Εθνικο Μετσοβιο Πολυτεχνειο
Τμημα Ηλεκτρολογων Μηχανικων Και Μηχανικων Υπολογιστων

Τομέας Τεχνολογίας Πληροφορικής & Υπολογιστών
Εργαστήριο Λογικής & Επιστήμης Υπολογιστών

Δημοπρασίες σε Μηχανές Αναζήτησης με Εξωγενείς Επιδράσεις

ΔΙΠΛΩΜΑΤΙΚΗ ΕΡΓΑΣΙΑ

του

Μάρκου Π. Επιτρόπου

Επιβλέπων: Δημήτρης Φωτάκης
Λέκτορας

Αθήνα, Ιούλιος 2012

ΕΘΝΙΚΟ ΜΕΤΣΟΒΙΟ ΠΟΛΥΤΕΧΝΕΙΟ
ΤΜΗΜΑ ΗΛΕΚΤΡΟΛΟΓΩΝ ΜΗΧΑΝΙΚΩΝ
ΚΑΙ ΜΗΧΑΝΙΚΩΝ ΥΠΟΛΟΓΙΣΤΩΝ
Τομέας Τεχνολογίας Πληροφορικής & Υπολογιστών
Εργαστήριο Λογικής & Επιστήμης Υπολογιστών

Δημοπρασίες σε Μηχανές Αναζήτησης με Εξωγενείς Επιδράσεις

ΔΙΠΛΩΜΑΤΙΚΗ ΕΡΓΑΣΙΑ

του

Μάρκου Π. Επιτρόπου

Επιβλέπων: Δημήτρης Φωτάκης
Λέκτορας

Εγκρίθηκε από την τριμελή εξεταστική επιτροπή την 23/07/2012.

....	....	....
Δημήτρης Φωτάκης	Ευστάθιος Ζάχος	Άρης Παγουρτζής
Λέκτορας	Καθηγητής	Επίκ. Καθηγητής

Αθήνα, Ιούλιος 2012

....
Μάρκος Π. Επιτρόπου
Διπλωματούχος Ηλεκτρολόγος Μηχανικός και Μηχανικός Υπολογιστών
Ε.Μ.Π.

ΕΘΝΙΚΟ ΜΕΤΣΟΒΙΟ ΠΟΛΥΤΕΧΝΕΙΟ
ΤΜΗΜΑ ΗΛΕΚΤΡΟΛΟΓΩΝ ΜΗΧΑΝΙΚΩΝ ΚΑΙ ΜΗΧΑΝΙΚΩΝ ΥΠΟΛΟΓΙΣΤΩΝ
Τομέας Τεχνολογίας Πληροφορικής & Υπολογιστών
Εργαστήριο Λογικής & Επιστήμης Υπολογιστών

Ευχαριστίες

Ολοκληρώνοντας αυτή την διπλωματική, θα ήθελα να ευχαριστήσω την τριμελή μου επιτροπή που αποτελείται από τους κ. Ζάχο, καθηγητή Ε.Μ.Π., κ. Παγουρτζή, επίκουρο καθηγητή Ε.Μ.Π., και κ. Φωτάκη, λέκτορα Ε.Μ.Π. για την στήριξη τους τόσο σε ακαδημαϊκό όσο και σε προσωπικό επίπεδο.

Ιδιαιτέρως θα ήθελα να ευχαριστήσω τον κ. Φωτάκη, επιβλέποντα αυτής της διπλωματικής, για την ενθάρυνση, την καθοδήγηση του και την συνεργασία μας κατά την διάρκεια αυτής της προσπάθειας. Θα ήθελα επίσης να τον ευχαριστήσω για την ευκαιρία που μου έδωσε να έρθω σε επαφή με ένα τόσο ενδιαφέρον και ευρύ πεδίο όπως είναι οι Αλγόριθμοι και οι εφαρμογές τους.

Τέλος, θα ήθελα να ευχαριστήσω όλα τα μέλη του Εργαστηρίου Λογικής και Επιστήμης Υπολογιστών για την βοήθεια που απλόχερα μου παρείχαν οποτεδήποτε την χρειάστηκα, για το ευχάριστο κλίμα που έχουν δημιουργήσει στο εργαστήριο και την συνεχή ανταλλαγή απόψεων και προβληματισμών.

Περίληψη

Σε αυτήν την διπλωματική εργασία εξετάζουμε μια κλάση δημοπρασιών, τις δημοπρασίες σε μηχανές αναζήτησης, γνωστές και ως δημοπρασίες λέξεων-κλειδί. Είναι ένας νέος τύπος ηλεκτρονικής διαφήμισης, όπου οι διαφημιστές ανταγωνίζονται για μια θέση δίπλα στα αποτελέσματα αναζήτησης. Οι κολοσσοί του διαδικτύου, Google, Yahoo!, Microsoft, πουλάνε χώρο για διαφήμιση με αυτόν τον τρόπο, δημοπρατώντας θέσεις στις μηχανές αναζήτησης. Τα κέρδη από αυτές τις δημοπρασίες καταλαμβάνουν μεγάλο ποσοστό των συνολικών κερδών τους. Παρουσιάζω το βασικό μοντέλο για αυτές τις δημοπρασίες και τους κύριους σχεδιαστικούς στόχους. Συγκεκριμένοι κανόνες διανομής των θέσεων και πληρωμών παραθέτονται και εξετάζονται ώς προς την ευστάθεια τους, την αποδοτικότητα τους και τα κέρδη. Νέα πειραματική ανάλυση δείχνει ότι το βασικό μοντέλο είναι πολύ απλουστευμένο, μην μπορώντας να αποτυπώσει εξωγενείς επιδράσεις μεταξύ των διαφημιστών. Παρουσιάζω νέα μοντέλα που μπορούν να εκφράσουν τέτοιες επιδράσεις, επεκτείνοντας το βασικό μοντέλο και αναλύοντας την ευστάθεια και αποδοτικότητα μηχανισμών.

Λέξεις Κλειδιά

Αλγοριθμική Θεωρία Παιγνίων, Σχεδιασμός Μηχανισμών, Δημοπρασίες σε Μηχανές Αναζήτησης, GSP, Εξωγενείς Επιδράσεις

Abstract

In this thesis, I consider a special class of auctions, Sponsored Search Auctions, also known as keyword auctions. It is a new form of online advertising, where advertisers compete for a place next to the web search results. Internet giants Google, Yahoo!, Microsoft sell online advertising by this form, who auction positions for sponsored search links. Search revenues accounts for a large percentage of their massive revenues. I introduce the basic setting in sponsored search auctions and the main design goals. Certain types of allocation and payments are presented, judging their stability, allocation efficiency and resulting revenue. Having new experimental analysis, it is proven that the basic setting is inadequate and externalities between advertisers exist. I present new settings to express such externalities, by expanding the basic setting, and analyze them in terms of stability and allocation efficiency.

Keywords

Algorithmic Game Theory, Mechanism Design, Sponsored Search Auctions, GSP, Externalities

Contents

Chapter 1

Mechanism Design

1.1 Game Theory

Game theory is the study of mathematical models of conflict and cooperation between intelligent rational decision-makers. It is the study of games where every agent acts rationally and tries to maximize her personal payoff. Many situations in real life can be modeled as a game and studied in depth. The game-theoretic model is used in many different sciences which involves conflicting parties, such as economics, computer science, political science and biology. Game theory studies the possible outcomes of such systems and analyzes their efficiency.

The most standard game to introduce game theory is the prisoner's dilemma. Two prisoners (agents) are interrogated and they have two choices, to confess or remain silent. These choices are the agent's strategies. According to the prisoners' choices a final verdict is taken which gives them prison sentences and servers as the outcome of the game. The two prisoners like or dislike the final decision and evaluate it, resulting in private valuations for each outcome. This game can be presented by two matrices, each depicting the prisoner's cost (prison sentence) for the 4 possible strategy combinations. It's a simple game with 2 agents and 2 strategies each.

The above game will be used to present the basic concepts of game theory and its goals. Specifically, I will introduce a way to define a possible outcome of a game, which must be formally stated, and present the context to evaluate efficiency.

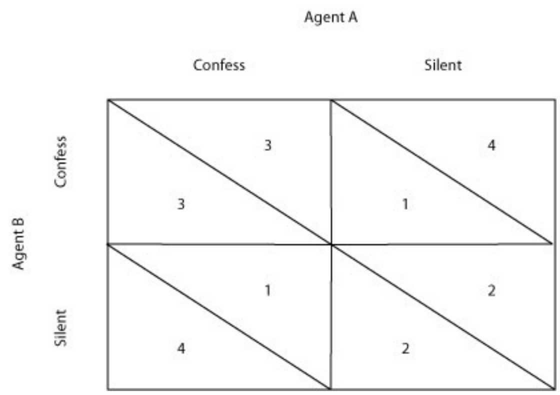

Figure 1.1: The prisoner's dilemma

1.1.1 Nash Equilibrium

Agents act according to their incentives, trying to maximize their own payoff. They choose the best strategy, assuming that they know the other agents' strategies. A moment occures that the agents have chosen their strategies and none has an incentive to change her strategy. This combination of strategies is known as a Nash equilibrium, depicting a notion of stability and an outcome of the game.

Formally, an agent's i strategy is chosen by a set of possible strategies, namely $\sigma_i \in \Sigma_i$. When mentioning a specific agent the other agents' strategies are modeled by a vector σ_{-i}, resulting in a complete vector of all strategies $\sigma = (\sigma_i, \sigma_{-i})$. Every agent i has a payoff for each outcome, namely $g_i(\sigma)$. A Nash equilibrium refers in such a vector which satisfies the following equation:

$$\forall i, \sigma'_i \in \Sigma_i : g_i(\sigma_i, \sigma_{-i}) \geq g_i(\sigma'_i, \sigma_{-i}).$$

Let's present another game, similar to the prisoner's dilemma and find a Nash equilibrium. The game needs no description but only two payoff matrices, depicting their payoff for each possible outcome. Generally these games are called bimatrix games.

One can see that there are two strategy combinations $(\sigma_A^{(1)}, \sigma_B^{(1)})$ and $(\sigma_A^{(2)}, \sigma_B^{(2)})$, where none has an incentive to change her strategy when knowing the other agent's strategy. Formally,

- $(\sigma_A^{(1)}, \sigma_B^{(1)})$:
  $g_i(\sigma_A^{(1)}, \sigma_B^{(1)}) = 5 > 1 = g_i(\sigma_A^{(2)}, \sigma_B^{(1)})$
  $g_i(\sigma_A^{(1)}, \sigma_B^{(1)}) = 5 > 1 = g_i(\sigma_A^{(1)}, \sigma_B^{(2)})$

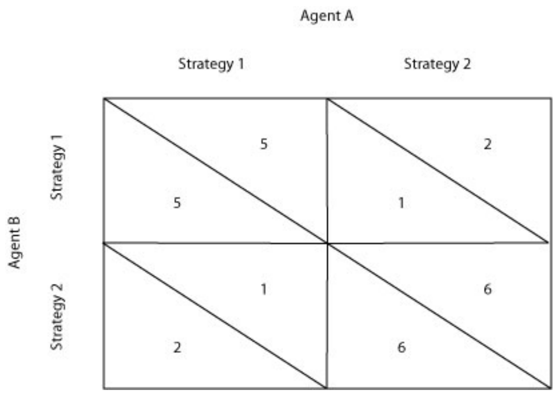

Figure 1.2: A bimatrix game

- $(\sigma_A^{(2)}, \sigma_B^{(2)})$:
  $g_i(\sigma_A^{(2)}, \sigma_B^{(2)}) = 6 > 2 = g_i(\sigma_A^{(1)}, \sigma_B^{(2)})$
  $g_i(\sigma_A^{(2)}, \sigma_B^{(2)}) = 6 > 2 = g_i(\sigma_A^{(2)}, \sigma_B^{(1)})$

1.1.2 Dominant Strategies Equilibrium

As shown in the above example an agent cannot determine her strategy without knowing the other agents' strategies. There is no "ideal" strategy. However, there are instances of games that the agents get the best payoff by choosing a specific strategy independently of the other agents' strategies. The example of the prisoner's dilemma stated above is such an instance.

Formally, a dominant strategy equilibrium satisfies the following equation:

$$\forall \sigma_{-i} \in \prod_{j \in [n]/i} \Sigma_j, \sigma_i' \in \Sigma_i : g_i(\sigma_i, \sigma_{-i}) \geq g_i(\sigma_i', \sigma_{-i})$$

In the prisoner's dilemma one can see that every prisoner, independently of the other prisoner's strategy, has an incentive to confess.

- If B confesses, A has an incentive to confess : $g_i("Confess", "Confess") = 4 < 5 = g_i("Silent", Confess")$
- If B remains silent, A has an incentive to confess : $g_i("Confess", "Silent") = 1 < 2 = g_i("Silent", "Silent")$

Simirarly, B has an incentive to confess, independently of the A's strategy.

The dominant strategies equilibrium is another concept which serves as a way to define the possible outcomes of a game. The Nash equilibrium may not be so much satisfactory to define such a state because an agent does not know the other player's strategies in many cases. However, the dominant strategies equilibrium defines a strategy for each agent which is a one-way road. The system will surely end up at that state.

The above notions of a dominant strategy equilibrium and a Nash equilibrium define the possible outcomes of a game. In the bibliography these notions are stated as solution concepts. Game theory also analyzes the efficiency of the solutions of a game. To do so an objective must be set to support the analysis. The most common objective is the social welfare maximization. The social welfare is the total payoff of all agents. I will describe this concept for the prisoner's dilemma as an introduction to the efficiency analysis.

In this specific example, I will use the total cost with the objective of minimizing it, because the matrices model costs instead of payoffs. The total cost is defined as $C(\sigma) = \sum_i g_i(\sigma)$. The above example has one solution ("Confess","Confess") with total cost $C("Confess", "Confess") = 3+3 = 6$. However, it can be seen that another strategy combination, though it isn't a solution, minimizes the total cost, namely $C("Silent", "Silent") = 2+2 = 4$. So the efficiency can be seen as the ratio between the total cost achieved under all possible solutions and the total cost achieved under an "ideal" strategy combination, which could be possible under a centralized guidance. The same analysis can be done when choosing the Nash Equilibrium as a solution concept.

1.2 What is Mechanism Design?

Research in computer science is concerned with protocols and algorithms applied at interconnected collection of computers. The emergence of the Internet motivated even more the study of such algorithms and protocols. However, the research before this breakthrough assumed a centralized authority which fully instructs each part of the network. This opposes the structure and rules of the Internet. The computers belong to different agents (persons and organizations) and it is unlikely to follow the designed protocols and algorithms. Their only concern could be to benefit the most. All this research seems inadequate to be realistically applied.

Mechanism design, a subfield of economics which similarly is interested in economic systems, seems to fill the gap. The research done in mechanism design assumes a game-theoretic setting, where the agents, the different members of society, act rationally in order to benefit the most. By borrowing ideas from mechanism design, it is possible to model the agents' rationality and construct algorithms which function in such settings. The new subfield of computer science was named **Algorithmic Mechanism Design**. In order to walk through, some basic ideas of mechanism design must be clarified.

A first approach to give some intuition for the difficulties of mechanism design will be to study the social choice problem. A single decision is to be made and every agent has private preferences for the final outcome. The social choice is the aggregation of the agents' preferences. The main difficulty of the social choice problem is that agents act rationally, declaring fake preference in order to mislead the mechanism and achieve a better outcome for theirselves. Which is the best social choice and how can we find it when agents lie? Condorcet's Paradox shows that the problem is not only difficult but also there are underlying difficulties which seem unapproachable. It also serves as an introduction to voting systems, which will be used to define some basic notions for mechanism design.

Condorcet's Paradox: Every mechanism involves a set of alternatives, i.e. the possible outcomes of the mechanism. Every agent has private preferences for the set of alternatives. These preferences can be modeled as private values for every alternative or an ordering of the alternatives etc. Suppose there are three alternatives A, B, C and three voters with private preferences:

- $A \succeq_1 B \succeq_1 C$
- $B \succeq_2 C \succeq_2 A$
- $C \succeq_3 A \succeq_3 B$

The social choice function is actually a function which takes as input the agents' private preferences and outputs the winning alternative. The most natural social choice is the majority. The alternative that the majority prefers must by chosen. However in such case it is obvious that all three alternatives are equally prefered. Furthermore, a voter who will decide not to support his prefered candidate, will have the opportunity to decide the outcome. For example, agent 3 may decide not to support alternative C and choose between A and B. His choice will simultaneously be the winning alternative.

Condorcet's Paradox shows that there are difficulties. An impossibility theorem will be presented that will form the setting and guide us to fully understand these underlying difficulties. As we will see an "ideal" mechanism may be a difficult task to design. But what is an "ideal mechanism"? Surely a system, which managed to extract all private orderings truthfully could succeed. L is the set of all possible private orderings. The following definition proposes the "ideal" social choice function:

Definition A social choice function f can be strategically manipulated by voter i if for some $\prec_1, \ldots, \prec_n \in L$ and some $\prec'_i \in L$ we have that $a \prec_i a'$ where $a = f(\prec_1, \ldots, \prec_i, \ldots, \prec_n)$ and $a' = f(\prec_1, \ldots, \prec'_i, \ldots, \prec_n)$. That is, voter i that prefers a' to a can ensure that a' gets socially chosen rather than a by strategically misrepresenting his preferences to be $\prec'_i$ rather than $\prec_i$. f is called incentive compatible if it cannot be manipulated.

However, there is a negative property that a mechanism should avoid. This is when f is a dicatorship, formally:

Definition Voter i is a dictator in social choice function f if for all $\prec_1, \ldots, \prec_n \in L, \forall b \neq a, a \succeq_i b \Rightarrow f(\prec_1, \ldots, \prec_n) = a$, f is called a dictatorship if some i is a dictator in it.

Thus, an "ideal" mechanism can be considered to be an incentive compatible mechanisms with no dictators in it. Unfortunately, Gibbard-Satterthwaite proved that there is no such mechanism.

Theorem 1.2.1 *Let f be an incentive compatible social choice function onto A, where $|A| \geq 3$, then f is a dictatorship.*

The theorem seems pessimistic and that the only way is to choose between dictatorships and strategically manupulated mechanisms. However, as any barrier, this theorem can be examined in depth to extract the real properties that make the task impossible. It makes clear that a mechanism must be enriched with additional features, in order to overcome this barrier.

1.3 Mechanism Design with money

A special class of mechanism design problems is those closely related with "money". In real world in most circumstances the mechanism designer (seller) seeks revenue for the goods distributed through the mechanism. So it is completely natural for a mechanism to identify these payments from the agents to the mechanism designer. By using money it is also much more easier to model an agent's preferences and quantify the difference for an agent between any two outcomes. This additional feature makes it possible to skip the imposibility theorems, mentioned before, and achieve mechanisms with the prefered properties. Basic notation will be presented for a general mechanism design model.

We assume that every agent has a utility function and intends to maximize it. The mechanism has a goal to produce an output (algorithm) and define the payments for every agent. The agents' utilities must form a Nash equilibrium, meaning that by changing their declared preferences they cannot meliorate their utility.

- n agents
- Agent's Private Valuations: $v_i \in V_i$
- Output: $f = f(v^1, \dots, v^n)$
- Payments: $p^i = p^i(v^1, \dots, v^n)$
- Agent's Utility: $u_i(v^1, \dots, v^n) = v_i(f(v^1, \dots, v^n)) - p_i(v^1, \dots, v^n)$

There is a special class of mechanisms, called incentive compatible mechanisms. An incetive compatible mechanism imposes an agent to truthfully declare her preferences. An agent admits her maximum utility when declaring her true value.

Definition A mechanism $(f, p_1, \dots, p_n)$ is called incentive compatible if for every player i, every $v_1 \in V_1, \dots, v_n \in V_n$ and every $v_i' \in V_i$, if we denote $a = f(v_i, v_{-i})$ and $a' = f(v_i', v_{-i})$, then $v_i(a) - p_i(v_i, v_{-i}) \geq v_i(a') - p_i(v_i', v_{-i})$.

1.3.1 Vickrey Auction

The well known Vickrey Auction is actually a single item auction with n agents bidding for the item. All agents have private valuations v_i for winning the item and zero valuation otherwise. The agents declare their values by bidding b_i. Ideally, the bids match their private valuations. The main goal of the mechanism is to generate the maximum social welfare. The natural social choice function f allocates the item to the agent bidding the most. The social welfare is equal to $SW(\mathbf{b}) = v_{f(\mathbf{b})}$. If the agents bid truthfully this function achieves its goal. However, an agent may win the item without having the maximum valuation by bidding the most. The social choice function accompanied with a payment scheme determines the agents' behaviour and eventually the auction's efficiency. Let's examine the above setting with the most natural payment scheme, ie the agent winning the item pays its bid. This is the single item first price auction.

The first price auction has a social choice function $f(\mathbf{b}) = \arg\max_i b_i$ and a pricing scheme $p_i = \begin{cases} b_i & \text{if } i = f(\mathbf{b}) \\ 0 & \text{otherwise} \end{cases}$. The agents have utility $u_i = \begin{cases} v_i - p_i & \text{if } i = f(\mathbf{b}) \\ 0 & \text{otherwise} \end{cases} = \begin{cases} v_i - b_i & \text{if } i = f(\mathbf{b}) \\ 0 & \text{otherwise} \end{cases}$. At first sight it is clear that an agent tends to bid lower that his true valuation in order to achieve a lower price.

Let's assume 2 players with valuations $v_1 > v_2$ and truthful bidding, namely $b_1 = v_1$ and $b_2 = v_2$. The item goes to player 1 and the agents have utilities $u_1 = v_1 - p_1 = v_1 - b_1 = v_1 - v_1 = 0$ and $u_2 = 0$. Agent 1 tends to lower her bid to $v_2 + \epsilon$ in order to get the item and achieve utility $u_1 = v_1 - (v_2 + \epsilon) \approx v_1 - v_2$. Thus, a first price auction supports strategic bidding. In this situation, strategic bidding cannot affect the mechanism's efficiency but in other cases there might be a problem. However, as an introduction, we will see a very natural way, eventually, to achieve truthfulness, which could seem a very demanding task at first sight.

Actually the idea of strategy-proofness is that at the state of truthful bidding, no agent has an incentive to change her bid, assuming that all others keep their bids as they are. No agent can achieve greater utility by changing her bid. As mentioned before, the problem with the first price auction appears as the winner bidder tends to pay the least he can to keep the item. This is actually the best second bid. Thus, it may be a good idea to impose a pricing according to the following bid. This also includes the obvious for a payment, which should not depend directly on the agent's bid. This is the single item second price auction or the Vickrey auction. It has a social choice function$f(\mathbf{b}) = \arg\max_i b_i$ and a pricing scheme $p_i = \begin{cases} b_{\arg\max_{j \neq i} b_j} & \text{if } i = f(\mathbf{b}) \\ 0 & \text{otherwise} \end{cases}$.

Let's consider that agents bid truthfully. There is the winner $f(\mathbf{v})$, who has positive utility and all other players with zero utility. We'll examine the two cases seperately to see if the agents have an incentive to change their bids.

- Winner: The advantage of such a mechanism is that an agent's utility won't change unless the allocation changes. Thus, the winner can only consider to lower her bid so that the item can be allocated to another agent. If this happens, her utility collapses to zero. Since at truthful bidding her utility is positive, the winner has no incentive to change her bid.

- Losers: Similarly, an agent can only consider bidding higher than the winner and gaining the item. However, such an allocation will generate negative utility, namely $u_i = v_i - p_i = v_i - v_w$, where v_w is the winner's valuation. Since the winner is the one with the maximum valuation $u_i < 0$. Thus, the agent has no incentive to change her bid.

The second price auction may now seem natural since it uses a natural social choice function and tries to impose the payments that an agent would use a strategy to go for. In the first price auction an agent would lower her bid to the least in order to gain the item. So in the second price auction they charge her this price from the beginning, so she would not consider lowering her bid. This gives an intuition of a mechanism which itself simulates the agent's strategic behaviour. A mechanism which manipulates the agents' bids before passing them into the natural mechanism.

1.3.2 Incentive Compatible Mechanisms

The following theorem makes use of this intuition to support the existence of incentive compatible mechanisms.

Theorem 1.3.1 *If there exists an arbitrary mechanism that implements f in dominant strategies, then there exists an incentive compatible mechanism that implements f. The payments of the agents in the incentive compatible mechanism are identical to those, obtained at equilibrium of the original mechanism.*

Now that the existence of incentive compatible mechanisms is theoretically proven, there is a question whether there is a standard method to design such mechanisms. In fact a family of incentive compatible mechanisms will be presented, which can be applied for every possible space of alternatives and social choice function. This is the Vickrey-Clarke-Groves mechanism. Actually it is a generalization of the mechanism proposed for the Vickrey auction. As mentioned before, a mechanism is a combination of a social choice function and a vector of payments for all agents. Let's consider the most natural social choice, that which maximizes the social welfare, namely $\sum_i v_i(a)$.

Vickrey-Clarke-Groves Mechanism: For all $v_1 \in V_1, \ldots v_n \in V_n$: $p_i(v_1, \ldots, v_n) = h_i(v_{-i}) - \sum_{j \neq i} v_j(f(v_1, \ldots, v_n))$, where $h_i : V_{-i} \rightarrow \mathcal{R}$, which does not depend on v_i. It is clear now that a user's utility keeps up with the social welfare. Since the social choice is the one with the maximum social welfare, a user gains when the mechanism works properly, so he acts truthfully.

Theorem 1.3.2 *Every VCG mechanism is incentive compatible.*

Proof Let's consider agent's i private value v_i and all other players' values v_{-i}. The agent must not have an incentive to falsely declare $v'_i \neq v_i$. Additionally, $a = f(v_i, v_{-i})$ and $a' = f(v'_i, v_{-i})$. If $a = a'$ the social choice does not change and the user has the same utility and no incentive to falsely declare his value. If $a \neq a'$, $v_i(a) - p_i(v_i, v_{-i}) = v_i(a) + \sum_{j \neq i} v_j(a) - h_i(v_{-i}) \geq v_i(a') + \sum_{j \neq i} -h_i(v_{-i}) = v_i(a') - p_i(v'_i, v_{-i})$. The inequality follows since a is the choice with the maximum social welfare. This completes the proof. ∎

However, the functions $h_{-i}(v_{-i})$ have not been defined. This function is independent of the agent's private valuation and must be manipulated in order to satisfy additional properties. The most natural properties of a mechanism are the following:

- Individually Rational: $v_i(f(v_1, \ldots, v_n)) - p_i(v_1, \ldots, v_n) \geq 0$
- No Positive Transfers: $\forall i : p_i(v_1, \ldots, v_n) \geq 0$

Applying these constraints in the Vickrey-Clarke-Groves mechanism the following equations must by satisfied for every agent i:

- Individually Rational: $\sum_i v_i(a) \geq h(v_{-i})$
- No Positive Transfers: $\sum_{j \neq i} v_j(a) \leq h(v_{-i})$

The Clarke pivot rule occurs, formally $h_i(v_{-i}) = \max\limits_{b \in A} \sum\limits_{j \neq i} v_i(b)$. So the VCG payment with the clarke pivot rule is the following

$$p_i(v_1, \ldots, v_n) = max_b \sum_{j \neq i} v_i(b) - \sum_{j \neq i} v_i(a)$$

, where $a = f(v_1, \ldots, v_n)$.

The above VCG mechanism was applied for a special social choice function, which maximizes the social welfare. However, this is a special incentive compatible mechanism applied into a specific social choice function. An incentive compatible mechanism's goal is, given a problem's properties and a desired social choice function, to define a vector of payments which admits strategy-proofness.

A problem's properties are in fact the set of alternatives, the valuation's behaviour over this set and the social choice function. It is challenging to start analyzing the ability to constuct an incentive compatible mechanism for different domains of preferences and social choice functions. Initially the case that the domain of preferences is unrestricted, namely $V_i = \mathcal{R}^A$, will be examined. In such situation it was shown that an incentive compatible mechanism can be imposed for a social choice function which is an "affine maximizer".

Definition A social choice function f is called an affine maximizer if for some subrange $A' \subset A$ for some player weights $w_1, \ldots, w_n \in \mathcal{R}^+$ and for some outcome weights $c_a \in \mathcal{R}$ for every $a \in A'$, we have that $f(v_1, \ldots, v_n) \in \arg\max\limits_{a \in A'} (c_a + \sum\limits_i w_i v_i(a))$.

The following theorem restricts us to consider only affine maximizers for the unrestricted case.

Theorem 1.3.3 *If $|A| \geq 3$, f is onto A, $V_i = \mathcal{R}^A$ for every i, and $(f, p_1, \ldots, p_n)$ is incentive compatible then f is an affine maximizer.*

The only incentive compatible mechanisms imposed for affine maximizers are a generalization of the VCG mechanism.

Theorem 1.3.4 *Let f be an affine maximizer. Define for every i,* $p_i(v_1, \dots, v_n) = h_i(v_{-i}) - \sum_{j \neq i} \frac{w_j}{w_i} v_j(a) - \frac{c_a}{w_i}$*, where* h_i *is an arbitrary function that does not depend on* v_i*. Then,* $(f, p_1, \dots, p_n)$ *is incentive compatible.*

It is natural to have difficulties in such a general case. However, by relaxing the $V_i = \mathcal{R}^{\mathcal{A}}$ constraint, more mechanisms may be applied. A well studied case are the single-parameter domains, where the domains V_i can be determined by a single parameter. We define for every agent i, a winning set $W_i \subseteq A$ such that for every $a \in W_i : v_i(a) = v_i$, and it's complement for which $v_i(a) = 0$.

Definition A single parameter domain V_i is defined by a $W_i \subset A$ and a range $[t^0, t^1]$. V_i is the set of v_i such that for some $t^0 \leq t \leq t^1$, $v_i(a) = t$, for all $a \in W_i$ and $v_i(a) = 0$ for all $a \notin W_i$.

The following two definitions are necessary for defining the class of incentive compatible mechanisms in single-parameter domains.

Definition A social choice function f on a single parameter domain is called monotone in v_i if for every v_{-i} and every $v_i \leq v'_i \in \mathcal{R}$ we have that $f(v_i, v_{-i}) \in W_i$ implies that $f(v'_i, v_{-i}) \in W_i$. That is, if valuation v_i makes i win, then so will every higher valuation $v'_i \geq v_i$.

Definition The critical value of a monotone social choice function f on a single parameter domain is $c_i(v_{-i}) = \sup_{v_i : f(v_i, v_{-i}) \in W_i} v_i$. The critical value at v_{-i} is undefined if $\{v_i | f(v_i, v_{-i}) \notin W_i\}$ is empty.

The following theorem fully determines the incentive compatible mechanisms in single-parameter domains.

Theorem 1.3.5 *A normalized mechanism* $(f, p_1, \dots, p_n)$ *on a single parameter domain is incentive compatible if and only if the following conditions hold:*

- *f is monotone in every* v_i*.*
- *Every winning bid pays the critical value. Formally, for every i,* v_i*,*v_{-i} *such that* $f(v_i, v_{-i}) \in W_i$*, we have that* $p_i(v_i, v_{-i}) = c_i(v_{-i})$*.*

This mechanism may be the first attempt to relax the assumption of unrestricted domain, but also serves as a complement to Arrow's theorem for the case of $|A| = 2$. The case of two alternatives falls into the category of single parameter settings, since only one parameter is able to express the valuation difference between the two alternatives.

1.4 Auctions - Examples

Auctions serve as the most practical and intuitive setting to study mechanism design. Technically, many mechanisms can be viewed as an auction. Furthermore, auctions are met in many instances of real life, so they are practically interesting. Two special classes of auctions follow in order to get well with the properties and challenges of auction design and mechanism design generally.

1.4.1 Combinatorial Auctions

In contrast with the single item auction, there may be a set of different items for sale. All n agents tend to purchase these m items and have private valuations for every possible combination of items.

Definition A valuation v is a real-valued function for each subset S of items, $v(S)$ is the value that bidder i obtains if he receives this bundle of items. A valuation must have "free-disposal", i.e., be monotone: for $S \subseteq T$ we have that $v(S) \leq v(T)$, and it should be "normalized": $v(\emptyset) = 0$.

Every agent must report her private valuations for every possible bundle. However, she can behave strategically to manipulate the mechanism and achieve better utility. The mechanism has to allocate the m items to the n agents with the most efficient method (social choice function) and price accordingly. The mechanism's efficiency is measured by the social welfare $SW = \sum_i v_i(S_i)$, where S_i is agent's i purchased set of items.

Such mechanisms are extremely demanding and there are many obstacles in the way. In this auctions the problems may start from the social choice function, that before was taken for granted. Since we talk for bundles of items, there is an exponential number of possible allocations. Consider that it is already timeworthy for an agent to declare her private valuations and the social choice problem might be hard. Furthermore, it is very challenging to design incentive compatible mechanisms or efficient mechanisms generally.

1.4.2 Digital Goods

Due to advance of technology and the Internet, a new type of auctions emerged. A company possesses audio files and can distribute them through the Internet. There are no constraints at the number of copies and downloads distributed. An auction is formated where the auctioneer sells an item, which is available in unlimited supply.

We assume that the auctioneer has no cost for each copy. There are different settings that this process may be applied. The auctioneer may be called to set a fixed price or set a price and as the agents buy copies she can manipulate the price. In such a setting a basic concept is the efficiency of the mechanism. A maximum social welfare objective seems trivial, as the auctioneer can distribute for free the audio file. In such a case all agents would buy the copy. However, such a mechanism is unrealistic since the auctioneer has no payoff. Thus, we introduce new objectives. The auctioneer tends to increase her revenue, ie the sum of payments in a setting where the agents' private valuations follow probability distributions, since she has no information of them. These auctions are suitable for the study of other objectives, such as the revenue maximization.

Chapter 2

Introduction in Sponsored Search Auctions

2.1 Motivation

Internet is the phenomenon of our age, which seems that it ushered in a new revolution as important as the industrial revolution. Long ago the creation of the Internet, the necessity of communication, imformation was known for their role in peoples' lives. Internet appears to be the breakthrough due to the direct communication of every two nodes in the whole world and the absence of limitations. At 2012, it is estimated that approximately 30% of the world population use the Internet, meaning up to 2,000,000,000 people. This new medium, due to its acceptance, created a new economy in terms of exchange and commerce. This was expected as the Internet has become a large portion of our society, where people can take actions similar to those in true life.

Google, Microsoft, Yahoo, Facebook are the pioneers of the Internet in our days. They provide services which attract the average user. Services which are trivial at first sight, however they are carefully and in depth developed and closely related to the main aspects of the Internet. Services which facilitate communication and imformation retrieval, such as communication clients, social networks and search engines. Google managed to claim over 1,000,000,000 unique users on its websites in a single month. It is quite an amazing stat if one considers the tremendous impact a company like that can have.

The most popular services mentioned are free. A single user can use them by only registering in the corresponding websites. This seems misleading, why such companies do not make money from the services they provide. However the answer is that they do make money indirectly, managing to maintain the popularity of their services. In order to monetize their free services they sell advertising. The area used in the websites for advertising may seem attractive for advertisers, if one considers the wide audience. In such way the companies manage to extract revenue and exploit their popularity. Examples of online advertising include contextual ads on search engine results pages, banner ads, blogs, Rich Media Ads, Social network advertising, interstitial ads, online classified advertising, advertising networks and e-mail marketing, including e-mail spam.

Auctions seem the predominant way to monetize these services. As mentioned before the advertising area is sold and it is limited in a sence. The most effective way of advertising until today are the sponsored search auctions. People use daily a search engine to gather information. It is the service that most users take for granted when they “surf” through the Internet. Sponsored search auctions is a clever way to monetize this service. When a user searches for a keyword, the search engine returns to him the most relevant and popular websites, in which the keyword appears. These are the so-called organic results of the search engine. Usually above or in the right of the organic results sponsored results appear. These results resemble the organic ones but they are the outcome of an auction. An auction which takes place every time this specific keyword is searched. In many occassions, sponsored results appear in websites, when a search engine reserves an advertising area.

Before defining the structure of Sponsored Search Auctions, it would be motivating to present their significance using stats. According to the Internet Advertising Revenue Report the internet advertising revenues in the United States totaled 31.7 billion dollars. Sponsored Search Auctions account for 46.5% of 2011 revenues and remains the most profitable advertising format. It is remarkable also that internet advertising has surpassed cable television advertising in revenues and Internet has become the most preferable advertising media before broadcast TV. It is clear now that internet advertising has become a huge industry and approximately half of its revenues come from Sponsored Search Auctions. Their significance is even more evident if one considers that up to 90% of Google’s revenues also come from Sponsored Search Auctions. Google is considered one of the pioneers of the Internet with a significant contribution to innovation.

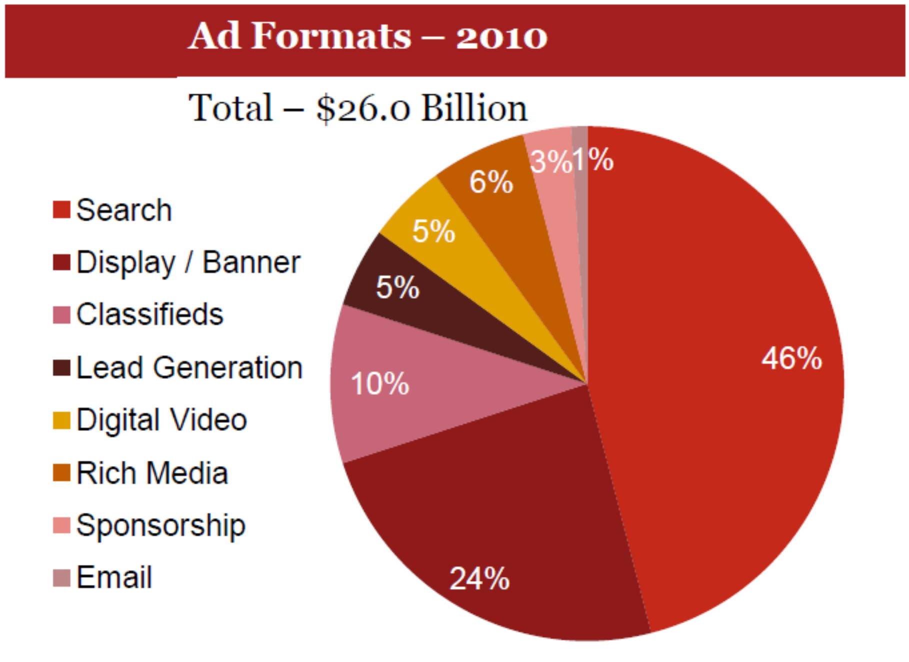

Figure 2.1: Internet Advertising

2.2 Structure

As mentioned before, Sponsored Search Auction refer to keyword auctions which take place every time a user searches for a keyword in a search engine. The search returns the main results which are the organic ones and another set of results called the sponsored results. Some times there is a set of sponsored results above the organic results and Microsoft Bing pioneered presenting sponsored results below the organic ones. The common format of a set of sponsored results is a listing of slots where every advertiser occupies an available slot. Each ad contains a title, which serves as a hyperlink to redirect the user to the ad's website, and a brief description.

The prevelance of sponsored search auctions as an advertising format is explained by the fact that seperate auctions run for each seperate search query. The user expresses her intent by searching for a query. This is indicative of the advertisements which could affect him the most and satisfy her desires.

Furthermore, the structure of these auctions impose the advertisers to

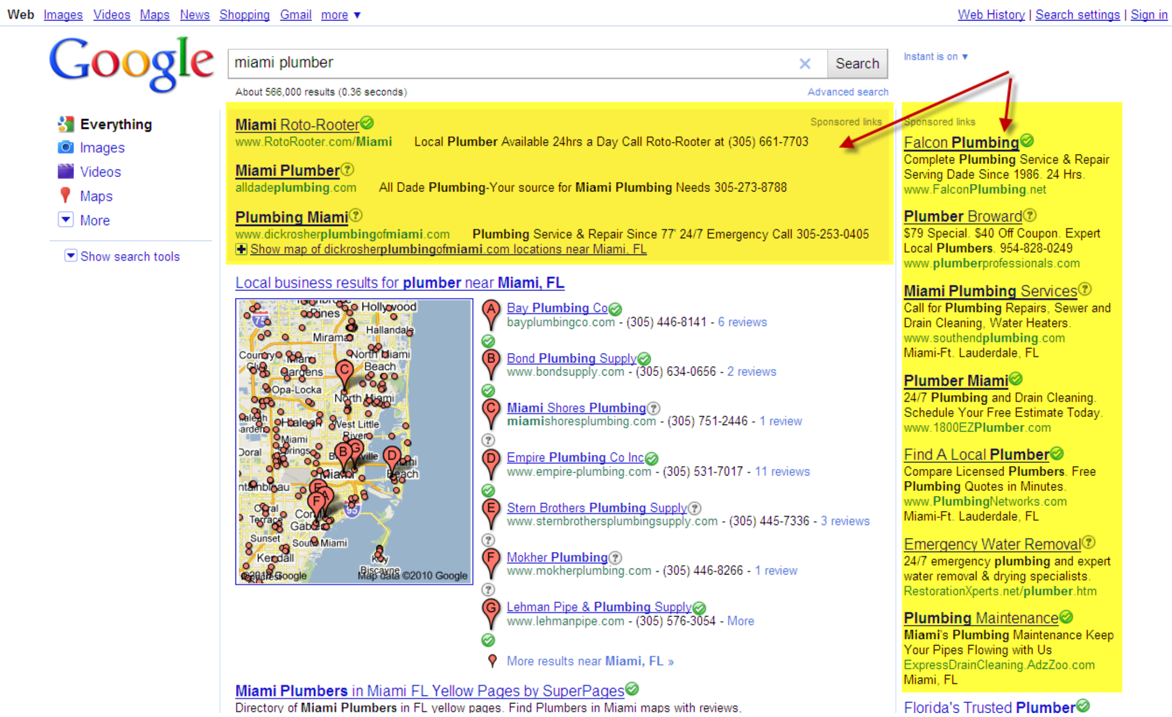

Figure 2.2: Sponsored Search Auctions format

choose a search query and submit single bids. The bids serve as the input of the auction. These bids express the maximum willingness of the advertisers to pay for each click. So the user pays per click. In the case of sponsored search auctions it is not clear of what is exactly sold. From the search engine's perspective, the unit sold is the impression made to the user who searches for a keyword. This suggests a payment per-impression. On the other side, it is natural for an advertiser to pay every time a user clicks on an advertisement and actually proceeds on a transaction. This suggests a payment per-conversion, which particularly is difficult to apply because of the difficulty to measure the actual effectiveness of the advertisement. Payment per-click actually serves as a middle ground between the natural payment methods of the two sides. Click fraud is brick wall which had to be overcomed. It is a technique used by rival parties who produced a large number of clicks, in order to charge a particular advertiser. Click fraud was considered unlawful and steps were made to be treated. These steps consist of methods to recognize such actions and restructure the rules of the auction to alleviate the bad effects.

Let's talk about bids. Advertisers submit single bids for a keyword auction. However an auction takes place each time a keyword is searched. This leaves the opportunity to an advertiser to change his bid between each two auctions. The model used suggests a continuous and infinitely repeated game, where each advertiser's profit is approximated in the long run and he may change bid accordingly.

Finally, despite the environment's special characteristics the advertisers submit single bids. It seems controversial as there are multiple positions with unique desirability for each position. For example two advertisers may prefer the first position from the second one, but a single type seems inadequate to fully determine their vague preferences. One bid per keyword may not be sufficiently expressive to fully convey preferences. On the other hand, the limitations put are not large enough to justify added complexity in the bidding language. A single bid accompanied with position-defined coefficients may do the job. At next chapters, consideration of externalities between advertisers will oppose our choice.

2.3 Notation

Sponsored search auctions admit K slots and n agents. The available slots are numbered from top to bottom. The unique structure of the game imposes a complete information setting, meaning that agents have full information of the other agents' bids. This makes sence if one considers that expertised advertisers can extract such information using statistical methods and by testing different bids and observing fluctuations in the allocation process. Each agent is indexed $i \in \{1, \dots, n\}$. An agent's preference can be expressed by his valuation v_i of a single click. A bid b_i is used to capture this variable and also expresses, as mentioned before, an agent's maximum willingness to pay for a click. Every time an auction takes place, the mechanism charges payments p_i.

In order to define a model for sponsored search auctions using a single parameter to express the agents' preferences, it is necessary to present a measure for the differences of different allocations of the slots to advertisers. Click through rate, which serves as a measure, is the probability that a specific ad is clicked under a specific allocation instance. It depends on the allocation rule to fully define click-through rate. It is also a very convenient way to express the desire for specific slots. For example, an advertiser may want to be allocated the first slot because the outcome will guarantee a high click-through rate. Intuitively this is true because advertisers yearn for a high probability of clicking their ad, in order to redirect the user at the corresponding website. Generally, we assume that click-through rates are known to the auctioneer and the agents.

We will use a basic model to analyze sponsored search auctions. Click-through rate may depend from a variety of factors. Any difference between two allocation instances may affect each agent's CTR. For a start, the model should satisfy two requirements. Firstly, an ad's click-through rate surely depends on it's quality. We will model it's quality as a probability to be clicked if a user views the corresponding slot. We call this probability relevance q_i. Furthermore, it appears that higher slots have a greater probability to be viewed. As a consequence, every slot has a probability λ_k to be viewed, satisfying $\lambda_1 \geq \lambda_2 \geq \ldots \geq \lambda_k$. Naturally, the click-through rate depends on both probabilities and we model it as $CTR_{k,i} = \lambda_k \cdot q_i$ for agent i occupying slot k. This is the seperable click-through rate model.

An advertiser i has a valuation v_i for every click and pays p_i. The mechanism allocates, according to the bids, every slot to an advertiser and we assume that advertiser i occupies slot k. So she has a probability $CTR_{k,i}$ to gain a click or equivalently gains $CTR_{k,i}$ clicks per period. The advertiser's profit or utility is $u_i = CTR_{k,i} \cdot (v_i - p_i)$. The agent wants to maximize this utility, acts selfishly and is willing to change her bid to achieve it. In order to overcome such a problem there is the need of a stable mechanism. The notion of stability in such circumstances is given by Nash. In our setting at a Nash equilibrium, no advertiser can change her bid, keeping other bids fixed, to maximize his utility:

$$\forall i, b'_i : u_i(b_i, b_{-i}) \geq u_i(b'_i, b{-}i)$$

. Only a mechanism, which guarantees at least a Nash equilibrium at every instance can be accepted and analyzed.

To evaluate the outcome of a mechanism we will use the social welfare. The social welfare is the sum of profits gained by all agents and the auctioneer. Thus, $SW = \sum_i u_i + \sum_i CTR_{k,i} \cdot p_i = \sum_i CTR_{k,i} \cdot (v_i - p_i) + \sum_i CTR_{k,i} \cdot p_i = \sum_i CTR_{k,i} \cdot v_i$. The social welfare measures the allocation efficiency of a given mechanism and actually doesn't depend on the payment but only on the allocation process. The maximization of the social welfare can be seen as a classical algorithmic problem, which is trivial in the basic case as we will see later. In our case the mechanism designer is the auctioneer, whose main goal is to maximize his revenue. Thus, another measure for evaluation is the revenue $R = \sum_i CTR_{k,i} \cdot p_i$, meaning the sum of all agents' payments. The mechanism designer aims to construct a mechanism, which maximizes the social welfare and the auctioneer's revenue. Generally, such a mechanism may not exist, but intuitively we could guess that a mechanism with a high enough social welfare will guarantee a high revenue for the auctioneer. Vice versa a mechanism with low social welfare cannot locate high payments.

Chapter 3

Allocation and Pricing in Sponsored Search Auctions

As mentioned in the previous chapter, a mechanism designer aims to construct a mechanism, which is stable and guarantees high social welfare and revenue for the auctioneer. Her goal is even more difficult, as the input of the mechanism may depict false values, clearly meaning the bids which should be identical with the private valuations. An ideal mechanism should push the agents to declare their true values. However this is not obligatory, as many mechanisms give rise to strategic behaviours which bring approximately good outcomes.

The mechanism can be divided into two parts. The allocation process which allocates the slots to specific agents and the pricing process which defines every agent's payment according to the allocation process. The allocation seems a simple maximization problem at first glance. However, it cannot be taken for granted because it may affect agents to falsely declare their private valuations. It is clear now that the process can be divided into two steps but the analysis requires the whole mechanism to evaluate. As I proceed, I will describe various allocation and pricing rules and combine them to analyze and evaluate the corresponding mechanisms.

3.1 Allocation

An allocation rule cannot be seen only as a simple maximization problem. As I said it may lead to strategic behaviours but also there is the need of simple allocation rules. A search engine must use simple rules to be easy for the user to understand them and effectively use them. Furthermore an allocation rule must meet a fairness requirement. For example, an advertiser may have very low quality and that leads to very low click-through rates. However it is unfair to never let her get a slot. An advertiser like that should have the opportunity to occupy a slot by bidding a little better that the others. With these in mind two allocation rules for the basic model have been proposed:

- **Rank by Bid:** Advertisers are ranked in the decreasing order of the submitted bids.
- **Rank by Revenue:** Advertisers are ranked in the decreasing order of the ranking scores, where the ranking score of an advertiser is defined as the product of the advertiser's bid and relevance ($r_i = q_i \cdot b_i$).

3.2 Pricing

Now that the basic allocation rules are clear, I will describe the basic payment rules. An analysis follows the description for every payment rule combining it with the rank by revenue rule. This is without loss of generality as the mechanism's behaviour remains the same with little differences. Rank by revenue achieves better results and that is the reason that it is currently used.

3.2.1 GFP

Generalized First Price payment rule was introduced in the first sponsored search auction design in 1997. It is a traditional payment rule, which reminiscent the traditional auctions. Every agent submits a bid, as a declaration of her willingness to pay, and if clicked pays this ammount. The format was used for several years but it had negative results and in 2002 a new rule was introduced to replace it. Gradually, the whole market abandoned the GFP pricing rule.

A basic property we seek from a mechanism is to be truthful. In our case, an agent's bid ideally should be indentical with her valuation. It is clear that GFP fails to acheive this. If an agent declares her true valuation and be clicked, she would pay the same ammount. This brings zero profit, which is unwanted for every agent. Thus, an agent would falsely declare a lower bid to occupy a slot and pay less. As mentioned before, not having a truthful auction is not basically a problem but becomes more complicated to achieve stable and effective mechanisms.

During the use of the GFP pricing rule, it was observed that it cultivated a bidding strategic behavour. Agents continuously changed their bids to maximize their utility. This starts to make sence if one considers that in this case not only the allocation but also the bid itself affects the payment. It seems natural now to analyze the mechanism in order to test it's stability. As mentioned before, we assume that a mechanism is stable if it has a Nash equilibrium. Presenting a simple example, it would be clear that the GFP mechanism, under the rank by bid rule without loss of generality, is unstable in the general case.

Example: Suppose there are two slots and three advertisers. All agents' qualities are excluded, meaning that $\forall i : q_i = 1$. I assume that the first slot receives 200 clicks per period and the second slot 100 clicks per period. Furthermore, the agents 1, 2, 3 have valuations \$10. \$4 and \$2 per click respectively. The agents' initial bids are \$2.02, \$2.01, \$2 respectively. Let's observe the agents' behaviour (Figure 3.1).

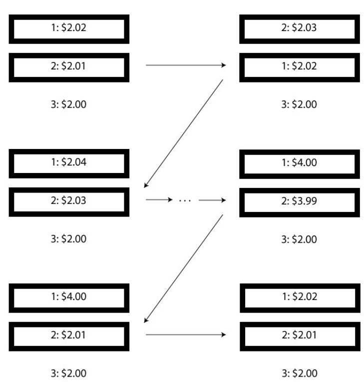

Figure 3.1: Example: First Price Auction

One can observe, that advertisers can use autobid systems to adjust their bids at any time, in order to achieve desired placement and avoid overbidding. It becomes now a matter of how quickly can one respond to changes and readjust her bid. This was observed empirically as the autobidders formed a distinctive "sawtooth" pattern by continuously readjusting their bids. This is undesirable generally. However, the failure of the generalized first price auction has gained importance as it generated revenue losses for the auctioneer.

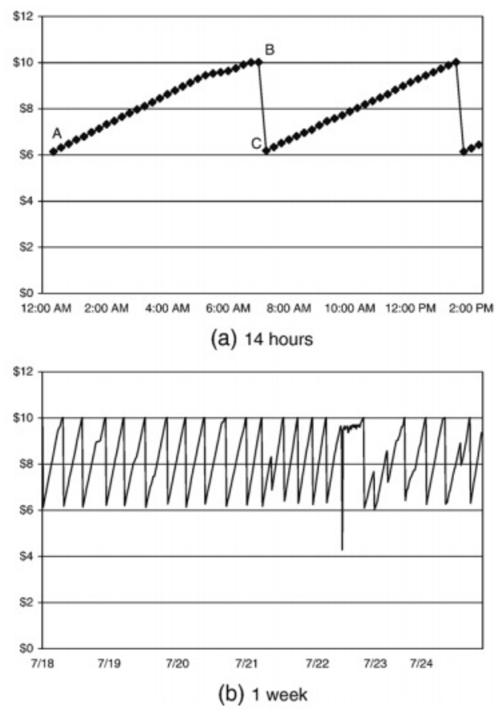

Figure 3.2: Sawtooth bidding pattern. (a) 14 hours (b) 1 week

Theoretically, revenue losses can be proven by combining the revenue generated from a first price auction and the revenue generated from a second price auction. A second price auction means that a bidder pays the bid of the agent located underneath her. As we see the second price auction, charges prices sustantially lower than the first price auction, but finally the revenue collected remains the same. Let's consider two agents a and b with private valuations $v_a > v_b$. Similarly with the example presented, the two bidders increase bids step by step from the minimum possible amount ϵ to v_b. Assuming that the agents spend equal time at each step the average revenue generated under the GFP rule is $\frac{v_b+\epsilon}{2}$. At the second price auction, if agents bid v_a and v_b, agent a occupies the first slot and pays v_b. This is clearly a Nash equilibrium. Thus, the first price auction fails to generate high revenue.

3.2.2 VCG

Having in mind the failure of GFP pricing scheme, it is natural trying to design a truthful auction. A truthful auction will be stable, as declaring the true valuations is a dominant strategy. It would be very convenient for the agents to participate in a truthful auction, as it would be no more necessary to determine a strategy but simply to declare her true valuation. This would simplify the bidding process significantly. After the design of a truthful auction, an analysis follows in terms of allocation efficiency and revenue.

The most natural choice seems to be a VCG mechanism which outputs the same ordering with the rank by revenue rule. A VCG mechanism admits a social choice function f, which is an affine maximizer, $f(v_1, v_2, \ldots, v_n) \in \arg\max_{a \in A}(c_a + \sum_i w_i v_i(a))$. It only remains to specify the weights $\{w_1, w_2, \ldots, w_n\}$ and the bias c_a. In our case, if allocation a gives slot j to advertiser i, $v_i(a) = CTR_{j,i} \cdot v_i = \lambda_j \cdot q_i \cdot v_i = \lambda_j \cdot r_i$. As the position multipliers are decreasing slot by slot, it is clear that an affine maximizer with $w_i = 1, \forall i$ and $c_a = 0, \forall a$ ranks the agents according to the rank by revenue rule. The rank by bid rule could also be simulated with a similar VCG mechanism with weights $w_i = \frac{1}{q_i}$. However, in this section, I would only analyze the rank by revenue case.

Without loss of generality, I assume that agents are indexed such that advertiser i occupies slot i. At the long run an agent pays $CTR_{i,i} \cdot p_i$. According to the VCG rule

$$CTR_{i,i} \cdot p_i = OPT_{-i} - \sum_{j \neq i} CTR_{j,j} \cdot b_j = \sum_{j=i}^{k}(CTR_{j,j+1} - CTR_{j+1,j+1}) \cdot b_{j+1} =$$

$$\sum_{j=i}^{k}(CTR_{j,i} - CTR_{j+1,i}) \cdot \frac{q_{j+1}}{q_i} \cdot b_{j+1}$$

$$\Rightarrow p_i = \sum_{j=i}^{k}(\frac{CTR_{j,i} - CTR_{j+1,i}}{CTR_{i,i}})\frac{q_{j+1}}{q_i}b_{j+1}$$

The above VCG mechanism was named from it's inspirators as laddered auction due to its recursive nature.

Theorem 3.2.1 *The laddered auction is truthful. Further, it is the unique truthful auction that ranks according to decreasing $r_i = q_i \cdot b_i$.*

Proof It is clear that changing a bid without any change on the allocation, an agent's utility remains the same. There is a change when the allocation changes.

Firstly, I will prove that the laddered auction is truthful. Consider an advertiser M, who occupies slot j and pays $p(j)$. If agent bids her true valuation v_M, she occupies slot x. r is the closest rank to x, which generates the maximum profit for the agent. if $r > x$ then the change to moving in slot $r-1$ is $(CTR_{r-1,x} - CTR_{r,x})(v_x - \frac{q_r}{q_x} b_r) \geq 0$. if $r < x$ then the change to moving in slot $r+1$ is $(CTR_{r+1,x} - CTR_{r,x})(v_x - \frac{q_r}{q_x} b_r) \geq 0$.

Secondly, I prove that the laddered auction is the unique truthful auction which ranks by revenue. Let's consider any auction $\mathcal{A}$, which also ranks by revenue. Agent M pays $p_{\mathcal{A}}(j)$, when occupying slot j. All other agents are indexed in the decreasing order of $w_i b_i$, excluding M.

Lemma 3.2.2 $p_{\mathcal{A}}(j) - p_{\mathcal{A}}(j+1) = (CTR_{j,M} - CTR_{j+1,M}) \frac{q_{j+1}}{q_M} b_{j+1}$

Proof

- Suppose $v_M = \frac{q_{j+1}}{q_M} b_{j+1} + \epsilon$: if M bids truthfully, she is ranked at slot j. The additional valuation of being ranked at slot j instead of slot j+1 is $(CTR_{j,M} - CTR_{j+1,M}) v_M$. In order to prevent agent M wanting slot j+1 the payment difference should be lower.

$$p_{\mathcal{A}}(j) - p_{\mathcal{A}}(j+1) \leq (CTR_{j,M} - CTR_{j+1,M}) \frac{q_{j+1}}{q_M} b_{j+1}$$

- Suppose $v_M = \frac{q_{j+1}}{q_M} b_{j+1} - \epsilon$: if M bids truthfully, she is ranked at slot j+1. The additional valuation of being ranked at slot j instead of slot j+1 is $(CTR_{j,M} - CTR_{j+1,M}) v_M$. In order to prevent agent M wanting slot j the payment difference should be lower.

$$p_{\mathcal{A}}(j) - p_{\mathcal{A}}(j+1) \geq (CTR_{j,M} - CTR_{j+1,M}) \frac{q_{j+1}}{q_M} b_{j+1}$$

$\blacksquare$

It is natural to assume that $p_{\mathcal{A}(k+1)} = 0$. Using the above lemma recursively the laddered auction pricing scheme is formatted. Assuming that our proposition is true for agent i occupying truthfully slot i, $p_{\mathcal{A}}(i+1) = \sum_{j=i+1}^{k} (CTR_{j,i} - CTR_{j+1,i}) \frac{q_{j+1}}{q_i} b_{j+1}$. Using the recursion of the lemma, $p_{\mathcal{A}}(i) = p_{\mathcal{A}}(i+1) + (CTR_{i,i} - CTR_{i+1,i}) \frac{q_{i+1}}{q_i} b_{i+1} = \sum_{j=i}^{k} (CTR_{j,i} - CTR_{j+1,i}) \frac{q_{j+1}}{q_i} b_{j+1}$.

$\blacksquare$

In this section I described a truthful auction for keyword auctions. By imposing such payment, every agent bids her true valuation and afterwards all agents are ranked accordingly. By ranking them in the decreasing order of $q_i b_i$, the social welfare is maximized. This is an ideal outcome in one sence. We do not know so far, how much revenue is generated for the auctioneer. An answer of this question will be given in the next section by comparing it with the second-price auction.

3.2.3 GSP

The drawbacks of the GFP scheme, have given rise to the idea that an advertiser will never want to pay one bid increment above the bid of the advertiser beneath her. In relation with the Vickrey auction format in 2002, Google introduced the Generalized Second Price auction. The allocation rules remain the same. Initially, Google used the rank by bid rule and afterwards changed it to the rank by revenue rule. Regarding the pricing, an agent pays per click an amount equal to the bid of the advertiser beneath her. This pricing scheme seems to constrain a strategic bidder behaviour.

An analysis of the GSP format will test its stability, its efficiency and the revenue generated. Before proceeding with the analysis, I consider it necessary to clarify some basic notion familiar with the GSP.

Stability In order to prove its stability, it is sufficient for the common sence to prove that the auction has a Nash equilibrium for all possible instances of agents' valuations. An analysis will be given for the rank by revenue rule, but the rank by bid case is entirely analogous.

Theorem 3.2.3 *There always exists an efficient complete information Nash equilibrium in pure strategies in the Generalized Second-Price Rank by Revenue slot auction.*

Proof The following lemma is essential for the proof:

Lemma 3.2.4 *Given an allocation $\sigma : [K] \rightarrow [n]$, there exists a Nash equilibrium profile of bids b leading to σ in a Second-Price Rank by Revenue slot auction if and only if*

$$(1 - \frac{\lambda_i}{\lambda_{j+1}}) r_{\sigma(i)} \leq r_{\sigma(j)}$$

for $1 \leq j \leq n-2$ *and* $i \geq j+2$.

The proof of the lemma can be found in the original paper. Actually, it presents the state of a Nash equilibrium as a linear system, with the Nash equilibrium constraints and the bids as variables. The equation above is produced by examining the existence of a solution for the linear system.

Regarding the lemma, let's consider a ranking σ according to the agents' true valuations. This ranking totally agrees with the conditions of the lemma, as $i > j : r_{\sigma_i} \leq r_{\sigma_j}$. Therefore, there exists a bid profile which leads to a Nash equilibrium. $\blacksquare$

Since the stability has been proven, meaning that there exists a Nash equilibrium at all instances, it is time to further analyze the game's Nash equilibria. I will mention a stronger class of equilibria, ie a stronger solution concept, called locally envy-free equilibria. These were the first steps made to justify the choice of the GSP scheme. The following analysis will be shown in terms of a simplified model were all agents have the same relevance q_i. Thus, an agent occupying slot i pays $p^{(i)} = \lambda_i b_{\sigma(i+1)}$.

Definition An equilibrium of the simultaneous-move game induced by GSP is locally envy-free if an agent cannot improve her payoff by exchanging bids with the agent ranked one position above her. More formally, in a locally envy-free equilibrium, for any $i \leq \min(n+1, K)$,

$$\lambda_i v_{\sigma(i)} - p^{(i)} \geq \lambda_{i-1} v_{\sigma(i)} - p^{(i-1)}$$

At the definition, a change with an agent above is mentioned because in the opposite direction the definition of a Nash equilibrium and a locally envy-free equilibrium are identical. Sometimes the definition of a Nash equilibrium is weak to prove stability in real life, as it admits that the rest agents keep their strategies stable. However, in our case, let's consider an agent increasing his bid to occupy a slot above her. The Nash equilibrium admits that the agent would bid more than the agent above her. That's not the case since as increasing her bid, the agent above her pays more and may be obliged to lower her bid. That means that occupying slot i-1 doesn't admit paying $b_{\sigma(i-1)}$, as the Nash equilibrium says. The extreme is when by a slight increase of agent's i bid, agent i-1 lowers her bid exactly beneath by bidding $b_{\sigma(i)}$. Thus, agent i finally pays $p^{(i-1)}$.

So, locally envy-free equlibria is a stronger class of equilibria than Nash equiblibrium and manages to ensure total stability in bids.

The following theorem is offered to justify the use of the GSP scheme as it proves the existence of efficient equilibria and the prevelance of GSP against the VCG scheme, in terms of revenue generated. A VCG mechanism would impose payments, which expresses the negative impact of an advertiser to the rest, namely $p^{V,i} = (\lambda_i - \lambda_{i+1})b_{i+1} + p^{V,(i+1)}$. For the theorem a strategy profile B^* is mentioned, which is also a locally envy-free equilibrium. All agents bid $b_i^* = \frac{p^{V,(i-1)}}{\lambda_{i-1}}$, $b_1^* = v_1$ and are indexed in the decreasing order of their bid.

Theorem 3.2.5 *Strategy profile B^* is a locally envy-free equilibrium of game Γ. In this equilibrium , each advertiser's position and payment are equal to those in the dominant-strategy equilibrium of the game induced by VCG. In any other locally envy-free equilibrium of game Γ, the total revenue of the seller is at least as high as in B^*.*

Proof Firstly, it must be proven that the allocation and payments, under B^*, coincide with those in the VCG mechanism. Regarding the allocation agents must be ranked in the decreasing order of their values, namely $b_i^* \geq b_{i+1}^*$:

$$\frac{p^{V,(i-1)}}{\lambda_{i-1}} \geq \frac{p^{V,i}}{\lambda_i}$$

$$\frac{(\lambda_{i-1} - \lambda_i)v_i + p^{V,i}}{\lambda_{i-1}} \geq \frac{p^{V,i}}{\lambda_i}$$

$$a_i(\lambda_{i-1} - \lambda_i)v_i \geq (\lambda_{i-1} - \lambda_i)p^{V,i}$$

$$\lambda_i v_i \geq p^{V,i}$$

which is true since agents pay less than they earn. It easy to check for the first agent:

$$b_1^* \geq b_2^*$$

$$v_1 \geq \frac{p^{V,1}}{\lambda_1}$$

$$\lambda_1 v_1 \geq p^{V,1}$$

Thus, allocation outcomes coincide. Under such a strategy profile, agent i pays $\lambda_i b_{i+1}^* = \lambda_i \frac{p^{V,i}}{\lambda_i} = p^{V,i}$. Thus, payments coincide.

Secondly, it must be proven that strategy profile B^* is a locally envy-free equilibrium. It is trivial to prove that strategy profile B^* is a Nash equilibrium. It remains to prove that agent i has no incentive to swap her bids with the agent above, namely $i-1$:

$$\lambda_i v_i - \lambda_i b^*_{j+1} \geq \lambda_{i-1} v_{i-1} b^*_j$$
$$\lambda_i v_i - p^{V,(i)} \geq \lambda_{i-1} v_{i-1} - p^{V,(i-1)}$$
$$\lambda_i v_i - \lambda_i v_i \geq \lambda_{i-1} v_{i-1} - \lambda_{i-1} v_{i-1}$$

Thus, in order to complete the proof, the revenue under this strategy profile must be compared with the revenue of the VCG mechanism.

The allocation in a locally envy-free equilibrium coincides with a stable assignment, proposed by Shapley and Shubik (1972). The core-elongation property informs us that there exists an allocation with the worst payments from all advertisers. I prove that strategy profile B^* admits the worst payments. In any locally envy-free equilibrium $p^{(k)} = \lambda_k v_{k+1} = p^{V,(k)}$. Let's assume that $p^{(i+1)} \geq p^{V,(i+1)}$ to prove by induction. All other payments should at least satisfy $p^{(i)} - p^{(i+1)} \geq (\lambda_i - \lambda_{i+1}) v_{i+1}$, otherwise agent $i+1$ would find it profitable to bid for slot i. Thus, $p^{(i)} \geq (\lambda_i - \lambda_{i+1}) v_{i+1} + p^{(i+1)} \geq (\lambda_i - \lambda_{i+1}) v_{i+1} + p^{V,(i+1)} = p^{V,(i)}$. Thus, the total revenue is at least $\sum_{i=1}^{k} p^{V,(i)}$ and there is no locally envy-free equilibrium with worse payments that those. This completes the proof. ∎

Allocation Efficiency As mentioned before, the VCG mechanism has a dominant strategy of truthful bidding. Truthful bidding and a rank by revenue allocation maximizes the social welfare guaranteed. However, due to its prevelance against the VCG mechanism in terms of revenue, the GSP mechanism is currently used. It was proven that the GSP mechanism is stable and has Nash equilibria, which acheive allocation efficiency. In order to fully clarify the allocation efficiency of a mechanism it is necessary to consider all possible Nash equilibria. The most popular way to measure such quantity is the Price of Anarchy.

The mechanism's input is a vector b of the agents' bids. I denote as $W(b)$ the social welfare generated under a Nash equilibrium of the game and OPT the optimal allocation of the game. Specifically, by labeling the agents in decreasing order of their values, namely $v_1 \geq v_2 \geq \ldots \geq v_n$, $OPT = \sum_i \lambda_i v_i$ and $W(b) = \lambda_i v_{\sigma(i)}$.

Definition The price of anarchy of a game over a given class of equilibria is defined as the worst ratio of the optimal social welfare over the social welfare of an equilibrium.

$$PoA = \max_b \frac{OPT}{W(b)}$$

At first sight, instances can be found that the price of anarchy can be arbitrarily large. Suppose an auction with two slots and two agents, with $\lambda_1 = 1$, $\lambda_2 = r$, $v_1 = 1$, $v_2 = 0$ and $b_1 = 0, b_2 = 1 - r$. According to these bids the second agent occupies the first slot and the first agent the second slot. This strategy profile is a Nash equilibrium, since changing a bid is of no interest. Specifically,

$$1 : \lambda_2(v_1 - 0) \geq \lambda_1(v_1 - b_2)$$

$$r \geq 1 - (1 - r)$$

$$r \geq r$$

$$2 : \lambda_1(v_2 - b_1) \geq \lambda_2(v_2 - 0)$$

$$0 \geq 0$$

The ratio in such an instance is $\frac{\lambda_1 v_1 + \lambda_2 v_2}{\lambda_1 v_2 + \lambda_2 v_1} = \frac{1}{r}$, which can be arbitrarily large.

However, such an instance is not realistic. It seems strange for an agent to bid over her valuation, as it is risky to have a negative utility. As a consequence, we admit only conservative bidders, namely $b_i \leq v_i$. The price of anarchy will be calculated over the class of conservative bidders equilibria. Such a hypothesis is in fact justified, because an agent doesn't gain by bidding over her valuation:

Lemma 3.2.6 *A bid $b'_i > v_i$ is dominated by $b'_i = v_i$.*

Proof There are two cases:

- $\exists j : b'_i > b_j > v_i$. Such case is not a Nash equilibrium, since agent i has negative utility.
- $\exists! j : b'_i > b_j > v_i$. In such a case, the allocation doesn't change and agent i has the same utility. ∎

It is time to analyze the price of anarchy and eventually calculate it, over the class of conservative bidders. The analysis will start with simple cases of 2 and three bidders and then will be generalized.

Theorem 3.2.7 *The price of anarchy over pure Nash equilibria of GSP auction games with 2 conservative bidders is at most 1.25.*

Proof The proof for the case of 1 slot is trivial. The price of anarchy is 1. Consider the case of two slots, with $\lambda_1 = 1$ and $\lambda_2 = r$. By normalizing the agents' valuation, I assume that $\lambda_1 v_1 + \lambda_2 v_2 = 1 \Rightarrow v_1 = 1 - rv_2$. The incovenient case is when bidder 2 occupies the first slot, namely $b_1 \leq b_2$. A Nash equilibrium and the conservative bidders hypothesis imply:

$$\lambda_2(v_1 - 0) \geq \lambda_1(v_1 - b_2) \geq \lambda_1(v_1 - v_2)$$

$$r(1 - rv_2) \geq (1 - rv_2) - v_2$$
$$[1 + r(1 - r)]v_2 \geq 1 - r$$
$$v_2 \geq \frac{1 - r}{1 + r(1 - r)}$$

Regarding the price of anarchy: $\frac{OPT}{W(b)} = \frac{\lambda_1 v_1 + \lambda_2 v_2}{\lambda_1 v_2 + \lambda_2 v_1} = \frac{1}{v_2 + r(1 - rv_2)} = \frac{1}{r + (1 - r^2)v_2} \leq \frac{1}{r + (1 - r^2)\frac{1-r}{1+r(1-r)}} = \frac{1 + r(1-r)}{r[1 + r(1 - r)] + (1 - r^2)(1 - r)} = 1 + r(1 - r) \leq 1.25.$ ∎

In order to prove the following bounds for the price of anarchy, the analysis will be given for a wider class than the class of pure Nash equilibria. Specifically, an upper bound will be found by testing weakly feasible assignments, a broader solution concept.

Definition An assignment σ is weakly feasible if for each pair of bidders i,j, it holds that $\lambda_{\sigma^{-1}(i)} v_i \geq \lambda_{\sigma^{-1}(j)}(v_i - v_j)$.

Clearly, the class of weakly feasible assignments are wider than the class of pure Nash equilibria, since the inequality of the weakly feasible assignments doesn't consider the payment at the current slot. Furthermore, a weakly feasible assignment is called proper if for any two slots $i < j$ with equal click-through rates, it holds that $\sigma(i) < \sigma(j)$. Generally, a weakly feasible assignment can be expressed with a graph $G(\sigma)$ which redirects every agent to the occupied slot, as it is evident from the image.

It is clear that such a graph will form cycles. In the optimal case, there will be no cycles, but otherwise there will be formed a number of cycles. In case that a single cycle is formed the graph is denoted as irreducible. In the opposite case, the graph will be called reducible.

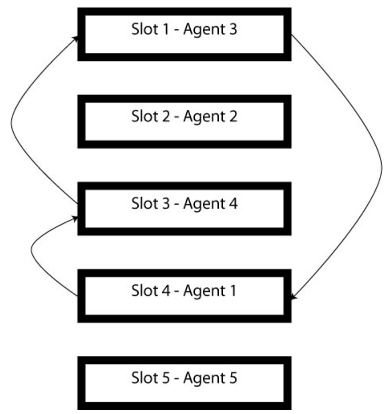

Figure 3.3: Graph G(σ)

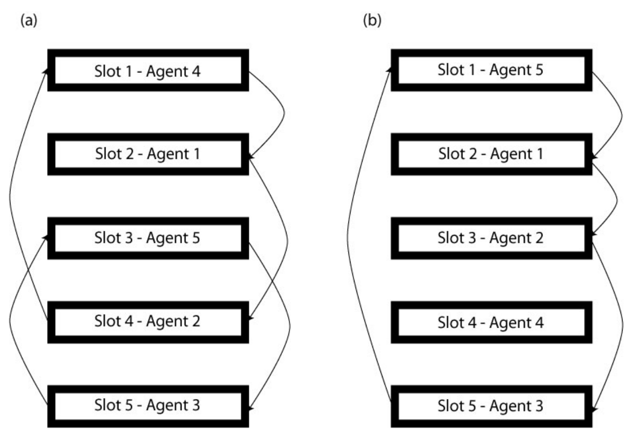

Figure 3.4: (a) Reducible Graph (b) Irreducible Graph

For each cycle, we can construct c auction subgames which correspond to the cycles. Regarding each cycle $l = 1, \ldots, c$, σ^l is the assignment of the l-th subgame. The following lemma is important for the analysis of the upper bounds.

Lemma 3.2.8 *If assignment σ is weakly feasible for the original GSP auction game, the σ^l is weakly feasible for the l-th subgame as well, for $l = 1, \ldots, c$. Then the efficiency of σ is at most the maximum efficiency among the assignments σ^l for $l = 1, \ldots, c$.*

Proof It is trivial to prove that σ^l is weakly feasible, since the inequality is satisfied for the remaining agents of the l-th subgame. Furthermore, the lemma is trivial to prove for the case of irreducible graphs and thus I will restrict my attention in reducible graphs. The total social welfare is the aggregation for all the c subgames, due to the linearity of the social welfare. Applying $\frac{a+c}{b+d} \leq \max(\frac{a}{c}, \frac{b}{d})$, it follows:

$$\frac{OPT}{SW} \leq \max_{i \in [c]}(\frac{OPT^{(i)}}{SW^{(i)}})$$

, where $SW^{(i)}$ and $OPT^{(i)}$ the social welfare and optimal social welfare of the i-th subgame. $\blacksquare$

The analysis of an upper bound for weakly feasible assignments with 3 conservative bidders follows. This upper bound also serves as an upper bound for the price of anarchy.

Theorem 3.2.9 *The price of anarchy over pure Nash equilibria of GSP auction games with 3 conservative bidders is at most 1.259134.*

Proof The agents are indexed in decreasing order of their values, $v_1 \geq v_2 \geq v_3$. Suppose there are 3 slots and the agents are ranked according to a proper weakly feasible assignment σ. In case of reducible graphs $G(\sigma)$ the upper bound drops to the case of two bidders, which is calculated to 1.25. Actually, in the case of three bidders this is the only case but such a proposition will be applied for n agents. In case of irreducible graphs, there are two subcases, evident in the following image.

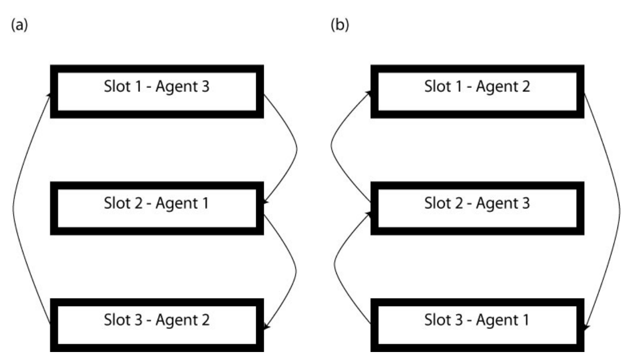

Figure 3.5: Two possible permutations

Without loss of generality, I restrict the analysis to the assignment $\sigma(1) = 3$, $\sigma(2) = 1$, $\sigma(3) = 2$. The agents' valuations are given by $v_2 = \lambda v_1$, $v_3 = \mu v_1$ and the click-through rates are given by $a_2 = \beta a_1$, $a_3 = \gamma a_1$.

$\frac{OPT}{W(\sigma)} = \frac{1+\beta\lambda+\gamma\mu}{\mu+\beta+\gamma\lambda} = \frac{1+\lambda-\mu\lambda+\mu-\frac{\mu^2}{\lambda}+\delta\lambda+\epsilon\mu}{1+\lambda-\mu\delta+\epsilon\lambda} \leq \frac{1+\lambda-\mu\lambda+\mu-\frac{\mu^2}{\lambda}}{1+\lambda-\mu}$

In order to find the maximum, I find the derivative with respect to μ and nullify it for $\mu \in [0,1]$. It follows, $\frac{OPT}{W(\sigma)} \leq \frac{\lambda^2+\lambda+2-2\sqrt{\lambda^3+1}}{\lambda} \leq 1.259134$. $\blacksquare$

An upper bound for the efficiency of the auction with n agents follows, in the same sence with the 3 conservative bidders.

Theorem 3.2.10 *The price of anarchy over pure Nash equilibria of GSP auction games with conservative bidders is at most* $\frac{61+7\sqrt{217}}{128} \approx 1.28216$.

The GSP mechanism proves to be quite impressive in terms of stability, allocation efficiency and revenue generated, since it combines all three factors. It's dominance against the GFP and VCG mechanisms was clarified and that's the reason it is currently used.

Chapter 4

The Cascade Model

4.1 A New Model - Experimental Analysis

Sponsored search auctions involve advertisers, bidding for a position in the sponsored list. Most times advertisers in the same list are competitors. It is believed that each one influences the rest. For example, an advertiser occypying the first slot may affect the click-through rate of the advertisers beneath her. Such an interaction between advertisers is called externalities, and are distinguished in positive and negative externalities. Surely such a concept must be simulated with a new model which revisits the formula of the click-through rate.

In 2008, Microsoft distributed clicking data from Microsoft Live to boost the research of the end user's clicking behaviour. Gomes, Immorlica and Markakis studied these data and proposed a new model which empirically seems closer than the separable model. The clicking data present the users' behaviour for specific keywords. From the data it is obvious that advertisers affect the ones beneath. It can be seen that the click-through rate of some ads vary analogously to the ads above. The original paper highlights these position externalities stating that a user may be tired from the poorly related ads or be satisfied from a specific ad. Let's define the new model, which is denoted as the cascade model.

The concept remains the same, meaning that an agent's profit is the product of a private valuation v_i and the click-through rate. However, in this case, the click-through rate's formula is different, in order to assimilate the characteristics of the new model. As mentioned before, in the cascade model we assume that the end user starts searching the slots from top to bottom. Thus, a continuation probability c_a is imposed, as the probability that a user will continue searching to the ads beneath after a specific slot a. A quality measure is also used, as being used in the seperable model, which is expressed as a quality probability q_a. The click-through rate of ad a_i, occupying slot i, is $r_{a_i} = q_{a_i} \cdot \prod_{j=1}^{i-1} c_{a_j}$. $C_i = \prod_{j=1}^{i-1} c_{a_j}$ serves as a cumulative continuation probability of the ads above.

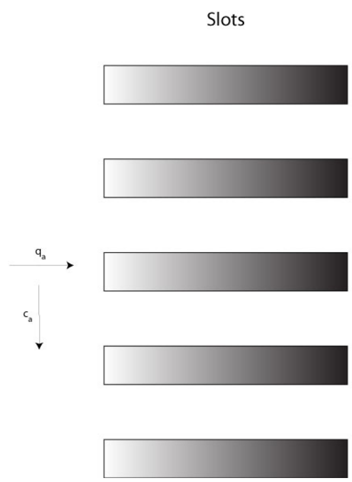

Figure 4.1: Cascade Model

4.2 Allocation

Let's study the allocation problem without considering any payments and the constraints imposed by the need of stability. The auctioneer has to choose k ads $a_1, a_2, \ldots a_k$, in order to maximize the social welfare $\sum_{i=1}^{k} [b_{a_i} \cdot q_{a_i} \cdot \prod_{j=1}^{i-1} c_{a_j}]$. At first glance, we see that an advertiser with high revenue $q_{a_i} \cdot b_{a_i}$ should occupy a high slot, but in the other hand if she has low continuation probability, by occupying a high slot, the remaining advertisers admit very low click-through rate. The following lemma makes clear how an optimal solution will look like.

Lemma 4.2.1 *Assume that the optimal solution places ad a_i in position i.*

Then, without loss of generality,

$$\frac{q_{a_1} b_{a_1}}{1 - c_{a_1}} \geq \frac{q_{a_2} b_{a_2}}{1 - c_{a_2}} \geq \dots \geq \frac{q_{a_k} b_{a_k}}{1 - c_{a_k}}$$

Proof The idea of the lemma is that if we are given the right set of k advertisers, we know their order. It can be observed that if two agents aren't in the right order the advertisers above them and beneath them are not affected, since the click-through rate is a simple product of the continuation probabilities. Thus, in the proof I will make use of this by comparing the welfare generated SW by two consecutive advertisers in the right order and in the opposite order SW'. The welfare of the rest advertisers remain the same.

- $SW = b_{a_i} q_{a_i} \prod_{j=1}^{i-1} c_{a_j} + b_{a_{i+1}} q_{a_{i+1}} \prod_{j=1}^{i} c_{a_j} = C_i \cdot (b_{a_i} q_{a_i} + b_{a_{i+1}} q_{a_{i+1}} c_{a_i})$
- $SW' = b_{a_{i+1}} q_{a_{i+1}} \prod_{j=1}^{i-1} c_{a_j} + b_{a_i} q_{a_i} c_{a_{i+1}} \prod_{j=1}^{i-1} c_{a_j} = C_i \cdot (b_{a_i} q_{a_i} c_{a_{i+1}} + b_{a_{i+1}} q_{a_{i+1}})$

$$SW - SW' \geq 0$$

$$C_i \cdot [b_{a_i} q_{a_i} + b_{a_{i+1}} q_{a_{i+1}} c_{a_i} - b_{a_i} q_{a_i} c_{a_{i+1}} - b_{a_{i+1}} q_{a_{i+1}}] \geq 0$$

$$b_{a_i} q_{a_i} (1 - c_{a_{i+1}}) \geq b_{a_{i+1}} q_{a_{i+1}} (1 - c_{a_i})$$

$$\frac{b_{a_i} q_{a_i}}{1 - c_{a_i}} \geq \frac{b_{a_{i+1}} q_{a_{i+1}}}{1 - c_{a_{i+1}}}$$

This completes the proof. ∎

Now that we know the right order when k advertisers are chosen, we restrict our attention in deciding the right set. A natural thought is to sort the agents according to the order of the lemma and start computing the optimal solution by adding every time a new agent i, for $i = n, \dots, 1$. Each time an optimal solution must be computed, we have the choice to add or not to add the new agent in our solution. If the agent is not added the optimal solution remains the same as it was without the new agent. If the agent is added, she will occupy the first slot. Thus, if we have n agents and k slots a recursive equation can be extracted, since the optimal solution is either the solution for n-1 agents and k slots, either the solution for agents $2, \dots, n$ and k-1 slots and the first slot occupied by the first agent. Analytically for ads $a, \ldots, n$ and slots $i, \ldots, k$ the recursive equation states:

$$A[a, i] = max(A[a+1, i], b_a q_a + c_a A[a+1, i+1])$$

An algorithm via dynamic programming, using the above equation, extracts the optimal solution. The algorithm sorts n agents in time $O(nlogn)$ and computes the matrix A in time $O(nk)$.

Theorem 4.2.2 *There is an algorithm with a running time of $O(nlogn+nk)$ which computes the optimal placement of n ads in k slots in the simple Cascade Model.*

4.3 Pricing

4.3.1 VCG Mechanism

At first glance, one would propose the Vickrey-Clarke-Groves mechanism to manage truthfull bidding. The VCG mechanism would accompany the above algorithm. Since the bidders report their valuations truthfully and the algorithm manages to extract the optimal solution, the mechanism admits maximum efficiency. The VCG mechanism sets standard payments which resemble the negative exteranality that the agent imposes to the rest. Clearly, the mechanism concerns the total payments. Formally, the agent occupying slot i pays:

$$C_i \cdot p_i = OPT_{-\sigma(i)} - (OPT - v_{\sigma(i)})$$

$$p_i = \frac{OPT_{-\sigma(i)} - (OPT - v_{\sigma(i)})}{C_i}$$

$$p_i = \frac{OPT_{-\sigma(i)} - (OPT - v_{\sigma(i)})}{(\prod_{j=1}^{i-1} c_{\sigma(j)}) \cdot q_{\sigma(i)}}$$

,where OPT_{-j} the social welfare of the optimal solution without agent j.

4.3.2 GSP Mechanism

The standard GSP mechanism, currently used, ranks the advertisers in the decreasing order of $q_i b_i$ and the advertiser occupying slot j pays $p_j = b_{\sigma(j+1)} \frac{q_{\sigma(j+1)}}{q_{\sigma(j)}}$. In contrast to the seperable model, the click-through rate's formula has changed and the mechanism must be analyzed again. Now, the advertiser's utility becomes $u_{\sigma(i)} = \prod_{j=1}^{i-1} c_{\sigma(j)} \cdot q_{\sigma(i)} \cdot (v_{\sigma(j)} - b_{\sigma(j+1)} \frac{q_{\sigma(j+1)}}{q_{\sigma(j)}})$.

The GSP mechanism is actually stable, since a Nash equilibrium always exists. Bidders are indexed in the decreasing order of $q_i v_i$. The following bids admit a Nash equilibrium.

$$b_s q_s = \begin{cases} v_1 q_1 & \text{for } s = 1 \\ \sum_{j=s-1}^{k+1} (\prod_{i=s}^{j} c_i) v_j q_j (1 - c_{j+1}) & \text{for } 1 > s \geq k \\ v_s q_s & \text{for } k > s \end{cases}$$

Since the stability of the GSP mechanism is guaranteed, it is important to test it in terms of efficiency. When testing the mechanism generally without any restrictions, the outcome is pessimistic, as in the basic model. Let's consider 1 slot and two players such that $q_1 v_1 = 0$, $q_2 v_2 = X$ and $b_1 > X$, $b_2 = 0$. In the optimal allocation player 2 occupies the slot generating social welfare $SW = a_1 X$. However, our case is a Nash equilibrium with social welfare $SW' = 0$. It is clear that the the price of anarchy is infinite.

Again the mechanism will be examined by resticting the analysis in the case of conservative bidders, $b_i < v_i$. The following theorem reveals the efficiency of the GSP mechanism.

Theorem 4.3.1 *The price of anarchy of GSP equilibria both against VCG and the best GSP equilibrium is k (the number of slots) in case of conservative bidders.*

Proof In order to prove the above theorem, I should prove an upper bound that the price of anarchy cannot overcome and then mention a case where the price of anarchy really touches this bound, so it is also a lower bound.

Upper Bound: We consider the player with the maximum revenue and analyze the case that she occupies the first slot and the case that she is sidelined. Let's assume that $q_x v_x = \max_i q_i v_i$. In case that she occupies the first slot, the efficiency is clearly $\frac{1}{k}$. In case that player x is sidelined, player y occupies the first slot and player x awarded slot j. Since we have a Nash equilibrium $C_j q_x (v_x - p_j) \geq q_x (v_x - b_y \frac{q_y}{q_x}$.

$$q_x v_x - q_y v_y \leq q_x (v_x - b_y \frac{q_y}{q_x} \leq C_j q_x (v_x - p_j) \leq C_j q_x v_x$$

$$q_x v_x - q_y v_y \leq C_j q_x v_x$$
$$q_x v_x \leq C_j q_x v_x + q_y v_y$$

The social welfare of the mechanism in such an equilibrium is greater than $C_j q_x v_x + q_y v_y$. So, an upper bound of $\frac{1}{k}$ is guaranteed.

Lower Bound: A specific instance of the game will be presented, and specific equilibria will be presented to justify the lower bound.

Table 4.1: Example

Player	1	2	3	...	k	k+1
V	X	$X-\delta$	$X-1-\delta$	...	$X-1-\delta$	$X-1-\delta$
c	0	$\frac{1}{1+\delta}$	1	...	1	1
q	1	1	1	...	1	1

At the table above I consider small positive constants ϵ, δ and large enough X.

Optimal Allocation: It is easy to check that the optimal allocation is $[k+1, k, \ldots, 4, 3, 1]$. The optimal allocation generates a social welfare $SW_{OPT} = kX - (k-1)(1+\delta)$.

Nash Equilibrium (NE1): All players bid their values except player 1 who bids $X - 1 - \delta$. The resulting allocation is $[2, 3, \ldots, k, 1]$ with a social welfare $SW_{NE1} = X - \delta + \frac{1}{1+\delta}((k-1)X - (k-2)(1+\delta)) \geq \frac{1}{1+\delta}kX - \epsilon X$.

Nash Equilibrium (NE2): All players bid their values except player 2 who bids $X-1$. The resulting allocation is $[1, 2, \ldots, k-1, k]$ with a socail welfare $SW_{NE2} = X$.

Both against the VCG and the GSP equilibria, there is an instance with a price of anarchy

$$\frac{\frac{1}{1+\delta}kX - \epsilon X}{X} \geq \frac{1}{1+\delta}k - \epsilon$$

∎

A similar theorem is given for the price of stability of the mechanism, which also reveals some thoughts for the efficiency of the mechanism. The theorem will be given without proof.

Theorem 4.3.2 *The price of stability of GSP equilibria against the VCG mechanism is k, and between $\frac{k}{2}$ and k in case of conservative bidders.*

The GSP mechanism's behaviour, in terms of efficiency, are reviewed in the following table:

Table 4.2: Efficiency

	PoA	PoS
Conservative bidders	k	k/2
Non-Conservative bidders	∞	[k/2,k]

Chapter 5

Modeling Externalities

While the new cascade model was presented, simultaneously research was conducted to better understand the end user's behavior towards a sponsored list. The cascade model explains some experimental observations which were made, but the situation seems much more complicated. By studying the user's data mentioned before, the cascade model is unable to go with the experimental analysis. Jeziorski and Segal made a more complete study over Microsoft's data, made significant observations and finally proposed a model for the end user's behavior. This model is much more general than the restrictive seperable and cascade models.

The first interesting observations made were the opposition of the seperable click-through rate model. The click-through rates of the advertisements of certain search strings were studied for a period of time. It is evident that the click-through rate of a certain domain in a certain position differs analogously to the advertisements in the other slots. For example, if a domain had a very competitive advertisement above it's click-through rate was low. The differences were statistically significant to extract safe conclusions for the presence of externalities. The externalities among advertisers seem to be negative due to the user's satiation. However, it is believed that there are also positive externalities but are not so evident from the experimental analysis due to the prevelance of the negative externalities. A user's satisfaction from a certain ad may raise her expectation about the relevance of the other ads.

From the data there is evidence that the cascade model is also unable to describe the end user's behavior. Namely,

- 46 % of users who click do not click in the sequential order of positions

- 57 % of users who click more than once do not "cascade", i.e., click on a higher position after clicking on a lower position.

5.1 Model

According to the seperable model an advertiser's click-through rate is the product of her relevance and a slot-specific multiplier. However, there is the need for a general model, which also accounts for the externalities between winning advertisers. The characteristics of the new model will be instilled in the click-through rate. The position-dependent multiplier will remain and the externalities will be instilled in the advertiser's relevance. The externalities between advertisers can be modeled with two Graphs G^+ and G^-. The allocation is given by a winning set S and a permutation function $\pi : [n] \Rightarrow [k]$. So a players click-through rate is given by the product $CTR_i(S, \pi) = \lambda_i \cdot Q_i(S, \pi)$.

So a player's utility is given by $u_i(S, \pi) = \lambda_{\pi(i)} \cdot Q_i(S, \pi) \cdot (v_i - p_i)$ and the resulting social welfare is $SW(S, \pi) = \sum_{i \in S} \lambda_{\pi(i)} \cdot Q_i(S, \pi) \cdot v_i = \sum_{i \in S} u_i(S, \pi) + \sum_{i \in S} \lambda_{\pi(i)} \cdot Q_i(S, \pi) \cdot p_i$. It remains to define the player's relevance according to the winning set and the permutation function π.

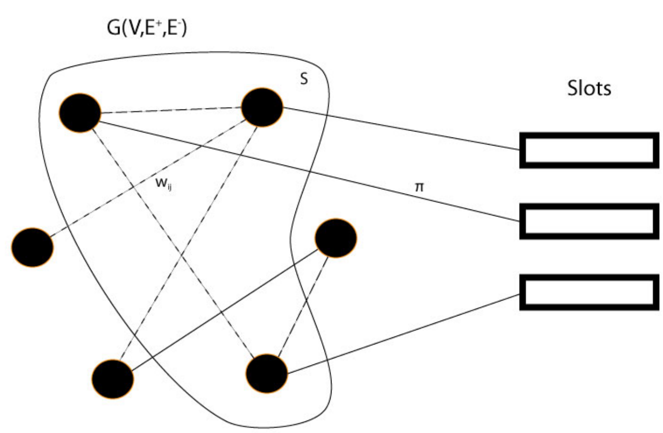

Figure 5.1: Winner Determination

A player has a probability q_i to be clicked independently of the other agents. However, his actual probability of being clicked also depends on the other agents and their relative position. Positive externalities can be expressed by the probability that a user is clicked, independently of others or after clicking an advertiser who influences positively. So, agent's i probability to be clicked is $Q_i^+(S,\pi) = 1-(1-q_i)\cdot \prod_{j \in N_i^+(S)} ((1-q_j)+q_j(1-w_{ji}(d_\pi(j,i)))) = 1-(1-q_i)\cdot \prod_{j \in N_i^+(S)} (1-q_j w_{ji}(d_\pi(j,i)))$. Additionally negative externalities are expressed by the probability that the user's attention is not distracted by other advertisers. So, $Q_i^-(S,\pi) = \prod_{j \in N_i^-(S)} ((1-q_j)+q_j(1-w_{ji}(d_\pi(j,i)))) = \prod_{j \in N_i^-(S)} (1-q_j w_{ji}(d_\pi(j,i)))$. The actual relevance is the product of the two probabilities: $Q_i(S,\pi) = Q_i^+(S,\pi)\cdot Q_i^-(S,\pi)$.

The model proposed is a general model, which includes many different cases of models. Facing this model, it is a priority to forget the game-theoretic aspects and verge it as an allocation problem. Specifically it would be interesting to determine the winning set S and permuation function π, which produces the optimal allocation. This is the Winner Determination problem. In the general case we denote the problem as MSW-E (Maximum Social Welfare with Externalities). Due to its generality, at first sight, the problem seems hard to compute. So it would be interesting to present specific subcases of the general model to analyze their behavior.

From the analysis of Microsoft's data, it seems that the users when browsing through the ads are affected towards an ad only by advertisers nearby it. So a window c can be proposed which includes a subset of successive ads which affect a specific ad. A new problem is proposed, similar to the initial one, where $w_{ji}(l) = 0$,if $|l| > c$. This is the MSW-E(c) problem.

Simpler instances of the above problems are the problem with positive externalities MSW-PE(c), meaning $E^- = \emptyset$ and the problem of forward-only positive externalities MSW-FPE(c), meaning $w_{ji}(l) = 0$, if $l < 0$. The forward-only externalities are inspired from the cascade model, where an ad is influenced by ads allocated above it.

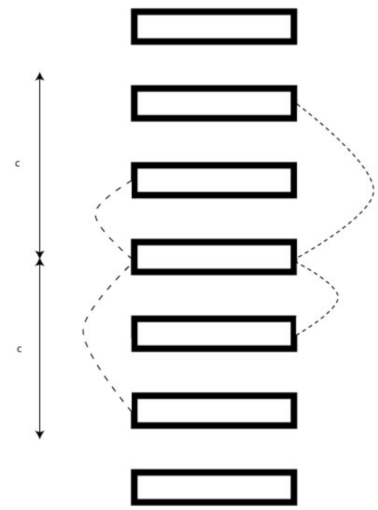

Figure 5.2: Window

5.2 Winner Determination

5.2.1 Computational Hardness

Theorem 5.2.1 *MSW-FPE(1) is NP-hard even in the special case of uniform position multipliers, valuation, and qualities.*

Proof A reduction to the Longest Path problem is given. The Longest path problem involves a directed graph G(V,E) and an integer $k \geq 2$ and a decision problem of giving a path with length k. Considering the Longest path problem, we construct an instance of the MSW-FPE(1) with valuations 1, relevance $\frac{1}{2}$, k slots and externalities graph G(V,E,∅) with $w_{ji}(1) = \frac{1}{2}$. There are two cases: a) An ad i is not influenced by the slot above it and generates welfare $q_i \cdot v_i = \frac{1}{2} \cdot 1 = \frac{1}{2}$ or b) an ad is influenced by slot j above it and generates welfare $1-(1-q_i)((1-q_j)+q_j(1-w_{ji}(1))) = 1-\frac{1}{2}(\frac{1}{2}+\frac{1}{2}\cdot\frac{1}{2}) = \frac{5}{8}$. An upper bound for the optimal allocation is k successive ads correlated one by one with welfare $\frac{1}{2}+(k-1)\frac{5}{8}$. If there existed a path with k nodes, there also exists an allocation with such a welfare. So if we knew the optimal allocation, we could also decide for the longest path problem. This completes the proof since the longest path problem is NP-complete.

Theorem 5.2.2 *MSW-FPE(1) is APX-hard even in the special case of uniform position multipliers, valuation, and qualities.*

To sketch the proof a reduction to the Traveling Salesperson problem with distances 1 and 2 is given. TSP(1,2) involves an undirected graph G(V,E) with edges of distance 1. All non-edges can be assumed to have distance 2. The goal is to find a path, which includes all vertices, with minimum total length. Given the graph G(V,E) we construct an instance of the MSW-FPE(1) problem. It can be proven that given a $(1+\epsilon)$-approximation of the MSW-FPE(1), a $(1+O(\epsilon))$-approximation can be extracted for TSP(1,2). Since TSP(1,2) is APX-hard, MSW-FPE(1) is also APX-hard.

5.2.2 Algorithms

The problem as stated is NP-hard to solve. However, by using the window there is not the need to choose from all possible permutations. With an algorithm based in color coding the number of permuations can be reduced.

Theorem 5.2.3 *MSW-E(c) can be solved optimally in* $2^{O(k)}n^{2c+1}log^2n$.

Trivially, the winner determination problem could be solved by enumerating all possible permutations and choosing the one which generates the maximum welfare. All possible permutations are $\binom{n}{k}$. So $O(n^k)$ time is needed. However the proposed algorithm clearly lowers the complexity and for $k = O(poly(logn))$ achieves polynomial complexity.

However, the computational time is exponential. In order to compute in polynomial time an approximation algorithm should be proposed. I would use an approach to the problem which makes a reduction to the 3-set packing problem. Before stating the theorem and proving it, a few words for the 3-set packing problem.

The 3-set packing problem involves sets with 3 elements at most and positive weights. The goal is to find the collection of disjoint sets with the maximum total weight. The problem is NP-hard but some approximation algorithms are known. The greedy algorithm achieves an approximation ratio of m, and there are two algorithms which give an approximation ratio of $\frac{2}{3}m$ and $\frac{m+1}{2}$ in quadratic and polynomial time respectively.

Theorem 5.2.4 *Given an* α*-approximation* $T(\nu)$*-time algorithm for Weighted 3-Set Packing with* ν *sets, we obtain a* $2\alpha c$*-approximation* $T(kn^2)$*-time algorithm for MSW-PE(c) with n ads and k slots.*

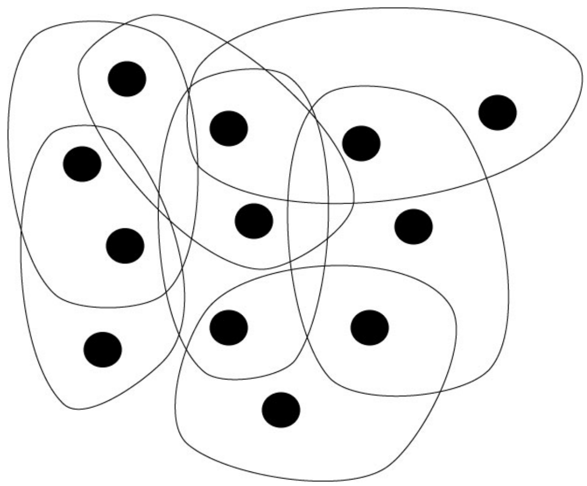

Figure 5.3: 3-set Packing Problem

Proof Initially, it will be proven that the welfare generated by the optimal allocation can be achieved by 2c instances of the weighted 3-set packing problem. As a consequence, the optimal welfare of the 3-set packing problem is a 2c-approximation of the optimal allocation's welfare. A description follows.

k slots are partitioned into $\frac{k}{2}$ pairs. Every pair is labeled, $p = 1, \ldots, \frac{k}{2}$. In every instance of the allocation problem, a pair is occupied by 2 ads i_1 and i_2. We construct all possible sets $\{i_1, i_2, p\}$ with welfare $W(i_1, i_2, p) = \max\{\lambda_p Q_{i_1}(i_1, i_2)v_{i_1} + \lambda_{p+1}Q_{i_2}(i_1, i_2)v_{i_2}, \lambda_p Q_{i_2}(i_2, i_1)v_{i_2} + \lambda_{p+1}Q_{i_1}(i_2, i_1)v_{i_1}\}$. We assume that an ad's relevance is only influenced by his partner in a pair, so $Q_{i_1}(i_1, i_2) = 1-(1-q_{i_1})(1-q_{i_2}w_{i_2 i_1}(1))$ and $Q_{i_2}(i_1, i_2) = 1-(1-q_{i_2})(1-q_{i_1}w_{i_1 i_2}(-1))$.

Lemma 5.2.5 *Let c by any positive integer. Given a list* $(1, \ldots, k)$*, there is a collection of 2c feasible 3-set packings such that for each pair* i_1, i_2 *of ads in* $(1, \ldots, k)$ *with* $|i_1 - i_2| \leq c$*, the union of these packings contains a set* $\{i_1, i_2, p\}$ *with* $p \leq \min\{i_1, i_2\}$.

Let's consider $W^{(j)}$ the welfare generated by the j-th feasible 3-set packing. Due to the lemma

$$\sum_{j=1}^{2c} W^{(j)} \geq \sum_{i=1}^{k} \lambda_i v_i \sum_{j=1}^{2c} Q_i^{(j)} \geq \sum_{j=1}^{k} \lambda_i Q_i^* v_i = W^*$$

In order to justify the second inequality let's focus on an ad i with $c < i < k-c$. It's relevance in the optimal allocation is $Q_i^* = 1-(1-q_i)\prod_{j=i-c, j\neq i}^{i+c}(1-$

$q_j w_{ji}(j-i))$. Simirarly it's relevance in the j-th set packing is $Q_i^{(j)} = 1 - (1-q_i)(1-q_j w_{ji}(-1))$. By repeatedly applying that for every $x, y, z \in [0,1]$, $(1-xy)+(1-xz) \geq 1-xyz$ we conclude that

$$\sum_{j=i-c, j\neq i}^{i+c} (1-(1-q_i)(1-q_j w_{ji}(j-i))) \geq 1-(1-q_i) \prod_{j=i-c, j\neq i}^{i+c} (1-q_j w_{ji}(j-i)) = Q_i^*$$

Given the optimal weighted 3-set packing, which is a 2c-approximation of the optimal allocation, the corresponding allocation also serves as a 2c-approximation, since we assume only positive externalities. So the additional externalities between corelated ads, which are in different pairs, can only increase the total welfare of the approximate solution.

Since the optimal weighted 3-set packing is a 2c-approximation of the optimal allocation, a α-approximation of the optimal weighted 3-set packing is a $2\alpha c$-approximation of the optimal allocation.

A more general theorem can be stated by generalizing to the m-set packing problem.

Theorem 5.2.6 *An $f(m)$-approximation $T(\nu, m)$-time algorithm for Weighted m-Set Packing with ν sets yields a $2f(2c+1)$-approximation $O(ckn^{2c} + T(kn^{2c}, 2c+1))$- time algorithm for MSW-PE(c) with n ads and k slots.*

5.3 Mechanisms

In order to achieve strategy-proofness the VCG mechanism can be applied when the allocation algorithm is monotone, since we consider a single-parameter setting. The exact algorithm for the winner determination problem, mentioned above, satisfies the monotonicity issue and can be paired with VCG payments. The VCG mechanism is well studied and very well known. In this section an analysis will be given for the GSP mechanism, since it is also the mechanism which is currently used. It will be examined in terms of stability, under the proposed model, and it's efficiency for some special cases.

5.3.1 GSP mechanism

The common use of the GSP mechanism is accompanied with the Rank by Revenue rule. Given a vector $\mathbf{b}$ the slots are allocated to advertisers according to $\phi_{\mathbf{b}} : \mathcal{K} \to \mathcal{N}$. Slot i is allocated to advertiser $\phi_{\mathbf{b}}(i)$. The payment rule imposes to advertiser i a price $\frac{(q_{\phi_{\mathbf{b}}(i+1)} \cdot b_{\phi_{\mathbf{b}}(i+1)})}{q_{\phi_{\mathbf{b}}(i)}}$. Thus, an advertiser $\phi_{\mathbf{b}}(i)$ has utility $u_{\phi_{\mathbf{b}}(i)}(\mathbf{b}) = \lambda_i \cdot Q_{\phi_{\mathbf{b}}(i)} \cdot (v_{\phi_{\mathbf{b}}(i)} - \frac{q_{\phi_{\mathbf{b}}(i+1)}}{q_{\phi_{\mathbf{b}}(i)}} \cdot b_{\phi_{\mathbf{b}}(i+1)})$. The GSP mechanism is tested in terms of stability and efficiency, under the assumption of conservative bidders. The proofs of the following theorems can be found in the original paper. The player's private valuation, the externalities and the position multipliers are carefully chosen in order to admit situations where there cannot be a stable allocation.

Theorem 5.3.1 *The strategic game induced by the Generalized Second Price Auction mechanism under the RBR rule and deterministic tie-braking does not generally have pure Nash equilibria in presence of forward positive externalities, even for 3 conservative agents and 2 slots.*

Theorem 5.3.2 *The strategic game induced by the Generalized Second Price Auction under the RBR rule and deterministic tie-braking does not generally have pure Nash equilibria in presence of bidirectional positive externalities, even for 3 conservative agents and 3 slots.*

Theorem 5.3.3 *There is an infinite family of instances of the strategic game induced by the Generalized Second Price Auction mechanism with unbounded Price of Stability, even with conservative bidders.*

In contrast to the optimistic results in terms of stability and efficiency, under the seperable model and the cascade model, when externalities are taken into consideration the GSP mechanism seems unable to guarantee a similar result.

5.4 Bidding Language Extension

In the present chapter, after it was experimentally shown that externalities between advertisers surely exist, we are trying to construct a model which includes such externalities. In the beginning, an advertiser's profit is defined as the product of her value per click times all clicks. The number of clicks were represented by the click-through rate. The main efforts to express externalities between advertisers were made by changing the click-through rate model, as shown. A user's attention can be distracted from a specific ad and be influenced to click or not to click other ads. Truly, it was experimentally shown that the total clicks of an ad are affected by the other allocated ads. This kind of externalities are called quantity externalities.

However, the externalities between advertisers are not present only in the clicking process. After a user click in several advertisement she chooses from them. If two ads are competitors, there is a small possibility for the user to buy both. So an agent's value per click may vary, depending on the allocated ads. This can be better undestood as the value per click can be seen as $v_{click} = v_{conv} \cdot Pr(conv|click)$. The probability that a user's click will be converted to a sale depends on the other clicked ads. This means that a single-dimensional representation of a user's value per click totally ignores any externalities. This kind of externalities are called value externalities.

Since the bids of these auctions express this value, in order to express such externalities it seems inevitable to extend the bidding language. The ideal model would require from every agent to report a value for every possible allocation of the agents to the slots. However, this would be extremely challenging from the user's perspective and it would dramatically increase the computational complexity of the allocation problem. There is a need for a middle solution. A bidding language which would not increase the bid's dimension a lot, but lets the agent to express her preferences. The research in such direction is a poor since a change in the bidding language would be difficult to be assimilated by the existing system.

Although, to have a better look in such direction I would present simple models which expand the current one. The models attempt to change the bid dimension as less as possible using clever ways. An analysis will propose possible mechanisms and test their stability, efficiency and revenue. At all times mechanisms which extend the current GSP mechanism would be prefered since they can extend the existing system. Subsequently, i will mainly describe such models and give a small intuition on how an analysis can be done.

5.4.1 Unit-bidder Constraints

Having in mind the prevelance of negative externalities it would be a great idea to let agents report their undesirable allocations. So, an advertiser can report her preferences with a two-dimensional vector for the desirable and undesirable allocations. So, an agent's bid is composed from a two-dimensional vector and a set of constraints. For example, an agent 1 reports that she doesn't want agent 2 above her. Formally, unit-bidder constraints can be represented by a set of triples (pos_i, j, pos_j). This means that bidder i, if allocated at slot pos_i, doesn't want bidder j at position pos_j. So a user can restrict possible allocations targeting on specific bidders. This model cannot fully define desirable and undesirable allocations but is quite close and extremely intuitive.

Such model also has the property to admit identity-specific and slot-specific constraints. Namely, a user can restrict a specific bidder being above her or restrict herself being below a specific slot. A user can also report that she want a specific bidder to be excluded from the sponsored list. This is quite intuitive if such a bidder is a major competitor. All these special constraints are actually subcategories of the general unit-bidder constraints.

A natural mechanism used for such a model is the expressive GSP (eGSP). The expressive GSP is an extension of the existing system and allocates slots to bidders greedily choosing every time from a set of bidders, in such a way that a constraint of the already allocated bidder is not violated. The allocated bidder is then charged the minimum bid she should have to gain the same place.

The approximation ratio of eGSP can be bounded but actually the stability analysis is worth mentioning. Bidders are called to report their constraints, meaning that they can falsely report them. This makes our job quite tricky because truthfulness cannot be achieved with standard ways. It can be shown that really this is extremely difficult, if not impossible, to achieve. So the analysis is done in the basis that bidders truly report their constraints but can falsely report their values. So we try to achieve semi-truthfulness. Even in this simplified version the task is challenging and only for special cases, such as for exclusion constraints, it can be achieved. See [6] for clearance.

5.4.2 Exclusive and non-Exclusive Display

Another simplified idea is that an advertiser may wish to be allocated a slot without any others nearby. This is called exclusive display. In such way she gets rid of negative externalities from competitors. For specific values of externalities and a user's values, the social welfare of the mechanism may be higher for exclusive display. Consider n agents having valuation ϵ if allocated together and an agent having valuation 1 if exclusively allocated. For small ϵ the exclusive display is optimal. Having this in mind a reasonal mechanism takes input from agents a two-dimensional bid (b, b') for exclusive and non-exclusive display.

So bidders report their values and a single bidder is allocated or multiple bidders. Two reasonable mechanism can be proposed. GSP_{2D} is a generalization of the GSP mechanism which charges the same if multiple bidders are displayed. However, if a bidder has higher valuation being exclusively displayed, so is done. She pays the social welfare achieved if multiple bidders were allocated. NP_{2D} is also a generalization of the GSP mechanism but is based on the next-price rule. The mechanism computes the optimal allocation and then charges every agent the minimum bid to retain her position. In such a two-dimensional setting, which is the next price. If an agent is exclusively displayed, she is charged the second maximum b or the minimum bid to retain the exclusive display. If an agent is displayed in a list, she is charged the maximum of the following advertiser's bid or the minimum bid to retain the non-exclusive display.

See [12] for stability and efficiency analysis and the revenue generated. Both mechanisms have equilibria and they are compared in terms of efficiency and revenue generated. They are also compared with VCG mechanism which achieves truthfulness. No mechanism is prefered since in some cases each one is better.

Bibliography

[1] Iab internet advertising revenue report. Technical report, 2011.

[2] G. Aggarwal, J. Feldman, S. Muthukrishnan, and M. Pál. Sponsored search auctions with markovian users. In *C. H. Papadimitriou and S. Zhang, editors, WINE 2008, volume 5385 of LNCS, pages 621-628.* Springer, Heidelberg, 2008.

[3] G. Aggarwal, A. Goel, and R. Motwani. Truthful auctions for pricing search keywords. In *Proc. of 7th ACM Conference on Electronic Commerce (EC)*, pages 1–7. ACM, 2006.

[4] A. Archer and E. Tardos. Truthful mechanisms for one-parameter agents. In *Proc. of the IEEE Symposium on Foundations of Computer Science (FOCS), pages 482-491*. IEEE, 2001.

[5] I. Caragiannis, C. Kaklamanis, P. Kanellopoulos, and M. Kyropoulou. On the efficiency of equilibria in generalized second price auctions. In *Proc. of the 12th ACM Conference on Electronic Commerce (EC).* ACM, 2011.

[6] F. Constantin, M. Rao, C.-C. Huang, and D. C. Parkes. On expressing value externalities in position auction. In *Proc. of the National Conference in Artificial Intelligence (AIII)*, 2011.

[7] A. Das, I. Giotis, A. Karlin, and C. Mathieu. On the effects of competing advertisements in keyword auctions. 2008.

[8] B. Edelman and M. Ostrovsky. Strategic bidder behavior in sponsored search auctions. *Decision support systems*, 43:192–198, 2007.

[9] B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second price auction: Selling billions of dollars worth in keywords. *American Economic Review*, 97(1):242–259, 2007.

[10] D. Fotakis, P. Krysta, and O. Telelis. Externalities among advertisers in sponsored search. In *Symposium on Algorithmic Game Theory*, 2011.

[11] A. Ghosh and M. Mahdian. Externalities in online advertising. In *Proc. of the World Wide Web Conference (WWW)*, 2008.

[12] A. Ghosh and A. Sayedi. Expressive auctions for externalities in online advertising. In *Proc. of the World Wide Web Conference (WWW)*, 2010.

[13] I. Giotis and A. Karlin. On the equilibria and efficiency of the gsp mechanism in keyword auctions with externalities. In *C. H. Papadimitriou and S. Zhang, editors, WINE 2008, volume 5385 of LNCS, pages 629-638*. Springer, Heidelberg, 2008.

[14] R. Gomes, N. Immorlica, and E. Markakis. Externalities in keyword auctions: an empirical and theoretical assessment. In *S. Leonardi, editor, WINE 2009, LNCS, pages 172-183*. Springer, Heidelberg, 2009.

[15] P. Jeziorski and I. Segal. What makes them click: empirical analysis of consumer demand for search advertising. Technical report, 2010.

[16] D. Kempe and M. Mahdian. A cascade model for externalities in sponsored search. In *C. H. Papadimitriou and S. Zhang, editors, WINE 2008, volume 5385 of LNCS, pages 585-596*. Springer, Heidelberg, 2008.

[17] E. Koutsoupias and C. H. Papadimitriou. Worst-case equilibria. In *Annual Symposium on Theoretical Aspects of Computer Science*, pages 404–413, 1999.

[18] V. Krishna. *Auction Theory*. Elsevier Science, 2002.

[19] S. Lahaie. An analysis of alternative slot auction designs for sponsored search. In *In Proc. of 7th ACM Conference on Electronic Commerce (EC)*. ACM, 2006.

[20] S. Lahaie, D. Pennock, A. Saberi, and R. Vohra. 28: Sponsored search auctions. In *N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani, editors, Algorithmic Game Theory*. Cambridge University Press, 2007.

[21] R. P. Leme and E. Tardos. Pure and bayes-nash price of anarchy for generalized second price auction. In *Proc. of the IEEE Symposium on Foundations of Computer Science (FOCS), pages 735-744*. IEEE, 2010.

[22] B. Lucier and R. Paes Leme. Gsp auctions with correlated types. In *Proc. of the 12th ACM Conference on Electronic Commerce (EC)*, 2011.

[23] N. Nisan. 9: Introduction to mechanism design (for computer scientists). In *N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani, editors, Algorithmic Game Theory*. Cambridge University Press, 2007.

[24] T. Roughgarden and E. Tardos. Do externalities degrade gsp's efficiency. In *Ad Auctions Workshop*.

[25] H. Varian. Position auctions. *International Journal of Industrial Organization*, 25:1163–1178, 2007.

[26] W. Vickrey. Counterspeculation, auctions and competitive sealed tenders. *Journal of Finance*, pages 8–37, 1961.