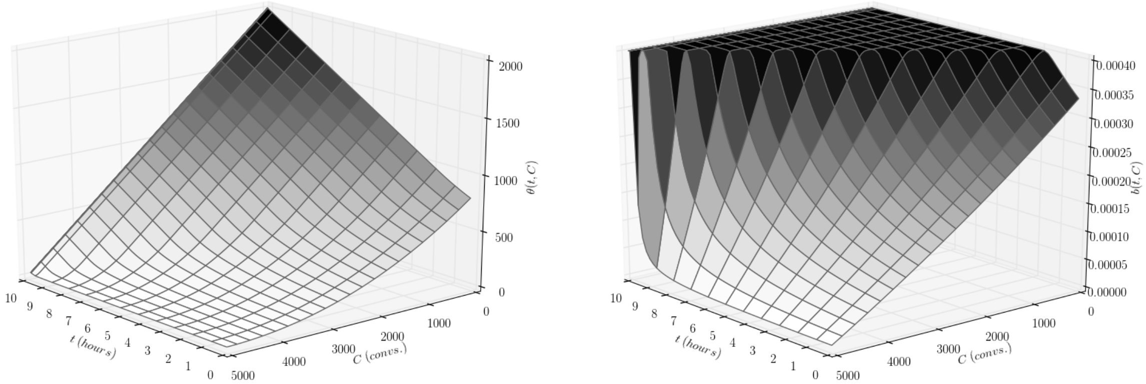

Figure 4: Left: the function $(c, t) \mapsto \tilde{\theta}_c(t)$. Right: the optimal bidding function $(c, t) \mapsto \tilde{b}^*_c(t)$.

Figure 5 shows the evolution of the number of impressions, the number of conversions, the cash spent and the bid level for a sample execution when one uses the optimal strategy.

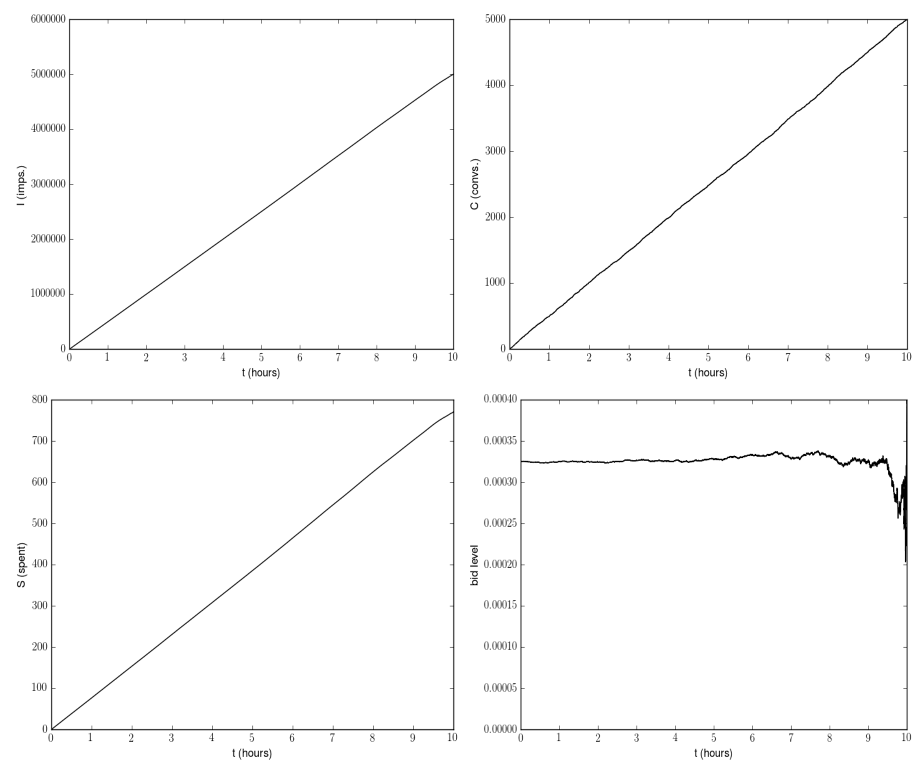

Figure 5: Evolution of the number of impressions, the number of conversions, the cash spent and the bid level for a conversion-based performance contract.

The indifference price $\tilde{\theta}_0(0)$ computed numerically is equal to \$771.61. By using Monte

Carlo simulations (10000 draws) with the optimal bidding strategy found numerically, we recover an average spent of \$758.91, and the standard deviation of the money spent is \$21.64. The distribution of the money spent is plotted in Figure 6. The average number of conversions obtained is 4997.11 with a standard deviation of 4.048 conversions.

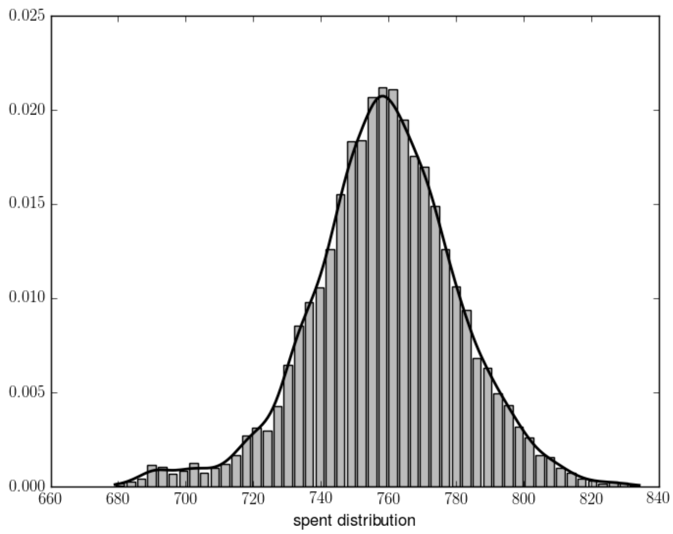

Figure 6: Distribution of the money spent over $[0, T]$ when the optimal bidding strategy is used – Monte Carlo simulations with 10000 draws.

As above, we see that, from the point of view of the ad trader (the intermediary), there is some risk in accepting to sign performance-based contracts. This is the reason why we believe that maximizing the expected value of the payoff is insufficient. In practice, agents are risk-averse and risk aversion should be taken into account, both in the pricing and in the bidding strategy.

4 Introducing risk aversion

In the previous section, we have only considered the case of a risk-neutral trader who minimizes his costs. In practice, one must take risk into account. For that purpose, we consider the case of a (risk-averse) trader who maximizes an expected utility criterion corresponding to a constant absolute risk aversion coefficient γ.

In our framework, this corresponds to the following optimization problem

$$\sup_{(b_t^1,\ldots,b_t^J)_t \in \mathcal{A}^J} \mathbb{E}\left[-\exp\left(-\gamma\left(-X_T + g(I_T^1,\ldots,I_T^J,C_T^1,\ldots,C_T^J)\right)\right)\right],$$

or equivalently to

$$\inf_{(b_t^1,\ldots,b_t^J)_t \in \mathcal{A}^J} \mathbb{E}\left[\exp\left(\gamma\left(X_T - g(I_T^1,\ldots,I_T^J,C_T^1,\ldots,C_T^J)\right)\right)\right]. \tag{4.1}$$

To solve this problem, we assume, in addition to the above assumptions (H1), (H2) and (H3) on the law of the *price to beat*, that:

$$\text{(H4)} \quad \forall j \in \{1, \ldots, J\}, \forall \phi > 0, \int_0^{+\infty} e^{\phi p} f^j(p) dp < +\infty.$$

4.1 Hamilton-Jacobi equation and reduction to a system of ODEs

The value function associated with this problem is the function:

$$U : (t, x, I, C) \in [0, T] \times \mathbb{R}_+ \times \mathbb{N}^J \times \mathbb{N}^J \mapsto \inf_{(b_s^1, \ldots, b_s^J)_{s \geq t} \in \mathcal{A}_t^J} \mathbb{E}\left[\exp\left(\gamma\left(X_T^{b,t,x} - g(I_T^{b,t,I}, C_T^{b,t,C})\right)\right)\right], \tag{4.2}$$

The associated Hamilton-Jacobi-Bellman equation is:

$$\begin{aligned} -\partial_t u(t, x, I, C) - \sum_{j=1}^J \lambda^j \inf_{b^j \in \mathbb{R}_+} \int_0^{b^j} f^j(p) \left[(1 - \nu^j)(u(t, x + p, I + e^j, C) - u(t, x, I, C))\right. \\ \left. + \nu^j (u(t, x + p, I + e^j, C + e^j) - u(t, x, I, C))\right] dp = 0, \end{aligned} \tag{4.3}$$

with terminal condition

$$u(T, x, I^1, \ldots, I^J, C^1, \ldots, C^J) = \exp\left(\gamma\left(x - g(I^1, \ldots, I^J, C^1, \ldots, C^J)\right)\right).$$

Eq. (4.3) is a non-standard integro-differential HJB equation in dimension $2J + 2$. In order to find a solution to this equation, we consider the following ansatz:

$$u(t, x, I^1, \ldots, I^J, C^1, \ldots, C^J) = \exp(\gamma(x + \theta_{I,C}(t))),$$

where $(\theta_{I,C})_{(I,C) \in \mathbb{N}^J \times \mathbb{N}^J}$ is a family of functions defined on $[0, T]$.

With this ansatz, it is straightforward to see that Eq. (4.3) becomes:

$$\begin{aligned} -\gamma \theta'_{I,C}(t) - \sum_{j=1}^J \lambda^j \inf_{b^j \in \mathbb{R}_+} \int_0^{b^j} f^j(p) \left[e^{\gamma p}\left((1 - \nu^j) e^{\gamma(\theta_{I+e^j,C}(t) - \theta_{I,C}(t))}\right.\right. \\ \left.\left. + \nu^j e^{\gamma(\theta_{I+e^j,C+e^j}(t) - \theta_{I,C}(t))}\right) - 1\right] dp = 0, \end{aligned} \tag{4.4}$$

with terminal condition $\theta_{I,C}(T) = -g(I, C)$.

For each $j \in \{1, \ldots, J\}$, it is straightforward to see that the minimum of

$$\int_0^{b^j} f^j(p) \left[e^{\gamma p}\left((1 - \nu^j) e^{\gamma(\theta_{I+e^j,C}(t) - \theta_{I,C}(t))} + \nu^j e^{\gamma(\theta_{I+e^j,C+e^j}(t) - \theta_{I,C}(t))}\right) - 1\right] dp$$

is reached at

$$b_{I,C}^{j*}(t) = \left(-\frac{1}{\gamma} \log\left(\eta_{I,C}^{j*}(\theta(t))\right)\right)_+,$$

where

$$\eta_{I,C}^{j*}(\theta(t)) = (1 - \nu^j) e^{\gamma(\theta_{I+e^j,C}(t) - \theta_{I,C}(t))} + \nu^j e^{\gamma(\theta_{I+e^j,C+e^j}(t) - \theta_{I,C}(t))}$$

Therefore, Eq. (4.4) can be written in two different manners:

1. by using the value at which the minimum of the integral is reached:

$$-\gamma\theta'_{I,C}(t) - \sum_{j=1}^{J} \lambda^j \int_0^{\left(-\frac{1}{\gamma}\log\left(\eta^{j*}_{I,C}(\theta(t))\right)\right)_+} f^j(p)\left[e^{\gamma p}\eta^{j*}_{I,C}(\theta(t)) - 1\right] dp = 0, \qquad (4.5)$$

2. or by using an integration by parts:

$$-\theta'_{I,C}(t) + \sum_{j=1}^{J} \lambda^j \eta^{j*}_{I,C}(\theta(t)) \int_0^{\left(-\frac{1}{\gamma}\log\left(\eta^{j*}_{I,C}(\theta(t))\right)\right)_+} F^j(p) e^{\gamma p} dp = 0, \qquad (4.6)$$

in both cases with terminal condition $\theta_{I,C}(T) = -g(I,C)$.

4.2 Solution and special cases

4.2.1 A general result

We now state a theorem in which we prove the existence of a solution to Eq. (4.4) and solve both the pricing problem and the strategic problem faced by the ad trader.

Theorem 4.1. *As in Theorem 3.1, we consider the set*

$$E = \left\{ (a_{I,C})_{(I,C)\in\mathbb{N}^J\times\mathbb{N}^J}, \|a\| = \sup_{(I,C)\in\mathbb{N}^J\times\mathbb{N}^J} \frac{1}{1+\|I\|_1+\|C\|_1}|a_{I,C}| < +\infty \right\}$$

There exists a unique function $\theta \in C^1([0,T],E)$, such that the family of functions $(\theta_{I,J})_{(I,C)\in\mathbb{N}^J\times\mathbb{N}^J}$ is solution of Eq. (4.4).

The function $u : (t,x,I,C) \mapsto \exp(\gamma(x+\theta_{I,C}(t)))$ is the value function U of Eq. (4.2) associated with the control problem (4.1).

The optimal strategy of the ad trader is given by

$$b^{j*}_{I,C}(t) = \left(-\frac{1}{\gamma}\log\left((1-\nu^j)e^{\gamma(\theta_{I+e^j,C}(t)-\theta_{I,C}(t))} + \nu^j e^{\gamma(\theta_{I+e^j,C+e^j}(t)-\theta_{I,C}(t))}\right)\right)_+.$$

The price of the contract is:

$$P^* = \frac{1}{\gamma}\log(u(0,0,0,0)) = \theta_{0,0}(0).$$

Proof. As above, by assumption, $(-g(I,C))_{I,C}$ is in the Banach space E.

Let us introduce

$$\eta^j : (a_{I,C})_{I,C} \in E \mapsto \left((1-\nu^j)e^{\gamma(a_{I+e^j,C}-a_{I,C})} + \nu^j e^{\gamma(a_{I+e^j,C+e^j}-a_{I,C})}\right)_{I,C}.$$

Let us introduce

$$G : (a_{I,C})_{I,C} \in E \mapsto \left(\sum_{j=1}^{J} \lambda^j \eta^j(a)_{I,C} \int_0^{b^{j*}_{I,C}} e^{\gamma p} F^j(p) dp \right)_{I,C},$$

where

$$b^{j*}_{I,C} = \left(-\frac{1}{\gamma} \log\left(\eta^j(a)_{I,C}\right) \right)_+.$$

The first step of the proof is to show that $\forall a \in E, G(a) \in \ell^\infty \subset E$.

For that purpose, let us notice that:

- if $\eta^j(a)_{I,C} \geq 1$ then $b^{j*}_{I,C} = 0$ and the j^{th} term of the sum in the definition of $G(a)_{I,C}$ is equal to 0,
- otherwise,

$$\begin{aligned} & \eta^j(a)_{I,C} \int_0^{b^{j*}_{I,C}} e^{\gamma p} F^j(p) dp \\ = \; & \eta^j(a)_{I,C} \int_1^{e^{\gamma b^{j*}_{I,C}}} F^j\left(\frac{1}{\gamma}\log(q)\right) dq \\ = \; & \eta^j(a)_{I,C} \int_0^{e^{\gamma b^{j*}_{I,C}}} F^j\left(\frac{1}{\gamma}\log(q)\right) dq \\ = \; & \frac{1}{\gamma} \eta^j(a)_{I,C} \int_0^{\frac{1}{\eta^j(a)_{I,C}}} F^j\left(\frac{1}{\gamma}\log(q)\right) dq \\ \leq \; & \frac{1}{\gamma}. \end{aligned}$$

Because $G(a) \geq 0$, we know therefore that

$$\|G(a)\|_\infty \leq \sum_{j=1}^{J} \frac{\lambda^j}{\gamma}.$$

Now, it is straightforward to see that $G : E \to E$ is a locally Lipschitz function.

As a consequence, Eq. (4.4) is a (backward) Cauchy problem for the function $\theta : [0, T] \to E$ which writes

$$\theta'(t) = G(\theta(t)), \qquad \theta(T) = (-g(I, C))_{I,C}$$

where $G : E \to E$ is a locally Lipschitz function. By Cauchy-Lipschitz theorem, there is a unique local solution θ to Eq. (4.4).

But, because $G(\theta(t)) \geq 0$ and $\|G(\theta(t))\|_\infty \leq \sum_{j=1}^{J} \frac{\lambda^j}{\gamma}$, the solution is global.

Now, by using a similar verification argument as in Theorem 3.1 (here the hypothesis (H4) is required), we easily prove that the value function of the control problem is indeed given by

$$u(t,x,I,C) = \exp(\gamma(x + \theta_{I,C}(t))),$$

and the optimal bidding strategy is given by

$$b^{j*}_{I,C}(t) = \left(-\frac{1}{\gamma} \log \left((1-\nu^j) e^{\gamma(\theta_{I+e^j,C}(t) - \theta_{I,C}(t))} + \nu^j e^{\gamma(\theta_{I+e^j,C+e^j}(t) - \theta_{I,C}(t))} \right) \right)_+ .$$

Finally, the indifference price of the contract, as defined in Eq. (2.1), is given by:

$$P^* = \theta_{0,0}(0).$$

□

Theorem 4.1 is a general result. As in the risk-neutral case of Section 3, the optimal bidding strategy and the price of the contract depend on a single function θ, which is the solution of an ordinary differential equation (in infinite dimension). As above, in many cases (see below), this equation boils down to a finite system of differential equations that can be solved easily, at least numerically on a grid.

4.2.2 Special cases

Simple inventory-based performance contracts

Let us consider first the special case where

$$g(i^1, \ldots, i^J, c^1, \ldots, c^J) = g_I(i^1 + \ldots + i^J).$$

In that case, by using the ansatz $\theta_{I,C}(t) = \tilde{\theta}_{I^1+\ldots+I^J}(t)$, Eq. (4.6) boils down to

$$-\tilde{\theta}'_\iota(t) + \int_0^{\tilde{b}^*_\iota(t)} \sum_{j=1}^J \lambda^j F^j(p) e^{\gamma(p + \tilde{\theta}_{\iota+1}(t) - \tilde{\theta}_\iota(t))} dp = 0, \tag{4.7}$$

with terminal condition $\tilde{\theta}_\iota(T) = -g_I(\iota)$, where

$$\tilde{b}^*_\iota(t) = \left(\tilde{\theta}_\iota(t) - \tilde{\theta}_{\iota+1}(t) \right)_+ .$$

In particular, the optimal bidding strategy is the same across all sources of inventory, and it only depends on the total number of impressions already purchased, and on the time to horizon:

$$\forall j \in \{1, \ldots, J\}, b^{j*}_{I,C}(t) = \tilde{b}^*_{I^1+\ldots+I^J}(t).$$

It is also noteworthy that when $g_I(\iota) = -\pi_I(\iota - \overline{I})_-$, Eq. (4.7) boils down to a triangular system of ordinary differential equations indexed by $\iota \in \left\{0, \ldots, \overline{I}\right\}$. This system is nonlinear, but the solution can be approximated very easily on a grid.

Simple conversion-based performance contracts

Let us consider now the special case where

$$g(i^1, \ldots, i^J, c^1, \ldots, c^J) = g_C(c^1 + \ldots + c^J).$$

In that case, by using the ansatz $\theta_{I,C}(t) = \tilde{\theta}_{C^1+\ldots+C^J}(t)$, Eq. (4.6) boils down to

$$-\tilde{\theta}'_c(t) + \sum_{j=1}^{J} \lambda^j \left((1-\nu^j) + \nu^j e^{\gamma\left(\tilde{\theta}_{c+1}(t) - \tilde{\theta}_c(t)\right)} \right) \int_0^{\tilde{b}^{j*}_c(t)} F^j(p) e^{\gamma p} dp = 0, \tag{4.8}$$

with terminal condition $\tilde{\theta}_c(T) = -g_C(c)$, where

$$\tilde{b}^{j*}_c(t) = \left(-\frac{1}{\gamma} \log \left((1-\nu^j) + \nu^j e^{\gamma\left(\tilde{\theta}_{c+1}(t) - \tilde{\theta}_c(t)\right)} \right) \right)_+.$$

In particular, the optimal bidding strategy is not the same across all sources of inventory. It depends on the total number of conversions already obtained, on the time to horizon, and on the probability of conversion associated with each source:

$$\forall j \in \{1, \ldots, J\}, b^{j*}_{I,C}(t) = \tilde{b}^{j*}_{C^1+\ldots+C^J}(t).$$

As above, it is noteworthy that when $g_C(c) = -\pi_C(c - \overline{C})_-$, Eq. (4.8) boils down to a triangular system of ordinary differential equations indexed by $c \in \left\{0, \ldots, \overline{C}\right\}$. This system is nonlinear, but the solution can be approximated very easily on a grid.

Inventory-and-conversion-based performance contracts

The third example we consider is that of payoff functions of the form:

$$g(i^1, \ldots, i^J, c^1, \ldots, c^J) = g_I(I^1 + \ldots + I^J) + g_C(c^1 + \ldots + c^J).$$

In that case, by using the ansatz $\theta_{I,C}(t) = \tilde{\theta}_{I^1+\ldots+I^J, C^1+\ldots+C^J}(t)$, Eq. (4.6) boils down to

$$-\tilde{\theta}'_{\iota,c}(t) + \sum_{j=1}^{J} \lambda^j \left((1-\nu^j) e^{\gamma\left(\tilde{\theta}_{\iota+1,c}(t) - \tilde{\theta}_{\iota,c}(t)\right)} + \nu^j e^{\gamma\left(\tilde{\theta}_{\iota+1,c+1}(t) - \tilde{\theta}_{\iota,c}(t)\right)} \right) \int_0^{\tilde{b}^{j*}_{\iota,c}(t)} F^j(p) e^{\gamma p} dp = 0, \tag{4.9}$$

with terminal condition $\tilde{\theta}_{\iota,c}(T) = -g_I(\iota) - g_C(c)$, where

$$\tilde{b}^{j*}_{\iota,c}(t) = \left(-\frac{1}{\gamma} \log \left((1-\nu^j) e^{\gamma\left(\tilde{\theta}_{\iota+1,c}(t) - \tilde{\theta}_{\iota,c}(t)\right)} + \nu^j e^{\gamma\left(\tilde{\theta}_{\iota+1,c+1}(t) - \tilde{\theta}_{\iota,c}(t)\right)} \right) \right)_+.$$

The optimal bidding strategy and the price of the contract can be obtained by solving a system of ODEs. The main points here are that the dimension of the indices is reduced from $2J$ to 2, and that there is (a priori) a mixed effect between ι and c.

It is also noteworthy that, when $g_I(\iota) = -\pi_I(\iota - \overline{I})_-$ and $g_C(c) = -\pi_C(c - \overline{C})_-$, Eq. (4.9) boils down to a system of ordinary differential equations indexed by $(\iota, c) \in \left\{0, \ldots, \overline{I}\right\} \times$

$\{0, \ldots, \overline{C}\}$. This system is nonlinear, but the solution can be approximated very easily on a grid.

Performance contracts with separable payoffs

The last example we consider is related to payoff functions of the form:

$$g(i^1, \ldots, i^J, c^1, \ldots, c^J) = g_I^1(i^1) + \ldots + g_I^J(i^J) + g_C^1(c^1) + \ldots + g_C^J(c^J).$$

In that case, by using the ansatz $\theta_{I,C}(t) = \tilde{\theta}^1_{I^1,C^1}(t) + \ldots + \tilde{\theta}^J_{I^J,C^J}(t)$, Eq. (4.6) boils down to

$$\forall j \in \{1, \ldots, J\}, \quad 0 = -\tilde{\theta}^{j\prime}_{I^j,C^j}(t)$$

$$+\lambda^j \left((1-\nu^j) e^{\gamma\left(\tilde{\theta}^j_{I^j+1,C^j}(t) - \tilde{\theta}^j_{I^j,C^j}(t)\right)} + \nu^j e^{\gamma\left(\tilde{\theta}^j_{I^j+1,C^j+1}(t) - \tilde{\theta}^j_{I^j,C^j}(t)\right)} \right) \int_0^{\tilde{b}^{j*}_{I^j,C^j}(t)} F^j(p) e^{\gamma p} dp, \tag{4.10}$$

with terminal condition $\forall j \in \{1, \ldots, J\}, \quad \tilde{\theta}^j_{I^j,C^j}(T) = -g_I^j(I^j) - g_C^j(C^j)$, where

$$\tilde{b}^{j*}_{I^j,C^j}(t) = \left(-\frac{1}{\gamma} \log \left((1-\nu^j) e^{\gamma\left(\tilde{\theta}^j_{I^j+1,C^j}(t) - \tilde{\theta}^j_{I^j,C^j}(t)\right)} + \nu^j e^{\gamma\left(\tilde{\theta}_{I^j+1,C^j+1}(t) - \tilde{\theta}_{I^j,C^j}(t)\right)} \right) \right)_+ .$$

The optimal bidding strategy for the source j only depends on the number of impressions and on the number of conversions associated with the source j. In other words, the bidding strategy of the different sources are independent. Therefore, everything works as if there were J contracts, one for each source, and the price of the contract is the sum of the prices of these J contracts:

$$P^* = \tilde{\theta}^1_{0,0}(0) + \ldots + \tilde{\theta}^J_{0,0}(0).$$

4.3 Numerical examples

We provide numerical examples to illustrate how the use of our risk-averse objective function affects the results obtained in the risk-neutral case in Section 3.3.

4.3.1 Inventory-based performance contracts

We consider again the contract "guaranteeing" 5000 impressions over a 10-hour period on a market with low liquidity and expensive inventory. The set of parameters remains the same as before: $\lambda = 1\ \mathrm{s}^{-1}$, $\mu = 200\ \$^{-1}$ and $\pi_I = \$0.001$. The risk-aversion parameter is taken to be $\gamma = 100$. The function $(\iota, t) \mapsto \tilde{\theta}_\iota(t)$ and the optimal bidding function $(\iota, t) \mapsto \tilde{b}^*_\iota(t)$ are plotted on Figure 7. Numerically, we obtain a price equal to $\tilde{\theta}_0(0) = \$1.8758$, to be compared to the lower price \$1.8251 obtained in the risk-neutral case.

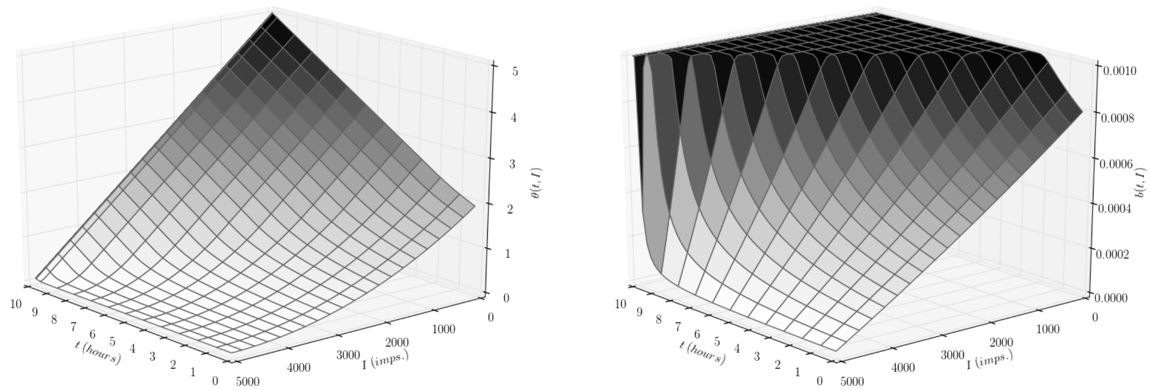

Figure 7: Left: the function $(\iota, t) \mapsto \tilde{\theta}_\iota(t)$. Right: the optimal bidding function $(\iota, t) \mapsto \tilde{b}^*_\iota(t)$.

Figure 8 shows the evolution of the number of impressions, the cash spent and the bid level for a sample execution when one uses the optimal strategy.

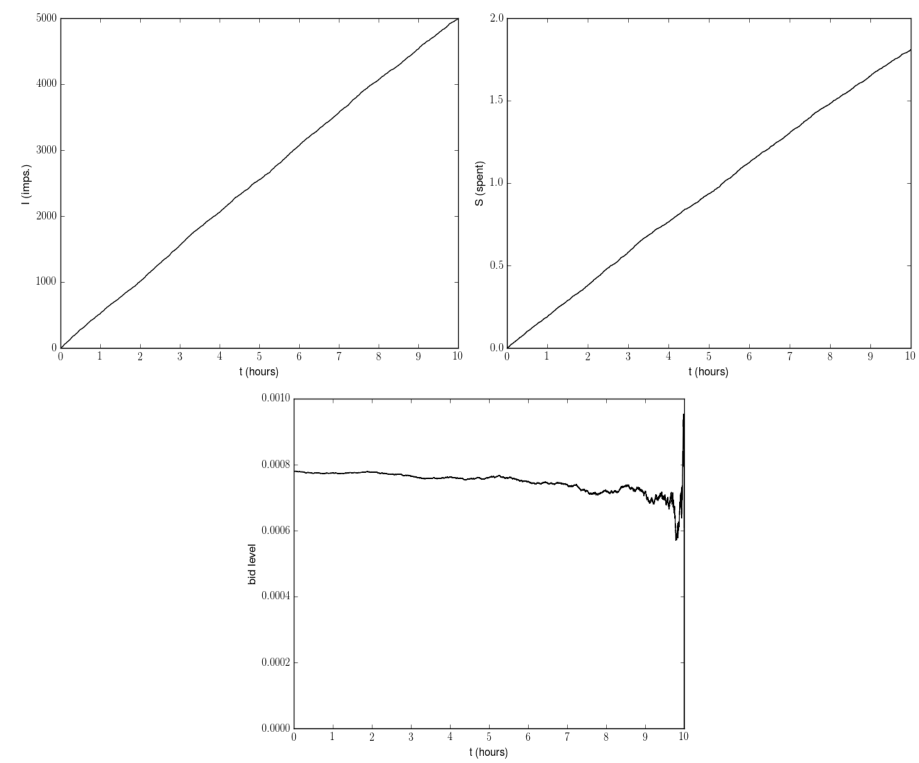

Figure 8: Evolution of the number of impressions, the cash spent and the bid level for an inventory-based performance contract.

Figure 9 shows the results of our Monte Carlo simulations (with 10000 trajectories). The average amount spent in this case is \$1.8274, with a standard deviation of \$0.0291. In this risk-averse situation the number of impressions purchased is in average 4999.11 with a standard deviation of 1.6722 impressions, closer to the target than in the risk-neutral case.

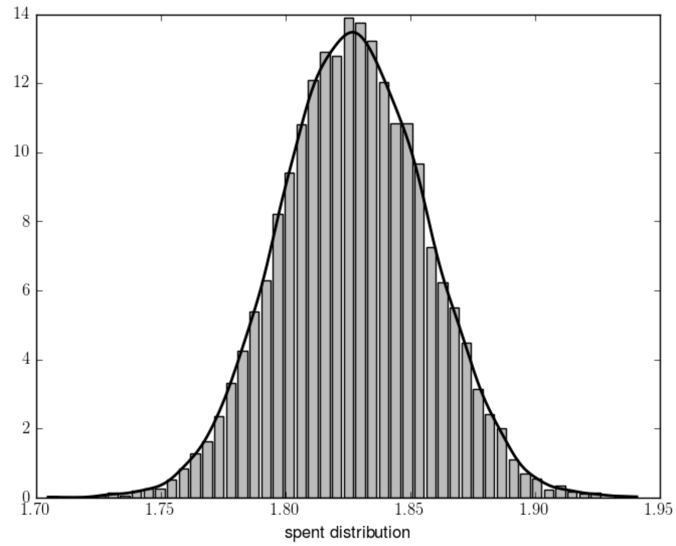

Figure 9: Monte Carlo simulations for the inventory-based performance contract in the risk-averse case.

4.3.2 Conversion-based performance contracts

As in the risk-neutral case, the second example we consider is that of a contract "guaranteeing" 5000 conversions over a 10-hour period on a market with high liquidity and cheap inventory. The set of parameters is the same as before: $\lambda = 500\ \mathrm{s}^{-1}$, $\mu = 1000\ \$^{-1}$ and $\pi_C = \$0.4$. The risk aversion parameter is taken to be $\gamma = 0.3$ (a smaller value than in the first example, because the value of the contract is far larger).

The function $(c,t) \mapsto \tilde{\theta}_c(t)$ and the optimal bidding function $(c,t) \mapsto \tilde{b}^*_c(t)$ are plotted on Figure 10. We see that, in the risk-averse case, the optimal bidding function is steeper than in the risk-neutral case: the ad trader chooses higher bids.

The price of the contract obtained numerically is $\tilde{\theta}_0(0) = \$868.25$, a higher value than in the risk-neutral case (\$770.61).

Figure 11 shows the evolution of the number of impressions, the number of conversions and the cash spent for a sample execution when one uses the optimal strategy. In the risk-averse case, the ad trader does not spend evenly throughout the day as in the risk-neutral case. Instead, he buys more at the beginning because he wants to reduce the risk he faces. In particular, the ad trader bids higher prices in order to buy inventory at a higher pace at the beginning, and this results in a higher value for the cash spent on average – this can be seen on Figure 12, which shows the results of a Monte Carlo simulation (10000 trajectories). The

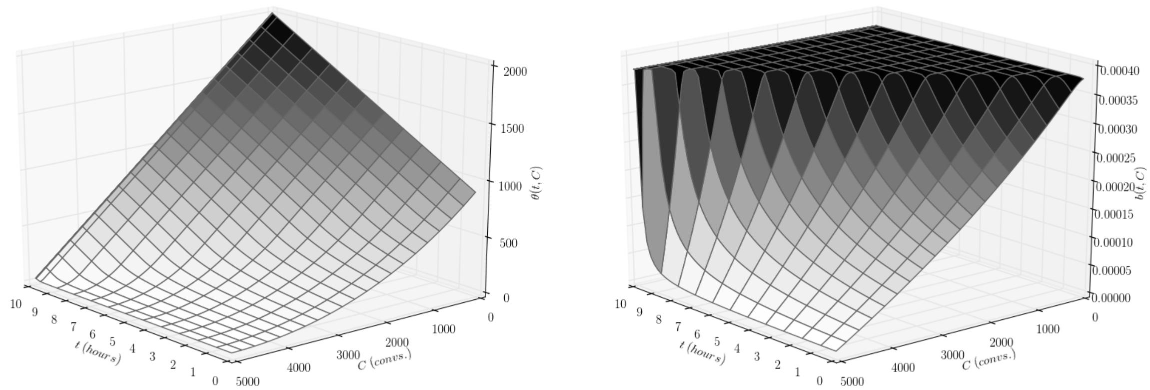

Figure 10: Left: the function $(c, t) \mapsto \tilde{\theta}_c(t)$. Right: the optimal bidding function $(c, t) \mapsto \tilde{b}^*_c(t)$.

average amount of money spent in the simulation is \$768.82, with a standard deviation of \$19.98. The average number of conversions in this case is 4999.58 with a standard deviation of 0.99 conversions: this is closer to the target than in the risk-neutral case.

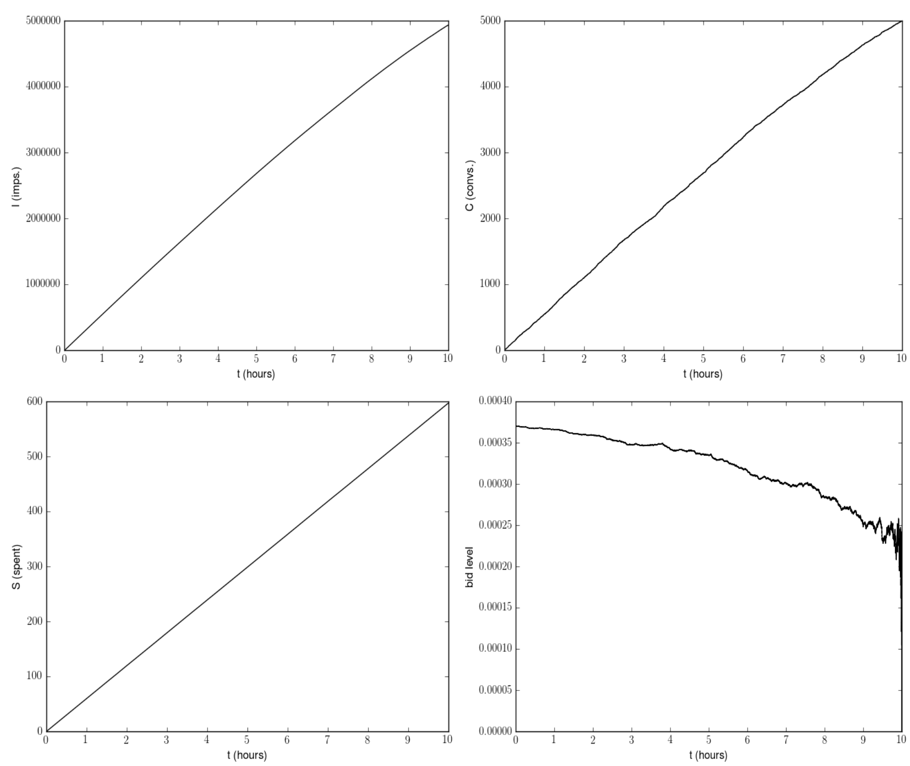

Figure 11: Evolution of the number of impressions, the number of conversions, the cash spent and the bid level for a conversion-based performance contract.

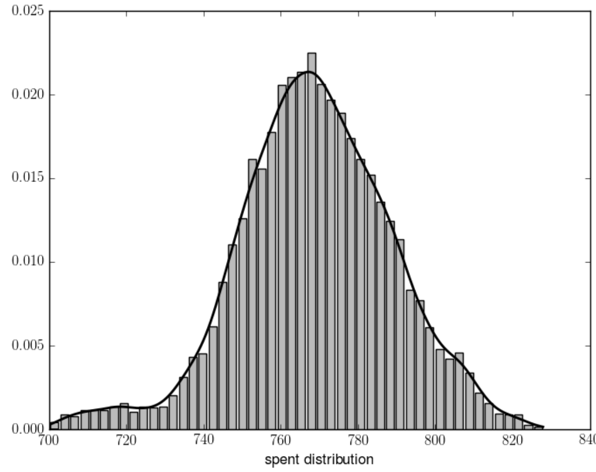

Figure 12: Monte Carlo simulations for the conversions-based performance contract in the risk-averse case.

5 Conclusion

In this article we have built a quantitative framework for the pricing and risk-management of ad-buying services, more specifically, in the case of performance-based contracts in real-time bidding (RTB). Our approach is based on a stochastic model for RTB auctions which incorporates the main features we encounter in real applications (continuous stream of auctions at random times, second-price auction dynamics, several possible performance goals, etc.).

Besides the optimization framework (based on stochastic optimal control techniques), the main contribution of this research paper is the introduction of a rigorous framework for defining and pricing performance-based contracts. Our framework relies on techniques from financial economics, in particular the indifference pricing approach, usually applied for the pricing of contingent claims. Another contribution is the use of Monte Carlo techniques for measuring risk.

Our approach opens the door to several lines of future research: financial and economic works about performance-based contracts (performance-elasticity of the different model parameters, more sophisticated models, dynamic model parameters, etc.), mathematical research about numerical analysis and Monte Carlo simulation issues [4, 12], and also research for the pricing of performance-based contracts when the parameters (or even the model) are unknown or uncertain – on this topic, reinforcement learning techniques (see [2, 13]) should be used, see [10].

References

[1] Asdemir, K., Kumar, N., & Jacob, V. S. (2012). Pricing models for online advertising: CPM vs. CPC. Information Systems Research, 23(3-part-1), 804-822.

[2] Azar, M. G., Gomez, V., & Kappen, H. J. (2012). Dynamic policy programming. The Journal of Machine Learning Research, 13(1), 3207-3245.

[3] Balseiro, S., & Candogan, O. (2015). Optimal contracts for intermediaries in online advertising. Available at SSRN 2546609.

[4] Brandimarte, P. (2014). Handbook in Monte Carlo simulation: applications in financial engineering, risk management, and economics. John Wiley & Sons.

[5] Chen, B., Yuan, S., & Wang, J. (2014, August). A dynamic pricing model for unifying programmatic guarantee and real-time bidding in display advertising. In Proceedings of the Eighth International Workshop on Data Mining for Online Advertising (pp. 1-9). ACM.

[6] Carmona, R. (2009). Indifference pricing: theory and applications. Princeton University Press.

[7] Fernandez-Tapia, J. (2015) An analytical solution to the budget-pacing problem in programmatic advertising. Technical report.

[8] Fernandez-Tapia, J. (2015). Statistical modeling of Vickrey auctions and applications to automated bidding strategies. To appear in Optimization Letters.

[9] Fernandez-Tapia, J., Guéant, O., & Lasry, J. M. (2016). Optimal Real-Time Bidding Strategies. To appear in AMRX, Oxford University Press, 2017

[10] Fernandez-Tapia, J., Guéant, O., & Lasry, J. M. (2016). Real-time bidding strategies with on-line estimation. Working paper.

[11] Guéant, O. (2016). The Financial Mathematics of Market Liquidity: From Optimal Execution to Market Making. CRC Press.

[12] Kushner, H., & Dupuis, P. G. (2013). Numerical methods for stochastic control problems in continuous time (Vol. 24). Springer Science & Business Media.

[13] Rawlik, K., Toussaint, M., & Vijayakumar, S. (2013). On stochastic optimal control and reinforcement learning by approximate inference. Robotics, 353.