

An Investigation of Pay Per Click Search Engine Advertising: Modeling the PPC Paradigm to lower Cost per Action

Alexandre Douzet

Author-Researcher

TheLadders.com

Co-Founder and Vice President, Marketing

alex@theladders.com

Abstract—This paper is aimed at increasing the general knowledge of the PPC search marketplace. The first piece will lay some background on the Pay-Per-Click (PPC) search engine market, analyzing differences in bidding and valuation between vendors, such as Google or Overture, and delivery methods like content or search listings. Next, I attempt to analyze the PPC paradigm, from the advertiser's perspective, treating keywords as a way of generating clicks, which will be used as an input into an order factory. We will look at both the refining of clicks from impressions, and the use of clicks in gaining orders. Of particular interest will be the question of whether all clicks provide an equal marginal benefit, as well as ways to improve the efficiency of refining keyword impressions into clicks. Finally, I will draw some conclusions on the nature of the PPC landscape, and offer some ideas on how to apply the learning from the model, to optimize a PPC campaign and lower Cost per Action.

I. INTRODUCTION

The pay per click (PPC) advertising model, particularly those implemented by Google, Yahoo, and other search engines are still in their relative infancy. Having been around for less than a decade, the amount of academic research that has gone into analyzing the systems at work is limited at best. In response to this, TheLadders.com has expended time, and capital in increasing the knowledge of this space.

While the PPC advertising marketplace is rather new, with little academic understanding, it bears large similarities to several very well understood business processes, complete with the models and decades of academic research that support them.

The most apparent feature of the PPC advertising marketplace is the auction component. Though the implementation differs from search engine to search engine they are all classical auctions, which have been studied significantly. The two primary auctions that take place in the PPC world are the sealed envelope or sealed bid auction and the English Auction. I will talk more about bidding strategy, in general, for each of these types of auctions, and how the auction type changes the dynamics of the auction.

From the advertiser's perspective, the PPC arena matches any finite natural resource dependant industry, be it power generation or steel production. Simplifying these industries, there are a few key points. First, there exists a finite

resource, be it iron ore, coal, or impressions on search pages. Second, the raw resource must be changed before it can be used, often with less than 100 percent efficiency. In the PPC arena, this is the conversion from impression to click. Finally, there is the manufacturing step, where a product is made. It is immaterial whether this step is melting pure iron ore to make steel, or turning curious clickers into paying customers. The same principles underlie both processes. Later I will examine each of these steps in more detail.

II. BACKGROUND INFORMATION ON THE PAY PER CLICK MARKETPLACE

Pay Per Click emerged onto the Internet in 1998. Originally introduced by GoTo.com, the idea was that sponsored versions of the search term would be auctioned off to the highest bidder. Beginning in 2000 Goto.com began selling their services to many of the largest search engines. Inktomi, yahoo, AltaVista, alltheweb.com, and others began using the paid results supplied by Goto.com [5]. By the end of 2001, under the company formally known as Goto, had a new name, Overture, and controlled nearly the entire PPC search market. It was for this reason that Yahoo purchased Overture in 2003 [5]. The notable exception was Google's pay per click model search model, which they developed in house,

In 2002, Google switched their model, in a similar move to the one made by Overture the year before, and began offering the Google paid search results on other search engines. During 2002 Google signed several major partnership deals, agreeing to serve sponsored search results for Earthlink, AskJeeves, and America Online [4]. Despite entering the game a little later, the revenue generated from Google's Pay-Per-Click program would help Google grow to a company with a market cap of well over 100 Billion dollars [5].

While smaller engines exist, and serve their own paid results, Google and Overture are, in 2005 the elephants in the living room of search engine marketing. Both companies would eventually offer a second model for companies to deliver advertising to the web in the form of content networks.

Similar to the way that Google and Overture provided smaller search engines with paid keyword listings,

eventually both would offer a way for websites to provide ads to the viewing public, based on the content of page being viewed. These content networks, though different from one another in ways that will be explored later, have are similar in the fact that they attempt to deliver context aware advertising, or in plain English, ads that directly relate to the web page being viewed at the time the ad is displayed.

III. SEARCH VS. CONTENT NETWORKS: COST AND TRAFFIC DIFFERENCES

In early 2005, a survey conducted on over 2000 adults, showed that in the general population, only one out of six people can tell the difference between paid search results, and those results that are organically delivered, resulting from the search itself [7]. This is an incredibly important bit of information. Furthermore, there is information showing that nearly 40 percent of those that know the difference between the two types of listings, actually prefer to click on paid results, saying they think the paid results are more relevant [4]. It is this population that paid search is based on. These people are looking for something in particular, and many are willing to look for their answer in paid listings, suggesting, that their clicks are out of a need to find something.

FIGURE 1: A COMPARISON BETWEEN CONTENT AND SEARCH LISTINGS

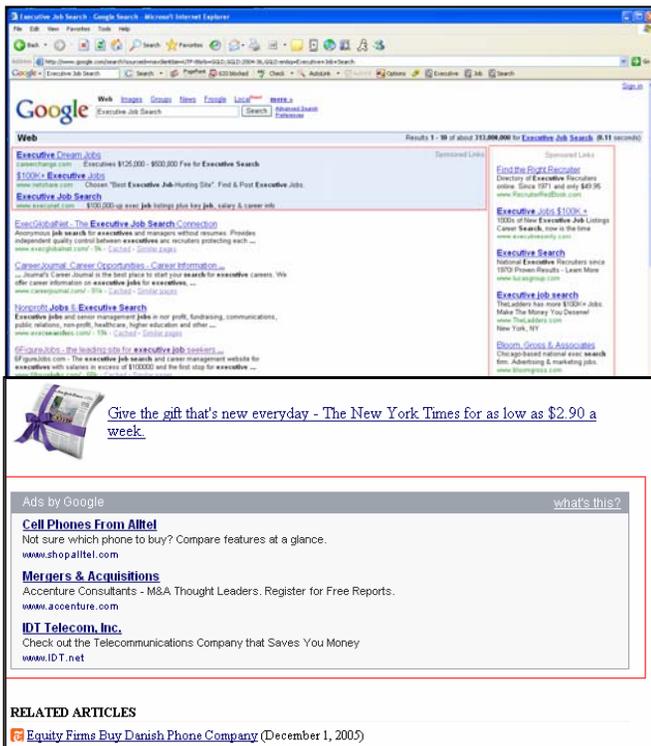


Figure 1 above, shows search listings on Google (top) and content listings on the NY Times website (bottom). The NY times ads are found at the bottom of articles.

While search listings are need based clicks, the clicks that come from the content side of the PPC networks are quite

different. Often appearing next to an article, under a blog, or adjacent to other content, these ads are not necessarily what the user is searching for. Most of the time, it is the article that the user is searching for. Because of this the clicks received on content networks are often out of curiosity. People have already found their main goal, but are interested and click. In addition, while there are sometimes as many as 20 or 30 words that will be auctioned in a search auction, more often than not, there are fewer than five spaces for ads on a content page, make it more costly to be seen on a content page. These differences are important, as they greatly affect the value of a click.

IV. AN EMPIRICAL STUDY OF POSITION AND EFFICIENCY OF CONTENT AND SEARCH:

During early 2005, TheLadders.com conducted an extensive study, aimed at gaining a better understanding of the search landscape, particularly the effect of position, and source on the volume of actions, of a portfolio of several hundred keywords. During this test, keywords were moved throughout the spectrum of costs and positions, allowing TheLadders.com to collect cost/benefit data at many price points.

A. Previous Academic Work

While test like this have been done before, they have been on a very small scale. The most notable, available work was done in February of 2003 at the London Business School [1]. The article in question looks only at the effect of bidding on the Exact Match (where the search term matches the auctioned term exactly) variety of the keyword “Real Estate.” Moreover, at the time, the content networks were just forming, not allowing the author to collect data on the different behaviors.

In addition, the fact that only one keyword is used in the testing is not sufficient to gather adequate results. There are routinely outlier words, which behave much differently than the general population, which is why TheLadders.com used thousands of keywords in its testing protocol. In addition, it has become commonplace for people to use not only exact matching, but other types of matching such as phrase or broad matching.

B. The Testing Process and Data Collection

The data presented in this study represents a subset of the data collected. While it changes the scale of the results, the shape of the curves, as well as the implications from the curves match those of the complete dataset. In this test, the action was the completion of a signup page, containing 6 fields.

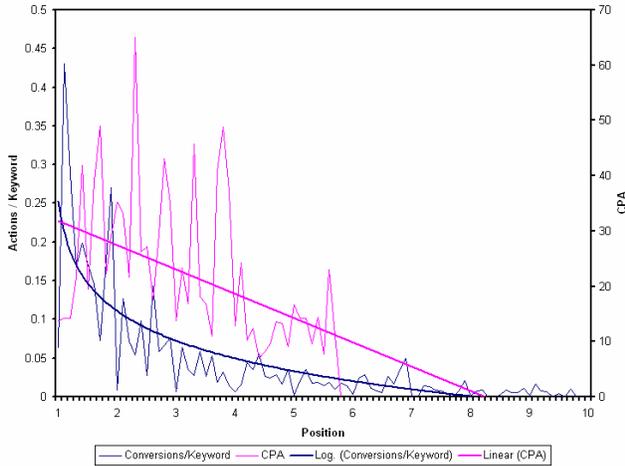
Actions were measured through tracking codes rather than through the use of pixel firings, as this is a more accurate method. Furthermore, despite the fact that data was collected on both Overture and Google campaigns, only data on Google will be on display in this paper. Again, this sampling does not effect the distribution of data, only the

quantity.

C. Data Analysis and Learnings from the Project

After the test ended, each keyword day was given an average position, based on Google’s reported information, from the Google Console. Position averages were rounded to tenths of a position. The main thing, for this paper’s purpose is the effectiveness of a click, and the quantity of a click based on position, and source.

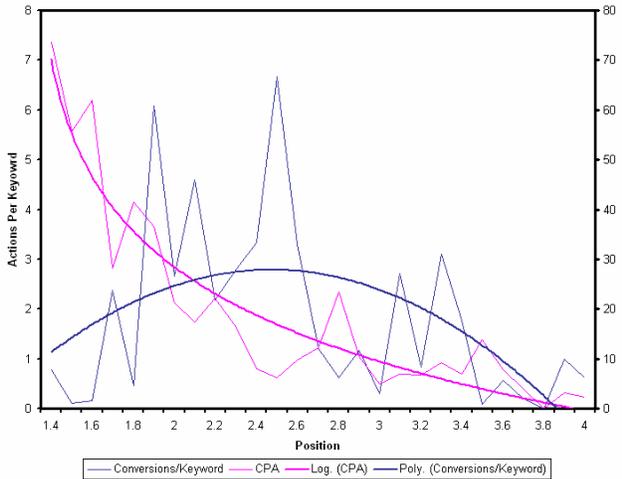
FIGURE 2: GOOGLE SEARCH MODEL



The Above Chart shows both the Cost per action, and the number of actions per keyword at each ranking level.

As you can see above, search behaves like one would expect, showing diminishing marginal returns from increased quality of goods, or in the specific case, increased position. Furthermore, it can be seen, that continuing to spend money will continue to increase the quantity of conversions, though at decreasing amounts. In addition, it can be seen that moving from one position to the next appears to be a linear function between positions one and position eight.

FIGURE 3: GOOGLE CONTENT MODEL



The Above Chart shows both the Cost per action, and the number of actions per keyword at each ranking level. Notice a very different quantity function than the search model.

Content listings however behave very differently. First of all, content listings have an increasing cost function; each additional step up to a higher position causes a greater increase in cost than the prior step. In addition, it appears as if there is, a volume maximizing position. This position, occurring between positions two and three generates the most volume of actions, though not necessarily the highest volume of clicks. In fact, one can tell that the highest volume of clicks come from the higher positioned keywords, where, the additional cost, can only come from additional clicks. These findings mean that search and content will have to be treated as separate entities, in order to maximize the value of our resources.

V. AUCTION TYPES IN PPC BIDDING

Before diving into the Convert Corp model, it is important to understand where keywords get their prices from, specifically, what results from the different types of auctions used in Google and Overture.

The first thing to realize is that there is, in fact, a new auction held, every single time a search happens. Secondly, unlike any auction you will encounter at Sotheby’s Google plays favorites. While price accounts for a large portion of the equation, Google also incorporates the likeliness that someone will click on your ad when choosing its position. While certainly something to be aware of, for the purpose of modeling the click through rate (CTR) adjustment is something that will be ignored.

A. A Second Price, Multiple Round, Sealed Bid Auction

The Google keyword model matches a model developed for timber auctions. The auctions start with a reserve price; all sealed bids must be greater than or equal to this reserve price. Each bidder places his bid with no information about what his competitors are bidding, or even how many competitors there are. In theory, with a collection of identical bidders, the same revenue is generated as in an open auction, because each bidder bids exactly their value for the auctioned good, no more, no less [8]. In practice, however, bidders are not all identical. They come with different strategies, and motives, and as such, the seller increases his revenue by between five and twenty percent by using a sealed bid auction. The variance is attributed to the competitiveness of the auction-taking place [8]. Because of the low barriers to entry and the relatively small price for a keyword, any given keyword auction is most likely to be highly competitive.

In all fairness to Google, the introduction of the second price option to the sealed bid auction adds a feature that attempts to make the auctions a more level playing field. In a second price auction, bidders do not have to pay their bids; they merely have to pay more than the next highest bid. This auction type is seen a lot online particularly in auction sites like Ebay [6]. The benefit of this is that even with the

gamesmanship that goes on with a group of heterogeneous bidders people are encouraged to bid their actual willingness to pay, since they will only have to pay the second place price.

One difference between the timber bid, and the keyword auctions is the repetitive nature of the keyword auction. Since hundreds of thousands of auctions per day take place, a bidder can gain knowledge of the pricing landscape through observing his or her results from prior auctions.

B. A second Price, Multiple Round, English Auction

The English auction is formally defined as an ascending auction, in which at all times, everyone knows how many bidders are left, and what the current bid is. Furthermore, once a bidder drops out, he is not allowed back in [3]. The English version is the type of auction one would expect to see in any movie, with an auctioneer increasing prices as bidders continue to show that they are still in the auction.

Overture's implementation of the keyword auction matches similarly to the English Auction. In the bidding interface, bidders can see exactly what bids are in place, for each keyword. Similar to the Google auction, an each auction occurs before each view of any given keyword knowledge from previous auctions can be applied to determine values of keywords.

C. Differences Between Keyword Auctions and Other Auctions

The biggest difference between the keyword auctions and the auctions in the rest of the world, is the fact that the auction is technically for first position (truly winning the auction), however, there exist many "runner up" items. If you have the fifth highest bid, you will have the fifth heterogeneous objects, each worth a different value. This is not a common thing.

The thought that you could be bidding on a Monet, and lose, and be told that you owe half of the winning bid, for a comic book, would not be acceptable in an auction house like Sotheby's. In a more standard auction house, each keyword-position would be auctioned separately.

VI. CONVERT CORP: MODELING THE PPC WORLD LIKE BIG INDUSTRY

So far, the discussion has been focused on individual pieces of the PPC marketplace, in complete isolation. Unfortunately, PPC advertising doesn't exist in isolation. To better understand the PPC world, and how it plays into business, it might be best to couch the PPC paradigm in something that has been studied for decades; production theory in a company requiring natural resources. As it will be seen, a great deal of intuition on how to deal with PPC campaigns can be learned by following this analogy through, from start to finish.

A. The Pieces of the Industrial Model

The best way to envision this marketplace is to consider something like a steel company. Before you can produce your widgets you must mine the iron ore out of the ground. Some mine sites might contain vast amounts of iron, and these sites will obviously cost more to obtain, than smaller mine sites. Following the mining process the ore must be processed and turned into steel.

The PPC market place is not much different. The mine in the PPC paradigm is the page you are attempting to show up on. The dirt in which the iron is suspended equates to the views of your ad. The iron ore itself is the clicks from the ads to your page. Finally, the steel you produce can be mapped to the actions you generate on your web page. In this model there are many places where increases in efficiency can be had. For any given mine, there is room for increased efficiency at the extraction level, taking the dirt out of the mine and the processing level. Furthermore, changing the size of the mine, can also give you an advantage, as it requires less capital to mine.

One thing that is nice about this model is the fact that resource economics have been studied for years. It is known, for instance, that you should keep producing steel, until your marginal cost matches the marginal benefit of steel in order to maximize benefit. It is also known how better capital and more efficient workers affect your business. We will use this in the coming sections.

B. Buying the Mine: Selecting the Position Where Your Keyword Appears

As it was shown before, search and content behave quite differently. As such, there will be different rules for selecting what position is ideal. In content, for instance, the decision is actually quite easy. Due to the fact that conversions actually fall at some level, the best position is the one where conversions/keyword is at a maximum. Granted this will be different for each keyword, but the value can be found, by starting at a low position, and increasing bids until a small decrease in conversions/keyword is witnessed. At this point, you will be close to the optimum "mine size" for your operation. In the data in this paper, that happened around position three, with a cost per action of just around ten dollars. A note of caution on this topic: despite this golden positions existence; if the position occurs at a price above your marginal benefit, than you should be in a lower position, where your marginal benefit equals your marginal cost from being at that position.

Search, however, behaves quite differently, acting like most normal goods. Utility increases, the more you buy, however at a diminishing rate. Seeing as this is the case, the only rule for determining what position you should be in is the marginal cost rule.

The dynamics of the PPC marketplace change quite frequently. Just like the invention of dynamite allowed for

more production out of a given mine, better creative can increase performance of the same position. Creative testing can lead to extremely unexpected results.

One particular test at TheLadders.com, investigated the use of the words “Search” and “Find.” The goal was to find out on which line of creative each word excelled on. It was found that the response rate of a creative with the word “Find” in the first line of creative had a response rate of 0.03% where as “search” in the same position had a 0.016 % response rate. With a z-score of just over 47, this data is significant at well over the 99.999% percent level of confidence. Similarly it was found, with a z-score of 51, “Search” in the second line had nearly twice the response rate of that of the opposite configuration, when looking at search terms, moving from a 0.018% to 0.035% response rate. This is not entirely surprising, as people on search engines don’t want to do more searching, they want to find things; after all they have already just done a search, why do another one?

FIGURE 3: SAMPLE KEYWORD ADVERTISEMENTS



Minor changes in creative can have a major impact. The only difference between the two ads, above, is the position of the words “Find” and “Search” yet there is a 100 % difference in the response rate of the ad.

By optimizing creative on a cost per action basis, (not just click through as is commonly done), you can significantly increase performance of your ads, essentially allowing your second position keyword to act as if it were actually running in the first position. Creative changes essentially act as artificial promotions (or demotions) in the ranking of your keyword.

C. Are All Clicks created Equal?

Just as there are different qualities of iron that can be mined, different positions can, and do generate different qualities of clicks. The processing of these clicks happens on the landing page, where the clicks are directed. It is possible, to dump all of your clicks to one page, though this would be akin to taking your iron with impurity X and refining it the same way you refine iron with impurity Y. It might work, but it certainly won’t be as efficient.

People coming from content words are in the exploration mode, while search clicks are looking for immediate answers. People searching for your company name

probably don’t need the same kind of page, as people who are looking for a fringe product you offer. Understanding this, and playing to it, can ultimately lower you cost per action, by generating more conversions out of the same number of resources.

At TheLadders.com, landing pages and flows are regularly tested through multivariate testing. So many things affect conversions such as graphics, text, form length, page layout, color, message along with many others that is difficult at best to come up with optimal landing pages without performing several rounds of multivariable testing. In addition, because of the nature of multivariate testing, several rounds of A/B testing can be accomplished all at once with smaller sample sizes. Therefore you not only save opportunity cost because of the quicker turn around but also you waist less advertising dollars allocated to testing (less time is endured with “losing” pages). These effects often offset the cost of renting the technology to use multivariate testing.

When TheLadders.com initially began using multivariate testing on one signup page (seen in appendix A), the page had a response rate of 3.47 percent. After one optimization flow made up of tree waves of testing the new signup page was put into service (seen in Appendix A). The new page saw response rate increase to just over 6.52 percent.

In resource economics, this increase in efficiency allows for either an increase in volume, by maintaining current levels of acquiring resources, or can allow a lowering of costs by reducing the amount of resources taken in, and producing the same amount of goods. So too, increases in efficiency on a landing page allow for the same benefits: more conversions or less cost (or a combination of the two).

VII. CONCLUSIONS

With some keyword accounts reaching into the hundreds of thousands of words, managers need to deal with millions upon millions of individual auctions a day. Obviously, it doesn’t make sense to manage the auctions individually; rather a law of averages is applied. Auctions of keywords are managed in groups, and updates to bids are made periodically.

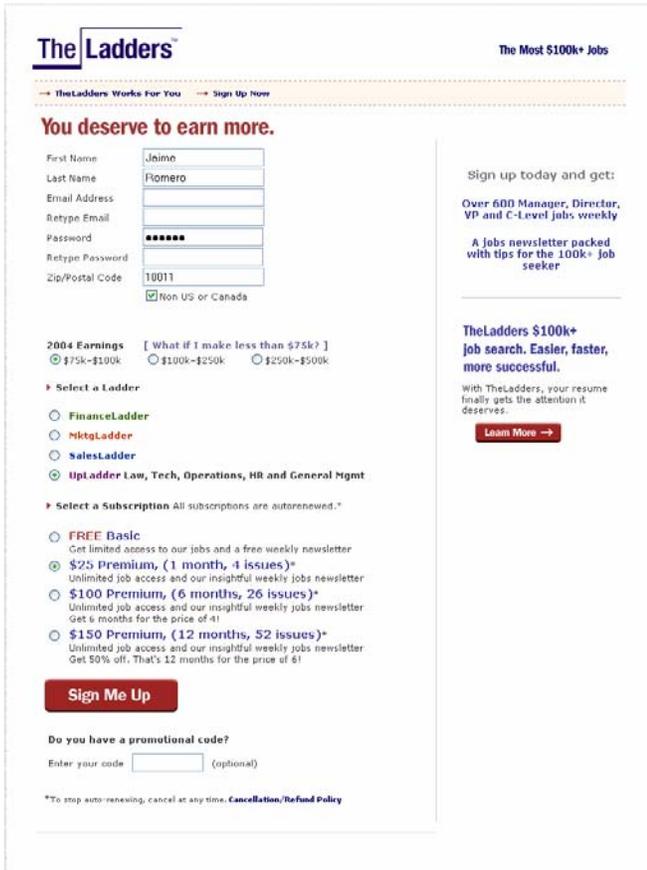
Here at TheLadders.com, knowing that the market place is based on the Internet traffic, it can change really quickly, and over the course of a quarter or two, the market, and the costs can change drastically. As such, we try to run tests to assess the look of the market place at least once or twice a year. Knowing the cost curves is imperative to being able to make good decisions about what position to be bidding for, as well as how you should be splitting up your spend between content and search.

References:

- [1] Ellam, Andrew, and Marco Ottaviani. Overture and Google: Internet Pay per Click (PPC) Advertising Auctions. Diss. London Business School, 2003. 10 Dec. 2005
<faculty.london.edu/mottaviani/PPCA.pdf>.
- [2] Levin, Jonathan, Susan Athey, and Enrique Seira. "13. A Theory of Auctions and Competitive Bidding." (2004). 10 Dec. 2005
<<http://ideas.repec.org/p/fem/femwpa/2004.142.html>>.
- [3] Milgrom, Paul R., and Robert J. Weber. "13. A Theory of Auctions and Competitive Bidding." *Econometrica* 50 (1982): 105-114.
- [4] Sullivan, Danny, and Chris Sherman. "Search Engine User Attitudes." *SearchEngineWatch*. 25 May 2004. 9 Dec. 2005
<<http://searchenginewatch.com/searchday/article.php/3357771>>.
- [5] Sullivan, Danny. "Yahoo to Buy Overture." *SearchEngineWatch*. 23 July 2003. 9 Dec. 2005
<<http://searchenginewatch.com/searchday/article.php/2234821>>.
- [6] Tucker, Joanne M. An Examination of Seller Pricing Options at Online Auctions. Diss. Shippensburg Univ., 2004.
- [7] "Users Confuse Search Results, Ads." *Wired.com*. 23 Jan. 2005. Associated Press. 9 Dec. 2005
<<http://www.wired.com/news/culture/0,1284,66374,00.html>>.
- [8] Perry, Motty, Elmar Wolfstetter, and Shmuel Zamir. "A Sealed Bid auction that Matches The English Auction" (2000). 10 Dec. 2005
<http://ideas.repec.org/p/ces/ceswps/_249.html>.
- [9] Cohen, David. Sr. Analyst, Research Assistant. *TheLadders.com*.

APPENDIX A

The Original Landing Page design is shown below
3.5% Response Rate



The revised page after multivariate testing:
6.5% Response Rate

