POLITECNICO
MILANO 1863

SCUOLA DI INGEGNERIA INDUSTRIALE
E DELL'INFORMAZIONE

EXECUTIVE SUMMARY OF THE THESIS

Convergence Analysis for Sponsored Search Auctions with Price Displaying

LAUREA MAGISTRALE IN COMPUTER SCIENCE AND ENGINEERING - INGEGNERIA INFORMATICA

Author: FEDERICO INNOCENTE

Advisor: PROF. NICOLA GATTI

Co-advisor: MATTEO CASTIGLIONI, DIODATO FERRAIOLI, GIULIA ROMANO

Academic year: 2021-2022

Abstract

Since the early 2000's online advertising has become one of the main revenue sources for companies, which have been using it to advertise their content and sponsor their goods. One of the main application fields for online advertising is search advertising, where the search engines submit to their users some sponsored contents next to the organic results. The way in which the displayed advertisements are selected is through a suitable auction mechanism. Mechanism design is a well studied sub field of game theory, and the problem of auctioneer's mechanism selection and advertisers' bidding strategy have a solid base in literature. However, very few studies have analyzed auctions in the context of e-commerce, namely *sponsored auction with price display*, where advertisements are associated to a price for the sponsored good and customers' clicks are influenced by the comparison of those selling prices. In this work, we initially formalize the extension of a well known bidding strategy for the generalized second price auction, the balance bidding, to the scenario with price displaying. Our main result consists in the proposal of a new auction mechanism based on GSP that is guaranteed to converge to its equilibrium if advertisers bid according to the extended bidding strategy when prices are fixed. We also study the efficiency of the equilibrium with respect to a parameter of the mechanism called *cut price*, proposing a randomized algorithm that guarantees a lower bound of the equilibria social welfare in expectation. Lastly, we provide some experimental results, in order to empirically analyze the average convergence time and the average efficiency for different auction settings.

1. Introduction

1.1. Goal

The goal of this thesis is to study a particular case of sponsored search auction we called *sponsored search auction with price display*, where advertisers aim to sell some good through their ads. This new framework introduces an externality among the ads, since when similar advertisements are shown along with their prices the users are induced to compare them to each other, and they are more likely to click on ads associated to a lower price. This preference will be modeled as a function, called *quality*, of the displayed selling prices, which determines the click probability of the ads. The problem of designing a suitable auction mechanism represents a challenge, since the allocation of an ad can, in principle, affects the click probability, and thus the utility, of all the agents taking part to the game. We will focus our attention on studying the convergence property of the auction, by proposing a mechanism that is guaranteed to converge to a Nash equilibrium when advertisers bid according to a myopic best response bidding strategy and the selling prices of each ad are constant over time. We will also study the efficiency in terms of social welfare of the reached equilibria with respect to the optimal allocation, and run some experiments to validate and empirically evaluate our results.

1.2. Related works

Edelman et al. [4] studied the equilibria of GSP auction, defining a notable class of equilibria called locally envy-free equilibria. They also showed that the untruthfulness of GSP can result in instability and bidding wars among the advertisers. On the basis of their work, Cary et al. [2] and Bu et al. [1] independently studied a myopic best response bidding strategy for GSP auction, which they called, respectively, Balance Bidding and Forward Looking. They proved that when advertisers bid according to that strategy, the auction is guaranteed to converge to a Nash equilibrium where the agents' utility and payments are the analogous to the one of VCG auction.
Castiglioni et al. [3] studied the sponsored search auction for price displaying, proposing both a VCG and a GSP indirect revelation mechanism for this class of auctions. They also studied the equilibria of the auction, analyzing the efficiency in terms of social welfare and auction revenue.

1.3. Original contribution

In this work, we study whether the extension to the price displaying setting of the bidding strategy studied by Cary et al. [2] and Bu et al. [1] can guarantee some convergence results. In particular, we first show that the GSP indirect revelation mechanism proposed by Castiglioni et al. [3] is not guaranteed to converge under this bidding strategy. We propose a new auction mechanism based on GSP called $\mathcal{M}_{ord}^{GSP}(p^*)$, dependent to a parameter p^* called *cut price*, that is guaranteed to converge to an equilibrium for some cut price selection policies when selling prices do not change over time. We study the efficiency in terms of social welfare of this equilibria, and propose a randomized algorithm to select the cut price in order to ensure an efficiency lower bound. We conclude the work by analyzing some experimental results to empirically evaluate the relationship between the convergence time and the efficiency with respect to some auction settings.

2. Background

2.1. Generalized second price

In a GSP auction, a set of N agents compete to be allocated in one of M available slots. Each slot is associated with a prominence $\lambda_j \in [0, 1]$ such that $\lambda_i > \lambda_j \, \forall \, i < j$, that represents the probability to observe the slot j. Each advertiser is associated with a value $v_i \in \mathbb{R}$ that he will collect every time his ad is clicked on. When an agent takes part to the auction, he submits a bid $b_i \in \mathbb{R}$ that represents that maximum amount he declare being willing to pay for a click on his ad, which can be updated any time. The mechanism assigns the ads to the slot in decreasing order of bid, and charges each allocated advertiser with an auction cost π_i equal to the next higher bid. GSP is a non truthful mechanism, which means that bidding $b_i = v_i$ is not a dominant bidding strategy, and thus advertisers need to elaborate complex bidding strategy to maximize their utility.

2.2. GSP bidding strategy

Cary et al. [2] and Bu et al. [1] both studied the same myopic best response bidding strategy for GSP, which they called, respectively, Balance Bidding (BB) and Forward Looking (FL). The idea on which this bidding strategy relies is that when an agent myopically targets the slot j^* that maximizes his utility, he can select his bid in the range $(\pi_i(j^*), \pi_i(j^* - 1))$, where $\pi_i(j)$ is the auction payment of player i for being allocated to the slot j. Thus, a tie breaking rule needs to be applied.

Definition 2.1 (Balance Bidding)**.** *The Balance Bidding strategy is the strategy for a player i that, given $\boldsymbol{b}_{-i}$, targets the slot j^* which max-*

imizes his utility

$$j^* = \arg\max_j \{\lambda_j(v_i - \pi_i(j))\}$$

and chooses his bid b_i' *as to satisfy the following equation:*

$$\lambda_{j^*}(v_i - \pi_i(j^*)) = \lambda_{j^*-1}(v_i - b_i').$$

If j^* *is the first slot, he (arbitrarily) bids* $b_i' = \frac{v_i+\pi_i(1)}{2}$.

The rationale behind BB strategy is that an agent wants to bid high enough to be allocated to his favourite slot, but not as high as to possibly regret being undertaken by another advertiser. If an agent cannot be assigned to any slot for any bid $b_i < v_i$, his best response is to bid $b_i = v_i$, since GSP is individually rational when agents do not overbid, and overbidding is a dominated strategy. The following result holds:

Proposition 2.1. *If all players follow the BB strategy in an auction with all distinct* λ*'s and asynchronous bidding, then the system converges to its unique fixed point. At this fixed point the revenue of the auctioneer (and the payment of each player) is equal to the one of the VCG equilibrium.*

3. Sponsored auction with price displaying

3.1. Model

In a sponsored search auction with price displaying (SSAPD), a set of N advertisers competes to be assigned to one of the M available slots, each one associated to a prominence $\lambda_j \in [0,1]$ representing the visibility of the slot, and such that $\lambda_i > \lambda_j \, \forall\, i < j$. Each ad competes for a single good and is associated with a selling price $p_i \in \mathbb{R}_{\geq 0}$, which is the price displayed on the ad. We will assume the selling prices as fixed in time for each ad: even if, in principle, an advertiser could change his good price at any time, in many practical scenarios this change is at most abrupt. For each ad a_i, we call the production cost of the sold good as $c_i \in \mathbb{R}$ and the buy probability of the item if the ad is clicked as $\alpha_i \in [0,1]$. Thus, the gain of the advertiser i is $\mu_i = \alpha_i(p_i - b_i)$. The pair (α_i, c_i) represents the type θ_i of the advertiser. In an indirect revelation mechanism, each advertiser submits to the auctioneer a bid b_i, which represents the declared gain of the advertiser for the ad a_i, and we say that the agent do not overbid if $b_i \leq \mu_i$. An allocation A is an assignment of the ads to the slot that is represented by a function $f : N \to M \cup \perp$ such that there is at most one ad per slot, and for all the ads not assigned to any slot in M it is $f(i) = \perp$. Each advertiser that joins to the mechanism is charged with a per-click payment $\pi_i \in \mathbb{R}$.

Since the ads' click probability is affected by the comparison of the users among the displayed ads' selling prices, we call $q_i \in [0,1]$ the probability (called *quality*) that the ad a_i is clicked, conditioned on its observation. We consider the quality as independent from the advertisers, and so a function of the selling price profile only. It will also be assumed that $q_i : \mathbb{R}^+ \times \mathbb{R}^+ \to [0,1]$, where $q_i(p_i, p_{min})$ denotes the player i's quality when his price is p_i and $p_{min} = \min_{k: f(k) \neq \perp} p_k$ is the minimum price among the displayed ads. Moreover, given p_{min}, $q_i(p_i, p_{min})$ is non strictly monotonically decreasing in p_i and, given p_i, is non strictly monotonically increasing in p_{min}. This assumption is due to the fact that users will compare the advertisements based on their prices, and their interest in a specific ad will decrease with the price difference with respect to the others.

We will refer as declared (expected) value to the value $\hat{v}_i(f, \boldsymbol{p}, b_i) = \lambda_{f(i)} q_i(p_i, p_{min}) b_i$ computed with the declared gain b_i in the allocation f, while the (true expected) value is $v(f, \boldsymbol{p}, b_i) = \lambda_{f(i)} q_i(p_i, p_{min}) \alpha_i(p_i - c_i)$. The expectation is referred to both the click probability $\lambda_{f(i)} q_i(p_i, p_{min})$ and the buy rate α_i. The social welfare of the allocation f with respect to the declared gains $\boldsymbol{b}$ is $s\hat{w}(f, \boldsymbol{p}, \boldsymbol{b}) = \sum_{i \in N} \hat{v}_i(f, \boldsymbol{p}, b_i)$, while the true social welfare is $sw(f, \boldsymbol{p}, \boldsymbol{b}) = \sum_{i \in N} v_i(f, \boldsymbol{p}, b_i)$.

3.2. SSAPD bidding strategy

We investigate whether the BB strategy mantains its convergence property in a SSAPD scenario. We start by formalizing the price dependent version of the bidding strategy.

Definition 3.1 (Balance Bidding for Price Displaying). *The Balance Bidding for Price Displaying (BBPD) strategy is the strategy for player* i *that, given* b_{-i} *and* $p = (p_1, ..., p_n)$*, targets the slot* $j_i^* = \arg\max_s \{\lambda_s q_i(p_i, p_{min})(\mu_i - \pi_i(s))\}$ *which maximizes his utility and chooses* *his bid b'_i for the next round so as to satisfy the following equation:*

$$\lambda_{j^*} q_i(p_i, p_{min})(\mu_i - \pi_i(j^*)) \\ = \lambda_{j^*-1} q_i(p_i, p_{min})(\mu_i - b'_i)$$

If j_i^* is the first slot, we (arbitrarily) choose $b' = \mu_i$. If an agent cannot be allocated to any slot without overbidding, he bids $b' = \mu_i$.

Starting from BBPD, we define an extended version of the bidding strategy.

Definition 3.2 (Extended Balance Bidding for Price Displaying)**.** *The Extended Balance Bidding for Price Displaying (EBBPD) is the strategy for player i that, given $\boldsymbol{b}_{-i}$ and $\boldsymbol{p} = (p_1, ..., p_n)$, targets the slot $j_i^* = \arg\max_s \{\lambda_s q_i(p_i, p_{min})(\mu_i - \pi_i(s))\}$ which maximizes his utility and chooses his bid b'_i for the next round as:*

$$b'_i = \max\{BBPD(\boldsymbol{b}_{-i}, \boldsymbol{p}), \bar{b}\}$$

where $\bar{b}$ is such that $\bar{b} = \min b_i : f^*(i) = j^*$.

3.3. Auction mechanisms for SSAPD

Castiglioni et al. [3] studied the equilibria efficiency of the GSP-based mechanism $\mathcal{M}_I^{GSP}$, where the allocation function f assigns the advertisers to the slots in order to maximize the declared social welfare, and the auction prices are charged in a second price manner: being k the number of allocated ads, $\pi_i = \lambda_{f(i)} q(p_j, p_{min}) b_j$ if $f(i) < k$, where j is such that $f(j) = f(i) + 1$, while $\pi_k = 0$ if $p_h < p_{min} \forall h$ such that $f(h) = \perp$, $\pi_k = \lambda_{f^*(k)} \max_{j: p_j \geq p_{min} \,\&\, f^*(j) = \perp}\{q_j(p_j, p_{min}) b_j\}$ otherwise. For all the agents such that $f(i) = \perp$ it is $\pi_i = 0$.

We show that the following result holds:

Proposition 3.1. *When advertisers bid according to EBBPD, the mechanism $\mathcal{M}_I^{GSP}$ is not guaranteed to converge to an equilibrium.*

We propose a new auction mechanism based on GSP called $\mathcal{M}_{ord}^{GSP}(p^*)$, where p^* is a parameter called *cut price*. The mechanism works as follows. Given the bids and the selling prices profiles, the ads in $N' = N \backslash \{a_i : p_i < p^*\}$ compete in the game, while the ones whose selling prices are lower than p^* are discarded. The remaining ads are sorted by $q(p_i, p^*) b_i$. We refer to such sorting function as σ, and to the first M ads according to σ as N_σ. As a last step, the ads in N_σ are assigned to the slots by the allocation function $f^* = \max_{f: f(i) \neq \perp \, \forall i \in N_\sigma} \sum_i \lambda_{f(i)} q_i(p_i, p_{min}) b_i$. We say that, for all agent $i \in N \backslash N_\sigma$ it is $f^*(i) = \perp$. Informally, the selected allocation is the one that maximizes the declared social welfare among the ones that assign all the ads in N_σ.

If k ads are allocated, the mechanism charges each ad such that $f^*(i) < k$ with a per click auction price of $\pi_i = \frac{q(p_j, p_{min}) b_j}{q(p_i, p_{min})}$, where j is such that $f^*(j) = f^*(i) + 1$. For the ad such that $f(i) = k$, instead, the payment is $\pi_i = \max_{j \in N': f^*(j) = \perp}\{q_j(p_j, p^*) b_j\} / q_i(p_i, p_{min})$.

A graphical representation of the mechanism $\mathcal{M}_{ord}^{GSP}(p^*)$ is shown in Figure 4.

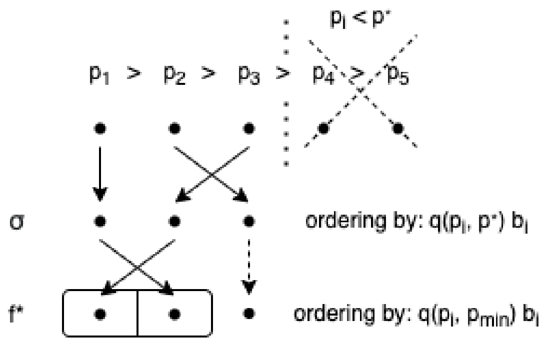

Figure 1: Graphical example of $\mathcal{M}_{ord}^{GSP}(p^*)$ mechanism workflow in a 5 players and 2 slots game. In the first step, the ads with a selling price lower than p^* are discarded. Then, remaining ads are sorted by σ. Lastly, the top 2 ads according to σ are allocated by f^*.

The following result hold for the mechanism $\mathcal{M}_{ord}^{GSP}(p^*)$.

Theorem 3.1. *The mechanism $\mathcal{M}_{ord}^{GSP}(p^*)$ is weakly budget balanced and individually rational in ex-post if agents do not overbid.*

The non overbidding assumption is reasonable, as overbidding is a dominated strategy for the advertisers.

3.3.1. $\mathcal{M}_{ord}^{GSP}(p^*)$ cut price

The cut price p^* is a parameter of the mechanism that affects the allocation for a given bid and price profile $\boldsymbol{b}$ and $\boldsymbol{p}$. The rationality behind this parameter relies on the fact that when the advertisers need to be allocated, they are selected accordingly to the ordering σ, and the agents with higher selling prices can be penalized with respect to the others, eventually affecting the social welfare. In a repeated auction, the cut price could, in principle, change for different instances of the mechanism. We will call a cut price selection strategy as *cut price dynamics*. Being the mechanism individually rational, a reasonable assumption is that every rational agent prefers to be allocated rather than not. Thus, if at any time an agent does not select a bid that grants him a slot, it means that he can not be allocated without overbidding. Based on this assumption, we define a particular cut price dynamic.

Definition 3.3. *Consider a SSAPD based on $\mathcal{M}_{ord}^{GSP}(p^*)$. At any time t, let $B \subseteq N'$ be the set of non discarded agents whose last submitted bid did not grant them a slot in the next iteration of the auction and $K = \{a_i : p_i = p^* \, \forall \, i \in N\}$ the set of agents such that their selling price is equal to the cut price. We call $\mathcal{D}$ the cut price dynamic that selects for the auction at time $t+1$ a cut price $p^{*\prime} = \min\{p_i : p_i > p^* \, \forall \, i \in N\}$ if $K \subseteq B$, or $p^{*\prime} = p^*$ otherwise.*

3.4. Properties of $\mathcal{M}_{ord}^{GSP}$

We summarize the main properties of $\mathcal{M}_{ord}^{GSP}$. Our main result concerns the convergence property of the mechanism.

Theorem 3.2. *A repeated SSAPD based on the mechanism $\mathcal{M}_{ord}^{GSP}(p^*)$ that implements $\mathcal{D}$ converges to its Nash equilibrium for all the initial cut price p^* when selling prices are fixed and the agents select their bids accordingly to EBBPD in a random and asynchronous way.*

We studied the equilibrium's efficiency in terms of social welfare of the proposed mechanism. The results are stated in the following.

Theorem 3.3. *The PoS of the social welfare for the mechanism $\mathcal{M}_{ord}^{GSP}(p^*)$ can be arbitrarily large for some cut prices p^* when the agents select their bids accordingly with the EBBPD strategy.*

Theorem 3.4. *In a SSAPD implementing $\mathcal{M}_{ord}^{GSP}(p^*)$ where agents select their bids accordingly with EBBPD, there is at least one cut price p^* such that the mechanism converge to an equilibrium that maximises the true social welfare.*

Thus, the social welfare of the mechanism equilibrium can span from arbitrarily bad to the optimal one for some initial cut price p^* when the advertisers bid accordingly to EBBPD. We propose a randomized algorithm over the initial cut price that ensures an efficiency lower bound in expectation.

Definition 3.4. *We define as $\boldsymbol{Part_{\bar{q}}}(N) = \{N_1, N_2, ..., N_k\}$ a partition over N such that*

$$N_i = \{a_j : q(p_j, \min\{p \in N_i\}) \geq \bar{p}\}$$

$$\bar{q} \leq q(p, p).$$

We will refer as $\boldsymbol{P}(Part_{\bar{q}}(N))$ to the set of selling prices such that

$$P(Part_{\bar{q}}(N)) = \{p : p = \min_{j \in N_i} p_j \, \forall \, N_i \in Part_{\bar{q}}(N)\}.$$

Given the defined partition over the agents, we state the following efficiency result.

Theorem 3.5. *In a SSAPD implementing $\mathcal{M}_{ord}^{GSP}(p^*)$ where the cut price dynamic is $\mathcal{D}$ and agents update their bids accordingly to EBBPD, if the initial cut price is selected from $P(Part_{\bar{q}}(N))$, then the social welfare at the equilibrium is in expectation at least $\frac{\bar{q}}{|Part_{\bar{q}}(N)|}$ of the optimal social welfare.*

This result is extended by the following bounding of the partition's cardinality.

Theorem 3.6. *For a general quality function $q(p, p_{min})$ and a given quality threshold $\bar{q}$, it is possible to compute a partition $\text{Part}_{\bar{q}}(N)$ such that*

$$|\text{Part}_{\bar{q}}(N)| \leq \left\lceil \frac{(\max_{i \in N} p_i - \min_{i \in N} p_i)}{\Delta p_{\bar{q}}} \right\rceil$$

where

$$\Delta p_{\bar{q}} = \min_{i \in N} p_i - \max\{\bar{p}_i : q(\bar{p}_i, p_i) \geq \bar{q}, \bar{p}_i \geq p_i\}.$$

Corollary 3.1. *For a k-Lipschitz quality function $q(p, p_{min})$ and a given quality threshold $\bar{q}$, it is possible to compute a partition $\text{Part}_{\bar{q}}(N)$ such that*

$$|\text{Part}_{\bar{q}}(N)| \leq \left\lceil \frac{(\max_{i \in N} p_i - \min_{i \in N} p_i)k}{q(p, p) - \bar{q}} \right\rceil$$

4. Experiments

We highlight some significant experimental results in order to analyze the convergence time and the equilibrium efficiency of the mechanism

$\mathcal{M}_{ord}^{GSP}$. The experiments are related to a sigmoid like quality function in the difference of the selling price and the minimum displayed price.

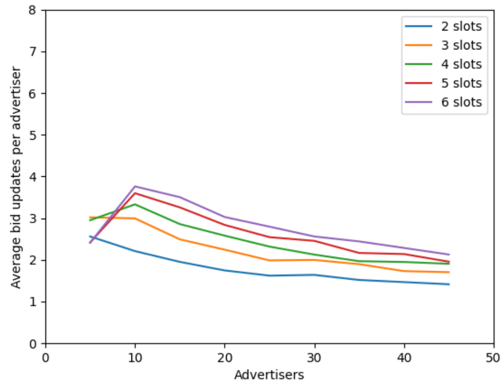

Figure 2: Average convergence time in terms of bid updates per agent for a different number of advertisers and slots.

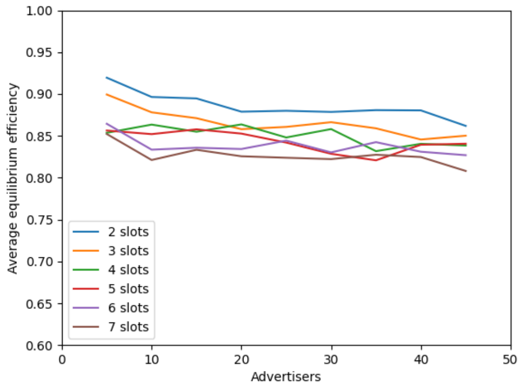

Figure 3: Average equilibrium efficiency with respect to the optimal allocation when selling prices are drawn from a normal distribution

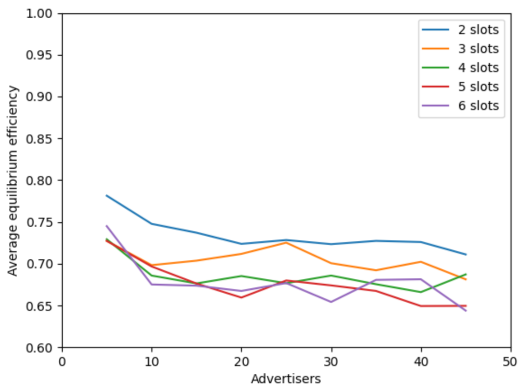

Figure 4: Average equilibrium efficiency with respect to the optimal allocation when selling prices are drawn from an uniform distribution

5. Conclusions

In this work we approached a particular case of sponsored search auctions called "sponsored search auctions with price display", which differs from the classic one by showing the users a price for each advertised ad, introducing a new externality among the customers. We formalized a bidding strategy for SSAPD as an extension of the GSP bidding strategy studied by Cary et al. [2] and Bu et al. [1]. We proposed a new GSP based allocation mechanism, namely $\mathcal{M}_{ord}^{GSP}(p^*)$, showing that it is guaranteed to converge to an equilibrium when agents bids according to the described bidding strategy and a parameter of the mechanism, called cut price, is properly updated. Then, we studied the social welfare efficiency of the convergence equilibria, proposing a randomized algorithm over the cut price in order lower bound guarantees. Lastly, we run some experiments to empirically evaluate the convergence time and the efficiency with respect to some auction settings.

References

[1] Tian-Ming Bu, Li Liang, and Qi Qi. On robustness of forward-looking in sponsored search auction. *Algorithmica*, 58(4):970–989, December 2010.

[2] Matthew Cary, Aparna Das, Benjamin Edelman, Ioannis Giotis, Kurtis Heimerl, Anna R. Karlin, Claire Mathieu, and Michael Schwarz. On Best-Response Bidding in GSP Auctions. NBER Working Papers 13788, National Bureau of Economic Research, Inc, February 2008.

[3] Matteo Castiglioni, Diodato Ferraioli, Nicola Gatti, Alberto Marchesi, and Giulia Romano. Efficiency of ad auctions with price displaying, 2022.

[4] Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 97(1):242–259, March 2007.