

Sponsored search auctions: an overview of research with emphasis on game theoretic aspects

**Patrick Maillé, Evangelos Markakis,
Maurizio Naldi, George D. Stamoulis &
Bruno Tuffin**

Electronic Commerce Research

ISSN 1389-5753

Volume 12

Number 3

Electron Commer Res (2012) 12:265-300

DOI 10.1007/s10660-012-9094-8



Your article is protected by copyright and all rights are held exclusively by Springer Science+Business Media, LLC. This e-offprint is for personal use only and shall not be self-archived in electronic repositories. If you wish to self-archive your work, please use the accepted author's version for posting to your own website or your institution's repository. You may further deposit the accepted author's version on a funder's repository at a funder's request, provided it is not made publicly available until 12 months after publication.

Sponsored search auctions: an overview of research with emphasis on game theoretic aspects

Patrick Maillé · Evangelos Markakis ·
Maurizio Naldi · George D. Stamoulis ·
Bruno Tuffin

Published online: 18 July 2012
© Springer Science+Business Media, LLC 2012

Abstract We provide a broad overview of the research that has been conducted until recently on the design of sponsored search auctions. We mainly focus on game theoretic and mechanism design aspects of these auctions, and we analyze the issues associated with each of the three participating entities, i.e., the search engine, the advertisers, and the users of the search engine, as well as their resulting behavior. Regarding the search engine, we overview the various mechanisms that have been proposed including the currently used GSP mechanism. The issues that are addressed include analysis of Nash equilibria and their performance, design of alternative mechanisms and aspects of competition among search engines. We then move on to the advertisers and discuss the problem of choosing a bidding strategy, given the mechanism of the search engine. Following this, we consider the end users and we examine how

P. Maillé (✉)
Institut Mines-Telecom, Telecom Bretagne, 2, rue de la châtaigneraie CS 17607,
35576 Cesson-Sévigné Cedex, France
e-mail: patrick.maill@telecom-bretagne.eu

E. Markakis · G.D. Stamoulis
Department of Informatics, Athens University of Economics and Business, 76 Patision St.,
Athens 10434, Greece

E. Markakis
e-mail: markakis@gmail.com

G.D. Stamoulis
e-mail: gstamoul@aueb.gr

M. Naldi
Dip. di Informatica Sistemi Produzione, Università di Roma Tor Vergata, Via del Politecnico 1,
00133 Roma, Italy
e-mail: naldi@disp.uniroma2.it

B. Tuffin
Inria Rennes Bretagne Atlantique, Campus Universitaire de Beaulieu, 35042 Rennes Cedex, France
e-mail: bruno.tuffin@inria.fr

user behavior may create externalities and influence the performance of the advertisers. Finally, we also overview statistical methods for estimating modeling parameters that are of interest to the three entities. In each section, we point out interesting open problems and directions for future research.

Keywords Sponsored search auctions · Keyword auctions · Search engines · Mechanism design · Game theory · Nash equilibrium

1 Introduction

Online advertising is a booming industry, accounting for a large percentage of the revenue generated by web services [51]. Online ads are essential to monetize valuable Internet services, offered for free to the general public, like search engines, blogs, and social networking sites; e.g. see [46, 90]. They have potential benefits for the advertisers, who can observe the results of their campaign within days or even hours; at the same time, they enhance the user experience by facilitating search and commerce decisions. The enhancement of the search experience provided by online advertising represents a key example of the welfare-increasing role played by search agents for time-constrained consumers [66].

Originally, the only form of online advertising available was in the form of banner advertisements (ads): the owner of a website and the advertiser would agree on a payment to display the ad a fixed number of times. In the last decade, the most popular advertising method has become sponsored search, which represents a very profitable market for search engines. The idea behind sponsored search is that, for queries with commercial interest (e.g., “digital camera”), Google, Yahoo!, Bing, and other search engines allow a certain number of ads to be displayed on the top or on the side of the search (organic) results. Typically, there are up to three links above the organic results (these are the mainline slots), and up to eight links on the right side the organic results (sidebar slots). The main advantage of such ads is that an advertiser is displaying his ad to users who have expressed interest for the specific keywords included in the query and are therefore more likely to be interested in his product.

The selection of the ads to be displayed is done by means of an auction, the main ingredients of which are described below. There are three different charging schemes that can be considered for the selected ads: (1) the *Pay-Per-Impression* (PPI) model, where each advertiser is charged every time his ad is displayed, (2) the *Pay-Per-Click* (PPC) model, where the advertiser is charged only when a user clicks on the ad, (3) the *Pay-Per-Transaction* (PPT) model, where the advertiser is charged when the click results in a conversion, i.e., a purchase by the user. The most popular model that is being used in almost all sponsored search auctions is the *Pay-Per-Click* model, and our survey will focus on this.

In order to design a sponsored search auction, we first need a rule that ranks the bidders and thus determines the allocations of the available slots to the ads. The ranking rule has to compute a score for each bidder, and rank bidders in decreasing order, according to that score. Throughout the history of sponsored search auctions, the score has varied from being simply the bid of each bidder to being a function of the bid and possibly of other parameters, most notably of the *Click-Through-Rate* (CTR).

The CTR of an ad is the probability that a user will click on the ad, and can be affected by the ad itself (due to the content of the text being displayed and/or the identity of the advertiser), the slot that the ad is occupying (higher slots typically receive more clicks), and several other factors, such as the presence of other competing advertisers. The use of CTR as a scoring parameter indicative of future revenue comes as a result of adopting the Pay-Per-Click model. The two most frequently used ranking rules are

- the rank-by-bid policy, where the bidders submitting the k largest bids win the k slots in the order of their bids, and
- the rank-by-revenue policy, where each bid b_i is weighted by a *quality score* w_i of advertiser i , which reflects the probability that a user will click on the ad of advertiser i . The rationale is that ranking only by the bid may lead to displaying ads with very low probability of attracting clicks and therefore lowering the total revenue of the search engine. On the contrary, the rank-by-revenue rule takes into account the expected revenue from each bidder. After sorting the advertisers by the product $w_j b_j$, the k highest advertisers get the k slots accordingly.

The ranking rule is complemented by the payment (pricing) rule, determining the amount that a bidder being allocated a certain slot for his ad will ultimately have to pay upon receiving a click. Back in 1997, when sponsored search auctions were launched by Overture (then GoTo; now part of Yahoo!), the allocation rule was ranking by bid and the payment rule was the “first price” one (i.e., “pay-your-bid”): any advertiser winning a slot would pay an amount equal to his bid. As this mechanism was gradually recognized to be unstable (it led to cycling bidding patterns and low revenues, see e.g. [31]), search engines switched, starting with Google in 2002, to the so-called Generalized Second Price (GSP) auction that we describe in the next section.

The definition of an auction mechanism is therefore the joint choice of a ranking rule and a pricing rule. Note, for example, that Yahoo! originally used first-price payments with bid-based ranking, then switched to GSP with bid-based ranking rule, and finally to GSP with revenue-based ranking rule. For more on the history of keyword auctions see [32]. For an earlier survey on sponsored search markets see [61]. In the present survey, and since the related literature is by now so extensive, we have decided to focus more on the game theoretic aspects of these auctions and the theoretical analysis of the corresponding games. We also present empirical and experimental observations and findings, serving either as motivation for the theoretical work or as an alternative means of extracting properties for the mechanisms.

The rest of the survey is structured as follows: we devote one main section to each of the interacting parties in sponsored search auctions, namely the search engine itself (Sect. 2), the advertisers who play the role of the bidders in the auction (Sect. 3), and the search engine users (Sect. 4). For each of them we discuss their interests and focus on modeling the important parameters that affect their overall benefit. Finally, Sect. 5 focuses on statistical techniques that can be applied by any of the involved entities to estimate unknown parameters, such as the valuation of other bidders (i.e., the price they are willing to pay to obtain a slot) or the CTRs.

We stress that the structure of our survey is inspired by the fact that sponsored search auctions take place within an ecosystem involving stakeholders with different

Table 1 The stakeholders involved in sponsored search auctions, with their objectives and strategic decisions

Type of actor	Objective	Strategic decisions
Search engines	Revenue from the auction	Number of slots to display Auction scheme ranking rule pricing rule charging scheme reserve price
Advertisers	Utility maximization (sales minus advertising charges)	Bidding budget Bid apportioning among search engines Bid level
Users	Relevance of the organic and sponsored search answers	Search engine selection

interests. Namely, search engines wish to attract both users and advertisers to maximize revenue, while advertisers submit bids in the hope of reaching users to finalize sales and users are sensitive to the quality of the results displayed by the search engines. These three types of actors interact, and their respective utility (payoff) criteria and strategic decisions are summarized in Table 1. Note moreover that there are several advertisers, as well as possibly several search engines in competition. Since each of these actors can reasonably be assumed to behave selfishly, game theory is the most appropriate tool to study their interactions. In the next sections, we concentrate on each of the three types of participants (search engines, advertisers, and users).

Before we proceed, we briefly recall the main principles of the most popular auction scheme currently in place, namely the GSP rule.

1.1 The GSP mechanism

We describe here formally the *Generalized Second Price* (GSP) mechanism, which is being used in practice by the major search engines.

First, we introduce some notation. Assume that there is a set $N = \{1, \dots, n\}$ of n advertisers, who compete for a set $K = \{1, \dots, k\}$ of k slots, where slot 1 indicates the slot on the top of the list and slot k is the slot on the bottom of the list. Typically, we have $n > k$. We will consider only the sidebar slots and ignore the slots on the top of the organic results. As already mentioned, we assume that Pay-Per-Click is employed. For each advertiser i , his valuation, v_i , expresses the maximum price per click he is willing to pay. When participating in the auction, advertiser i submits a bid b_i that may differ from the actual valuation. The vector of the advertisers' bids, $\mathbf{b} = (b_1, \dots, b_n)$, will usually be referred to as a *bidding profile* or *strategy profile*. Submitting a bid b_i guarantees to advertiser i that he will not be charged a price higher than b_i per click. Moreover, as already mentioned, an important parameter in the context of sponsored search auctions is the click-through-rate (CTR), interpreted as the probability that a given ad will be clicked when displayed. The CTR can be decomposed in different parts. We assume that, for each slot s , there is a slot-dependent parameter θ_s , denoting the probability that a user will click on an ad on slot s ; this

is often referred to as the CTR of the slot. Also, for every bidder i , there is a bidder-dependent parameter q_i , which is the CTR of bidder i , i.e., the probability that a user will click on an ad of bidder i . This is a crucial parameter for the advertisers. From their point of view, one of the challenges they face is to design and monitor their ad (while conforming to the rules imposed by the search engine) so as to maximize the parameter q_i , which has led to the use of various marketing and behavioral approaches [85].

In this survey, with the exception of Sect. 4, we assume that the overall CTRs are *separable*: the probability that an ad of advertiser i , occupying slot s , receives a click is $q_i\theta_s$. Note, however, that the exact meaning of the term CTR will follow from the context. The statistical estimation of the CTR values from observations of user behaviors is addressed in Sect. 5.2. For the rest of the sections, we consider the CTR values as given.

The GSP pricing rule then charges each bidder the minimum bid value he could have offered to be assigned the same slot. In practice, if an advertiser obtains slot s , his charge equals the bid that would have had him exactly tie with the advertiser in slot $s + 1$ in the current ranking, plus possibly a very small constant ϵ .¹

Given a ranking rule, let (s) denote the index of the bidder who is ranked in the s -th position according to the rule. Depending on the ranking rule selected, GSP auctions thus give the following:

- With a rank-by-bid policy, the winner of a slot $s \leq k$ is charged a price of $b_{(s+1)} + \epsilon$. Hence the winner of slot s pays the bid of the person who obtains the slot right below him, if $s < k$. If $s = k$, then the winner of slot s pays the highest losing bid.
- With a rank-by-revenue rule, each bidder that gets a slot pays again the amount that would be necessary to bid to keep his current position. Hence, for the advertiser who obtained slot s , the payment p_s should satisfy the inequality $w_{(s)}p_s \geq w_{(s+1)}b_{(s+1)}$. The minimum price resulting from that inequality is

$$p_s = \frac{w_{(s+1)}b_{(s+1)}}{w_{(s)}} + \epsilon,$$

where $w_{(s)}$ is the quality weight of the advertiser who obtains slot s . In the remaining of our survey we will take $\epsilon = 0$, as this is not an important parameter of the mechanism. This rule was introduced by Google in 2002; at that time the quality weight was taken to be equal to the estimated CTR of the advertiser who won slot s . In light of our previous discussion on CTRs, this would be equal to the bidder-dependent parameter $q_{(s)}$. This approach can be generalized by considering the quality weight to be equal to a power of the CTR. This form of functional dependence makes it easier to remove irrelevant advertisements, i.e., advertisements with low CTR though accompanied by high bids [65] (as irrelevant ads diminish the trust of customers and are therefore undesirable). At present, however, Google's quality score does not depend just on the CTR but also on other qualities of the advertiser, including also the text of the ad. The exact method of determining the

¹The exact value of ϵ is provided by every search engine for every currency. It usually equals $\epsilon = 0.01$.

Table 2 Example of slot allocation and pricing

Bidder	Bid	Quality score	Rank by bid	Rank by rev	Price by bid	Price by rev
1	9.62	0.07	5	4	–	–
2	10.47	0.05	2	5	10.45	–
3	10.45	0.09	3	1	9.68	9.68
4	10.64	0.08	1	3	10.47	8.42
5	9.68	0.09	4	2	–	9.46

quality score is not publicly available. In 2007 the revenue-based ranking rule was adopted also by Yahoo! and Microsoft Live (now Bing).

Another important quantity in the context of sponsored search auctions is the welfare produced by an allocation of slots to the advertisers, as defined below.

Definition 1 Given a strategy profile \mathbf{b} , and a ranking rule, the social welfare of the resulting allocation is $SW(\mathbf{b}) = \sum_{j=1}^k \theta_j q_{(j)} v_{(j)}$, where (j) denotes the index of the bidder occupying slot j . An allocation is called efficient if it achieves the maximum possible social welfare.

The impact of the ranking rule can be grasped by considering the simple example in Table 2, where 5 bidders compete for 3 slots. For simplicity, in this example the slot dependent CTRs $(\theta_1, \theta_2, \theta_3)$ are all assumed to be equal to 1.

It should be also noted that there are several details involved in the practical application of the above rules; e.g., bids for ads that appear to be of low relevance or quality can be excluded, or be subject to higher reserve prices; see [59].

2 The search engine's view and interests

In this section, we focus on topics that are of interest to the auctioneer, i.e., the search engines. We start with an analysis of the Nash equilibria of the GSP mechanism, and a study of their properties (Sects. 2.1 and 2.2). We then discuss the design of alternative mechanisms (Sect. 2.3), the impact of reserve prices on revenue (Sect. 2.4) as well as aspects of competition between search engines (Sect. 2.5).

2.1 Analysis of Nash equilibria

As sponsored search auctions are essentially games among advertisers, the ideal situation for the search engine is to ensure that the advertisers have no incentive to misreport their valuations. This would eliminate the possibility of potential manipulations of the mechanism by the advertisers. However, simple examples demonstrate that neither the rank-by-revenue nor the rank-by-bid GSP mechanisms are truthful, and that bidders can be better off by shading their valuations. See, e.g., the example in [32].

Given the absence of truthful dominant strategies, a natural approach is to identify the set of Nash equilibria of such games. A Nash equilibrium is defined [36] as a stable outcome of the game, i.e., a situation where no player can improve his payoff utility by a unilateral strategy change. Here we consider the bidding game, where the players are the bidders and the strategies are the bids they submit. Hence, an outcome is a Nash equilibrium if no player has an incentive to improve his utility by targeting a different slot than the one he is currently occupying. The set of Nash equilibria depends on the rules of the game, which are given by the auction scheme implemented by the search engines. The set of equilibria is typically large, and there has been published a stream of articles that focus on a subset of Nash equilibria that are called “symmetric Nash equilibria” (SNE), which we define below. Symmetric Nash equilibria have specific properties of interest for the search engine and for the advertisers, and therefore they could be the equilibria that the search engine would prefer to attain; see [32, 60, 87]. Below, we present the main research results related to such equilibria.

We first set up the model in which these equilibria are studied. The equilibrium analysis was performed for the rank-by-bid rule in [32, 87] (but this is generalized in [60]), where the CTR of advertiser i for position s is assumed to be the same for all advertisers, and to depend only on the position slot s , i.e., the bidder-dependent part of CTR, q_i , is assumed to be the same for all bidders. We start by presenting this analysis for the rank-by-bid rule and we later incorporate the bidder dependent CTRs. Denoting the CTR for slot s by θ_s , assume that $\theta_1 > \theta_2 > \dots > \theta_k > 0$. Following the GSP principle, the price paid by the advertiser occupying slot $s \leq k$ is $p_s = b_{(s+1)}$ (since $\epsilon = 0$, as in the previous section) and the total value that the outcome has for him, i.e., his utility, in the game-theoretic vocabulary, is $(v_{(s)} - p_s)\theta_s = (v_{(s)} - b_{(s+1)})\theta_s$. To simplify the notation in the analysis that follows, for a given instance of the auction we renumber the bidders so that v_s is the valuation of the bidder assigned slot s . Hence his utility would then be $(v_s - p_s)\theta_s$.

We assume in this section a one-shot and simultaneous game with complete information (if the game is played repeatedly, the Folk Theorem [36] would lead to a very large potential set of equilibria). At a Nash equilibrium, no advertiser would have an incentive to obtain a different slot. Recalling the GSP pricing rule, we can express this formally:

Definition 2 A bid vector is a Nash equilibrium if for every slot s and for the advertiser at this slot, it holds:

$$\begin{aligned} \theta_s(v_s - p_s) &\geq \theta_j(v_s - p_j) \quad \forall j > s, \\ \theta_s(v_s - p_s) &\geq \theta_j(v_s - p_{j-1}) \quad \forall j < s. \end{aligned}$$

The index $j - 1$ in the last equation comes from the fact that the ordering is changed if advertiser s changes his bid to target a higher slot.

A restricted class of equilibria is considered in [32, 87], called *symmetric Nash equilibria* (SNE) in [87] and *locally envy-free equilibria* in [32]:

Definition 3 A symmetric Nash equilibrium is a set of bids that satisfies:

$$\theta_s(v_s - p_s) \geq \theta_j(v_s - p_j) \quad \forall j, s. \tag{1}$$

Definition 3 just considers the inequality of Nash equilibria for $j > s$, but applied to all positions. The rationale behind this notion becomes clearer if we look at pairs (j, s) such that $s = j + 1$. If the bidder at slot s starts raising slightly his bid so as to increase the payment of the bidder above him, then bidder j can underbid him as a retaliation, and essentially this means that they will have swapped their bids. The right hand side expresses the payoff of bidder s if bidders j and s swap their bids. Symmetric Nash equilibria capture the notion that there should be no incentives for such swapping of bids between any pair of players.

It is straightforward to verify that if a bid vector satisfies the inequalities (1), then it will be a Nash equilibrium [87]. Hence the class of SNE is a subclass of the set of Nash equilibria. The following key properties are satisfied by SNE and can be desirable by the search engine [87]:

- At an SNE, there is monotonicity in the valuations of the winning bidders, i.e., the value v_s of the bidder assigned to slot s is decreasing in s .
- To check if a bid vector is an SNE, it suffices to examine, for every slot s , only the inequalities that concern the slots $s - 1$ and $s + 1$. This also justifies the term “locally-envy-free” that was introduced by [32].
- There is an SNE maximizing the search engine’s revenue among all possible Nash equilibria.

The fact that we only need to verify the inequalities of SNE for neighboring slots allows for more explicit characterizations of the bidding vectors. Since the advertiser at position $j + 1$ does not want to move one slot up, that is, $(v_{j+1} - p_{j+1})\theta_{j+1} \geq (v_{j+1} - p_j)\theta_j$, and the one at position j does not want to move one slot down, that is, $(v_j - p_j)\theta_j \geq (v_j - p_{j+1})\theta_{j+1}$, we have

$$v_{j-1}(1 - \gamma_j) + b_{j+1}\gamma_j \geq b_j \geq v_j(1 - \gamma_j) + b_{j+1}\gamma_j, \tag{2}$$

where $\gamma_j = \theta_j/\theta_{j-1} \leq 1$. We thus obtain recursive upper and lower bounds for the bids:

$$\begin{aligned} b_j^U \theta_{j-1} &= v_{j-1}(\theta_{j-1} - \theta_j) + b_{j+1}\theta_j \\ b_j^L \theta_{j-1} &= v_j(\theta_{j-1} - \theta_j) + b_{j+1}\theta_j, \end{aligned} \tag{3}$$

whose solutions are $b_k^U \theta_{k-1} = \sum_{j \geq k} v_{j-1}(\theta_{j-1} - \theta_j)$ and $b_k^L \theta_{k-1} = \sum_{j \geq k} v_j(\theta_{j-1} - \theta_j)$. The upper bound corresponds to the case where advertiser k bids the amount in the upper bound in (2) with $k = j$, while the lower bound is when he bids the lower bound of (2). It is rather advised in [87] that advertisers bid according to the lower bound; thus, an advertiser would make a profit if he moves up in the ranking.

Finally, it is interesting to compare the GSP mechanism with the classical Vickrey-Clarke-Groves (VCG) auction, which is truthful and where each bidder pays for the externality that he is causing to the other bidders, that is, the loss of utility that is due to his participation in the auction. For more on the VCG mechanism, see [22, 48,

88]. In [32], a particular SNE is constructed for GSP, in which the slot assignment and the payments coincide with the allocation and payments of the VCG mechanism when all bidders declare their true valuations. If $v_1 > v_2 > \dots > v_n$, then this SNE is defined recursively as follows:

$$b_j^* = \begin{cases} 2b_2^*, & j = 1, \\ \gamma_j b_{j+1}^* + (1 - \gamma_j)v_j, & 2 \leq j \leq k, \\ v_j, & k < j \leq n, \end{cases} \quad (4)$$

where $\gamma_j = \theta_j / \theta_{j-1}$. Note that b_1^* can actually be any quantity greater than b_2^* since it does not affect the price of any slot. Except for this degree of freedom, this Nash equilibrium does not involve over-bidding. That is, the bids of the rest of the players do not exceed the corresponding advertisers' valuations.

Following [18], we will refer to this as the *VCG equilibrium* of GSP. This SNE was shown in [32] to be the worst SNE for the search engine in terms of revenue, and the best for the advertisers in terms of their utility. In other words, the engine revenue under GSP is always better than when using the truthful VCG mechanism, which provides an explanation of why the GSP is adopted instead of the well-known VCG auction. The findings of [32] are summarized below.

Theorem 1 *The bidding vector \mathbf{b}^* defined by (4) is an SNE. In this equilibrium the assignment and the payments are identical to the dominant strategy equilibrium of the VCG mechanism. Furthermore, in any other SNE, the revenue is at least as high as the revenue of \mathbf{b}^* .*

The proof of Theorem 1 is based on viewing these games as *assignment games*, which were introduced by Shapley and Shubik in [83]. The construction is based on results of Leonard [62] and Demange et al. [25] concerning stable assignments in such games. For more details on the proof we refer the reader to [32] as well as to the nice exposition in [30], Chap. 15.

It should be noted that the comparison of the auctioneer's revenue under the VCG and the GSP mechanisms is further investigated in [37]. As mentioned above, it is established in [32] that the revenue under the VCG dominant strategy equilibrium constitutes a lower bound to the revenue of any SNE of the GSP mechanism. In [37], Fukuda *et al.* extend the comparison between the sets of SNE of the GSP and the VCG mechanisms, motivated by the fact that it has been observed that bidders do not play the dominant strategy of VCG in reality. In particular, they prove that the lower bound for the revenues of the two sets of equilibria is still the revenue under the VCG dominant strategy equilibrium, and that the maximum revenue attainable under the SNE of VCG is the same with GSP. In the sequel, the authors of [37] investigate the revenues of the two mechanisms experimentally. It appears that GSP in general produces somewhat higher revenues than VCG, although both mechanisms come close to the lower bound. On the other hand, the efficiency (i.e., the social welfare) attained in the experiments under VCG was higher than that under GSP, although an improvement due to repetition was observed for both mechanisms.

Finally, related to the above comparisons of VCG and GSP auctions is the work of Babaioff and Roughgarden, who in [10] derive conditions under which a payment

rule combined with the rank-by-bid policy shares the same properties of VCG that GSP does. To this end, they derive necessary and sufficient conditions for a payment rule so that the resulting mechanism both has a full-information Nash equilibrium that is identical to the dominant-strategy VCG outcome, in terms of allocations and payments, and admits an ascending implementation such as that introduced in [32] (and described in the following paragraph). The latter property applies as a consequence of certain monotonicity properties and of the “upper triangular” property, whereby the price paid for slot j is a function only of the bids $b_{(j+1)}, \dots$, i.e., those that are lower than $b_{(j)}$. The authors of [10] also formalize the intuitive fact that among the payment rules with the aforementioned properties, GSP is the simplest one.

Convergence to equilibria An interesting issue in equilibrium analysis is whether (and how) the advertisers would eventually converge to an equilibrium, using an iterative process. We revisit these issues in Sect. 3, where we deal with convergence related to repeated runs of the single-shot game with sealed bids. For now, we point out that one possibility in order to understand further the GSP mechanism is to imagine a process such as the *Generalized English Auction*, introduced in [32], as an analogue of the standard English auction that helps us understand the Vickrey auction. In the Generalized English Auction the price increases linearly and continuously from zero and advertisers decide when to quit the auction. Their bid is then taken as the price level at that specific moment, and the allocation is decided when all advertisers have “announced” their bid. The bidding game among advertisers can be studied as a Bayesian game. If we assume that valuations are random variables, independent and identically distributed, it can be shown [32] that there exists a unique perfect Bayesian Nash equilibrium where the price $p_i(j, h, v_i)$ at which advertiser i quits the English auction depends on his valuation v_i , the number j of remaining advertisers in the auction and the history $h = (b_{(j+1)}, \dots, b_{(n)})$ of advertisers that have already dropped out: $p_i(j, h, v_i) = v_i - (v_i - b_{(j+1)})\theta_j/\theta_{j-1}$. In other words, advertiser i is better off before the price reaches the level at which he is indifferent between paying $b_{(j+1)}$ at the $(j + 1)$ -th slot and paying p at the j -th slot. It turns out that the position and the payoff of each advertiser in this unique perfect Bayesian equilibrium are the same as in the dominant-strategy VCG equilibrium that we defined earlier.

Incorporating bidder-dependent CTRs The whole analysis above remains almost the same if we integrate the quality scores, q_i , to handle different CTRs among advertisers. In [60], the CTR of advertiser i for position s is assumed to be separable of the form $q_i\theta_s$, as explained in Sect. 1.1. The model is also generalized to the rank-by-revenue rule, so that advertisers are ranked in decreasing order of their score $w_i b_i$ for the auction, where a weight w_i is associated to advertiser i . The authors focus on weights of the form $w_i = q_i^d$, where the exponent d is a parameter that can vary in the interval $(-\infty, +\infty)$. Note that this family of ranking schemes includes the rank-by-bid rule (for $d = 0$) as well as the case where $w_i = q_i$ (for $d = 1$). With these new definitions, the utility of advertiser i , when he occupies slot s becomes $q_i\theta_s(v_i - p_s)$, and the payment p_s would be $b_{(s+1)}w_{(s+1)}/w_i$. For ease of notation, as we did with the rank-by-bid rule, let us renumber the bidders so that v_s is the valuation of the person occupying slot s . The generalization of a symmetric Nash equilibrium yields

then the inequalities

$$q_s \theta_s \left(v_s - \frac{w_{s+1}}{w_s} b_{s+1} \right) \geq q_s \theta_j \left(v_s - \frac{w_{j+1}}{w_s} b_{j+1} \right) \quad \forall j \neq s.$$

It is assumed in [60] that advertisers play the smallest SNE, i.e., the bid according to the lower bound obtained from the generalization of (3), the goal being here to maximize a lower bound of revenue. This gives the recursion

$$\theta_s w_{s+1} b_{s+1} = \sum_{j=s}^k (\theta_j - \theta_{j+1}) w_{j+1} v_{j+1}.$$

The goal of the search engine is to determine the weights $\{w_i\}_{i \in N}$ that will maximize the expected revenue. Determining them in full generality can be a difficult problem, but focusing on the parametric family $w_i = q_i^d$, as explained earlier, leads to a more tractable problem, since it is essentially reduced to finding the optimal value for the parameter d . The authors then first show that the efficiency, defined as the sum of revenues of the search engine plus those of the advertisers, is maximized for $d = 1$, while the relevance, defined as the total CTR, is increasing with d . These findings imply that, if the auctioneer imposes bounds on the efficiency and relevance loss he is willing to tolerate, this will derive upper and lower bounds for the value of the parameter d . The revenue curve can then be plotted within the allowable range for d and the optimal value of d can be selected. The authors of [60] complement their analysis with simulations, where the valuation v_i and the CTR effect q_i of each bidder are taken from the same joint density. The joint distribution of these two parameters is inferred from real data in Yahoo! auctions, taking into account their correlation. It is then shown that ranking by bid yields a higher revenue than ranking by revenue if the correlation between value and CTR effect is positive, while it is the opposite if the correlation is negative. Using the optimal value of d (which can be determined by the revenue curve) can result in a significant revenue increase.

2.2 Social inefficiency under the GSP mechanism

In this section we continue the analysis on Nash equilibria of the GSP mechanism but we take a different direction. We focus on the question of whether the equilibria of the mechanism lead to near optimal social welfare. Surprisingly, even though inefficiency of equilibria has been studied in many other contexts in game theory (e.g., congestion games), this was not given much attention in sponsored search auctions until recently.

Before we proceed, we give some relevant definitions. For simplicity, we ignore the quality score of each bidder, however the results can be easily generalized. Recall that, according to Definition 1, for a given bidding profile $b = (b_1, \dots, b_n)$, the social welfare associated to b is $SW(b) = \sum_{j=1}^k \theta_j v_{(j)}$, which depends on the vector of bids through the slot allocation rule. On the other hand, if $v_1 > v_2 > \dots > v_n$, the optimal social welfare would be $OPT = \sum_{j=1}^k \theta_j v_j$. The two metrics SW and OPT differ whenever $v_{(j)} \neq v_j$, for some indices j , which means that the order of bids differs from that of the valuations. The usual way for capturing the inefficiency of Nash equilibria in games, known as *Price of Anarchy*, is by considering the worst possible case [57]:

Definition 4 The Price of Anarchy of the game induced by the GSP mechanism is:

$$PoA = \sup \frac{OPT}{SW(b)},$$

where the supremum is taken over all bid-profiles that constitute Nash equilibria.

Simple examples showing that Nash equilibria may fail to be efficient are easy to obtain even for two slots, see for example [78]. The first formal analysis on the Price of Anarchy was given by Lahaie in [59]. An upper bound of $(\min_{i=1, \dots, k-1} \min\{\gamma_{i+1}, 1 - \gamma_{i+2}\})^{-1}$ was obtained, where we assume that $\gamma_{k+1} = 0$ (recall that $\gamma_i = \theta_i / \theta_{i-1}$, as defined in Sect. 2.1). For arbitrary CTRs, this may lead to very high inefficiency. However, for geometrically decreasing CTRs, with decay parameter δ , i.e., $\theta_i = 1/\delta^i$ and $\gamma_i = 1/\delta \forall i$, the bound becomes $(\min\{1/\delta, 1 - 1/\delta\})^{-1}$. In the experimental work of [34], it was observed, using various empirical datasets, that click-through data fit well with the exponential decay model with $\delta = 1.428$, thus implying a price of anarchy of at most 3.336. Hence, this can be seen as a positive result that for datasets fitting this model, the inefficiency is not arbitrarily high.

Improved upper bounds for general CTRs have been obtained recently in a series of works [16, 17, 68, 78]. Motivated by the observation that all known examples of very high inefficiency occurred at equilibria that involved over-bidding, the authors of [78] considered only Nash equilibria among conservative bidders, i.e., bidders that never bid above their valuation. This is a reasonable assumption as bidding above your valuation can be dominated by other strategies. The authors obtained an upper bound on the Price of Anarchy of 1.618 for pure Nash equilibria, an upper bound of 4 for mixed Nash equilibria, and a bound of 8 for Bayesian equilibria. These were later improved by [16, 68] and [17], resulting in upper bounds of 1.282, 2.310, and 2.927, for pure, mixed and Bayesian equilibria respectively. The bound of 2.310 also holds for the class of coarse correlated equilibria, which is an interesting class of equilibria, since it consists essentially of the points of convergence of regret-minimizing algorithms. These are algorithms where players adjust their strategy over time and their average regret for their choices tends to 0 (for more see [91]). It is not yet known whether these upper bounds are tight and also whether these can be improved if one focuses on special cases of CTR distributions, such as geometric decreasing assumptions. It is an interesting open question to tighten the bounds and have a complete picture on the inefficiency of Nash equilibria. A concrete question here is whether the combination of conservative bidding and geometric CTRs can yield improved upper bounds. As for lower bounds, we do know that the Price of Anarchy is at least 1.259, see e.g. [17], for all the concepts presented in Table 3.

The inefficiency of equilibria has also been studied experimentally in [86]. There, the authors considered various assumptions on the preference profiles of the bidders and conducted simulations with each class of preferences. Their main experimental findings are that the currently used, rank-by-revenue rule was more efficient than both the generalized first price auction (pay-your-bid) and the rank-by-bid second price rule. In fact, for some of the preference profiles they considered, the rank-by-revenue rule was approximating efficiency quite well.

Table 3 Known upper bounds on the loss of efficiency incurred by the GSP auction scheme

	Pure equilibria	Pure locally stable profiles	Mixed equilibria	Coarse correlated equilibria	Bayesian equilibria
General case	$(\min_{i=1,\dots,k-1} \min\{\gamma_{i+1}, 1 - \gamma_{i+2}\})^{-1}$	–	–	–	–
Conservative bidders	1.282	$(1 - \max_i \gamma_i)^{-1}$	2.310	2.310	2.927
Conservative bidders and geometric CTRs: $\gamma_i = 1/\delta$	1.282	$\delta/(\delta - 1)$	2.310	2.310	2.927

Finally, social inefficiency has been recently analyzed with respect to locally aware bidders in [69]. Given the fact that learning all of your competitors' bids entails a cost in time, effort, budget, and other factors, [69] focuses on bidders who are making only local moves, i.e., they are aware only of the price of the slot right above and below them in the current configuration. The *local stability ratio* is then defined as the analogue of the Price of Anarchy for such locally stable configurations, where no local move is profitable (a relaxation of the notion of Nash equilibrium). The authors obtain upper bounds on the local stability ratio, which imply that, for the case of conservative bidders and geometrically decreasing CTR distributions, the inefficiency is no more than the bound of Lahaie [59] for Nash equilibria. As with the rest of the results outlined above however, no tight lower bounds are yet known and it is still an open problem to resolve whether the upper bounds are the best possible.

Note here that a lower bound for the Price of Anarchy for any of the above concepts is obtained by simply exhibiting an instance of a GSP auction along with a Nash equilibrium of the appropriate inefficiency. When we allow over-bidding, it is quite easy to construct such examples. For conservative bidders however, lower bounds still remain elusive, see e.g. the discussion in [16].

The known efficiency results for all the equilibrium concepts described in this section are summarized in Table 3. This line of research falls within the recent initiative of analyzing inefficiency of mechanisms that do not possess truthful dominant strategies; see [11] for an analysis of combinatorial auctions along these lines. In our context, the loss of efficiency in most cases is only some small constant factor, which implies that even if bidders play strategically, the final allocation may not be far from the optimal one, thus mitigating efficiency losses arising from the adoption of the GSP by the search engines due to the associated higher revenues. Regarding the potential for improving these results, although by now tight results for the price of anarchy have been obtained for several other classes of games, the context of sponsored search auctions seems to require different technical arguments. Therefore, determining whether the upper bounds presented here are tight is a challenging problem for future work.

2.3 Truthful auction mechanisms

In this section we discuss the issue of designing alternative mechanisms that do not give incentives to the bidders for misreporting their true valuation, contrary to the GSP mechanism.

In [3], Aggarwal *et al.* deal with the design of mechanisms in which bidding the true valuation for a keyword is a dominant strategy for each bidder. This would render optimal bidding for advertisers simpler than in the standard mechanisms already overviewed. Indeed, a bidder would only have to determine his actual valuation, without having to take into account how the others would bid. In particular, the authors of [3] assume that there is already a mechanism in place that ranks bidders in decreasing order of $w_j b_j$, where again b_j is the bid of bidder j and w_1, \dots, w_n is a set of given and fixed weights. The exact problem analyzed in the article is as follows: given this ranking rule, what is the truthful auction mechanism that produces the *same* rankings as the original rule? Of course, due to this restriction, the only remaining degree of freedom is the payment rule. The authors first show by means of counterexamples that the standard GSP rule does not lead to truthful bidding, while under direct ranking of bids (i.e., $w_j = 1$ for all j), then the famous VCG auction is not always attainable, even when modified through some weighting. Indeed, it is not possible to find a set of bid-independent weights for VCG that would produce the desired ranking. In fact, this inapplicability of weighted VCG is further extended to other cases, yet holds only for non-separable click-through-rates (see Sect. 1.1). The authors then introduce the “laddered auction”, according to which, the bidder ranked at position s pays the sum of two terms accounting for: a) the clicks that this bidder would have received if ranked at position $s + 1$, at the price he would have paid in that position, b) the extra clicks due to being ranked at position s , for an amount equal to the minimum bid necessary to maintain that position. Therefore, the payment per click p_s is expressed as follows:

$$p_s = \sum_{j=s}^k \frac{CTR_{s,(j)} - CTR_{s,(j+1)}}{CTR_{s,(s)}} \frac{w_{j+1}}{w_i} b_{(j+1)}, \tag{5}$$

where $CTR_{i,(j)}$ is the CTR of bidder i when his ad is displayed on slot j . It is then established that the laddered auction is truthful. Next, the authors of [3] compare the revenues under the laddered auction and those under the standard GSP auction in equilibrium. In particular, for separable click-through rates, they construct a deterministic equilibrium (with respect to bids of GSP) that yields the same revenues as the laddered auction.

In [41], Goel *et al.* propose a different mechanism in an attempt to extend the existing model of Pay-Per-Click. Their mechanism is based on a hybrid scheme, where each bidder is asked to submit two bids: a per-impression bid and a per-click bid, indicating the maximum amount he is willing to pay for being displayed and for receiving a click, respectively. The authors of [41] first investigate *myopic* bidders, i.e., bidders that try at every time step to optimize some function of the expected revenue, given prior distributions on each bidder’s CTR. In the case of a single slot, they propose a VCG-based pricing scheme and show that their mechanism is truthful in

expectation, when bidders are risk-neutral. They then propose two generalizations for auctions with multiple slots. The first one is based on a generalization of the GSP scheme, and is not truthful. The second generalization is based on VCG and the ladder auction of [3], and achieves truthfulness. Finally, *semi-myopic* settings are also explored, where bidders are trying to maximize expected utility over a time horizon.

As a concluding remark, it should be noted that, despite the importance of truthful bidding as a property facilitating efficiency, such mechanisms are not currently used in practice. This is mainly due to the fact that search engines prefer simple mechanisms (so that advertisers are not discouraged to participate) and they also aim to maximize their revenue rather than social welfare (recall Definition 1). However, we still believe that it is an important research direction to investigate further the design of alternative, more sophisticated mechanisms that achieve desirable properties, such as truthfulness, and exploit more information from the bidders' side, i.e., via submission of more parameters as in [41], while maintaining simplicity at the same time. Obtaining tradeoffs between these aspects still remains to be explored.

2.4 The impact of reserve prices on revenues

It is well-known in the theory of optimal auctions that introducing reserve prices in a mechanism can lead to increased revenues for the auctioneer, or even to revenue maximization, as established in the pioneering works by Myerson [74], by Riley and Samuelson [82], and in several other works that followed. In [77], Ostrovsky and Schwarz presented the first investigation of the impact of reserve prices on the revenue of GSP auctions. In particular, they report the results of related simulated auctions, whose various parameters (number of bidders, the moments of bidders' valuation distribution etc.) were selected on the basis of a Yahoo! auctions dataset. The reserve price was computed according to the theory of optimal auctions. This was subsequently personalized on a per advertiser (bidder) basis, according to the quality score of each of them, since this score is taken into account by the bidders' ranking mechanism. The authors compare the resulting revenues to those of the case of a fixed reserve price of \$0.1. The results reveal that the introduction of the aforementioned reserve prices does have a positive overall effect on the average revenue per keyword, which the authors estimate at 2.7 %. However, this effect is not positive in all cases. For example, for keywords with low search volumes, or with low values of reserve prices, the effect was negative. Clearly, the impact of reserve prices in sponsored search auctions is an important topic deserving further investigation.

2.5 Competition among search engines

One of the main issues that has been mostly ignored in the adwords literature is the fact that the analysis and optimization of parameters are performed when dealing with a single search engine (i.e., in the case of a monopoly). But there exist in practice several such search engines among which advertisers can choose, the two most important examples being Google and Yahoo!. The behavior of search engines as a reaction to this competition for advertisers requires a thorough investigation and can lead to different (equilibrium) situations than in a monopoly. Only a few articles deal with this topic, for instance [8, 45, 67].

In [67], two search engines in competition for advertisers are considered. The goal is to choose the best auction rules given that advertisers will go to the engine that best serves their interest, and to understand the impact of ranking policies in a competitive environment. To simplify the analysis, it is assumed that each engine offers a single slot. There are two classes of advertisers, one with a high expected CTR q_h and the other one with a low CTR q_l (with $q_l < q_h$ regardless of the engine). Let β be the probability that a given advertiser is in the class with high CTR. Valuations v_i are independent and taken from a cumulative density function F . The bid b_i submitted by an advertiser i is the amount he is going to pay per click if he wins the slot (that is, a first price auction), but it is also shown that a second price strategy would lead to the same expected payoffs and revenues for advertisers and search engines respectively. Two potential ranking rules are considered: either the search engine ranks according to bid, or according to expected revenue. For the different possible pairs of ranking policies employed by the two search engines, the advertiser game is analyzed, where an advertiser chooses (exactly) one auction and places his bids; a Nash equilibrium in mixed strategies is then obtained. Essentially, when the two engines adopt the same ranking rule, advertisers are indifferent between the two auctions (going with probability $1/2$ to each of them). If the engines do not implement the same ranking rule, the advertiser equilibrium depends on the proportion β of high-quality advertisers: if $\beta \geq 1/2$, i.e., advertisers are more likely to be of high-quality, all low-quality advertisers go to the price-only auction, while high-quality ones go to the price-only auction with probability $(2\beta - 1)/(2\beta)$ and to the quality-adjusted auction with probability $1/(2\beta)$; if $\beta < 1/2$ (and under some assumptions on F), all high-quality advertisers participate in the quality-adjusted auction where low-quality advertisers choose the price-only auction with a probability that depends on their valuation and on q_l/q_h , this probability being 1 above a threshold v^* . The existence of a Nash equilibrium in the ranking game is then discussed (when search engines try to maximize their revenue) depending on the value of β and the CTR ratio q_l/q_h . Few trends are extracted from the study. For instance, being the only quality-adjusting (resp. price-only) engine gives a market advantage when the number of high quality advertisers is high (resp. low) with β close to 1 (resp. 0). Also, competition produces incentives to adopt the quality-adjusted (i.e., rank-by-revenue) rule even if this is not the optimal strategy in the case of a monopoly. This could explain why Yahoo! moved from rank-by-bid to rank-by-revenue due to the competition with Google.

The impact of competition is also analyzed in [45]. The model considers a double auction with advertisers on one side and slot sellers on the other. The goal is to study the efficiency and incentive compatibility properties depending on whether the separability assumption (i.e., the fact that the valuation is the product of the CTR at a given position, independent of the advertiser, and the per-click valuation of the advertiser) applies or not. It is shown that if separability is not assumed, the VCG mechanism has to be applied (for all participants from both sides of the auction) to obtain a truthful and efficient mechanism. But in that case, the market maker will potentially run a budget deficit.

The VCG payment rule is also applied in [8], when there are two search engines in competition, each engine offering a potentially different number of slots, and having different CTRs (θ_k at position k for the first engine, and θ'_k for the second one). At

each auction, slots are allocated by decreasing order of the bid. As a first result, it is possible that, if players can participate in both auctions at the same time, a consistently more popular auction (if $\theta'_k > \theta_k \forall k$) yields a smaller revenue. This comes from the VCG pricing rule: whatever the “performance” of the auction, an advertiser pays only for the loss he creates on the system, not the value itself. If on the other hand players have to choose between auctions (this being their only strategy choice, advertisers submitting their real value due to the incentive compatibility property), there is a unique equilibrium such that with a given probability $q(v_i)$ player i chooses to join the first auction, while with probability $1 - q(v_i)$ he chooses the second one. In that case, the more popular auction always gets a higher revenue.

To summarize, certain models for competition among search engines have already been published in the literature, however these pertain to rather special cases of the problem. Therefore, we feel that the relevant research is in its early stages. Thus, defining the most “robust” (in terms of revenue) mechanisms requires more investigation. We view this as an interesting and promising research direction from the point of view of the search engines.

3 The advertisers: bidding strategies and their properties

Once the search engine has chosen a mechanism, it is then the advertisers' turn to play the game. Hence, given the GSP mechanism and its variants, the main question that the advertisers face is to decide what they should bid. As already discussed in Sect. 2, it has been observed that truthful bidding is not a dominant strategy under the GSP mechanism. This gives rise to strategic behavior by the bidders in order to increase their utility, as was also established empirically in [31], where the authors studied a Yahoo! dataset of auctions from 2002 and 2003. Furthermore, if we view the process as a repeated game, it is not always the case that the game will converge to a better state for all players, even when bidders try to profit by lying. It is known that a plethora of Nash equilibria exist [32, 87], and it is not a priori clear whether any of these equilibria are actually reached in real keyword auctions.

All these issues make the bidding decisions much more complex: advertisers often end up assigning their bidding campaign to consultants or other companies, specializing in such campaigns, see e.g., [4, 28]. In this section we review some of the proposed bidding schemes and study their properties.

3.1 Greedy bidding strategies in auctions for a single keyword

Most often auctions for the same keyword are performed repeatedly. Thus, a natural approach to bidding is to use the past as a prediction for the future. Hence, if an advertiser assumes that the bids of the other players in the next round will remain fixed, the best choice for him is to bid so as to win the slot that maximizes his own utility (or to bid so as not to win if winning leads to negative utility). Hence we can define the class of *greedy bidding strategies* as all the strategies in which an advertiser i chooses a bid for the next round so as to maximize his utility, assuming that the vector of bids $b_{-i} = (b_1, b_2, \dots, b_{i-1}, b_{i+1}, \dots, b_n)$ remains as it was in the previous round.

In most instances, there is a range of bids that achieve maximum utility given the bids of the other players. Specifying further how to choose a bid within this allowed range gives rise to various greedy strategies. For example, suppose that the utility of advertiser i is maximized when he acquires slot s . One way to bid then is to submit the smallest possible value needed to acquire slot s , given b_{-i} . This is usually referred to as *altruistic* bidding as in that case the bidder who wins slot $s - 1$ will pay the smallest possible amount. An alternative line of reasoning is that, since bidders are usually business competitors, one should try and push the other bidders' payments as high as possible. This can be achieved by submitting the maximum bid that will guarantee slot s rather than slot $s - 1$. This is referred to as *competitor busting*. Finally, a more balanced approach is to bid somewhere in the middle so as to still push prices up but without running the risk of paying more than expected if one of the other bidders changes his bid. In order to define those strategies more precisely, let $p_s(i)$ be the price that player i has to pay when he bids so as to win slot s , given b_{-i} . In [18], the following greedy strategies were introduced and studied:

1. **Balanced Bidding (BB)**. In this scheme, bidder i first targets the slot s_i^* that maximizes his utility, i.e., $s_i^* \in \arg \max_s \{\theta_s(v_i - p_s(i))\}$. Given the desired slot, he then chooses his bid b for the next round so as to satisfy:

$$\theta_{s_i^*}(v_i - p_{s_i^*}(i)) = \theta_{s_i^*-1}(v_i - b).$$

The intuition is that player i should bid high enough so as to push the prices paid by his competitors up but at the same time it should not be the case that the utility of i decreases if the competitor right above him decides to bid just below b and i ends up at the higher slot $s_i^* - 1$.

2. **Restricted Balanced Bidding (RBB)**. This scheme is based on the same intuition as BB except that bidder i looks only at slots with no higher CTR than the slot he currently has. Hence, if his current slot is s_i , then he first targets the slot s_i^* that belongs to $\arg \max_s \{\theta_s(v_i - p_s(i)) : s \geq s_i\}$. Given s_i^* , he then chooses his bid b according to the same equation as in BB.
3. **Altruistic Bidding (AB)**. In this scheme, bidders are trying to not overcharge other players by bidding just what is necessary to get the slot they desire. Hence the slot s_i^* is selected just as in BB but then the bid b is chosen equal to $\min\{v_i, p_{s_i^*}(i) + \epsilon\}$ for some small $\epsilon > 0$.
4. **Competitor Busting (CB)**. This is the opposite of AB, and therefore bidders are simply trying to push prices as high up as possible. Again s_i^* is selected as in BB but then the bid b is set to $\min\{v_i, p_{s_i^*-1}(i) - \epsilon\}$. This strategy has been observed in practice and is also referred to as *anti-social* or *vindictive* bidding [14, 92].

Unfortunately, not all of the above strategies converge to some steady state and cycles may appear. For example, the schemes AB and CB do not always have a steady state. For more on this, see [18, 69]. However, BB and RBB do have nice convergence properties. In particular, both the BB and RBB processes have a unique fixed point, which is precisely the VCG equilibrium of GSP as defined in (4) of Sect. 2.

The positive convergence results as obtained in [18] are summarized as follows:

Table 4 Summary of results for four different greedy bidding strategies

Strategy	Fixed point	Convergence	Number of rounds	Efficiency loss
BB	Unique (efficient VCG equilibrium)	For asynchronous with random bid updating order, also for synchronous but with ≤ 2 slots	General CTRs: $O(n^{n+2^k})$ Geometric CTRs: $O(n^{n+k^3})$	None
RBB	Unique, same as above	Both for synchronous and asynchronous	General CTRs: $O(k2^k)$ Geometric CTRs: $O(k^3)$	None
AB	None	–	–	$(1 - \gamma)^{-1} + \gamma^{-1}$
CB	None	–	–	$(1 - \gamma)^{-1} + \gamma^{-1}$

Theorem 2

1. Both BB and RBB have a unique fixed point, at which players bid according to the VCG equilibrium.
2. RBB converges to the VCG equilibrium in both the synchronous (all bidders updating their bids simultaneously) and asynchronous (bidders updating their bids one by one) models.
3. BB converges to the VCG equilibrium in the synchronous model with 2 slots and in the asynchronous model when players bid in random order.

In general, RBB does not converge in polynomial time (in the number of bidders and the number of slots), but it does so when the CTRs are geometrically decreasing. As for BB, in the cases where it converges, the currently known theoretical upper bounds on the number of rounds that are required seem prohibitively high, see Table 4. However, the simulations presented in [18] reveal that these strategies typically converge quite fast. Finally for the synchronous model with at least 3 slots and for the asynchronous model where the order of updating is not random but a priori fixed, examples have been obtained that demonstrate the non-convergence of BB. For more details we refer the reader to [18].

Apart from the convergence to equilibrium under the assumptions of Theorem 2, other properties of BB have also been studied. In particular, the performance of BB and other greedy strategies are analyzed in Bayesian settings in [75, 89]. In [89], simulations are performed to evaluate greedy bidding strategies, under incomplete information. Their experimental results reveal that BB seems to be the most stable strategy, and, when all players follow BB, the game ends near an equilibrium, i.e., the additional gain from a deviation tends to be low. In [75], an extensive simulation study is conducted, assuming four different probability models for the distribution of click valuations, to assess the mismatching between the slot that advertisers aim for and what they actually obtain, and to evaluate the expected utility of advertisers. The study shows that an advertiser is typically not assigned the slot he was aiming for. In most cases the advertiser would get a larger profit if he were assigned a lower slot than the actual one. In fact, with lower slots advertisers get fewer clicks but also pay less, so that the profit may be larger with lower slots. The overall consequence is that in most cases advertisers get a higher-positioned slot than the optimal one and

pay more for something that will lead to lower profits. Such results are also proven to hold true (through the application of known results in the theory of order statistics) for the case of truthful bidding.

The properties of AB and CB have also been studied further, especially since CB is often encountered in practice. The pricing mechanism embedded in GSP appears to be prone to the Competitor Busting phenomenon, since an advertiser may raise his bid, thereby increasing the price paid by his competitor for the next higher slot, while suffering no consequences as to the price he is paying himself. In [47] a new pricing rule, named Penalized Second Price (PSP), has been proposed to alleviate CB. According to this rule the price paid by each advertiser is a linear combination of his own bid as well as of the next lower bid. With PSP an advertiser pays the consequences of his own aggressive strategy. In [76] it was shown that, when aggressive bidders playing CB do not have a budget advantage over bidders playing BB, the CB strategy is not beneficial, since it does not lead to aggressive bidders getting more slots over time. In [64], cooperative and vindictive bidding as well as existence of equilibria are studied for games where the utility of a bidder can express various levels of malicious behavior towards the other players. Finally in [69] the social welfare of configurations that are steady states with respect to AB and CB is considered and the authors obtain upper bounds on the inefficiency of such configurations.

To summarize, the properties of the strategies discussed are presented in Table 4. It should be noted that in the last column, the upper bounds in the efficiency loss for AB and CB hold for profiles that are steady states with respect to these two strategies. Even though AB and CB do not always converge, as we have already mentioned, the simulations in [18] and [69] show that in the majority of the cases, convergent states were found. Given the current literature, the main conclusion is that the balanced bidding strategies possess desirable properties both theoretically and experimentally. We feel however that the analysis of best response bidding strategies is far from complete. Even though there is a whole interval of bids that allow a player to obtain his best response slot, we are not aware of any further analysis of such strategies. Another interesting direction is to consider different bidding dynamics. Dynamic behavioral models have been analyzed successfully in other contexts, such as congestion games and load balancing games and some potential approaches in our context would be to study the performance of imitation dynamics [21], or no-regret algorithms [19]. For the latter, some recent progress has been made in [17] with regard to the social welfare achieved by such dynamics, as already pointed out in Sect. 2.2. Finally, it would also be interesting to investigate convergence to alternative solution concepts, other than Nash equilibria, such as convergence to sink equilibria, see e.g., [43]. An initial step has been taken in [15], where convergence to forward looking Nash equilibria has been considered.

3.2 Taking budgets into account

So far we have not taken into account budget considerations for the advertisers. In practice, advertisers can be required by the search engine to submit a budget, with the option of having it renewed at the end of a certain period. Of course, this budget can be so high that it does not constitute an actual constraint. Most companies, however, have to come to terms with the ensuing advertising costs and specifying a high cap

may not be sustainable for a long term period. In such cases, the advertiser should try not to follow extremely aggressive strategies as he may exhaust his budget before the renewal time.

In [13], Borgs *et al.* consider a model where, if the actual payment of the advertiser for the clicks on his ad exceeds a certain threshold (i.e. the advertiser's budget), then his utility collapses and becomes $-\infty$; each advertiser's valuation and budget are taken as private information. It is then proved that, under such hard budget constraints, it is not possible to design a truthful mechanism that allocates all the slots to different bidders, even in the case of two bidders and two slots. They also design an asymptotically optimal (in terms of revenue) mechanism that may not allocate all slots. Furthermore, in [7], Ashlagi *et al.* consider another private information model with budgets, whereby, if the actual payment of the advertiser exceeds his budget, then his utility vanishes to 0. These authors develop a modification of the Generalized Ascending Auction of [32], whose ex-post equilibrium outcome maintains the nice properties of the original design (see Sect. 2), despite the fact that the original design was not applicable to the case of budgets. Finally, a weakly dominant bidding strategy is considered in [80], where all bidders with budget constraints are led to state their true budget rather than understate their own valuations.

All the aforementioned works regarding budget considerations assume bidding is not repetitive; that is, advertisers submit their bids once, slot allocations and payments per click are determined, and then advertisers pay for all subsequent clicks accordingly. Contrary to this assumption, Drosos *et al.* propose in [29] a strategy for repetitive bidding for a single keyword auction based on dynamic programming. The objective is for bidders to carefully avoid overspending the available budget. Simulations are also conducted in order to evaluate the performance of this strategy, and comparisons are made with the balanced bidding protocol. The conclusions made so far reveal that the available budget can have an impact on how one should bid and can be particularly helpful for bidders who are not within the highest valuation range.

With this in mind, an interesting question that arises is to design bidding strategies that take this extra dimension into account. Despite its obvious applicability, these issues have not been extensively explored and we believe it is a topic worth further investigation.

3.2.1 Budgets for multiple keywords

A yet more realistic model is to assume that a budget needs to be split among several keywords. In practice advertisers select a set of keywords and participate in all the corresponding auctions. For example a company that sells digital devices may wish to appear on queries for laptop, digital camera, mp3 player, etc. Hence for a set of relevant keywords, each advertiser i should specify his bid b_{ij} on each keyword j . At the same time, the bids should be such that the resulting payments should not exceed the total budget of the advertiser.

From an optimization viewpoint, there has been a series of papers on designing algorithms for various settings regarding such budget-constrained bidders. For revenue maximization of the search engine see among others [72], where online algorithms are designed and their competitive ratio is analyzed, i.e., the ratio of the optimal value

of the objective function to that obtained by the algorithm. The question of maximizing the profit of a single advertiser is studied in [20, 33, 73]. These works concern either stochastic models where the advertiser has some information about the other bidders' behavior in the form of some distributions for the cost of obtaining a certain slot [33, 73], or online models [20], where the bids of the other advertisers are known and in each round the bidder has to decide which slot to target.

From a game theoretic viewpoint, we are only aware of [12], where a bidding strategy is proposed and is proved to converge in some cases to a market equilibrium. That is, the prices attained are such that the seller (i.e., the search engine here) sells his entire supply, while demand in these prices equals supply. In the model of [12], every advertiser has a budget which is renewed at the beginning of every round (e.g., daily). Advertisers need to choose simultaneously a bid b_{ij} for every keyword j and the search engine selects the winners of each auction taking into account that no advertiser can pay more than his budget. The authors propose a natural bidding heuristic that is based on equalizing the marginal *return-on-investment* (ROI) across all keywords. To this end, they also add a random perturbation in order to avoid cycles that may appear when all bidders use this heuristic. It is proved that when everybody adopts the perturbed ROI heuristic, the system converges to its market equilibrium in the case of the first price mechanism with a single slot. In the case of the second price mechanism on a single slot, experiments reveal that the system converges, but no theoretical results have been obtained. It is an interesting open problem to obtain theoretical results for the second price mechanism on one slot and more generally for the GSP mechanism in the case of multiple slots.

4 User models and externalities among bidders

In this section, we focus on two interrelated topics: models for user behavior and externalities among bidders. The literature that we have discussed so far has ignored the behavior of the end users and is based on the assumption that CTRs are separable: the CTR of a bidder i in slot s is the product of two quantities, the first expressing the quality of the bidder and the second the quality of the slot he occupies ($q_i \cdot \theta_s$). Most other articles are also based on that assumption, thus defining CTR as a function of the bidder i and the slot s even when not assuming separability. Such assumptions however are not always justified. As an example, if a user searches for a commercial product and decides to click first on the ads on the top of the list, he may not end up clicking on the last ad if he finds what he was looking for before reaching the bottom slot. Hence, the CTR of an advertiser is clearly dependent on the search behavior of the users and some recent works have focused on developing models of user behavior that are consistent with empirical observations.

Apart from the user behavior, the CTR is also crucially dependent on the quality of the other advertisers that are present. Advertisers offering similar products create positive or negative externalities to their competitors, depending on the satisfaction that a user receives by clicking on their ad. This calls for the design of new auction mechanisms. Externalities in general settings of auctions have been studied before in the economics literature, admittedly though not to a great extent. The earliest work

that we are aware of is by Jehiel *et al.* [52], where the value of a loser depends on the identity of the winner. See also [53] for a follow up work of these authors on the topic. In the context of online advertising and in relation to sponsored search auctions, the first work appeared in [38], where a model was presented for online lead generation. In their setting, advertisers receive leads about potential customers whom they can contact and offer quotes about their service. After seeing the quotes the user selects the advertiser with the lowest quote.

In this section, we discuss models that have been proposed specifically for sponsored search auctions. We start in Sect. 4.1 with user models that have been studied recently in the literature, and especially the sequential search model, along with the externalities that they create. Then we move on, in Sect. 4.2, to discuss alternative models for externalities along with the corresponding mechanisms.

4.1 User behavior models in sponsored search auctions

Regarding user behavior, one line of research has focused on identifying user intentions, as clearly not all users are interested in clicking on the ads or making a purchase. In [6], the authors use click-through data and learning techniques to classify search queries into commercial/non-commercial and navigational/informational. This approach allows for better predictions of CTRs, for a given query with particular intentions. For more on detecting user intentions, see also [5].

A different approach is introduced in [54], where a game theoretic model is presented. End users are viewed as rational agents in a game played under uncertainty (here uncertainty refers to the fact that users do not know the value of clicking on an ad). Each user then decides sequentially on which ad to click on so as to maximize his expected utility under uncertainty. The authors also provide an empirical investigation based on a dataset of Microsoft Live from 2007 and estimate the parameters of their model.

The majority of the remaining works on user models has focused on the so called *sequential search* model and its variants, motivated by the experimental work of [24], which in turn was inspired by the eye-tracking experiments of [55], as described below. The main elements of this model are that the users (i) browse the sponsored links from top to bottom and (ii) they make clicking decisions slot by slot. After reading each ad, users decide whether to click on it or not and, subsequently, decide whether to continue browsing the sponsored list or to simply skip it altogether.

The basic sequential search model and its variants The first model for ordered search was introduced and studied empirically in [24]. This formed the baseline for the more general version that we present here, which was introduced independently in [56] and [2].

Formally, we assume that there is an intrinsic quality q_i of each advertiser i , specifying the probability that a user will click on i when he reaches the slot where i 's ad is displayed. Furthermore, there is also a *continuation probability* c_i that specifies the probability that the user continues to the next slot after looking at i 's ad, (and possibly clicking on it). Finally, v_i is the valuation of advertiser i for a click. Suppose now that the slots $1, \dots, k$ contain the ads a_1, \dots, a_k respectively. Then the user will behave as follows:

1. He starts by looking at ad a_1 of slot 1 and clicks on it with probability q_{a_1} .
2. Independently of whether ad a_1 was clicked or not, he continues to the ad a_2 with probability c_{a_1} , otherwise the process ends with probability $1 - c_{a_1}$.
3. He repeats steps 1 and 2 for the following slots a_2, a_3, \dots , till the process terminates.

The focus on such an ordered search model is motivated by various reasons. First, as the work of [24] demonstrates, position bias is present in organic search. In particular, [24] compares a sequential search model with four other models (including the separable model) and concludes that sequential search provides the best fit to the click logs they have considered. Secondly, sequential search is further advocated as a natural way to browse through a list of ads by the eye-tracking experiments of Joachims *et al.* [55], where it is observed that users search and click in a top down manner. Moreover, as the value per click of each advertiser tends to be correlated with its relevance, ordered search is a good heuristic for users (see [9]).

Under this model, which is also referred to as the cascade model, the willingness to click on an ad changes as a user collects new information through his search, and hence the decision about whether to continue reading ads naturally depends on the click history of the user. Hence the CTR of ad a_s , placed on position s is:

$$R_{a_s} = q_{a_s} \cdot \prod_{j=1}^{s-1} c_{a_j}.$$

Other variations have also been proposed, introducing more parameters and generalizing the basic model. These involve

1. Adding slot-dependent CTRs. This was studied in [56] and allows for the presence of an additional parameter θ_s , the probability of clicking an ad at slot s .
2. Splitting the continuation probability in two parameters. This was studied in [44] and assumes that there is a different continuation probability when a user clicks on an ad and a different parameter when the user looks at the ad and decides not to click on it and continue to the next ad.
3. Considering the dependence of CTR on history of clicks. This was also studied empirically in [44] and is based on having the parameter q_a of advertiser a depend on the clicking history of the user.
4. Allowing multiple ad slates. This was introduced in [56] and captures the fact that nowadays sponsored links are displayed both on the right hand side but also on the top of the search results. As a result, there can be different groups of users, depending on whether they first scan the top results and then the ones on the right hand side or vice versa. This can be further generalized by allowing different groups of users to scan the ads in different orders.
5. Considering the Pay-Per-Transaction model instead of the usual Pay-Per-Click model. This was studied in [58], where further comparisons between the VCG and the GSP mechanisms were investigated as well as issues of robustness to manipulations under this model.

The results that have been obtained so far can be split into two categories: algorithm design for finding the optimal allocation in the basic model and its variants,

and equilibrium analysis of the GSP and related mechanisms under this type of user behavior. These are overviewed below.

The winner determination problem This is the problem of finding the allocation of slots to advertisers that achieves the highest social welfare for the advertisers. In [2] and [56] it was established that the efficient allocation (i.e., the one that achieves optimal social welfare for the advertisers) can be found in polynomial time by means of dynamic programming.

Theorem 3 *The winner determination problem can be solved in polynomial time in the sequential search model.*

As is noted in [56], the same is true in the variant where there are two types of continuation probabilities, since in the dynamic programming algorithm the two probabilities act cumulatively. In [56], the problem is also studied for two more variants of the basic model. The first one is the case of multiple slates, where a polynomial time approximation scheme is presented. It is still an open problem to determine whether the winner determination problem is NP-hard. The second variant is in the presence of slot-dependent CTRs. In this case, a 4-approximation algorithm is established, as well as a quasi-polynomial time approximation scheme. Again, it is still not known if this variant is NP-hard. Determining the exact complexity of the winner determination problem in these variants is an interesting open problem. It is also interesting to note here that all the algorithms for these variants are based on Knapsack-related problems.

Equilibrium analysis Beyond the algorithmic question of finding the optimal allocation, it is of natural interest to study how the equilibria of the GSP mechanism are affected by the user behavior or investigate mechanisms that take into account the user behavior and the continuation probabilities. In both such cases, the properties of the CTRs arising as a consequence of such behavior by the users should be taken into account. The first equilibrium analysis under the more general model that also includes slot-dependent CTRs was obtained in [40]. The authors proved that pure Nash equilibria still exist. However, in contrast to the usual models presented in Sect. 2 we cannot guarantee that there exist equilibria that implement the optimal allocation along with the VCG payments. In particular, it is proved that the social welfare of an equilibrium can be as far as a factor k away from optimal (k being the number of slots) for equilibria where bidders never overbid and it can be arbitrarily far from optimal if there are no restrictions on the bids. In [44], the implementation of efficient or revenue-maximizing allocations is studied for various scoring rules in the generalization of the basic model that allows for two types of continuation probabilities. A scoring rule is simply any ranking scheme in which the ranking of the bidders is performed according to the product $w_i b_i$, where w_i is a weight that depends only on advertiser i . The authors of [44] identify a profile of bidding strategies that constitutes a revenue-maximizing and efficient equilibrium if and only if the scoring rule used by the search engine has a particular form that depends on both q_i and the continuation probabilities. Namely, the weight in this rule should be a multiple of $q_i/(1 - c_i)$.

Interestingly, this is the same ranking rule by which the winners should be ranked in [2, 56] for solving the efficient allocation problem (in a non-strategic environment). They also extend the negative result of [40] showing that no scoring rule can implement an efficient equilibrium where advertisers pay their VCG payments for all valuations and search parameters. Finally, in [27] the rule of ranking by the weight $q_i/(1 - c_i)$ is investigated further. A particular pure equilibrium is constructed and its efficiency properties are studied.

An interesting open question here is to obtain a complete characterization or a better understanding of the set of Nash equilibria under these user behavior models for any scoring rule. It is also interesting to see what happens with regard to revenue maximization when the weights of the scoring rule are not given by $q_i/(1 - c_i)$. In addition, more experimental analysis with real data, following [44], would be very valuable for providing further validation to this model. Finally, it would be nice to generalize the model of sequential search. A candidate abstract setting for this would be to think of the user as moving in a Markov chain, so that a user can not only visit the next ad in the list but also jump with a certain probability to the other ads.

4.2 Models of bidder externalities

In this section, we study alternative models, where the externalities among bidders are not derived explicitly from the user behavior, but are nevertheless motivated by such considerations.

Recently, in [35] a quite general model has been presented, where externalities are modeled by a social context graph. The graph specifies two disjoint sets of edges, E^+ and E^- ; an edge from bidder i to bidder j indicates a positive (resp. negative) externality if $(i, j) \in E^+$ (resp. $(i, j) \in E^-$). Each edge also has a weight which depends on the distance between the two advertisers if they are displayed on the same impression. Hence, the closer the advertisers, the stronger the effect, whether positive or negative. An additional parameter of the model is a constant c , indicating that there are no effects between advertisers who are at a distance higher than c . Essentially, this implies that the underlying assumption about the users is that they are allowed to browse a bounded scope section of consecutive ads in the list (at most c) but no other restriction is made on the order in which they visit the ads. Hence an advertiser cannot influence other advertisers who are far away in the list of impressions. The authors of [35] show that the winner determination problem is NP-hard in this model, unlike the sequential search model, and provide a polynomial time approximation algorithm and an exact algorithm, which is polynomial when the number of slots is relatively small. Finally, they also study game theoretic aspects and revenue considerations, where some negative results are obtained.

Another type of externality is considered in [39]: the value of a click is supposed here to depend on exclusivity, i.e., it is larger when the ad is the only one displayed, as it is more likely that a click will be converted into a sale. In such a context, the authors suggest to use two-dimensional bids: one value stands for the case when only that ad is displayed, and the other value of the bid corresponds to the classical auction schemes when several ad slots are used. The auctioneer then has to decide whether to display one or several ads, how to allocate the slots, and to compute payments.

The authors study two GSP-inspired mechanisms with two-dimensional bids; the first one coincides with the one-dimensional GSP scheme when the outcome is a multiple ad display, while the second one extends the “next-price” GSP rule according to which each participant pays the minimum price necessary to keep its position. The authors consider equilibria where losers bid at least their true valuation. For the former scheme, at any such equilibrium the revenue is at least half what a VCG scheme would give, and efficiency is at least $1/3$ of the optimal. For the latter scheme, if bidders do not play dominated strategies and losers bid at least their true value, then any equilibrium has efficiency larger than half the optimal; moreover there exists an equilibrium yielding as much revenue as VCG.

Also, in [23]. Constantin *et al.* introduce a model of negative externalities of the values per click. In particular, each bidder can submit a set of constraints on his position relatively to that of certain other bidders (for example, he may insist on being allocated a higher slot than a certain competitor). The authors assume that the bidder will pay his bid b_i provided that all his submitted constraints apply under the allocation, otherwise he will pay 0. The authors mostly focus on the case where each constraint submitted by a bidder is related to the position of one more bidder. They investigate a greedy winner determination algorithm applicable under such constraints and show that it is not possible to achieve truthfulness on the declaration of both the bidder's value and his constraints, even under VCG-type payments. On the other hand, a GSP payment rule would achieve truthful declaration of the constraints, provided that bidders have downward-monotonic value externalities (that is, if a slot is not acceptable under certain conditions, lower slots are not acceptable either). The authors also investigate other forms of value externalities.

To summarize our discussion of Sects. 4.1 and 4.2, one can see that even though in the early history of sponsored search user behavior was not taken into account, it has by now evolved to an important research dimension. Undoubtedly user behavior creates externalities that can be observed in practice and have motivated the models that have been proposed so far. The sequential search model has attracted the most attention so far, however it is quite challenging to reach a conclusion as to which model captures in a better way the real life scenarios and user behaviors. Apart from [24, 44, 54], the rest of the works discussed here propose theoretical models without providing any experimental findings. The nature of this topic also makes it more time consuming to reach well established experimental conclusions, as one needs to observe users' behavior over substantial time periods. Hence, it remains for future works to provide better evidence on this matter and establish what the best model is.

5 Statistical learning techniques

Keyword auctions are held repeatedly, and the participating set of bidders may be largely overlapping from one auction round to the next one. In such repeated games the information gathered on the bidders' behavior may be exploited in later rounds. Hence both the auctioneer and the bidders have an interest in such information. In fact, each bidder may use the information on other bidders' bids to modify his own bid, in order to get a more preferable slot or to get the same slot for less. A particular

role is played by the auctioneer, which, among all the stakeholders, has the largest information dataset: knowing the bids submitted by each bidder and the ads clicked on by the customers, he is in the best position for (and has the largest interest in) learning two very important quantities revealing the bidders' and users' preferences, namely the values attached by bidders to clicks, and the ads and slots preferred by the customers. In this section we deal with both issues, reviewing the most important techniques employed for those purposes.

5.1 Learning advertisers' valuations

The knowledge of the value that advertisers attach to clicks is crucial for the proper management of keyword auctions. Auctions allow for differential pricing, whereby the seller (the search engine) can extract the largest possible income from the sale, by having the prospective buyers declare what they are willing to pay. In truthful mechanisms (such as the VCG one) bidders are induced into declaring their valuation. But in pricing mechanisms such as the GSP, the bidding strategies lead each bidder to submit bids that are lower than his valuation, hence retaining a surplus margin. The knowledge of the true valuation would allow the search engine to further exploit the willingness to pay by the bidders, e.g., to estimate a revenue-maximizing reserve price. However, in non-truthful pricing mechanisms, the bids are observed whereas the valuations are not and have then to be estimated.

The characteristics of the valuations are typically described by resorting to either of two paradigms:

- *Common Value*;
- *Independent Private Value (IPV)*.

In the *Common Value* case all the items of the auction have the same value for all the bidders, who however have incomplete information about it and then try to estimate it. In the *IPV* case the value of the item is different for each bidder, but all the values can be considered as independent random variables drawn from the same probability distribution [71]. In the context of keyword auctions the *IPV* paradigm has been adopted, e.g., in [32, 75, 79]. Under the *IPV* assumption the problem is then the estimation of the probability distribution of valuations.

The problem of estimating the distribution of valuations has been investigated mainly for first price auctions. As far as the authors know, no approach has been proposed for the context of GSP auctions. Hence, in the following we briefly overview the two main approaches proposed in the literature for first-price auctions. They are due respectively to Guerre et al. [49] and to Marmer and Shneyerov [70].

In the work of Guerre et al. [49] the bidders are assumed to adopt a Bayesian Nash equilibrium strategy, whereby the bid b_i submitted by the generic bidder i (among N bidders) is determined by its private valuation v_i and the cumulative distribution function of valuations (cdf) $F(v)$ (under the *IPV* paradigm) as follows

$$b_i = v_i - \frac{1}{[F(v_i)]^{N-1}} \int_{p_0}^{v_i} [F(u)]^{N-1} du, \quad (6)$$

where p_0 is the reserve price, i.e., the minimum accepted bid. This relationship can be inverted to provide the individual valuation as a function of the individual bid and the

probability distribution and density of bids. By repeating this inversion procedure for a number of bids, we obtain a sample of valuations (pseudo-values), each pertaining to an individual bid. This sample can finally be employed to get a nonparametric estimate of the probability density function of private values (namely through the kernel method, see [84]).

A similar approach is proposed by Marmer and Shneyerov [70]. They consider again a first-price sealed-bid auction, with the same hypotheses as Guerre. However, they avoid the use of pseudo-values, and arrive at the estimation of the probability density function of valuations by using the non-parametric estimators of the pdf, cdf, and quantiles of bids. By exploiting the monotonicity of the inverse bidding strategy exploited by Guerre, they introduce the following relationship between the quantile function of valuations $Q(\tau)$ and that of bids $q(\tau)$

$$Q(\tau) = q(\tau) + \frac{\tau}{(N - 1)w(q(\tau))}. \tag{7}$$

The estimate of the pdf of valuations is then obtained by estimating first the quantile function of bids, using then expression (7) to get the quantile function of valuations, from which we can finally obtain the cdf and the pdf of valuations.

To summarize, two approaches have been proposed in the literature for learning advertisers' valuations in First-Price auctions. It should be noted though that both of these approaches can be also adopted for GSP auctions, if an expression linking valuation and bids is available and invertible. In fact, the steps involved in both methods after the inversion procedure are quite general and do not rely on any assumption on the auction pricing method. It is therefore an interesting open question to successfully apply these techniques to GSP mechanisms.

5.2 Click-through-rate estimation

The CTR is a measure of the interest of customers for a given ad or a measure of how often they click on a given slot. For the case of a bidder's ad, if we indicate by x the number of times the ad is clicked on and by y the number of impressions on which that ad appears, the CTR is measured as the ratio

$$CTR = \frac{x}{y}. \tag{8}$$

The methods we present here can be applied both to the bidder-dependent $CTR(q_i)$ of a bidder i , as well as to the slot-dependent $CTR(\theta_s)$ of a slot s . Clearly the variables x, y in the above estimation can be adjusted according to the case we are interested in. From now on, we focus on the estimation of an ad's CTR (i.e., of a bidder).

As reported in [46], the typical average CTR is around 2 %. With such quite low values, the estimate will be characterized by a large variance. In the simple example reported by Richardson et al. [81], if the true CTR is 5 %, we need 1000 impressions to have an estimated CTR within ± 1 % of the true value with an 85 % confidence level.

Operationally, the ratio (8) can be measured in three different ways:

1. Setting a time interval T and measuring the impressions and the clicks taking place within that interval (average over a fixed time window);
2. Setting a limit number of impressions and measuring the number of clicks observed till we reach that limit number (average over a fixed impression window);
3. Setting a limit number of clicks and measuring the number of impressions needed to get that limit number (average over a fixed click window).

In addition to these straightforward estimates, Immorlica et al. [50] introduce an exponential discounting estimate that is the weighted average of the clicks observed over all the impressions, so that the weights favor the most recent impressions. If we indicate by α the weighting parameter, and by x_i the indicator variable taking value 1 if the i th most recent impression resulted in a click and 0 otherwise, this estimator is

$$CTR = \frac{\sum_i x_i \exp(-\alpha i)}{\sum_i \exp(-\alpha i)}. \tag{9}$$

Obviously the most recent impressions have a better weight under this scheme. In fact, in [50] it is shown that all the above estimators fall in the general class of estimates of the form:

$$CTR = \frac{\sum_i x_i \delta(i, t_i, c_i)}{\sum_i \delta(i, t_i, c_i)}. \tag{10}$$

In the above formula, $\delta(\cdot)$ is a decreasing function in all three parameters, where recall that index i corresponds to the i th most recent impression. The parameter c_i is the number of impressions that received clicks between impression 1 and impression i , and t_i is the time elapsed between impression 1 and impression i .

As stated in [50], all the above methods provide an estimate arbitrarily close to the true CTR for an appropriate setting of parameters (e.g., a large enough number of impressions in the average-over-impressions method). Though none appears more preferable on the basis of its accuracy, the picture changes when we consider other desirable properties, such as fraud-resistance. This is the capability to maintain a correct estimate of the CTR, while a competitor generates clicks on an ad with the sole intent of increasing the payment of the advertiser holding that ad. Not all the estimators in the class defined above satisfy this additional requirement. As shown in [50], under some natural assumptions on the function $\delta(\cdot)$, the subfamily of the so-called click-based estimators, which includes the average-over-clicks method, exhibits fraud-resistance. Examples are also provided showing that methods not belonging to this particular subfamily may fail to be fraud-resistant.

Taking a different approach, one drawback of the estimators defined above is that they fail to highlight the relationship between the CTR and its main determinants. This is the purpose of the estimator proposed in [26], where the CTR is considered as a function of the ad itself, the ad's position (the slot), and the page on which the ad appears. The estimator is then obtained through the maximum likelihood approach. The real-time applicability on a massive scale of the estimator proposed in [26] has not been investigated yet and it remains to be seen in practice. On the other hand, the

family of estimators defined by (10) appears quite simple to implement and with a minimal computational effort.

Despite all these different available methods, we may still not have a satisfactory estimate of the CTR. For example, we can have the following two cases:

- The advertiser and the publisher (the search engine in our case) have different estimates of the CTR;
- The advertisement is relatively new and the estimation of its CTR is based on a short time sample.

In the first case there is a discrepancy between what the advertiser expects to pay and what the publisher expects to receive. In order to reduce the effects of divergence in valuations, Goel *et al.* [42] have introduced contract auctions, which generalize the classical second price auction. In particular, they propose an impression-plus-click pricing mechanism, in which advertisers pay a fixed amount per impression plus an additional amount if their ad is clicked.

In the second case, the publisher faces conflicting requirements when trying to efficiently allocate the ad space and simultaneously estimate the CTR. Hence, he has to strike a balance between exploring (i.e., showing an ad to get a better estimate of its CTR) and exploiting (i.e., showing ads that have the best performance, according to its current estimates of the CTRs). In [63] it is shown that an advertiser has an incentive to increase his bid by some amount when the search engine has not done enough exploration, which the authors call the value of learning.

To summarize, the CTR is an important parameter to estimate, for reasons of predicting the revenue from ads and of avoiding fraud. We have overviewed several methods of estimating the CTR, and discussed the specificities associated with this parameter, e.g. due to the lack of an adequate number of observations. There still remain some interesting open questions, particularly regarding how to incorporate the main determinants of the ad in the estimation of its CTR.

6 Conclusions

We have presented an overview of research that has been conducted in sponsored search auctions mainly in the last five years. Our overview has focused more on game theoretic aspects and strategic considerations of the interacting entities. We believe this is a promising area for future research as can be also evidenced by the annual workshops on ad auctions (see e.g. [1] for the latest one).

Apart from theoretical analysis, we have also included empirical considerations in our survey. It is undoubtedly very important to perform experimental analysis with empirical data, which however are rarely available publicly. Future research should both incorporate to a greater extent the recent findings of work in theoretical analysis, and study additional empirical datasets, which will hopefully become publicly available.

Acknowledgements This work has been partly funded by the Euro-NF Network of Excellence, through the specific research project AMESA. The authors thank the anonymous referees for their constructive comments.

References

1. 8th ad auctions workshop. <http://sites.google.com/site/adauctions2012/>.
2. Aggarwal, G., Feldman, J., Muthukrishnan, S., & Pál, M. (2008). Sponsored search auctions with Markovian users. In *Workshop on Internet and network economics (WINE)* (pp. 621–628).
3. Aggarwal, G., Goel, A., & Motwani, R. (2006). Truthful auctions for pricing search keywords. In *ACM conference on electronic commerce (EC)* (pp. 1–7).
4. Apex Pacific. <http://www.apexpacific.com/index.html>.
5. Ashkan, A., & Clarke, C. L. A. (2009). Characterizing commercial intent. In *Conference on information and knowledge management (CIKM)* (pp. 67–76).
6. Ashkan, A., Clarke, C. L. A., Agichtein, E., & Guo, Q. (2009). Classifying and characterizing query intent. In *European conference on information retrieval research (ECIR)* (pp. 578–586).
7. Ashlagi, I., Braverman, M., Hassidim, A., Lavi, R., & Tennenholtz, M. (2010). Position auctions with budgets: existence and uniqueness. *B.E. J. Theor. Econ. Adv.*, 10(1). doi:10.2202/1935-1704.1648
8. Ashlagi, I., Monderer, D., & Tennenholtz, M. (2011). Simultaneous ad auctions. *Mathematics of Operations Research*, 36(1), 1–13.
9. Athey, S., & Ellison, G. (2011). Position auctions with consumer search. *Quarterly Journal of Economics*, 126(3), 1213–1270.
10. Babaioff, M., & Roughgarden, T. (2010). Equilibrium efficiency and price complexity in sponsored search auctions. In *Proc. of 6th workshop on ad auctions*, Cambridge, MA, USA, June 2010.
11. Bhawalkar, K., & Roughgarden, T. (2011). Welfare guarantees for combinatorial auctions with item bidding. In *ACM-SIAM symposium on discrete algorithms (SODA)* (pp. 700–709).
12. Borgs, C., Chayes, J., Etesami, O., Immorlica, N., Jain, K., & Mahdian, M. (2007). Dynamics of bid optimization in online advertisement auctions. In *International conference on world wide web (WWW)* (pp. 531–540).
13. Borgs, C., Chayes, J., Immorlica, N., Mahdian, M., & Saberi, A. (2005). Multi-unit auctions with budget-constrained bidders. In *ACM conference on electronic commerce (EC)* (pp. 44–51).
14. Brandt, F., & Weiß, G. (2001). Antisocial agents and Vickrey auctions. In *8th international workshop on agent theories, architectures and languages (ATAL)* (pp. 335–347).
15. Bu, T., Deng, X., & Qi, Q. (2008). Forward looking Nash equilibrium for keyword auction. *Information Processing Letters*, 105(2), 41–46.
16. Caragiannis, I., Kaklamanis, C., Kanellopoulos, P., & Kyropoulou, M. (2011). On the efficiency of equilibria in generalized second price auctions. In *ACM conference on electronic commerce (EC)* (pp. 81–90).
17. Caragiannis, I., Kaklamanis, K., Kanellopoulos, P., Kyropoulou, M., Lucier, B., Leme, R. P., & Tardos, E. (2012). On the efficiency of equilibria in generalized second price auctions. [arXiv:1201.6429](https://arxiv.org/abs/1201.6429).
18. Cary, M., Das, A., Edelman, B., Giotis, I., Heimerl, K., Karlin, A., Mathieu, C., & Schwarz, M. (2007). Greedy bidding strategies for keyword auctions. In *ACM conference on electronic commerce (EC)*, San Diego, California, USA.
19. Cesa-Bianchi, N., & Lugosi, G. (2006). *Prediction, learning and games*. Cambridge: Cambridge University Press.
20. Chakrabarty, D., Zhou, Y., & Lukose, R. (2008). Budget constrained bidding in keyword auctions and online knapsack problems. In *Workshop on Internet and network economics (WINE)* (pp. 566–576).
21. Chung, C., & Pyrga, E. (2009). Stochastic stability in Internet router congestion games. In *Symposium on algorithmic game theory (SAGT)* (pp. 183–195).
22. Clarke, E. H. (1971). Multipart pricing of public goods. *Public Choice*, 11, 17–33.
23. Constantin, F., Rao, M., Huang, C.-C., & Parkes, D. (2011). On expressing value externalities in position auctions. In *AAAI*.
24. Craswell, N., Zoeter, O., Taylor, M., & Ramsey, B. (2008). An experimental comparison of click position-bias models. In *ACM international conference on web search and data mining (WSDM)* (pp. 87–94).

25. Demange, G., Gale, D., & Sotomayor, M. (1986). Multi-item auctions. *Journal of Political Economy*, 94, 863–872.
26. Dembczyński, K., Kotlowski, W., & Weiss, D. (2008). Predicting ads' click-through-rate with decision rules. In *Proc. 17th international conference on world wide web (WWW)*.
27. Deng, X., & Yu, J. (2009). A new ranking scheme of the GSP mechanism with Markovian users. In *Workshop on Internet and network economics (WINE)* (pp. 583–590).
28. DoubleClick. <http://www.google.com/doubleclick/>.
29. Drosos, D., Markakis, E., & Stamoulis, G. D. (2012). *Budget constrained bidding in sponsored search auctions*. Working paper.
30. Easley, D., & Kleinberg, J. (2010). *Networks, crowds and markets: reasoning about a highly connected world*. Cambridge: Cambridge University Press.
31. Edelman, B., & Ostrovsky, M. (2007). Strategic bidder behavior in sponsored search auctions. *Decision Support Systems*, 43, 192–198.
32. Edelman, B., Ostrovsky, M., & Schwarz, M. (2007). Internet advertising and the generalized second-price auction: selling billions of dollars worth of keywords. *American Economic Review*, 97(1), 242–259.
33. Feldman, J., Muthukrishnan, S., Pal, M., & Stein, C. (2007). Budget optimization in search-based advertising auctions. In *ACM conference on electronic commerce (EC)* (pp. 40–49).
34. Feng, J., Bhargava, H. K., & Pennock, D. M. (2007). Implementing sponsored search in web search engines: computational evaluation of alternative mechanisms. *INFORMS Journal on Computing*, 19(1), 137–148.
35. Fotakis, D., Krysta, P., & Telelis, O. (2011). Externalities among advertisers in sponsored search. In *Symposium on algorithmic game theory (SAGT)* (pp. 105–116).
36. Fudenberg, D., & Tirole, J. (1991). *Game theory*. Cambridge: MIT Press.
37. Fukuda, E., Kamijo, Y., Takeuchi, A., Masui, M., & Funaki, Y. (2010). Theoretical and experimental investigation of performance of keyword auction mechanisms. In *Proc. of 6th workshop on ad auctions*, Cambridge, MA, USA.
38. Ghosh, A., & Mahdian, M. (2008). Externalities in online advertising. In *Proc. 17th international world wide web conference (WWW)* (pp. 161–168).
39. Ghosh, A., & Sayedi, A. (2010). Expressive auctions for externalities in online advertising. In *Proc. 19th international conference on world wide web (WWW)*, Raleigh, NC, USA, April 2010 (pp. 371–380).
40. Giotis, I., & Karlin, A. (2008). On the equilibria and efficiency of the GSP mechanism in keyword auctions with externalities. In *Workshop on Internet and network economics (WINE)* (pp. 629–638).
41. Goel, A., & Munagala, K. (2009). Hybrid keyword search auctions. In *Proc. 18th international conference on world wide web (WWW)* (pp. 221–230).
42. Goel, S., Lahaie, S., & Vassilvitskii, S. (2009). Contract auctions for sponsored search. In *Workshop on Internet and network economics (WINE)*, Rome, Italy (pp. 196–207).
43. Goemans, M. X., Mirrokni, V. S., & Vetta, A. (2005). Sink equilibria and convergence. In *Symposium on foundations of computer science (FOCS)* (pp. 142–154).
44. Gomes, R., Immorlica, N., & Markakis, E. (2009). Externalities in keyword auctions: an empirical and theoretical assessment. In *Workshop on Internet and network economics (WINE)* (pp. 172–183).
45. Gonen, R. (2009). Characterizing optimal syndicated sponsored search market design. In *Proc. of the 5th workshop on ad auctions*, Stanford, CA, USA, July 2009.
46. Green, D. (2003). Search engine marketing: why it benefits us all. *Business Information Review*, 20(4), 195–202.
47. Grillo, A., Lentini, A., Naldi, M., & Italiano, G. F. (2010). Penalized second price: a new pricing algorithm for advertising in search engines. In *8th annual conference on communication networks and services research (CNSR 2010)*, Montreal, 11–14 May 2010 (pp. 207–214).
48. Groves, T. (1973). Incentives in teams. *Econometrica*, 41(3), 617–631.
49. Guerre, E., Perrigne, I., & Vuong, Q. (2000). Optimal nonparametric estimation of first-price auctions. *Econometrica*, 68(3), 525–574.
50. Immorlica, N., Jain, K., Mahdian, M., & Talwar, K. (2005). Click fraud resistant methods for learning click-through rates. In *Workshop on Internet and network economics (WINE)* (pp. 34–45).
51. Internet advertising bureau. <http://www.iab.net>.
52. Jehiel, P., Moldovanu, B., & Stacchetti, E. (1996). How (not) to sell nuclear weapons. *American Economic Review*, 86(4), 814–829.

53. Jehiel, P., Moldovanu, B., & Stacchetti, E. (1999). Multi-dimensional mechanism design for auctions with externalities. *Journal of Economic Theory*, 85(2), 258–294.
54. Jeziorski, P., & Segal, I. (2009). *What makes them click: empirical analysis of consumer demand for Internet search advertising*. Working paper.
55. Joachims, T., Granka, L., Pan, B., Hembrooke, H., & Gay, G. (2005). Accurately predicting click-through data as implicit feedback. In *ACM SIGIR* (pp. 154–161).
56. Kempe, D., & Mahdian, M. (2008). A cascade model for externalities in sponsored search. In *Workshop on Internet and network economics (WINE)* (pp. 585–596).
57. Koutsoupias, E., & Papadimitriou, C. H. (1999). Worst-case equilibria. In *STACS* (pp. 404–413).
58. Kuminov, D., & Tennenholtz, M. (2009). User modeling in position auctions: Re-considering the GSP and VCG mechanisms. In *International conference on autonomous agents and multiagent systems (AAMAS)* (pp. 273–280).
59. Lahaie, S. (2006). An analysis of alternative slot auction designs for sponsored search. In *Proc. seventh ACM conference on electronic commerce (EC)*, Ann Arbor, Michigan, USA (pp. 218–227).
60. Lahaie, S., & Pennock, D. (2007). Revenue analysis of a family of ranking rules for keyword auctions. In *Proc. ACM conference on electronic commerce (EC)*, San Diego, California, USA, June 2007 (pp. 50–56).
61. Lahaie, S., Pennock, D., Saberi, A., & Vohra, R. (2007). Sponsored search auctions. In N. Nisan, T. Roughgarden, E. Tardos, & V. Vazirani (Eds.), *Algorithmic game theory* (pp. 699–716). Cambridge: Cambridge University Press. Chap. 28.
62. Leonard, H. (1983). Elicitation of honest preferences for the assignment of individuals to positions. *Journal of Political Economy*, 91, 461–478.
63. Li, S.-M., Mahdian, M., & McAfee, R. P. (2010). Value of learning in sponsored search auctions. In *Workshop on Internet and network economics (WINE)* (pp. 294–305).
64. Liang, L., & Qi, Q. (2007). Cooperative or vindictive: bidding strategies in sponsored search auction. In *Workshop on Internet and network economics (WINE)* (pp. 167–178).
65. Linden, G., Meek, C., & Chickering, M. (2009). The pollution effect: optimizing keyword auctions by favoring relevant advertising. In *Proc. of the 5th workshop on ad auctions*, Stanford, CA, USA.
66. Ling, M. C., & Lawler, K. (2001). Internet advertising, game theory and consumer welfare. *Electronic Commerce Research*, 1(1/2), 169–181.
67. Liu, D., Chen, J., & Whinston, A. (2008). Competing keyword auctions. In *Proc. of 4th workshop on ad auctions*, Chicago, IL, USA, July 2008.
68. Lucier, B., & Leme, R. P. (2011). GSP auctions with correlated types. In *ACM conference on electronic commerce (EC)* (pp. 71–80).
69. Markakis, E., & Telelis, O. (2010). Discrete bidding strategies and their social inefficiency for locally aware bidders. In *Workshop on Internet and network economics (WINE)* (pp. 523–530).
70. Marmer, V., & Shneyerov, A. (2009). Quantile-based nonparametric inference for first-price auctions. SSRN eLibrary.
71. McAfee, R., & McMillan, J. (1987). Auctions and bidding. *Journal of Economic Literature*, 25(2), 699–738.
72. Mehta, A., Saberi, A., Vazirani, U., & Vazirani, V. (2007). Adwords and generalized online matching. *Journal of the ACM*, 54(5).
73. Muthukrishnan, S., Pal, M., & Svitkina, Z. (2007). Stochastic models for budget optimization in search-based advertising. In *Workshop on Internet and network economics (WINE)* (pp. 131–142).
74. Myerson, R. (1981). Optimal auction design. *Mathematics of Operations Research*, 6(1), 58–73.
75. Naldi, M., D'Acquisto, G., & Italiano, G. F. (2010). The value of location in keyword auctions. *Electronic Commerce Research and Applications*, 9(2), 160–170.
76. Naldi, M., Pavignani, A., Grillo, A., Lentini, A., & Italiano, G. F. (2011). The competitor busting strategy in keyword auctions: who's worst hit? In *Proceedings of the ninth annual conference on communication networks and services research (CNSR 2011)*, Ottawa, Ontario, Canada, 2–5 May 2011 (pp. 153–160).
77. Ostrovsky, M., & Schwarz, M. (2010). Reserve prices in Internet advertising auctions: a field experiment. In *Proc. of 6th workshop on ad auctions*, Cambridge, MA, USA, June 2010.
78. Leme, R. P., & Tardos, E. (2010). Pure and Bayes-Nash price of anarchy for generalized second price auction. In *Foundations of computer science (FOCS)* (pp. 735–744).
79. Pardoe, D., Stone, P., Saar-Tsechansky, M., & Tomak, K. (2006). Adaptive mechanism design: a meta-learning approach. In *Eighth international conference on electronic commerce (ICEC)*, August 2006 (pp. 92–102).

80. Ravi, R., Hafalir, I., & Sayedi, A. (2009). Sort-cut: a pareto-optimal and semi-truthful mechanism for multi-unit auctions with budget-constrained bidders. In *Proc. of the fifth workshop on ad auctions*, July 2009.
81. Richardson, M., Dominowska, E., & Ragno, R. (2007). Predicting clicks: estimating the click-through rate for new ads. In *International conference on world wide web (WWW)* (pp. 521–530).
82. Riley, J., & Samuelson, R. (1981). Optimal auctions. *American Economic Review*, 71(3), 381–392.
83. Shapley, L., & Shubik, M. (1972). The assignment game. I. The core. *International Journal of Game Theory*, 1, 111–130.
84. Silverman, B. (1986). *Density estimation for statistics and data analysis*. London/Boca Raton: Chapman & Hall/CRC Press.
85. Spiekermann, S., & Paraschiv, C. (2002). Motivating human-agent interaction: transferring insights from behavioral marketing to interface design. *Electronic Commerce Research*, 2, 255–285.
86. Thompson, D. R. M., & Leyton-Brown, K. (2009). Computational analysis of perfect-information position auctions. In *ACM conference on electronic commerce (EC)* (pp. 51–60).
87. Varian, H. (2005). Position auctions. *International Journal of Industrial Organization*, 25, 1163–1178.
88. Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16(1), 8–37.
89. Vorobeychik, Y., & Reeves, D. M. (2007). Equilibrium analysis of dynamic bidding in sponsored search auctions. In *Workshop on Internet and network economics (WINE)* (pp. 155–166).
90. Yen, B. P.-C., & Ng, E. O. (2002). Migrating procurement onto the Internet. *Electronic Commerce Research*, 2, 113–134.
91. Young, H. P. (2004). *Strategic learning and its limits*. London: Oxford University Press.
92. Zhou, Y., & Lukose, R. M. (2007). Vindictive bidding in keyword auctions. In *International conference on electronic commerce (ICEC)* (pp. 141–146).



Patrick Maillé graduated from Ecole Polytechnique and Telecom ParisTech, France, in 2000 and 2002, respectively. He has been with the Networks, Security, Multimedia department of Telecom Bretagne since 2002, where he obtained his Ph.D. in applied mathematics in 2005, and is now an associate professor. He spent 6 months at Columbia University in 2006 to work on nonatomic routing games. His research interests are on game theory and economic concepts applied to telecommunication ecosystems: resource pricing, routing, consequences of user selfishness on network performance.



Evangelos Markakis received his diploma from NTUA in 2000 and his Ph.D. from Georgia Tech in 2005. Since 2009, he has been a Lecturer in the department of Informatics of the Athens University of Economics and Business. He is a member of the theoretical computer science group of his department and his research interests include design and analysis of algorithms, approximation algorithms, combinatorial optimization, mechanism design and algorithmic game theory. He has served several times on program committees of international conferences on algorithmic game theory and electronic commerce (e.g. EC, WINE, SAGT, WWW) and as a reviewer in many journals related to theoretical computer science. He has also participated in basic research projects funded by EU programs, such as ESF/COST. In particular he was a member of the management committee of a COST action on “Algorithmic Decision theory” and was also a member of the coordinating team in a project funded by the network of excellence EURO-NF.



Maurizio Naldi graduated cum laude in 1988 in Electronic Engineering and then received his Ph.D. in Telecommunications Engineering from the University of Rome at Tor Vergata. After graduation he pursued an industrial career, first at Selenia, and then in the Network Planning Departments of Italcable, Telecom Italia, and Wind Telecomunicazioni, where he was appointed Head of the Traffic Forecasting & Network Cost Evaluation Group. Since 2000 he holds a position at the University of Rome at Tor Vergata, where he is now Aggregate Professor. His current research interests lie in the fields of network economics and computational advertising.



George D. Stamoulis (<http://nes.aueb.gr/users/gd.html>) received the Diploma in Electrical Engineering (1987, with highest honors) from the National Technical University of Athens, Greece, and the M.S. (1988) and Ph.D. (1991) degrees in Electrical Engineering from the Massachusetts Institute of Technology, Cambridge, Massachusetts, USA. From 1995 to 2000 he was an Assistant Professor at the Computer Science Department of the University of Crete, Greece, and a member of the Telecommunications and Networks Division of the Institute of Computer Science, Foundation for Research and Technology Hellas (ICS-FORTH). Since 2001 he is a faculty member, currently Professor, in the Department of Informatics of Athens University of Economics and Business (AUEB). His research interests are in economic and business models for networks, reputation mechanisms for electronic environments, auction mechanisms for bandwidth and digital goods, telecommunications regulation, and demand response in electricity consumption.

He has published over 55 articles in scientific journals such as IEEE/ACM Transactions on Networks, Telecommunications Systems, Computer Communications, IEEE Transactions on Communications, Journal of the ACM, and IEEE Transactions on Information Theory, as well as in conferences such as INFOCOM, ITC, ACM SIGMETRICS, GLOBECOM, ICQT, IFIP networking and GP2PC. He has also collaborated with the Greek Regulatory Authorities for Communications and Power on auction design as well as with Greek telecom companies on service pricing.



Bruno Tuffin received his Ph.D. degree in applied mathematics from the University of Rennes (France) in 1997. Since then, he has been at INRIA in Rennes. He spent eight months as a postdoc at Duke University in 1999. His research interests include developing Monte Carlo and quasi-Monte Carlo simulation techniques for the performance evaluation of telecommunication systems, and developing new Internet-pricing schemes and telecommunication-related economic models. He has published more than one hundred papers on those issues. He has also led or participated into several French and European projects, and co-organized several conferences. He is currently Associate Editor for INFORMS Journal on Computing, ACM Transactions on Modeling and Computer Simulation and Mathematical Methods of Operations Research. He has written or co-written two books devoted to simulation: "Rare event simulation using Monte Carlo methods" published by John Wiley & Sons in 2009, and "La simulation de Monte Carlo" (in French), published by Hermes Editions in 2010.