LIBERA UNIVERSITÀ INTERNAZIONALE DEGLI STUDI SOCIALI

DIPARTIMENTO DI IMPRESA E MANAGEMENT

CATTEDRA: MATEMATICA

TESI DI LAUREA

Auction theory applied to online ads: mechanisms and strategies

Laureando

Andrea Longhi

Relatore

Marco Dall'Aglio

Anno accademico 2020/21

Index

Introduction

Online advertising is a digital marketing technique that uses the Internet to target promotional messages to consumers, through devices like personal computers, tablets and mobile phones, and in various forms, among which very popular are the clickable banners that link users to the advertisers' contents. There are three main actors involved: the publisher who sells ads slots in his own web page, search engine or social network platform, the advertiser who buys slots and provides the ads to be displayed, and the users who visit web pages and become potential buyers of the products being sponsored.

Ads slots were initially sold on a pay-per-impression basis, through flat contracts directly negotiated between advertisers and web companies. Later on, in parallel with the increasing demand for on line ads, auction mechanisms began to be adopted to assign ads positions, spanning from the most attractive, typically in the upper part of a page, to the cheaper and less attractive, down to the bottom part of a page. An efficient auction is designed in a way that the so called *social welfare* is maximized, meaning that the utility of buyers (advertisers) and sellers (publishers) is overall maximized through a suitable assignment of slots-to-ads.

Objective of this work is to provide an in-sight of the most popular ads auction types, analyzing the theory behind them, which has a strong derivation from on the strategic game theory.

The first chapter focuses on the assignment problem: how to find the best possible resource-to-task allocation in a way that a certain function is maximized? One of the most popular methods to solve the assignment problem is the Hungarian algorithm, which is described together with a practical example.

After giving a recall of the main concepts of auction theory and its derivation from the strategic game theory, the second chapter describes the four basic auction types in two different cases: single-item auctions and multiple-item auctions.

The spread of sponsored links, i.e. paid ads for which advertisers buy the keywords for which they want to appear in a certain web page, is one of the reason to explain the development of so many different algorithms for implementing on-line ads auctions. Chapter 3 provides an analysis of the most popular methods and strategies used by the larger internet actors, like Google and Facebook, for online ads auctions, including a description of relevant models, main performance indicators and the current effort by auction designers of modelling the effect of “externalities” and building algorithms for the prediction of the winner’s utility in complex multi-ads scenarios.

Chapter 1
The assignment problem and its solution using the Hungarian Algorithm

1.1 The assignment problem

Suppose that in a real situation there are a number of tasks to be done and a number of available resources, like for example people or machines, each of them being in principle able to perform any of these task: the *assignment problem* becomes significant when the available resources show different “performance”, also referred as *cost*, when facing different task. In this context, the term *cost* may mean time, efficiency or any other measurable parameter that is worth to be minimized and thus each specific resource-task assignment result in general in a correspondent cost.

In such scenario, the assignment problem consists in covering as many tasks as possible with the minimum achievable total cost, which is defined as the sum of all the resource-task assignment costs. In addition there are two constraints:

- each resource cannot perform more than one task
- each task cannot be performed by more than one resource.

Let’s make a practical example of an assignment problem: a beer delivery firm receives 5 simultaneous orders for delivery to five customers (e.g. restaurants) located in different areas of the same city. Being the time to deliver a key mission of the firm, the problem translates in assigning 5 riders (i.e. delivery agents) to each of the customers in a way that the total time (cost) of all deliveries is minimized.

The assignment problem is operations research problem, also known as maximum-weight bipartite matching problem and then it can be described using the graph theory.

A *graph*[1] is a structure amounting to a set of objects in which some pairs of objects have a certain relation; the objects are called vertices or nodes or points, while the related pairs of vertices are called edges or links or lines, which on turn can be directed or undirected[2].

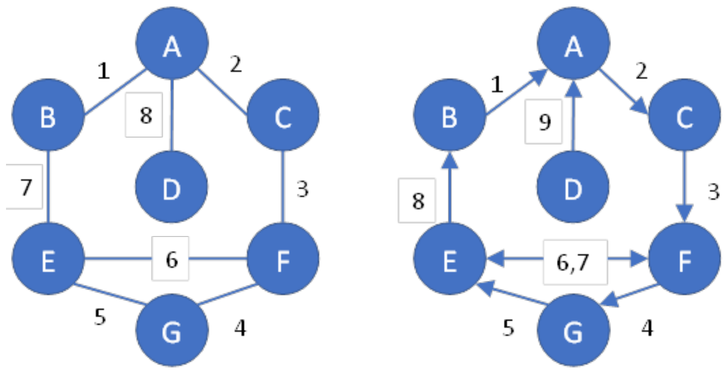

Figure 1: Example of an undirected graph with 7 vertices and 8 edges (on the left) and a directed graph with 7 vertices and 9 edges (on the right). Source: own elaboration.

A particular type of graph is the *bipartite graph*, in which all vertices are grouped in two different disjoint sets V and W, called the parts of the graph, in a way that every edge connects a vertex in V to one in W[3].

The graph in Figure 1 cannot be a bipartite graph because of the "triangle" composed by vertices EFG: when removing for example the edge number 6, a bipartite graph can be build, as depicted in Figure below.

---

[1] The word "graph" was first used in this sense by J. J. Sylvester in 1878 in a direct relation between mathematics and chemical structure.. Source Wikipedia
[2] https://en.wikipedia.org/wiki/Graph_(discrete_mathematics)
[3] A bipartite graph does not contain any odd-length cycles. Source: https://en.wikipedia.org/wiki/Bipartite_graph

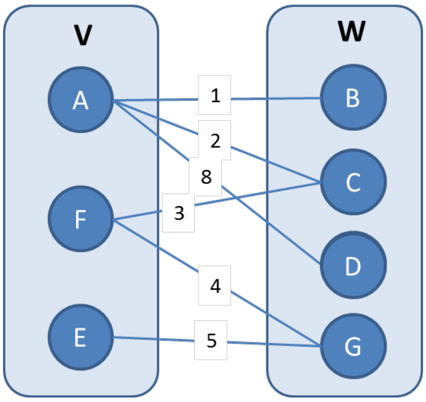

Figure 2: Example of a bipartite graph obtained from the graph in Figure 1 by removing edge number 6. Source: own elaboration.

A bipartite graph is called a *complete bipartite graph* when the parts V and W are disjoint sets and every node in V are connected with all nodes in W.

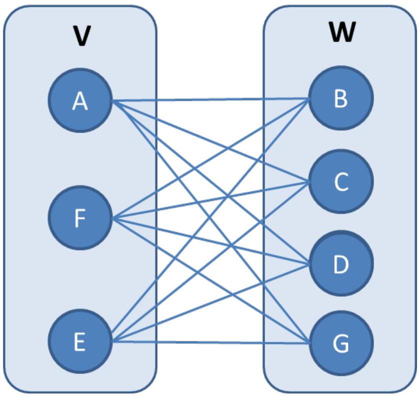

Figure 3: A complete bipartite graph. Every vertex in U are connected with all vertices in V. Source: own elaboration.

1.1.1 Formal definition of the assignment problem

Let define (V, W, Z) a complete bipartite graph[4] with:

V, W the set of nodes

$Z = VxW$ the set of edges

[4] R. Colini-Baldeschi, J. Mestre, O. Schrijvers, C. A. Wilkens, "The Ad Types Problem"

$C: Z \rightarrow R^{+}$ the edges weight functions (costs)

The goal of the assignment problem is to find a bijection $f: V \rightarrow W$ that allows to minimize to cost function:

$$\sum\nolimits_{v \in V} C(v, f(v))$$

Or, in other words, that allows to find a matching *M* of maximal total minimizes the sum of weights of the edges.

Very often in the computation of an assignment problem, the weight function is represented by square matrix C. In such cases the cost function can be re-stated as:

$$\sum\nolimits_{v \in V} C_{v,f(v)}$$

Therefore, the assignment problem can be represented in the form of a VxW cost matrix *C* of real positive elements[5], as described below.

$$\begin{array}{cc} & \textbf{Tasks} \\ \textbf{Resources} & \begin{bmatrix} C_{11} & C_{12} & \cdots & C_{1w} \\ C_{21} & \cdots & \cdots & C_{2w} \\ \cdots & \cdots & \cdots & \cdots \\ C_{v1} & C_{v2} & \cdots & C_{vw} \end{bmatrix} \end{array}$$

Figure 4: An assignment problem with v resources and w tasks to be done, represented through a VxW Cost Matrix C. Source: own elaboration.

[5] https://www.engineeringenotes.com/project-management-2/operations-research/assignment-problem-meaning-methods-and-variations-operations-research/15652

The assignment problem is often referred as linear problem when the cost function to be optimized and all the constraints are linear.

A rough-and-ready solution for the assignment problem consists in calculating the cost of each resource-task pairs and manually implementing the assignments to minimize the cost functions. Such procedure is obviously very inefficient, considering that in problems for example with *n* resources and *n* tasks, there will be factorial of *n* different assignments, therefore the computation time could become unsustainable.

Fortunately, a number of algorithms available in literature are able to solve the problem in a time that is polynomial in *n*, among them the so called Hungarian Algorithm, which will be described in section 1.3.

1.2 Balanced and unbalanced assignment

If the two parts *V* and *W* of a graph have equal cardinality, then the graph is called a *balanced* bipartite graph.

Coming back to the example of the food delivery firm and consider two alternative scenarios:

- suppose that, after having received 5 order, there are exactly 5 riders available. This is a *balanced assignment problem* and as mentioned in the previous section, its solution is the combination of riders-customers that results in the lowest total cost;
- conversely, suppose now that there are 6 riders available, but still 5 customers. With the cardinality of V higher than that of W, the assignment problem becomes unbalanced. One way to solve an *unbalanced assignment problem* is to define an additional "dummy" task, a sort of "void task" that have a cost equal to 0 when it is assigned to one of the available riders. This way, the problem is re-conducted to a balanced problem and then it can be solved using the same algorithms available for the balanced problems. Of course, similar adjustments can be done when there are more deliveries to be done than available riders (i.e. cardinality of W higher than

that of V): for example, a rider could serve two customers in a single trip in a way that two tasks can be grouped in only one (eventually considering an additional penalty for the inefficiency due to the double task), and also in this case the problem is reduced to a balanced problem.

1.2.1 Algorithms for balanced problems

There are a number of polynomial-time algorithms[6] available to solve a balanced assignment problem. Among them, one of the first methods to be introduced was the Hungarian algorithm, a *global* algorithm[7] that will be described later in section 3.1.

Differently from the global methods, *local algorithms* aim at finding local augmenting paths, instead of full augmenting paths. Even if local algorithms have in general worse asymptotic runtime than global algorithms, they are known to be more manageable and work better in many practical cases. Local algorithms are often also referred as *auction algorithms*, *push-relabel algorithms*, or *preflow-push algorithms*[8].

Some local algorithms require that a perfect matching is possible, otherwise they run to infinite loops without achieving a solution. In such cases the solution consists in adding dummy edges with very large weights that exceed the weights of all existing matchings (in order to avoid to find the dummy edges in the solution), thus extending the graph to become a *complete* bipartite graph.[9]

1.2.2 Algorithms for unbalanced problems

In unbalanced problems, the bipartite graph has in general one part with n vertices, another (smaller) part with $r < n$ vertices and a maximum cardinality s of the possible matching in the graph. When the graph admits a one-sided-

[6] An algorithm is said to be polynomial-time if its execution time is bounded at the top by a polynomial expression in the size of the input to the algorithm, i.e., $T(n) = O(n^k)$ for some constant k. Source: Wikipedia
[7] Global algorithms focus on finding the minimum or maximum over the entire given set, while local algorithm are devoted to finding local minima or maxima.
[8] https://en.wikipedia.org/wiki/Assignment_problem
[9] https://en.wikipedia.org/wiki/Assignment_problem

matching of size r, we have $s = r$ (perfect matching); in such cases, the objective is to find a matching of size s with the minimum possible cost.

Moreover, an unbalanced assignment can be adjusted to become, and therefore be managed, as a balanced assignment . A possible way to do this is to add $n - r$ new vertices to the smaller part of the graph and connect them to the larger one with zero-cost edges: this modification could be not easily manageable as it requires a number of $n * (n - r)$ edges to be added.

One more efficient alternative way to reduce an unbalanced problem represented by a graph G into a balanced one is the so called *doubling technique*. It consists in building a graph G' by taking two copies of G, a "forward copy" G_f and a "backward copy" G_b, and flipping G_b so that each side of G' has the same number of $n + r$ vertices. To achieve the balancing we need to add two kind of edges (see Figure below):

- *Large-to-large*: zero-cost edges to connect each vertex of larger part of G_f to the corresponding vertex in G_b;
- Small-to-small: very-high-cost edges to connect each vertex of the smaller part of G_f to the corresponding vertex in G_b.

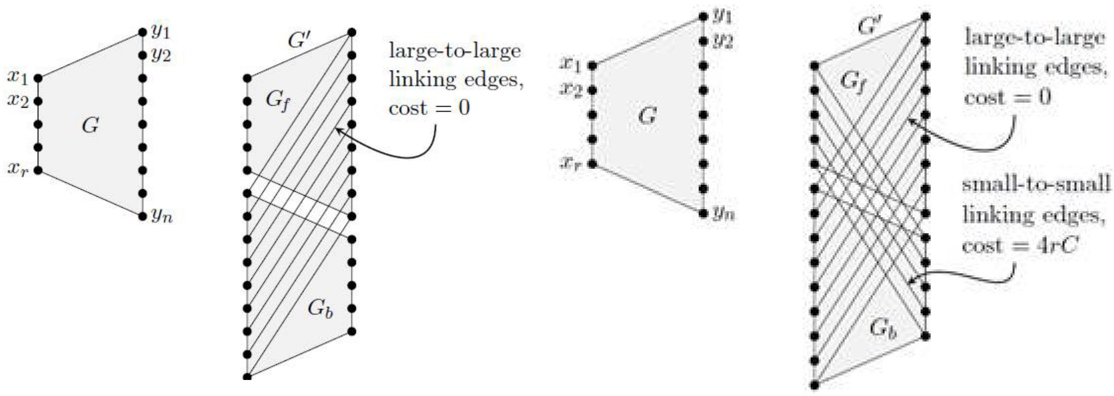

Figure 5: Graphical representation of the doubling technique. Source: L. Ramshaw, R.E. Tarjan, "On Minimum-Cost Assignments in Unbalanced Bipartite Graphs".

The resulting graph always has a perfect matching of size $n + r$ by adding at most $n + r$ new edges. The drawback of this doubling technique is that it becomes very inefficient when the two parts of the graph start to have very different cardinalities, until there is no speed gain when $n \ll r$. [10]

1.2.3 Other methods and approximation algorithms

There exist many other methods to solve the assignment problem, each of them characterized by their own complexity, which has decreased over the years. However, despite decades of research, the complexity of such algorithms has been always bounded by $O(m\sqrt{n})$ where m and n are the number of edges and vertices respectively. Duan and Pettie[11] demonstrated that such barrier for the maximum weight matching problem can be bypassed by approximation algorithms which runs in linear time for any fixed error bound.

1.3 Solving assignment problems with the Hungarian Algorithm

The Hungarian [12] algorithm is a combinatorial global optimization method developed by Harold Kuhn in 1955. It is used to find a maximum weight matching in a bipartite graph or in other words, to solve the assignment problem, in polynomial time. Indeed, the computational complexity of the algorithm was $O(n^4)$ in its first version, and decreased down to $O(n^3)$ with slight modifications.

As described in section 1.1, the assignment problem is represented by means of a bipartite graph, in which the vertices are the elements to be associated, the edges are the possible choices of pairs and each arc has an associated cost. The algorithm works improving step by step the matching along augmenting paths, alternating paths between unmatched vertices.

The algorithm starts from an initial partial solution of the problem, composed by a trivial primal solution (empty matching) and a trivial dual solution, and then it

---

[10] Lyle Ramshaw, Robert E. Tarjan, "On Minimum-Cost Assignments in Unbalanced Bipartite Graphs"
[11] R. Duan, S. Pettie, "Linear-Time Approximation for Maximum Weight Matching"
[12] It is called "Hungarian" as it is based on previous works by the Hungarian mathematicians Dénes König and Jenő Egerváry. Source: Wikipedia

iteratively increases the cardinality of the matching while improving the value of the dual solution until the value of the primal solution equals that of the dual. Therefore, at each iteration, the goal is to find a path that increases the number of elements in the solution.

1.3.1 Step by step procedure

If we consider again the graph $A = (V, W, Z)$ introduced in section 1.1, where V and W represent the sets of elements to be associated, Z is the set of edges and $C: Z \rightarrow R^{+}$ the weights of the edges, the primal and dual pair problems can be represented with[13]:

$$maximize \sum\nolimits_{(i,j)\in Z} v_{ij}\, x_{ij} \qquad\qquad minimize \sum\nolimits_{i\in V} u_i + \sum\nolimits_{j\in W} p_j$$

$$subject\ to \begin{cases} \sum\nolimits_{j} x_{ij} \leq 1 & \forall i \in V \\ \sum\nolimits_{i} x_{ij} \leq 1 & \forall j \in W \\ x_{ij} \geq 0 & \forall (i,j) \in Z \end{cases} \qquad\qquad subject\ to \begin{cases} u_i + p_j \geq v_{ij} & \forall (i,j) \in Z \\ u_i \geq 0 & \forall i \in V \\ p_j \geq 0 & \forall j \in W \end{cases}$$

The algorithm operates the following steps:

1. It starts from the following primal and dual solutions:

Primal solution (empty)	Trivial feasible dual solution
$M = \emptyset$	$\begin{cases} u_i = 0 & for\ all\ i \in V \\ p_j = max_{(i,j)\in Z} v_{i,j} & for\ all\ j \in W \end{cases}$

2. In each iteration, the algorithm identifies the set of tight edges

$$T = \{(i,j) \in E : u_i + p_j = vij\}$$

[13] R. Colini-Baldeschi, J. Mestre, O. Schrijvers, C. A. Wilken, “The Ad Types Problem”

and builds an alternating BFS[14] tree B, named Hungarian tree in (V, W, Z) out of the free vertices in W

3a. If the alternating tree B contains an augmenting path K, the cardinality of M is increased by K

3b. if no such path is available, the dual solution is updated by reducing the dual value of $W \cap B$ and increasing the dual value of $V \cap B$ by the same amount until a new edge become tight. This update maintains feasibility while reducing the value of the dual solution and makes at least one new edge tight, which in turn allows us to grow the alternating tree further.

4. at the end of the algorithm a matching is found whose weight equals the value of the dual feasible solution.

During the algorithm execution the dual solution shall always be feasible and the edges in the matching M shall always be tight.

1.3.2 Matrix interpretation

In practical cases, the Hungarian method is often used in its matrix representation. Given an $n \ x \ n$ square matrix[15] representing the cost matrix of the assignment problem, the algorithm can be expressed through the following steps (see also Figure 5):

- Step 1: find the smallest element in each row of the given cost table and then subtract that from each element of that row
- Step 2: find the smallest element in each column and then subtract that from each element of that column. Now there is at least one zero in each row and column.
- Step 3: if a complete assignment is still not possible, identify those rows and/or columns such that all zeros in the matrix are covered with as few rows and/or columns as possible. There exist different procedure to do this,

[14] BFS is the acronym for breadth-first search, a search algorithm for graphs that starting from a node called "source" allows to search for the path to another chosen node and connected to the source node.

[15] When the matrix is not a square matrix, zero-cost dummy rows or columns shall be added to reduce the problem in a balanced problem

of which one is described in Figure 6;

- Step 4: select the smallest value *m** from the cells not covered by any line identified on step 3 and:
    - Subtract *m** from cells not covered by the identified lines;
    - Add *m** to all cells in the intersection between two identified lines;
- Step 5: repeat Step 4 and 5 until a complete assignment is possible

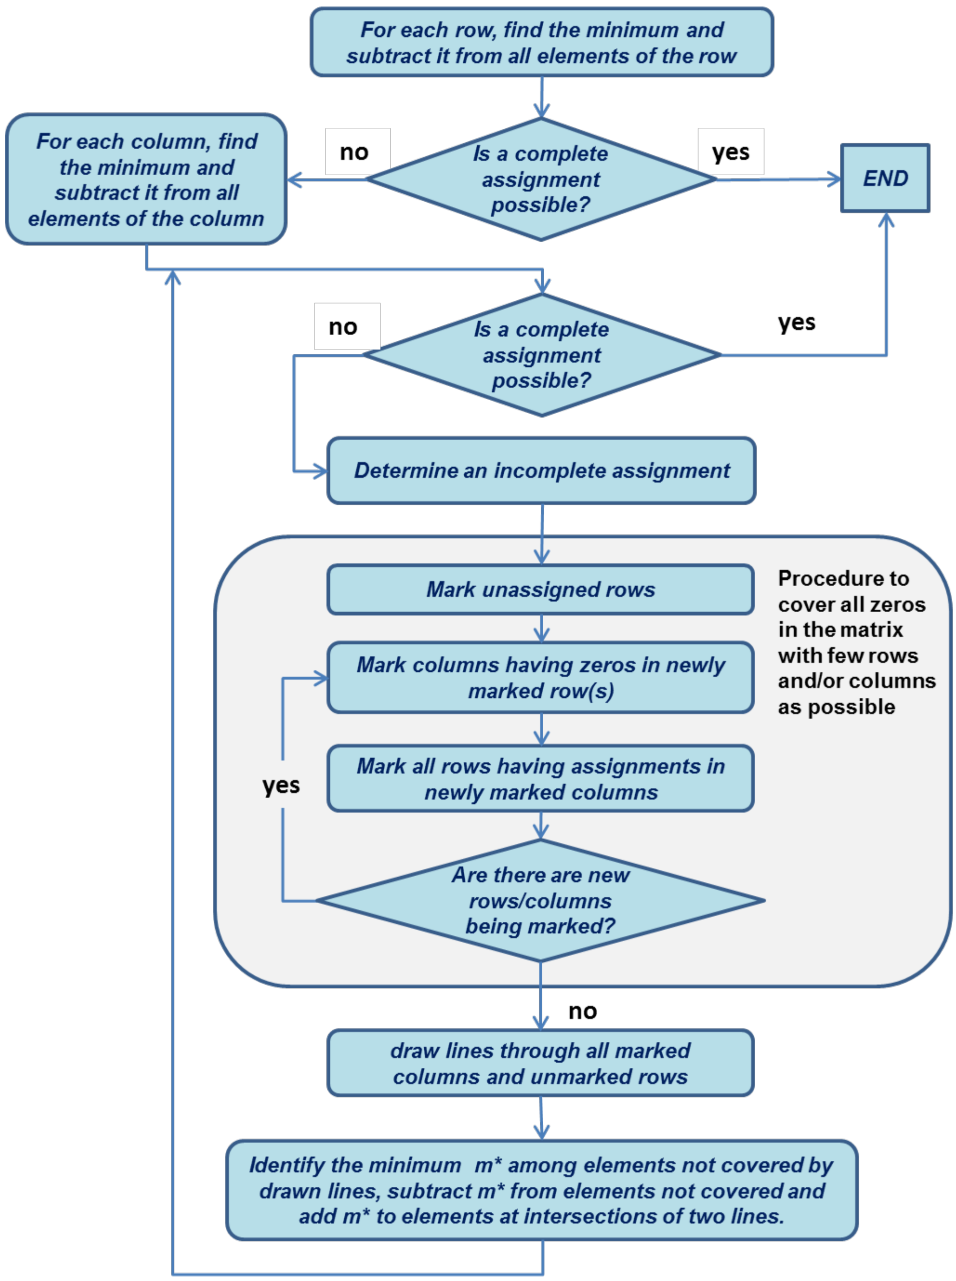

Figure 6: Matrix interpretation of the Hungarian algorithm. Source: own elaboration

In order to make an assignment, and therefore to determine if a complete assignment is possible, the following procedure can be used:

a) Scan rows until a row with exactly one unmarked zero is obtained and assign that zero by making a square around it;
b) For each zero that becomes assigned, eliminate (Strike off) all other zeros in the same row and/ or column;
c) Repeat step (a) and (b) for each column also with exactly single zero value all that has not been assigned;
d) If a row and/or column has two or more unmarked zeros and one cannot be chosen by inspection, then choose the assigned zero cell arbitrarily
e) Continue this process until all zeros in row column are either enclosed (Assigned) or struck off (x)

1.3.3 Practical example

Let's consider again the example described in section 1.1, in which a food delivery firm receive five orders for delivery to five customers located in different areas. The problem here consists in assigning riders to customers in a way that the total time spent by the riders is minimized.

		Time to delivery to restaurants (min)				
		A	B	C	D	E
Riders	Marc	22	30	28	12	10
	Peter	18	40	25	15	13
	Noemi	16	18	32	24	26
	Susan	8	17	6	17	11
	Tom	26	21	8	36	28

Figure 7: Matrix representation of the assignment problem in the examined example. Source: own elaboration

We can apply now the steps of the Hungarian Algorithm to solve the assignment problem. First, we perform step 1 and step 2, obtaining the following two matrices respectively.

12	20	18	2	0
5	27	12	2	0
0	2	16	8	10
2	11	0	11	5
18	13	0	28	20

12	18	18	0	0
4	25	12	0	0
0	0	16	6	10
2	9	0	9	5
18	11	0	26	20

Figure 8: Outcome of step 1, 2 of the algorithm. Source: own elaboration

From the second one, it is clear that a complete assignment is not possible since row 4 and row 5 have a zero value in the same column. We then proceed finding the minimum number of rows and column needed to cover all the zeros in the last matrix.

			*		
	12	18	18	0	0
	4	25	12	0	0
	0	0	16	6	10
*	2	9	0	9	5
*	18	11	0	26	20

			*		
	12	18	18	0	0
	4	25	12	0	0
	0	0	16	6	10
*	2	9	0	9	5
*	18	11	0	26	20

Figure 9: Outcome of step 3 of the algorithm. Source: own elaboration

It is found that with 3 rows (rows 1, 2 and 3) and 1 column (column 3) all the zeros are covered with a minimum number of four lines.

Now we proceed with step 4:

- finding the minimum m* = 2 among the cells which are not highlighted in the last matrix
- subtracting 2 from all the not-highlighted cells
- adding 2 to all the cells which are on the intersection between two highlighted lines, i.e. the dark blue cells in the last matrix.

The result of step 4 is represented in the following matrix, in which it is clear that a complete assignment is possible: green cells represent the assigned column to each row.

			*		
	12	18	20	0	0
	4	25	14	0	0
	0	0	18	6	10
*	0	7	0	7	3
*	16	9	0	24	18

Figure 10: Outcome of step 4 of the algorithm. Source: own elaboration

Finally, it is possible to draw the bipartite graph of the problem, representing the assignments just found and indicating, for each edge, the associated cost as per the original cost matrix.

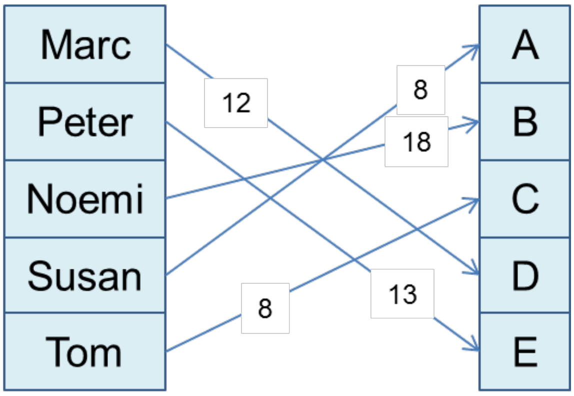

Figure 11: Bipartite graph with complete assignment that minimize the cost function after the algorithm implementation. Source: own elaboration

The sum of the five costs, each representing the cost associated to a rider-customer pair, is 59. It means that, if the deliveries are done according with the above assignment, the total delivery time is 59 minutes, corresponding to an average time of 11.8 minutes, which is the minimum among all the possible assignments.

Chapter 2
Auction Theory and applications

2.1 The auction theory and its derivation from the game theory

An auction can be defined as a bid-based buying and selling process in which items are awarded by the bidder who makes the highest offer.

To find the historical origins of the auctions[16] we shall go back in time until the fifth century BC, when Herodotus described in the first book of the *Histories*, where it is told that Babylonians used to combine marriages by means of auctions. In the following centuries the use of auctions rapidly spread, for example for trading slaves or war spoils, however they finally obtained a diffusion in the form know today from the eighteenth century, when Sotheby's and Christie's auction houses were founded.

Today an auction can be held between various suppliers for the purchase of goods or services, or to obtain the right of use of the electromagnetic spectrum, or between airlines for the hourly availability for landing at airports, or even among sellers willing to buy an advertising space on a website or on a social network platform.

By definition, an auction is a resource allocation system characterized by a set of rules governing the exchange between economic agents, where specific procedures are defined to determine the price that must be paid for awarding an asset, a service or a right[17].

The auction theory is an applied branch of *game theory*, which, in turn, studies the individual decisions of a subject in conflict situations or strategic interaction

---

[16] The term *hasta* which is the Latin translation of auction, came into use in ancient Rome: before proceeding in a public sale in which the interest of the state was involved, a thick and long stick of wood, i.e. the hasta was planted in the ground to declare the presence of the state.

[17] Source: Wikipedia

with other opponents aimed at maximizing the gain of each subject: following a feedback mechanism, the decisions of each player can influence the results achievable by the others and vice versa. Game theory aims to search for competitive or cooperative game solutions, using mathematical models.

2.1.1 Strategic games

The main purpose of the Game Theory is to study situations in which a number of players interact pursuing common, different or even conflicting objectives, or, in other words, the players' strategies. In game theory, a *game* (sometimes referred as *strategic game*) is a set of interactions among players, which plan their moves (strategies) once and for all, and such decisions are made simultaneously.

The following statements describe the fundamentals of a strategic game:

- each player plans a set of moves, which form the player's strategy;
- all players must be aware of the rules of the game, intended as the relation between strategies and corresponding outcomes;
- when all the strategies have been deployed and completed, each player receives a *pay-off*, that can be for example a remuneration (win) or a penalty (loss).

Thus, a strategic game can be described by the following four elements:

1	N players	$1, \cdots, N$
2	m_i strategies for each player i	$S_i = \{S_{i1}, \dots, S_{im_i}\}$
3	Each player has a *payoff* associated to each of the possible combinations of the player's strategies	
4	Set of game rules that link strategies to payoffs	

Strategic games are classified according to different criteria, among them:

- <u>*Cooperative or competitive games:*</u> in *cooperative games* different players share common objectives and can cooperate to gain higher profit than that they would gain by playing individually. Binding agreements ensure that winnings are distributed in proportion to the role played by each player and according to his strategy. Conversely, in *competitive games* (also referred as *non-cooperative games*) players are not allowed to enter into binding agreements;

- *Full information or incomplete information games:* In *full information games* every player knows completely the context, for example the number of opponents, their potential strategies and payoff, even if the actions of the other players are not necessarily known, as they could take place simultaneously; conversely, in incomplete information games, the players are not fully aware of the context;
- *Static and dynamic games:* in *static games*, actions implemented by a player are independent from the actions of other players; static games are also called simultaneous games. On the other hand, when players can decide to make an action based on knowledge of choices made by opponents the game is referred as dynamic game;
- *Finite and infinite games:* Finite games are characterized by a limited number of possible game situations, differently from the infinite games, those in which there are an infinite number of potential game scenarios;
- *Zero-sum and non-zero sum games*: in *zero-sum games* the sum of the pay-offs, positive or negative, derived from the selected strategies, is always zero, meaning than players are in a “total opposition”; In non-zero-sum games the sum of the pay-offs can be positive or negative.

Basically the rational behavior of each player in non-cooperative games is such as to always pursue the most advantageous strategy for himself (strategy of the maximum).

A strategy is called "dominant" if choosing such strategy a player obtains the maximum result whatever the choice of the others players. In general, there are two kinds of dominant strategies:

- *strictly dominant strategies* (or simply dominant strategies), when a player always gains the maximum utility, regardless of the other player’s strategy;
- *weakly dominant strategies* are those strategies that provides at least the same utility when combined with all the other player’s strategies, and in

some cases, i.e. for some of the other player's strategies, even strictly greater utility.

A *pure strategy* occurs when there exist a complete definition of how a player will play a game, while a mixed strategy is an assignment of a probability to each pure strategy, allowing for a player to randomly select a pure strategy[18].

If there is a dominant strategy common to all players, then the game has an equilibrium, which is famous as the *Nash equilibrium*[19]. John Nash has shown that any finite game with *n* players admits at least a point of equilibrium in mixed strategies.

A classic example of full information game is the "Prisoner dilemma"[20], in which two criminals, both accused for the same crime, are locked in two different jail cells, preventing them from communicating to each other. Each of them are given two choices, to collaborate (confessing the crime) or not collaborate (not confessing), and each combination of choices (strategies) results in a different outcome (pay-off) for the two prisoners, according to the following table.

		Prisoner B			
		Cooperate		Not cooperate	
Prisoner A	Cooperate	6	6	0	7
	Not cooperate	7	0	1	1

Figure 12: Matrix representation of the prisoners dilemma Source: own elaboration from Wikipedia

It is noted that best result for the two prisoners together is not to cooperate (both obtain 1 year in prison instead of 6), however, since they don't know what the

[18] Source: Wikipedia

[19] If each player has chosen a strategy and no player can increase their own expected payoff by changing their strategy while the other players keep theirs unchanged, then the current set of strategy choices constitutes a Nash equilibrium. John Forbes Nash Jr. (June 13, 1928 – May 23, 2015) was an American mathematician who made fundamental contributions to game theory, differential geometry, and the study of partial differential equations. Source: Wikipedia.

[20] Prisoner's Dilemma was proposed in the 1950s by Albert Tucker as a game theory problem and It was also used to explain the arms race of the 1950s by the USA and USSR (i.e. the two prisoners) during the Cold War.

other will choose, the best strategy is to collaborate, that leads to a risks of 0 or 6 years instead of a risk of 1 or 7 years. The strategy “not collaborate” is said to be strictly dominated by the strategy “collaborate”. It is eliminating the strictly dominated strategies that the Nash equilibrium is reached.

A relevant example is provided by the famous “Munich agreement”, occurred during 1938, when British Prime Minister Chamberlain needed to decide if making or not concession[21] to Hitler, giving Germany the region of Sudetenland and getting in return form Hitler a commitment not to proceed with the war. Chamberlain's preferences were obviously to avoid war, possibly without offering concessions, therefore the worst possible payoff was represented by Hitler continuing the war after having been granted the said concessions. The information was incomplete as Chamberlain did not know Hitler's preferences and could not assign any payoffs to Hitler in every possible situation.

In such cases a solution can be provided by the Harsanyi method[22], that works introducing in the game a new element that represents the expectations of the players regarding random events. In the Chamberlain case, the incomplete information about Hitler thoughts was transformed in an incorrect information by considering two potential scenarios: one based on Hitler who remains faithful to the agreements, with a certain assumed probability p^*, the other considering the opposite situation, with probability $1 - p^*$. This way it became possible to assess a dominant strategy common to the two negotiating parts, therefore an equilibrium, which is called, in the case of a Bayesian game, a Bayes-Nash equilibrium[23].

2.1.2 Auctions

If the goal of the game theory is to study the behavior of players in strategic games, the objective of the auction theory is to analize the behavior of participants to

---

[21] The concession consisted in allowing Germany to annex Czechoslovakia’s Sudetenland, after that they had just conquered Austria and it was common opinion that they were considering a similar action against Czechoslovakia

[22] Harsanyi (1967–68) proposed a method for transforming uncertainty over the strategy sets of players into uncertainty over their payoffs. The transformation appears to rely on an assumption that the players are rational, or, indeed, that they are rational and that there is common belief of rationality

[23] Unfortunately, the history tells that notwithstanding the concessions were granted, Germany entered Prague the following spring and invaded Poland a year later, starting World War II.

auction markets and characterize their theoretical properties. This mean that most of the findings from the game theory can be applied to the auction theory.

Generally an auction involves three categories of agents: sellers, auctioneers and bidders (buyers). Sometimes, without loss of generality, the bidders can be potential sellers (e.g. willing to sell a service) instead of buyers, like for example in the public tenders for a service to citizens.

The auction mechanism can be described by the following key elements:

1	Set of n bidders	$N = \{1, \dots, n\}$
2	Set of elements ("signals") that concur in the evaluation of the bidders reserve price	$V_i = \{V_{i1}, \dots, V_{im}\}$ $i \in \{1, \dots, n\}$
3	Set of bids, i.e. actions implemented by the bidders	$B_i = \{B_{i1}, \dots, B_{im}\}$ $i \in \{1, \dots, n\}$
4	Set of bidders' strategies, i.e. functions that transform signals in to bids	$S_i: V_i \rightarrow B_i$ $i \in \{1, \dots, n\}$
5	Utility (pay-off) for each agent	$U_i = \{U_{i1}, \dots, U_{im}\}$ $i \in \{1, \dots, n\}$
6	Set of auction rules that link strategies to pay-offs: - Items to be auctioned - Rules for entering the auction and presenting bids - Rules for winner determination - Rules for payment	

A reserve price is a minimum prices that a seller is willing to accept from a buyer or a maximum price that a buyer is willing to pay for awarding a bid. In an auction, the players are typically not required to disclose the reserve price.

Bid functions maps the agents' *value* to a bid price. The payoff of each player under a combination of strategies is the expected utility (or expected profit) of that player under that combination of strategies[24].

In general it can be said that sellers/auctioneers are interested in making profit, while bidders want to win the bid at the price possibly lower than their reserve price. It is important to remark the difference between *market value and auction value:*

- the market value of an item is in general regulated by the laws of supply and demand. For example, if there is an high demand for that item, its market value increases and if item is hard to find, buyers are willing to pay more for it. But at an auction, items are rarely sold at their market prices;
- the specific context of auctions in general induces bidders to give an item an "auction value" which is different from the market value. This can happen when a bidder is willing to spend more money for an item in the spirit of winning. For example, an auction can be characterized by the presence of so called "externalities" among bidders, i.e. situations in which a bidder will do anything to avoid that his competitors win. The effect of externalities in ads auction will be described later on in section 3.3. Of course, also the sellers (or auctioneers) may consider auction values different from market values, to get higher profit.

One of the available general definitions of an auction is the following[25]: supposing there are n agents (participants to an auction), each of them with some value function v_i over outcomes $\mathrm{v}_i\colon \Omega \to \mathbb{R}$, where Ω is a finite set of outcomes, an auction $\mathcal{A} = (\chi, \mathcal{P})$ can be described with the following two functions:

$$\chi\colon (\Omega \to \mathbb{R})^n \to \Omega$$

$$\mathcal{P}\colon (\Omega \to \mathbb{R})^n \to \mathbb{R}$$

The allocation function χ takes the bids of the n agents and selects the outcome from Ω, while $\mathcal{P}$ represents the payment function.

---

[24] Source: Wikipedia.

[25] R. Colini-Baldeschi, J. Mestre, O. Schrijvers, C. A. Wilkens, "The Ad Types Problem", (2019).

Usually a goal of the auctioneer is to obtain the outcome ω^* corresponding to the maximal total value[26]:

$$\omega^* = argmax_{w \epsilon \Omega} \sum_{i=1}^{n} v_i(\omega)$$

ω^* is called the social welfare and represents the sum of values of all agents.

Referring as b_i the bid of the i-th bidder and as b_{i-1} the bids of all other agents except I, each agent shall focus on their utility, which can be expressed by:

$$u_i(b_{i,} b_{i-1}) = v_i(\chi(b_{i,} b_{i-1})) - \mathcal{P}(b_{i,} b_{i-1})$$

Therefore, the auctioneer, on the basis of bid from other agents b_i, will report

$$b_i^* = argmax_{b_i} u_i(b_{i,} b_{i-1})$$

The pricing function might be used by the auctioneer to "push" bidders, to make offers according to their *true valuations*, in what is often referred as agents incentivization.

In general, an auction is *incentive-compatible* when participants can achieve the best result by setting a strategy in line with their own preferences. More specifically, there exist different levels of incentive-compatibility, that range between the following two extreme situations:

- the *Bayesian-Nash incentive-compatibility (BNIC)*, that occurs if there is a Bayesian Nash equilibrium when all participants act according to their true preferences. In other words, if all the other participants act truthfully, then it is also best (or at least not worse) for you to do the same;
- the *dominant-strategy incentive-compatibility (DSIC)*, that occurs when being truthful leads to a weakly-dominant strategy, which means that acting truthfully leads to the best result (or at least not worse), regardless of what the other participants do. In a DSIC mechanism, strategic considerations cannot help any agent achieve better outcomes than

[26] To be noted that the auctioneer doesn't have direct access to the valuation function v_i. Source: R. Colini-Baldeschi, J. Mestre, O. Schrijvers, C. A. Wilkens, "The Ad Types Problem"

bidding their true value.

Therefore in a DSIC auction, each agent *i* maximizes its utility when bidding through its (true) value v_i, regardless of the bids b_{i-1} of all other agents other than *i*,:

$$v_i \epsilon \ argmax_{b_i} u_i(b_{i,} b_{i-1})$$

Vickrey, Clarke and Groves demonstrated that the above result is always achievable, meaning also in non DSIC auction, by charging agents their externality[27].

2.2 The four basic auction types

The majority of studies related to auction theory[28] are focused on the following four "basic" auction types:

- *English auctions,* also called o*pen ascending-bid auctions,* in which participants make increasing bids, until the price reach a value they are not prepared to pay. The auction goes on until there are no more raising bids and at this point the item is awarded by the highest bidder at the last offered price. Sometimes the item is sold only if the bidding reaches a reserve price set by the seller;
- *Dutch auctions*, also called *open descending-bid auctions,* in which the starting price is sufficiently high to deter all bidders, and is progressively lowered until a bidder is prepared to buy at the current price, thus winning the auction at that price. These auctions are generally very fast and are used to sell perishables, like in fish or flowers markets;
- *First-price sealed-bid auction* (FPSB), in which bidders place their bid in a sealed envelope. The envelopes are opened by the auctioneer and the winning bid is the highest one. The winning bidder shall pay exactly the price

[27]R. Colini-Baldeschi, J. Mestre, O. Schrijvers, C. A. Wilkens, "The Ad Types Problem", (2019).

[28] Typical examples of DSIC mechanisms are majority voting between two alternatives, and second-price auction, while typical examples of mechanisms that are not DSIC are plurality voting between three or more alternatives and first-price auction.

he put in the sealed envelope;

- *Second-price sealed-bid auctions* (SPSB), also called Vickrey auctions[29], in which bids are delivered in sealed envelopes and simultaneously opened by the auctioneer. The highest bid wins, however, differently from FPSB, the price to pay equals the second-highest bid.

In English auctions and SPSB auctions (Vickrey auctions), dominant strategies lead to a unique equilibrium: in fact, the asset is assigned to the bidder who offered a price such that all other bidders are excluded, and pays that price in the case of English auction and a price corresponding to the second highest bid in the case of Vickrey auctions. Such equivalence is described in details in the following pages.

As said, one of the main seller's goals is maximizing his revenues and he could easily do it if he could know all the evaluations of the potential bidders, thus offering the item at a price equal to maximum evaluation. In general this is not possible and the same kind of problem involves each bidder, who cannot know the other bidders' evaluations. In other words, the so-called *asymmetry of information*, that happens when the auction participants do not have the same set of information, is the main cause of limitation of the utility.

For this reason, the oral character of English auctions provides an advantage for the bidders, who can exploit information coming from other bidders and observe at any time the current price, thus being allowed to dynamically update his evaluation accordingly. Conversely, FPSB auction requires complex analysis that shall involve bidders potential valuation, but also assumptions on how bidders could estimate other bidders' valuations.

There is another important auction classification in terms of the *value model* that is used to assign, or to guess, the value that is given to an auctioned item:

- *Indipendent Private Values auctions*: each participant has a private individual value of the item, being such value independent from both the

[29] This type of auction was first described academically by Columbia University professor William Vickrey in 1961 though it had been used by stamp collectors since 1893. Source: WIkipedia

other bidders' values and from other external factors[30]. Private Values auctions have an important feature, formalized by the "Revenue Equivalence Theorem" [31], consisting in guaranteeing the same revenue expected by the seller for all the four basic type of auction listed above.

- *Generalized Private Values auctions.* It is a generalization of the private value model, in which it is assumed that bidders' valuations depend not only on private information, but also on unknown information that are supposed to be available for example from the auctioneer or others participants. Bidders are therefore willing to change their private valuation if and when they get other information during the auction;
- *Common Value Model.* The item has a fixed and identical value for all participants, but none of them know it as it can only be estimated on the basis of private information[32]. The development of models and methods for analyzing common value situations responds to the need to explain a singular phenomenon that emerged in real application, like for example auctions for the assignment of rights for exploitation of particular natural resources. In such situations, it can happen that the winning company obtains very low or even zero profits, due to the lack of information from other bidders' values (this phenomenon, is also known as the "winner's curse") [33].

English auctions are by definition the ones characterized by the higher amount of information available, mainly as a result of the living interactions between bidders. It is important to remark that in general "public auctions" generate higher revenues

---

[30] As an example, different art collectors can assign different values for the same artwork, according to their tastes which in general should not be affected by the influence of other bidders

[31] In 1996 Vickrey won the Nobel Prize in Economics with his "Revenue Equivalence Theorem", which states that *the expected earnings for the seller is the same for both the auction in a sealed envelope at the best price and for the auction at the second price* (that is, the Vickrey auction)

[32] It is the case, for example, of an auction for a company branch: bidding companies, even if operating in the same sector, can have different signals about conditions and profitability they will find themselves when acquiring that asset.

[33] The phenomenon called "winner's curse" occurs when winning bidders obtain zero profit (or very lower than expected), the reason being linked, in most cases, to a lack, or incorrect use, of the available information during an auction.

than “secret auctions”, as more information is disseminated and, as a consequence, bidders can change their valuations and, if necessary, adapt its offer to the highest bid.
Therefore, the best policy for the auctioneer to get bidders *incentivized* to offer at levels close to their evaluation is to disseminate the maximum amount of information during the auction. For this reason, from the seller's point of view, the preference is for an ascending open auction model.

In English auctions where the assumption of independent private values is valid, each participant has the dominant strategy that can be formulated as: *to increase the last offered price by the allowed minimum raise, until its reserve price is not exceeded*. If every bidder sticks to this strategy, the item is awarded at a price equal to the second reserve price, plus the minimum raise.

$$P^* = R_2 + \delta b$$

Where P^* is the winning price, R_2 is the second reserve price and δb is the minimum raise admitted in the auction.

Thus, there is a substantial equivalence among English auction and Vickrey's auction, since in the latter, on equal terms, the winning price would be $P^* = R_2$, therefore differentiating from English auction just for the minimum raise δb.

In both cases the item is always assigned to the bidder with the higher *value*, who gets a pay-off U^* equal (or very close in the case of English auction) to the difference between his reserve price and the second reserve price.

$$U^* = (R_1 - R_2) - \delta b$$

As seen from the sellers point of view, U^* is also its portion of profit, also referred as *surplus*, that is left on the winning bidder’s hands.

In single-item auctions, the bidders’ valuation are always single-parameter functions, while in multiple-item auctions those functions might be multiple-parameter. Assuming in general there is a certain function $x_i: \Omega \rightarrow \mathbb{R}^+$ that maps the space of possible outcomes Ω to a quantity of item that the bidder i receives, the single-parameter valuation function can be expressed as $v_i(w) = v_i * x_i(w)$, where v_i is multiplied by the function x_i : therefore in single-item auctions, where an agent has a value v_i associated to the auctioned item, $x_i(w) = 1$ if agent i obtains the good in outcome ω and it is 0 otherwise[34].

2.3 Multiple-item auctions

In auctions with multiple items, the goods could be either tendered one by one, through a sequence of single-item auctions, or tendered simultaneously, in a way that each participant can bid for one or more units (simultaneous auction). The two methods are equivalent when a set of identical goods are offered, while they have significantly different properties when goods have different characteristics.

It has been found empirically that sequential auctions are often characterized by decreasing item by item price trend. This phenomenon, called "afternoon effect", can be explained both with the bidders aversion to risk (the need to obtain at least one good could result in high bids in the first few auctions), and as a result of bidders' strategic behaviors: for example, a bidder with high valuation might start with low bids to induce lower valuation to other participants and thus awarding goods in subsequent auctions at lower prices, when the competition is less likely to be challenging, as some competitors are out. As a consequence, this kind of auction has an overall inefficient result, determined by the effect of distortion between bidders' values and their actual bids.

Simultaneous auctions are therefore more efficient, however they present two main issues:

- *"Synergies" among items*: bidders could be interested in awarding a set of goods, rather than a single asset. The valuation of the single item could even be zero or, in an opposite situation, an item could be valued less if purchased together with others. If this cases it becomes difficult to identify an efficient auction model.
- *Asymmetry between bidders*: some bidders could be interested to buy all the goods and others in buying only some goods, but totally disinterested

[34] R. Colini-Baldeschi, J. Mestre, O. Schrijvers, C. A. Wilkens, "The Ad Types Problem", (2019)

in others.

To solve the above issues, it is sometimes necessary to find support in “external” criteria, on a case-by-case basis, that are independent of the auction mechanism.

While Vickrey's original paper mainly considered auctions where only a single, indivisible item is being auctioned, the Vickrey-Clarke-Groves (VCG) auction is a generalization of Vickrey's auction for the case of multiple goods.

2.3.1 Vickrey, Clarke and Groves (VCG) auctions

We now refer to an auctions where:

- a set of identical products are being sold;
- bidders participate in the auction by declaring the maximum price they are willing to pay to receive N products;
- bidders can also declare more than one bid, since their targeted pay per unit price may vary depending on how many items they gets;
- bidders cannot see other people's bids since they are sealed;
- once all the bids are made, the auction is closed.

At this point, the auction system evaluates the possible combinations of bids and select the combination C^* that meets the following conditions:

- It maximizes the total sum of bids;
- the total amount of items shall not be exceeded;
- the combination includes no more than one bid from each bidder.

Bidders who made a bid included in C^* receive the won products, according to the quantity specified in their bid; however, instead of paying exactly the amount of their initial bid, they will pay a part of it, which corresponds to the *marginal harm* caused by their bid to the other bidders.

We can indicate with $A = \{a_1, \dots, a_m\}$ the m items being auctioned, with $B = \{b_1, \dots, b_n\}$ the n bidders, and with V_B^A the sum of the values each winner assigns to the won items, being the value zero when no item is won. V_B^A is named *the social value* of the VCG auction.

Suppose now that a bidder b_i made a bid $v_i(a_j)$ for the item a_j. If b_i wins the bid, the amount P^* that he will pay is:

$$P^* = V_{B\backslash\{b_i\}}^{A} - V_{B\backslash\{b_i\}}^{A\backslash\{a_j\}} \quad (1)$$

P^* is the social cost of the winning that is incurred by the other bidders and it is the difference between the two terms:

- $V_{B\backslash\{b_i\}}^{A}$ which is the social welfare achievable by the n-1 bidders without the presence of the bidder b_i but with all items available
- $V_{B\backslash\{b_i\}}^{A\backslash\{a_j\}}$ which is the social welfare achievable by the n-1 bidders without the presence of the bidder b_i, but with the item a_j no more available, given the bidder b_i got it.

If a winning bidder made a bid with its true value $v_i(a_j)$ he obtains the higher possible utility:

$$v_i(a_j) - P^* \quad (2)$$

An important characteristic of this kind of mechanism, is that the only situation in which the price paid by the winners is the same as their initial bid is when the second best combination has the same total bids amount of the best combination C^*. Only in such a case the seller's revenue is maximized, while in all other cases the buyers payment is lower and the seller's revenue is not maximized.

To make a practical example of VCG auction, let's suppose there are 3 bidders, namely B_1, B_2 and B_3, competing for 2 separate side-by-side positions on a web page:

- B_1 and B_2 wish to get only one position and are willing to pay 20€ and 12€;
- B_3 wants to get the package of two positions (he could wish to promote two different products of his own) and are not interested in awarding only one slot. He values the package 30€.

If the three advertisers bid truthfully, the VCG auction will result in the assignment of the two slots to B_1 and B_2, since the sum of their bids, i.e. 32€, is higher than B_3's bid for the entire package, i.e. 30€.

We can now apply (1) and (2) to calculate payments and utilities reported in Figure 13.

	$V_{B\setminus\{b_i\}}^{A}$	$V_{B\setminus\{b_i\}}^{A\setminus\{a_j\}}$	$P^* = V_{B\setminus\{b_i\}}^{A}$ - $V_{B\setminus\{b_i\}}^{A\setminus\{a_j\}}$	$U = v_i(a_j) - P^*$
B_1	30€	12€	18€	2€
B_2	30€	20€	10€	2€
B_3	32€	0€	0€	-

Figure 13: Example 1 of VCG action outcome. Source: own elaboration

If we consider the same auction, but with the only difference in B_3 value, that is now 35€ (thus higher than the sum of B_1 and B_2 values), the two slots are assigned to B_3 with the payment and utility as reported in Figure 14.

	$V_{B\setminus\{b_i\}}^{A}$	$V_{B\setminus\{b_i\}}^{A\setminus\{a_j\}}$	$P^* = V_{B\setminus\{b_i\}}^{A}$ - $V_{B\setminus\{b_i\}}^{A\setminus\{a_j\}}$	$U = v_i(a_j) - P^*$
B_1	35€	35€	0€	-
B_2	35€	35€	0€	-
B_3	32€	0€	32€	3€

Figure 14: Example 2 of VCG action outcome. Source: own elaboration

Therefore, in VCG auctions bidders have incentives to bid their true valuations, and making truthful bids corresponds to adopt a weakly-dominant strategy.

It has to be remarked that the optimal social welfare can be degraded in case of bidder collusion (e.g. single bidder who makes multiple bids under different names); however, also in case of collusion, the VCG outperforms the generalized second-price auction both in terms of seller's revenues and in terms of assignment efficiency.

The following section will introduce the phenomenon of the on-line auctions and in particular those in which items are being sold, while Chapter 3 will focus auctions for on-line advertising.

2.4 Online auctions for buying and selling items

With the diffusion of the internet-based service and applications, part of the traditional buying and selling activities has been progressively moved on digital platforms, leading to the development of the so called electronic commerce, or simply *e-commerce*.

Today, at least one part of an e-commerce transaction is typically handled through the Web. Initially the most common examples of online sales concerned the purchase of online books (such as Amazon) or music (download from the iTunes Store) and some payment services or purchase of tickets for events; today, thanks also to the large number of users who access the internet, it is possible to complete more or less any type of transaction online. There are three areas of e-commerce: online retailing, electronic marketplaces, and online auctions. This chapter will focus on online auctions and specifically on online auction for advertising.

Online auctions (or electronic auction or e-auction) are by definition auctions held over the internet and, like “physical” auctions, they can be of one of the four basic types defined in chapter 2, or a combination between two basic types, or even other less common types.

The potentiality offered by the internet has enabled online auctions to overcome some limits of traditional auctions, such as the need of a physical location for the auction and the need of physical and simultaneous presence of the bidders.

Maybe the best known site in the world for online auctions is eBay, founded on 1995 by Pierre Omidyar and today capable of a turnover of about 25 billion dollars. When you list an item for sale in an eBay auction, you choose a starting price and whether your auction should run for 1, 3, 5, 7, or 10 days and interested buyers place bids; when the auction ends, you sell to the highest bidder;. when selling certain type of items, like for example vehicles, you can add a reserve price to make sure you get the price you want for your item[35]. It can be said that eBay

[35] Source: eBay web site

auctions are on-line time-limited English auctions. eBay's revenue consists of commissions that are charged to the seller to start an auction, and of additional fees for optional advanced listing upgrades and services.

Chapter 3
Methods and strategies for online ads auctions

Online advertising, or simply on-line ads, has developed in recent years thanks to the rapid growth of internet and of the number of social networks users.

Since the second half of the 1990s began to spread the use of *banners*, i.e. advertising spaces sold by the web site owner to any person who wanted to promote any good or service. This form of advertising was in most cases present on search engine sites and did not follow an auction logic: the spaces on the web page were sold at a price determined by the characteristics of the banner, for example its size and position, and by the number of appearances on the website.

This kind of advertising mechanism is often called *pay per impression (PPI)*, meaning that advertisers shall pay any time the ads is *impressed*, i.e. visualized. Such kind of mechanism

Advertising was no longer sold on a basis of thousand showings. Every time a sponsored link was clicked by a consumer, the advertiser had to pay an amount equal to the advertiser's last bid. The highest bid was shown in the highest position, following a descending order of bids.

3.1 Auctions for on-line advertising

In more recent years, the mechanism behind on line advertising has progressively moved from one-to-one agreements among sellers (e.g. website owner) and advertisers to auction mechanisms for assigning one or more potential available ads spaces to a number of potential buyers.

Such progressive transformation was driven in particular by:

- the growing demand for on-line advertising;
- the increase of number of users accessing the internet and the social

networks, and therefore of potential viewers of the ads;

- the development of suitable algorithms for efficient management of on-line auction mechanisms;
- the increasing digital computational speed, which today allows to manage complex and simultaneous operations on web platforms in nearly real time.

On-line ads auctions can be implemented in various ways

There are today plenty of on-line platforms, search engines and social networks that implement auction mechanisms to sell online ads and, in general, each of them can design its own auction mechanism, which is customized on specific objectives.

However, most of the implemented ads auctions can be conceptually grouped within the so called *position auctions*, in which advertisers can make their bid to get assigned a suitable space to promote their products or service. The following section provides an in-depth of such kind of auctions.

3.1.1 Position Auctions

The first search engines used the so called meta-tags, to record a list of keywords and link them to relevant web pages: once one or more recorded keywords were put in the search field, the engine displayed a list of result sorted, as best as possible, from the most relevant down to the less relevant from those that matched some way with the keywords. In addition to being inefficient due to the need to manually enter keywords and update them continuously, such method was also unprotected against webmasters who started to insert “dummy keywords” into their pages, in order to increase their potential appearance as resulting pages, even when their site had no matching at all with the typed keywords.

In a context characterized by increasing demand for online ads in a strong competitive scenario, from the second half of the 90’s the bigger internet companies (first examples were Yahoo!, Microsoft and Google), began to cooperate with eminent economists and researchers in order to study proper algorithms and mechanisms to make their ads auctions more profitable; as a result, the most popular search engines began to sell the “best” positions of their search results, alternating them with those really matching with the keywords[36].

To make a practical example of how this mechanism works today we can suppose a user types the words “car” in order to make a research on such topic. As shown in Figure 15, the web page returns relevant ads (and corresponding locations on a map) just above the search results. Two ads are immediately visible (best positions), while other are displayed only if the user scrolls left to right such ads (less ranked positons).

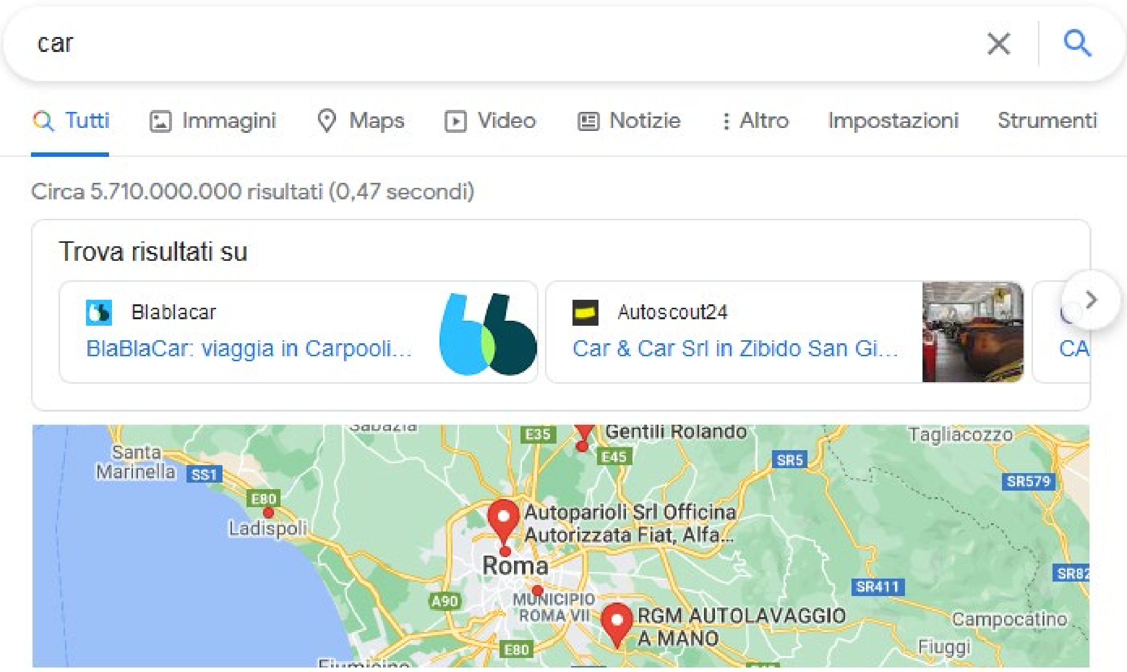

Figure 15: Sponsored results from Google page.

Generally the design of a position auction is based on the following steps:

- an advertiser defines a set of keywords relevant to the item it wishes to sell;
- on the seller’s side (e.g. the web platform owner), there could be a number of spaces available to accommodate ads on a certain web page, each of them having different characteristics in term, for example, of size, position in the page and more in general of relevance for a potential viewer;

[36] Still today many search engines, like e.g. Google, uses to place in the highest positions, results that, although relevant to the search, are flagged as "promotion".

- for each defined keyword, each advertiser places a bid indicating the amount that it is willing to pay if a user clicks on its ad; also, the bid shall be based on advertiser's preferences, such as desired type of advertising, frequency of appearance, maximum sustainable cost and even target of buyers to whom the ads has to be shown. By bidding on specific keywords of interest, advertisers can target their potential customers and also they can receive useful feedback on how certain keywords are efficient for their business;
- when a user's search query is launched, a set of ads that someway match the typed keywords is displayed. These ads are ranked by bids (or a function of bids) and the ad with the highest bid receives the "best position", i.e those more "profitable" from a marketing point of view; even if best positioning strictly depends on the web page layout, very often the best positions are those the one on the top of the first page, down to those on the bottom and even on the following pages, being of course the higher positions mostly likely to be clicked on.
- If the user clicks on an ad, the advertiser is charged an amount that depends on the bid of the advertiser below it in the ranking.

Practically, the original pay per impression (PPI) has been replaced by mechanisms in which advertisers, once they win a bid and their ads is shown, pay on the basis of the "success" of the ads, in terms of number of interactions (e.g. clicks) it receives.

It is important to clarify the difference among three kind of events that can occur once an ads is going to be shown:

- *impression*: already mentioned in the introduction to this chapter, an impression is simply the event of appearing an advertising on some user's screen. It doesn't imply that such user has looked at it;
- *click*: this event is triggered when a user just clicks on an ad, regardless if other actions will follow the click;
- *conversion*: when a click is converted into a determined action, a conversion occurs. A conversion can be for example an item purchased on the web page addressed after the click, or a phone call to the advertiser's

company for asking more information about a sponsored item, or a form duly filled by a user, and so on. Conversions could also be of more complex type, such as for example a conversion which is triggered by users who keep visiting a web page for a certain minimum time.

A key difference between traditional auctions for buying and selling items and most of the on-line ads auctions is on the criteria to determine the winner. In the former the only criteria is the offered price, i.e. who made the higher bid wins. In the latter there is an additional factor, often called *quality of bid*, that is how relevant an ad is for the target viewers, or in other words, the probability that the ads will generate an interaction (e.g. by an item, watch a video and so on) by targeted users. This performance will be further explained in following sections.

3.1.2 Google Ads case

Google Ads is the Google's main platform for advertising and the most used by large companies and professionals, since it offers a large range of positioning and targeting, with a high level of customization. There exist simplified versions, like Google Ads Express, that offer automated, less flexible, tools for supporting advertising choices, and therefore they are mostly used by small companies or individuals with limited resources. [37]

The following rules regulates the *cost per click* (CPC) mechanism offered today by Google Ads[38]:

- advertisers only pay for clicks on their ads;
- a maximum cost-per-click bid, called *max CPC*, shall be set by the advertiser. It represents the maximum amount that it is willing to be paid for one click on the ad;
- if an ad is selected, the *actual CPC* is the final amount that is actually charged for a click. Actual CPC is upper bounded by the *max CPC*;
- advertisers can chose between manual bidding, i.e. directly choosing the bid amounts, and automatic bidding, that is Google sets bids to get the

[37] Source: https://performance-ppc.com.
[38] Source: Google web site.

advertiser the maximum number of clicks within its budget.

The ads are selected according to an *ad rank* that is based on the combination of several parameters, including:

- bid, expressed by the advertiser as maximum cost per click (max CPC), that is the maximum amount an advertiser will pay for a click on its ad. Such indicator can also be monitored by the advertiser, who can use it to have indication of the probability of success of its ads and eventually try to improve it;
- expected click-through-rate;
- ad relevance;
- landing page experience.

To say it differently, the success of an ad is determined mainly by the CPC and other three parameters that together represent the ad quality, that provides a measure of how relevant and attractive an ad and the website it links are to the users who will see it.

Each advertiser's Ad Rank is then used compared with *Ad Rank thresholds* to determine where the ad appears and, if it is the case, in which position: even when an advertiser makes a bid lower than competitors, he could win a higher position if its ad has higher quality.

All ads that are eligible for a certain position (i.e. with ad rank above the relevant threshold) participate to an SPSB[39]-like auction (that is multiple-item auction when more than one position is available), in which the amount to be paid by the winner is not just the second price, but it is the actual CPC needed to achieve an ad rank immediately above the second ranked ads.

[39] Second Price Sealed Bid (see section 2.2).

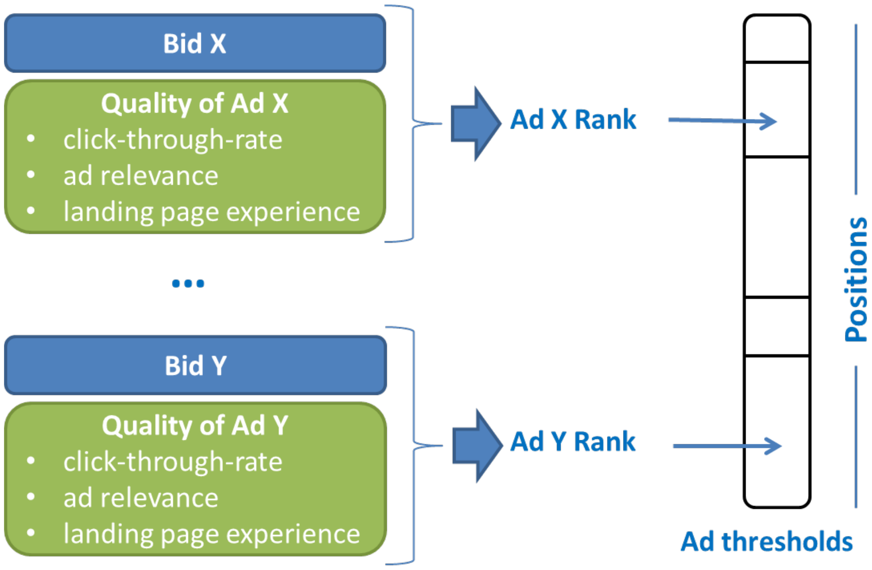

Figure 16: Graphical representation of Google Ads position auctions. Source: own elaboration.

To make a practical example, we can suppose there are six advertisers $A_1,\ldots,A_6$ with ad ranks of 10, 20, 30, 40, 50 and 60 respectively, three available ad positions above the search results (the most desirable position), for which the threshold is 55 and other three available positions below the search results, for which the threshold is 15. The outcome of the auction would be:

- A_6 has a rank of 60 and wins a top side position, being the only ad overpassing the first threshold at 55. Having no competitors, he pays a CPC corresponding to its reserve price;
- A_5 (rank 50) wins the first available bottom side position (threshold at 15) and pays a CPC just enough to beat A_4 (rank 40);
- A_4 (rank 40) wins the second available bottom side position and pays a CPC just enough to beat A_3 (rank 30);
- A_3 (rank 30) wins the third available bottom side position and pays a CPC just enough to beat A_2 (rank 20);
- A_2, although with a rank higher than the second threshold, will not award a position since all have been won by A_3, A_4 and A_5;
- A_1 will not award a position since its rank is below the second threshold.

In order to be as much as possible competitive, for an advertiser it is fundamental:

- to define the goals of its ads campaign and relevant strategies, as for example:
    - when an advertiser wishes to increase awareness of its brand, focusing on impressions may be the best strategy;
    - if the objective is just to increase traffic to a website, a cost-per-click (CPC) bidding strategy may be the right one;
    - if the main goal is that customers take a direct action on the landing page, the best would be to focus on conversions;
    - when the objective is to increase views or interactions by running video ads, a cost-per-view (CPV) or cost-per-thousand impressions (CPM) bidding could be the right choice;
- to make a good ads design, according to the defined bid strategy;
- to have a method for measuring the probability of success of an ad (attractiveness to users), which on turn influences the probability of winning a good position in the auction. At this aim, the Google Ads package includes a bid simulator which can estimate the number of clicks and corresponding cost that a certain ad would receive.

In addition to search engines, important players in the world of the on-line ads auctions are the social network companies, which make large investments to design efficient ads auction mechanisms, being advertisement the biggest contributor of their revenues. Facebook ads is a relevant example described in following section.

3.1.3 Facebook Ads case

Like Google, Facebook manages its ads through an auction system: candidate ads are sorted by an algorithm that takes into account several variables. Every time a user connects to Facebook, the system can access a very large set of candidate advertisements and choose those who maximize a total value that comes from a combination of three main factors[40]:

- *bid*: the amount an advertiser is willing to pay to achieve the desired result;

[40] Source: Facebook web site.

- *relevance of the ad*, which is an indicator composed by:
    - quality of the ad;
    - estimated action rates.

The estimated action rate provides an indication of the likelihood that a particular user will interact with a particular ad, leading the advertiser to its desired result.

Like in the Google case, the relevance of an ad is an important factor in Facebook auctions, since ads with very high relevance often cost less and perform better than ads with lower relevance.

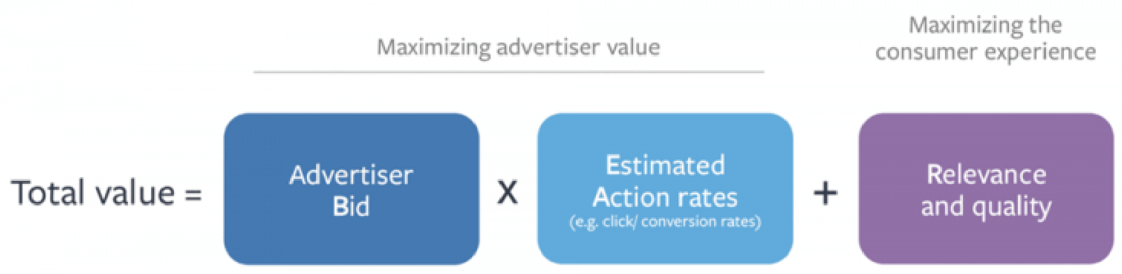

Figure 17: The three main factor for determining the total value of an ads. Source: https://trucchifacebook.com

Advertisement with the highest total value wins the auction.

When an ad campaign is not matching the advertiser's expectations, a key point is how the probability of winning could be maximized. To do this, different strategies could be adopted:

- manually increasing the offer to make it in line with the actual return;
- increasing the budget at a level at least equal to the amount of the offer chosen, daily.
- extending the target in order to reach more customers (Facebook offers dedicated guides to achieve such objective);
- activating more positioning options, in order to let Facebook set up the best one according with your ads campaign characteristics, thus increasing the probability to award a position.

3.1.4 Ads performance indicators

There exist several metrics for analyzing both the potentiality and the actual result of an ad campaign. The most popular are described in the following:

- *click through rate (CTR):* The CTR is an indicator of the percentage of clicks resulting from a number of impressions. As an example, if an ad has been impressed 100 times and clicked 20 times, the CTR is 20%;
- *click-to-open rate* (CTOR) is similar to CTR, but it measures the number of clicks by different users. Thus, again with 100 impressions, if one user, and no others, clicks 10 times on the same ad, the CTOR is 1%, while the CTR is 10%. Consequently, the CTOR is more reliable than the CTR in evaluating the success of online marketing campaigns. Since advertisers typically pay more for a high click-through rate, getting many clicks with few purchases is undesirable for advertisers;
- the *conversion rate* is simply the ratio between number of conversions and number of clicks, e.g. if 100 clicks generate 5 conversions, the conversion rate is 5%;
- the *Return on Ad Spend (ROAS)* is an indicator of the profitability of the ads campaign. For example, if the cost sustained for a campaign was 100€ and only one product was sold for 20€, the ROAS was 20%.

Sometimes it could be dangerous to rely on only one metric and founding its ads design on as many metrics as possible, and studying the relation among them, could be a winning strategy.. For example, an ad could have a very high CTR, but a poor conversion rate, which would result in unsuccessful campaign when the goal is to have conversions. Or, in another kind of situation, an ad could have very few click but most of them leading to a conversion; also in this case, a global view at different metrics could help in properly scale the campaign to more significant size.

3.2 Development of models for position auctions

Since the 90', position auctions for on-line advertising began to be conceived and implemented, mostly by the larger web search engines, who captured the increasing demand from advertisers and the growing number of customers accessing the internet. Different models have been proposed since then, starting from the simpler ones based on English auction to the more complex used today by most of the larger search engines and social network companies.

The first on-line ads were paid, in most of the cases, on the basis of the number of times the ads appeared on a certain page, thus advertisers bought flat fees, typically for a fixed number of 1000 impressions..

By 1997, the internet company Overture (later acquired by Yahoo!) introduced the Generalized First-Price Auctions (GFP) as a new model of selling on-line ads. In the GFP:

- each advertiser bids for a particular keyword and declares his budget for a click (payment on a per-click basis), thus, instead of paying to show his ad to everyone visiting a page, an advertiser can target his customers by choosing keywords relevant to his products.
- every time a customer clicks on a link, its owner is automatically billed the amount of his most recent bid.
- ads are positioned in descending order of bids, thus giving more visibility to the highest bids.

Such mechanism has the drawback of being very unstable as advertisers may update their bid very frequently to respond to other agents' bid[41]. So, for example, when 4 agents compete for 3 positions and have a value of 8€, 6€, 4€ and 2€ respectively for a click on their own ad, if agent 2 (value 6€) bids 4.01€ to win at least a position, agent number 1 (value 8€) may just raise the bid to 4.02€ in order to get the highest position; on turn agent 2 may again raise his bids and so on. The most probable consequence is that some advertisers shall invest money into

---

[41] B. Edelman, M. Ostrovsky, M. Schwarz "Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords".

bidding robots to speed up their raises, resulting in decreasing the social welfare, being such situation inefficient for both advertisers and search engine.

To fix such inefficiencies, the on line ads market started to develop second-price structures. The generalized second price (GSP) auction is one of the most popular today. In the first and simpler implementation of GSP:

- each advertiser submits a bid for positioning his ad;
- ads slots are assigned on the web page in descending order of the bids' amount;
- the advertiser who wins the first (i.e. best) position pays, for each click on his ad, a price equal to the second advertiser's bid, plus an increment established in the auction rules;
- in the same way, the second advertiser pays the bid of the third advertiser plus an increment, and so on till there are available positions and bidders who compete for them.

Hence, similarly to VCG mechanisms, in GSP auction the amount that a bidder shall pay (which is as per-click payment in the particular case of ads position auctions) does not directly depend on his bid. Edelman, Ostrovsky and Schwarz showed that, although GSP looks similar to VCG, its properties are very different: in particular, unlike the VCG mechanism, GSP generally does not have an equilibrium in dominant strategies, and truth-telling is not an equilibrium of GSP[42].

VCG and GSP auction mechanisms are both very popular (for example Google mostly uses GSP, while Facebook mostly uses VCG) and both have their pros and cons. GSP is tailored to the specific characteristics of the online ads market, while VCG is in general more computationally complex, since it requires determining the harm each bidder causes to the others, and therefore its complexity grows with increasing number of bidders. Conversely, calculating the winners in GSP is quite easier, as it simply involves matching the highest bidders to the best ad position.

---

[42] B. Edelman, M. Ostrovsky, M. Schwarz "Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords".

Varian and Harris [43] provided three reasons to explain Google's preference for GSP over VCG: "*We thought very seriously about changing the GSP auction to a VCG auction during the summer of 2002. There were three problems: 1) the GSP auction was growing very rapidly and required a lot of engineering attention, making it difficult to develop a new auction; 2) the VCG auction was harder to explain to advertisers; 3) the VCG auction required advertisers to raise their bids above those they had become accustomed to in the GSP auction. The combination of these issues led to shelving the VCG auction in 2002*."

3.3 The effects of externalities in on line ads auctions

In the previous sections we have seen that the most popular metrics for determining the value that an advertiser obtains when he win an online ads auction are the CTR (click-through-rate), the CTOR (click-to-open rate) and the conversion rate. It follows that also most of the ads auction models are built using such indicators as the advertiser's utility that is worthy to maximize.

This type of model is well representative of situations where only one ad space is available on a web page, and therefore the winning ad is the only one being shown on a certain web page.

Let's consider now a different case, in which two or more ads are impressed together in the same page, and suppose:

- that one of those ads is particularly attractive than it can capture all the attention of the viewers, thus precluding the other ads on the same page to be considered;
- or that the impression of a low quality ad can lead viewers to quit the page, thus, again, causing a negative impact to other advertisers having their ads on such page;

The above situations are examples of effects which cannot be controlled with models that only consider the CTR and the conversion rate. Such kind of effect is

[43] H. R. Varian, C. Harris, "VCG in Theory and Practice", Google, Inc. (December 2013).

commonly referred as *externality* and occurs when, in an auction, the utility to a winner has a certain dependence from the other winners.

Several studies have been conducted to model the effect of externalities and to build algorithms for the prediction of the winner's utility in complex multi-ads scenarios.

Ghosh and Mahdian[44] initiated the study of externalities in online advertising and proposed several models to describe them, primarily in the context of on-line lead generation advertising, i.e. the commercialization of selected customers' personal information (leads) to companies that can make business with such leads, by directly contacting the potential customer[45].

By analyzing the problem of modeling externalities in online ads, they studied the so called *winner determination problem*, proposing models which are based on the preferences of the users.

The utility to an advertiser who buys a lead depends:

- on the number of advertisers who buy the same lead: with a limited number of buyers, the value to each advertiser is in general higher, since there is less competition around the potential customer;
- on which other advertisers obtain the lead: if a competing advertiser offers a similar service with better quality (or less price) the value of the lead decreases much more than facing with a competitor who is offering a similar service with lower quality (or higher price) or even a different service.

The model is built on such premises and works as follows:

- there are n advertisers $1,...,n$, each of them with a private value v_i for capturing the business of the user (i.e. the value advertiser i obtains when

[44] A. Ghosh, M. Mahdian, "Externalities in Online Advertising", WWW 2008 / Refereed Track: Internet Monetization - Online Advertising April 21-25, 2008 Beijing, China

[45] The lead generation, or pay-per-lead, is a very popular advertising strategy. To make just two examples: mutuionline.it provides financial institutes with interested to obtain a mortgage; edilportale.com provides professionals with personal information of people who require domestic maintenance services. In both cases, potential customers are required to fill dedicated forms containing both personal information and detailed specifications about the required service.

he is selected by a user);

- there is a user, who can choose at most one among a number of advertisers, on the basis of his perception of the quality q_i of the advertiser *I*;
- there exist a best "outside option" of quality q_0. Outside options means that the user could receive a better quote through another medium and finally choose such external option;
- all quality parameters q_i are random variables, to take into account the fact that users do not all make the same choices among the advertisers. However, in general the q_i's need not be independent, since the choices of users are often dictated by the same general principles[46];
- the user picks the advertiser with the highest q_i among a set *S* of advertisers provided that q_i is higher than q_0 ;
- such winning advertiser derives a value of v_i, while all other advertisers derive a value of 0.

It results that the expected value deriving from a selection of a set S of advertisers is:

$$v(s) = \sum_{i \in S} v_i * \Pr\left[q_i \geq q_j \; \forall j \in S \cup \{0\}\right]$$

From the above expression it is clear that $v(s)$ does not depend on the random values q_i while it depends on their relative ordering.

The winner determination problem in this model consists in choosing a set *S* of advertisers to maximize $v(s)$. For this problem Ghosh and Mahdian proposed an

[46] For example, knowing that a user perceives a large financial group as better in quality then a small regional bank, increases the likelihood that he also prefers another large financial group to another small regional bank.

approximation algorithm[47] with an approximation factor that is logarithmic in the ratio of the maximum to the minimum advertiser's bid.

Let's now consider the effect of externalities in the context of the keyword advertising, that is for example one of the main business of Google.

Gomes, Immorlica and Markakis[48] defined a model in which users perform *ordered search*, as follows:

- they browse sponsored links from top to bottom on a certain web page;
- after having red each ad, they decide whether to click on it or not;
- if not, they decide whether to continue looking at the sponsored list or quit the page.

Accordingly, they defined:

- *continuation probabilities,* which are the probabilities (for each ad) that a user continues searching through the sponsored list after clicking or not on some ad. In general users stop browsing the page either because they found what they needed or because they were disappointed with the previous links, thus continuation probabilities capture *position externalities*, consisting in the negative effect that the ads on the top cause to the click-through rates of the ads in the bottom;
- *conditional click-through rates*, that is the probability, for each ad, of a click conditional on the user's previous clicking history. Such parameter capture *information externalities*, since it is possible to assess how the information collected by the user by clicking on one given link impacts the click-through rates of the other links he could read.

Each advertiser's score is obtained from their bid multiplied by a weight which depends on each advertiser's characteristics. Advertisers are then ranked by their

[47] For example, knowing that a user perceives a large financial group as better in quality then a small regional bank, increases the likelihood that he also prefers another large financial group to another small regional bank.

[48] R. Gomes, N. Immorlica, E. Markakis, "Externalities in Keyword Auctions: an Empirical and Theoretical Assessment"

score and the available slots are assigned in decreasing order of scores. Finally, each advertiser pays (per click) the minimum bid necessary to keep his position.

The value of acquiring a position is strongly affected by quality and the position of the other ads, thus externalities imposes among advertisers influence the click-through rates.

Most of the sponsored search auctions models cannot capture the externalities that one advertiser imposes on the others, since they are based on the assumption the CTRs are separable:

$$CTR_i = q\_adv_i * q\_pos_i$$

meaning that the CTR of an ads *i* is a product of a quantity q_adv expressing the quality of the ads and a quantity q_pos_i expressing quality of the occupied position.

Gomes, Immorlica and Markakis deviated from the separable model and studied a new model that integrates the users' search behavior and the advertiser's bidding behavior, to take into account externalities in sponsored search auctions. Their work was conducted through the following steps:

- from an empirical point of view:
  - they retrieved recorded data from Microsoft Live[49] (three months of impression and clicking data) to assess how the CTR of an ad is affected by the user's click history and by the other ads which appear in the same page;
  - they used the retrieved data to estimate their *ordered search model*, a model of users' behavior that assumes ordered search. In particular, the model assumes that users choose which ad to click by analyzing one link at a time and that they browse the ads from top to bottom;
  - they assessed impacts on the CTR of an ad caused by both user's click history and other competing ads;
  - then they compared their model with the separable click-through rate model in terms of clicking predictions;

[49] The brand Microsoft Live included services offered by Microsoft like "Microsoft Office Live", "Live Anywhere", "Xbox Live", "Live Mesh" and others. Source: Wikipedia.

- as one of the results of this empirical steps, they found that users that click one sponsored link are more likely to keep browsing the sponsored list than users that don't make clicks at all;
- as a second result, they found that the new empirical model is shown to have more predictive power than the separable click-through rate model.

- From a theoretical point of view, they focused on scoring rules and in particular on how the choice of a scoring rule affects the set of complete information Nash equilibrium of the GSP auction. They analyzed a specific GSP equilibrium that maximizes the search engine's revenue among all pure strategy Nash equilibria. As an important result they found that it is not possible to find a scoring rule which implements an efficient equilibrium with VCG payments[50] for all profiles of valuations and search parameters.

[50] VCG payments charge each advertiser the welfare difference imposed on the others.

Conclusion

With the spread of the digital technologies, part of the traditional buying and selling activities has been progressively moved on digital platforms. Online ads auctions is a relevant example: the development of powerful algorithm running on the web has enabled search engines and social network to progressively adopt auction mechanisms to sell their ads slots to advertisers.

In fact, if originally the ads were sold on a pay-per-impression basis through one-to-one agreements among web publishers and advertisers, in more recent years, the mechanism behind on line ads has moved to complex auction mechanisms based on assigning sponsored slots according with a combination of two main parameters, the advertisers' bids and the quality of ads, both concurring to the so called ads rank.

VCG and GSP auctions are two relevant examples of implementation of such mechanisms, being utilized by internet giants like Facebook and Google. Both have pros and cons: GSP is tailored to the specific characteristics of the online ads market, while VCG in general is more computationally complex. The theory behind those algorithms finds its strong derivations from strategic games theory: this is the reason why eminent economists like Varian, Edelman and many others started to collaborate at different levels with ads auction providers to study and design increasingly powerful mechanisms.

The new challenging is represented by the modelling of the effect of externalities in online ads auctions, which occurs when the utility to a winner has a certain dependence from the other winners. Several studies have been conducted to model such effect, to build suitable algorithms for the prediction of the winner's utility in complex multi-ads scenarios and the estimation of the winner determination problem, proposing also models which are based on the preferences of the users (web visitors), in addition to those of the advertisers.

Bibliography

Colini-Baldeschi, R., Mestre, J., Schrijvers, O., & Wilkens, C. A. (2020, December). The ad types problem. In *International Conference on Web and Internet Economics* (pp. 45-58). Springer, Cham.

Edelman, B., Ostrovsky, M., & Schwarz, M. (2007). Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American economic review*, *97*(1), 242-259.

Varian, H., & Harris, C. (2013). VCG in theory and practice. *Rapport technique, Google inc*.

Ghosh, A., & Mahdian, M. (2008, April). Externalities in online advertising. In *Proceedings of the 17th international conference on World Wide Web* (pp. 161-168).

Gomes, R., Immorlica, N., & Markakis, E. (2009, December). Externalities in keyword auctions: An empirical and theoretical assessment. In *International workshop on internet and network economics* (pp. 172-183). Springer, Berlin, Heidelberg.

Ramshaw, L., & Tarjan, R. E. (2012). On minimum-cost assignments in unbalanced bipartite graphs. *HP Labs, Palo Alto, CA, USA, Tech. Rep. HPL-2012-40R1*.

Duan, R., & Pettie, S. (2014). Linear-time approximation for maximum weight matching. *Journal of the ACM (JACM)*, *61*(1), 1-23.

Varian, H. R. (2007). Position auctions. *international Journal of industrial Organization*, *25*(6), 1163-1178.

Quattrocchio, L. M. (2016). La teoria delle aste: la matematica delle procedure competitive.

Web References

wikipedia

Google web site

Facebook web site

www.engineeringenotes.com

https://performance-ppc.com

https://www.italiaonline.it

ihttps://www.hpa4u.com

https://trucchifacebook.com