ASSESSING ADVERTISING BIDS IN GENERALIZED-SECONDARY AUCTIONS FOR IMPROVED SELECTION OF SEARCH ENGINE KEYWORDS

Sean L. Humpherys, West Texas A&M University
Neil R. Meredith, West Texas A&M University

ABSTRACT

As internet users query search engines, relevant advertisements accompany the search results. Advertisers bid for limited ad spots in a generalized-secondary auction employed by search engines. This paper uses publically-available data from Google to test strategies advertisers may use when deciding where and how much to bid for a keyword phrase. A predictive model of suggested bid price may provide knowledge of when a keyword phrase is too costly to bid on. Alternative strategies are discussed that focus the bidder on customer conversion rate and profitability instead of seeking to maximize click through rate. **JEL Classification:** D44, M37

INTRODUCTION

Each day billions of people around the world search for information on search engines. For example, in December 2014, 26.7 million searches were performed using the term "Xbox" and related variants using the Google search engine (Google AdWords, 2015). Search engines provide users a free service by listing relevant websites that match the keywords searched for by internet users. Search engines, such as Google, Yahoo, and Microsoft's BING, use a variety of algorithms in an attempt to increase the relevance of search results to the keyword phrase queried by the internet user. The more relevant and useful the search results, the more the prevalent the network externality becomes where people perceive the search engine as useful and continue to use it instead of others. The search results page shows a list of websites relevant to the terms used by the internet user. Some of the results are organic, which means being listed without charge because of relevance to the keyword phrase and perceived value according to the internet community, i.e., page-rank. Lines of research have been dedicated to how search engines use algorithms and how to improving algorithms to increase the relevance and value of the search engine results page to that of the internet user's query (Baeza-Yates, Hurtado, & Mendoza, 2007; Berry & Brown, 2005; Xie, Wang, & Goh, 1998).

The aim of this paper is to create a predictive model that can help online advertisers select keyword phrases to pursue and how much to pay for advertising on them. Internet forums are replete with advice and heuristics, i.e. rules of thumb, regarding the "best" strategies for selecting keyword phrases. With publically available data, the validity of various heuristics is tested and the accuracy of a predictive model for keyword selection is evaluated.

The organization of manuscript is as follows: First, a brief literature review is put forth. The second section offers a discussion of three proposed hypotheses. The next section describes the methodology. The fourth section shares results. The fifth section is a discussion of the results with recommendations. The final sections look at limitations and conclusions.

LITERATURE REVIEW

Search engines monetize their search results by placing ads to appear in or alongside organic search results. For instance, Google generates advertising revenue by auctioning 11 ad spots on search page results. As of 2015, Google places three paid ads above its organic search results. Eight ad spots are placed in a panel to the right of the organic search results (see Figure 1). Yahoo and Microsoft's Bing follows a similar presentation of ads. In 2014, Google collected $59.6 billion in advertising revenues (Google Investor Relations, 2015). Ads are typically text-only and include a one-sentence title and a brief one or two sentence description. Text-based ads mimic the appearance of organic search results. Advertisers may include a call to action in the title, such as "Try our auto dialer free", and use the description for more detail about the offer, product, or service, e.g., "Reach more live people per hour. Free trial software". Clicking the ad takes the internet user to the advertiser's website for more details and is called a *click through*. The rate at which users responds to an ad by clicking on the ad is called the *click through rate*. The click through rate is computed as the ratio of the number of times an ad appears to the number of times a prospect clicks on the ad.

Advertisers choose a venue for their ads such as a search engine or social network. Subject to a constraint of scarce advertising dollars, advertisers also face a difficult choice of what keywords to place ads upon. Moreover, the level of competition for an ongoing keyword auction combined with meeting a target click through rate can further complicate marketing decisions. The ability to place ads on relevant keywords can result in customer web traffic and can increase sales revenues (Agarwal, Hosanagar, & Smith, 2011). For example, Yan et al. (2009) demonstrate that selecting keyword phrases based on behavioral models can increase click through rate by 670%.

Many search engines use a generalized-secondary price (GSP) auction to sell online advertising spots. Building on the second-price auction from Vickrey (1961), the GSP auction for online advertising spots, as Edelman, Ostrovsky, & Schwarz (2007) explore, is an auction where the highest winning bidder pays the second highest price to place an ad at the highest position on a search results page. Advertisers only bid on keyword phrases that are relevant to their products or services. For example, a telephony company that sells a predictive auto dialer may want to place an ad on keyword phrases such as "predictive auto dialers", "auto dialers", "auto dialer software", "buy auto dialers", "compare auto dialers", or "robo dialers".

Advertisers bid for ad spots by auctioning for relevant keywords they believe customers are searching for. Rather than paying a fixed price for ad spots, search engines auction the ad spots to maximize revenue. While auctions can result in more revenue and efficient allocation of resources, where the value of a good or service is better aligned with its price, as Vickrey (1961) and Klemperer (1999) denote, the GSP auction can result in inefficiency as Edelman et al. (2007), Varian (2007), and Caragiannis, Kaklamanis, & Kanellopoulos (2015) suggest. Moreover, Edelman et al. (2007) demonstrate that GSP does not have an equilibrium in dominant strategies nor does disclosing to the search engine provider your true maximum bid, known as truth-telling, achieve equilibrium.

For GSP auctions, an advertiser submits a hidden, maximum bid price to a search engine. The maximum bid price is used to place the 11 top winning bids on a search engine results page. However, an advertiser is charged only when an internet user clicks an ad, which yields *cost per click*. Clicking on an ad takes a user to a web page with typically a single product sales offer called a landing page. Doubleclick Ad Research (2015), the display advertising arm of Google, reports an average click through rate of 0.06% for all ad formats and ad placements between 2010 and through 2014.

Predictive Models to Aid Online Advertisers Select Which Keywords to Target

The process of bidding for ad spots is largely reactive. Advertisers select potential keywords related to their product, place bids for ads, which may include bidding on thousands of keyword phrases at a time, and monitor the results. Advertisers must place a high enough maximum bid to be included in ad spots. Advertisers seek to align keyword phrases with the interest and behaviors of the ideal customer for the ad's target audience (Yan et al., 2009). Only after spending money does the advertiser acquire results useful for adjusting bid amounts and keyword selection.

To aid advertisers and web content providers, Google and other search engines make available search query data. Given a keyword phrase, Google presents a list of related keyword phrases derived from actual searches by internet users and related data. Google Support (2015) defines the data as follows:

> *Average monthly searches*: The average number of times people have searched for the exact keyword based on the location and Search Network targeting that you've selected. By default, the number of searches for the term is averaged over a 12-month period.
>
> *Competition:* The number of advertisers that showed on each keyword relative to all keywords across Google. Note that this data is specific to the location and Search Network targeting selected. In the 'Competition' column, a reviewer can see whether the competition for a keyword is low, medium, or high.
>
> *Suggested bid:* A suggested bid is calculated by taking into account the costs-per-click (CPCs) that advertisers are paying for this keyword for the location and Search Network settings selected. The amount is only a forecast, and the

actual cost-per-click may vary. (Google Support, 2015)

Table 1 presents an example of data available from Google. When a keyword phrase is searched for in Google AdWords relevant keyword phrases are also presented. An advertiser may create a simple heuristic to allocate scare advertising dollars. For example, a search volume of 3,600 per month for the term "auto dialer" seems to suggest that an ad on that keyword phrase would be seen by more potential customers than "auto dialer software", which only has 880 searches per month. Yet interestingly, the latter has a higher suggested bid ($54.71 compared to $29.46, see Table 1). This data source presents some research opportunities that might validate or discredit heuristics currently used by online advertisers or lead to a predictive model useful for advertisers.

HYPOTHESES

H1. A Relationship Exists Between Average Monthly Search Volume, Competition, and Suggested Bid.

As the first hypothesis, the relationship between average monthly search volume, suggested bid, and competition is explored. All things being equal, the higher a keyword phrase's search volume, the more opportunities exist for potential customers to view an ad. Higher keyword search volumes should be valued more by online advertisers than lower keyword search volumes. Basic supply and demand principles suggest that given the same average monthly search volume, a change in competition score should correlate with a change in bid price. Because the actual bid prices are not publically available, suggested bid price is used as a proxy given Google's claim that suggested bid price is based on historical bid prices for a particular keyword phrase.

H2. Average Monthly Search Volume, Competition, and the Interaction of the Two Can Be Used To Predict Suggested Bid Price.

The second hypothesis is aimed at evaluating bidding strategies. It is common for internet marketers to share online advertising strategies and keyword auction bid strategies. Paid training courses and commercial software claim to teach strategies on how to select keywords, how to bid in GSP auctions, and how to maximize click through rate using data from Google AdWords and other search engines. For example, in an online forum of internet marketers, one marketer provides his strategy regarding Google AdWords data.

> Quite simply, what I was taught regarding that tool was (from the keyword academy); when analyzing keywords to target, focus on the low and medium [competition] and forget about the high. So basically what I've done is to find the best keyword phrase with the lowest competition and the most monthly searches. (Denny, 2012)

Internet marketers seem to favor correlations between the Google AdWords data. For the second hypothesis, it is argued that search volume and competition score are predictors of suggested price. The higher the search volume the more people are searching for that keyword phrase. This indicates more opportunity to have an ad viewed by prospects and may results in higher competition to place ads. Higher competition should result in higher bid prices. The interaction of average search volume and competition is included in the second hypothesis to account for the assumption that for a given search volume an increase in competition should result in an increase in bid price. If correlations are present, is an optimal bidding strategy is considered if it exists.

The evaluation of the second hypothesis will validate to the community that Google's suggested bid price is consistent with the intuitive concept that desirability of a keyword phrase (as determined by a higher monthly search volume) and competition drive the expected price of an ad. A negative finding on this research question may call into question the rules of thumb internet marketers use to select keyword phrases and to set maximum bid prices for the GSP auction.

H3: The Positive Difference Between the Google's Suggested Bid Price and Our Model's Predicted Bid Price Signals an Opportunity to Purchase a Keyword at Below Market Value.

The third hypothesis considered is to model bid price and compare it to Google's suggested bid. If a substantive difference exits, the placement of where to commit advertising dollars is evaluated. A machine learning algorithm can create a *predicted bid price*, comparable to Google's suggested bid price. The predicted bid price can be based on thousands of data points from related keyword phrases (i.e., an ad group) rather than calculated based on the variability within a single keyword phrase, as is Google's suggested bid price. The argument in favor of a model that incorporates data from multiple, related keyword phrases is based on a prior observation that internet user's search behavior and intent to purchase is similar to that of other internet users when searching in related keyword phrases (Yan et al., 2009). Google combines related keyword phrases into ad groups. The data from an ad group could be sampled as a training set for the machine learning algorithm. More data has the potential to increase the power of a predictive model. In other words, statistical sampling from an ad group of thousands of related keyword phrases should result in a model that is more compelling than from sampling data in only one keyword phrase.

Subject to external validation, identifying a difference between the model's predicted bid price and Google's suggested bid price may signal an opportunity to purchase a keyword at below market value. For example, assume the predictive model reports a predictive price of $50.00 for "auto dialer software" instead of Google's suggested bid of $54.71. The predictive model suggests that an advertiser may only have to pay $50 instead of $54.71 for a given keyword (positive difference of $4.71). A bidder using the predictive model can place a bid of $54.71 to win an auction but knows from the predictive model that a true price of $50 will likely result. With this knowledge, a bidder can set an advertising budget to either bid on more keywords than would have previously been possible or to reduce the advertising budget. Other bidders without the predictive model may evaluate Google's $54.71 as the expected cost per click and choose an alternative keyword phrase to bid on or reduce the number of bids placed on alternative keyword phrases.

Three assumptions must be met to support the position this paper takes. First, it is assumed that suggested bid price is a good approximation of the actual bid price, as implied by Google's public definition of the term. Second, internet user behavior is similar across related keyword phrases (Yan et al., 2009). Third, that sampling of data from multiple, related keyword phrases (an ad group) produces a more accurate model than sampling from one keyword phrase. Validating the model with a goodness of fit measure can demonstrate the model meets the assumptions and support our argument of buying opportunities where differences exist.

A human might be good at managing several hundred word auctions, but when several thousands of keywords need monitoring, the task becomes more difficult. Building software that includes a predictive model or algorithm can aid the internet marketer.

METHOD

To evaluate the hypotheses, data is collected from Google AdWords. Eight keyword phrases are randomly selected as phrases from a variety of industries to start the inquiry (see Table 2) and Google returned a set of related keyword phrases (N = 5,269). Pearson correlations can then be used to investigate relationships between these variables to evaluate the first hypothesis.

Modeling Suggested Bid Price

Exploring the second and third hypotheses requires an accurate predictive model of suggested bid price. Suggested bid price is used as the dependent variable. For the independent variables, average monthly search volume, competition score, and the interaction of the two are employed. When a user downloads the data from Google AdWords, the categorical data (low, medium, high) is presented as a continuous variable 0.00 through 1.00.

Machine learning is a study of patterns in data to create models useful in creating predictions using computational algorithms with the ability to learn from the data. Supervised machine learning is a branch of machine learning where a learning algorithm is presented with a training dataset that includes relevant data (i.e., independent variables) and the correct answers (i.e., dependent variable). Supervised refers to learning by giving the algorithm the correct answer to learn from. From the training data, the learning algorithm creates a model using a variety of statistical or alternative modeling algorithms. For this study, a learning algorithm that outputs a linear regression model is chosen. Next the learning algorithm evaluates the accuracy of the model by making predictions on a subset of the data called the testing dataset. Comparing the predicted value to the known value in the testing dataset generates accuracy measures. As a common practice, a 10-fold cross-fold validation technique is used where 90% of the data is used for training and 10% is used for testing. The algorithm is run 10 times, randomly selecting a 90% training set and a 10% testing set. Results are averaged to give an idea of accuracy across the 10 validation runs. Averaging also prevents a model from overfitting the data and can give a more accurate description of what the model might do when presented with a future data.

Machine learning algorithms have several advantages to alternative statistical techniques. They can handle large sample sizes and tens of thousands of variables.

They can compute useful accuracy measures in addition to the typical coefficient of determination (R-Squared). Learning algorithms can reduce the chance of overfitting. Overfitting is where the model exactly predicts each data point but is too exact to accurately fit other samples from the sample population.

For this study, WEKA (Hall et al., 2009), an open-source software with numerous, built-in machine learning algorithms, is used. WEKA is freely available at http://www.cs.waikato.ac.nz/ml/weka/. Within WEKA, a learning algorithm was selected that outputs a linear regression model and uses feature reduction to automatically remove variables that do not improve the accuracy of the model. Other learning algorithms were attempted that either made accuracy worse or resulted in the same accuracies.

RESULTS

To evaluate the hypotheses, 5,269 keyword phrases were gathered from Google AdWords. Basic descriptive statistics are viewable in Table 2. Average monthly search volumes can reach over 100,000 per month per keyword phrase. Keyword phrases with the highest monthly search volume in this dataset are “iPhone” and “iPhone 5” at 3,350,000 and 2,740,000 queries per month, respectively. Of the top keyword phrase with over 100,000 searches per month ($n = 61$), 53 come from the ad groups of iPhone, smart phone, women clothes, and shoes.

Regarding suggested bid price, the greater the value of a product (e.g., iPhone, telephony software, etc.), the greater the maximum bid. Suggested bid is a cost-per-click measure. Companies advertising for the keyword phrase “phone system business” have paid up to $157.11 *for a single click* on an ad, assuming Google’s suggested bid represents at least an approximation of historical bid data. Of the keyword phrases in our sample, 91 suggested bids over $50.00 are in the technology category of telephony or auto dialer. One exception is the keywords “Canadian small business accounting software” with a suggested bid of $97.40.

First Hypothesis Results: Relationship Between Variables

The first hypothesis of this paper tests whether or not there exists a relationship between average monthly search volume, competition, and suggested bid. It is expected that online marketers desire keywords with higher traffic because an ad could be viewed by more prospects. A higher search volume should correlate with a higher competition score or higher suggested bid. Table 3 presents the Pearson correlations between average monthly search volume, competition, and suggested bid. All tests of significance use an alpha of .05.

Even though two coefficients have statistical significance, the correlations are very small. The correlation between average monthly search volume and suggested bid is .005 with $R = -.039$. The observed statistical significance may simply be a random chance event in a very large sample ($n = 5,269$). A scatter plot of average monthly search volume and suggested bid shows the skewed distribution of the data with many extreme outliers (see Figure 2). Fitting a regression line to the data results in an R-Squared of .001.

Filtering the data to keywords with less than 2,000 average monthly searchers strengthens the Pearson correlation to -.189 ($n = 3,539$, $p < .001$) for suggested bid to average monthly search volume and .184 ($n = 3,539$, $p < .0001$) for suggested bid to competition. Competition to average monthly search volume has a Pearson correlation of -.109 ($n = 3,539$, $p < .0001$). This subset shows that for queries under 2,000 per month, a stronger relationship exists than in the larger dataset, although it remains weak overall.

The cutoff of 2,000 queries is arbitrary but was decided upon after comparing several scatter plots and zooming into the concentration of plots near the axis (see Figure 3). Fitting a regression line to the subset of data results in an R-Squared of .036, which does not account for much variation.

Figure 3 shows a quantization of data around certain whole numbers. It is assumed that average monthly search volume should be a continuous integer with an equal probability of settling on any integral value. The structure in Figure 3 appears unnatural. It seems possible that Google is not reporting integer data but has some algorithm to quantize the data before reporting it to the public. Regardless of the reason, the quantization may lose information inherent in the data and may help contribute to low correlations.

Regarding the relationship between competition and suggested bid, the correlation is .156 with statistical significance less than 0.001 ($n = 5,269$). Figure 4 shows a scatter plot of the data. There does not seem to be a significant relationship. In essence, the evidence appears to reject the first hypothesis that there is a relationship between average monthly search volume, competition, and suggested bid.

Second Hypothesis Results: Predicting Suggested Bid Price

For the second hypothesis, it is contended that average monthly search volume and competition are predictors of suggested bid price. Because the average of the suggested bid price varies greatly by ad group, a linear regression models is created for each ad group separately. This decision is made because a product in one ad group has little potential for being sold when advertised to a different market (e.g., shoes being sold in the auto dialer market).

An interaction variable is added by multiplying average monthly search volume with competition. The interaction is generated because it is argued that higher search volumes on a keyword should be more desirable to advertisers and, therefore, may increase competition for those keywords and result higher bid prices. Table 4 presents the results of the linear regression model.

The linear regression machine learning algorithm in WEKA first uses a M5 feature selection method to prune variables that may weaken model accuracy. Specifically, the M5 method steps through the variables removing the variable with the smallest standardized coefficient until no improvement is the estimated error given by the Akaike information criterion (Akaike, 1973; Hall et al., 2009). Average monthly search volume and the interaction term are automatically pruned in all but one ad group. This is consistent with the previous observation of a weak correlation among these variables. Evidence appears to be lacking to support the second hypothesis that average monthly search volume and competition are predictors of suggested bid.

Evaluating Model Accuracy

In addition to the R-Squared value, a machine learning algorithm can provide other measures that reflect on accuracy of the model. The models used in this study predict suggested bid price (labeled “predicted bid price”). The *mean absolute error* gives an average of how different the predictive bid prices are from the actual suggested bid prices. The values are absolute to account for the predictive bids being greater or less than the suggested bid prices. The greater the mean absolute error, the greater the divergence. For the ad group iPhone, the mean absolute error is $1.25, indicating that on average the predicted value is plus or minus $1.25 different from the suggested bid price (see Table 5).

Having a mean absolute error difference can indicate two conditions. First, the model may not represent the data well. Second, the model may be a good fit and the difference indicates exploitable discrepancies in the bid price. A lower predicted price may indicate a keyword phrase that is undervalued by the market given the search volume and competition. The discrepancies may be an opportunity to bid for a higher-than-average search volume in a lower-than-average competitive auction. If this second scenario is accurate, then an online advertiser could concentrate the ad budget on keyword phrases that are predicted to be lower than the suggested bid price. As a result, an advertiser could reduce ad expenses or expand the number of keyword phrases which can be bid on for advertising.

To evaluate the goodness of fit for the model, the *relative absolute error* is reported (see Table 5). A relative absolute error is calculated by comparing the mean absolute error for the linear regression to an alternative model’s error, which in Weka is the mean of the suggested bid price. It is then considered whether the linear regression model predicts the suggested bid price better than by taking the mean of the suggested bid price. If it is better, then the linear regression model may be useful. If it is not, then the linear regression model is not useful. Parsimony would favor the mean model if both were equally accurate.

The lower the relative absolute error, the better the linear regression model predicts the suggested bid price compared to the mean model. Hypothesis testing desires a lower percentage in relative absolute error. A relative absolute error of exactly 100% means both the linear regression model and the mean model produce equally accurate predictions. A value greater than 100% indicates the linear regression model does worse than simply taking the mean of suggested bid price. For the ad group iPhone, the linear regression model’s relative absolute error is 101.8% — worse than simply using the mean to predict the suggested bid price.

Third Hypothesis Results: Comparing Google’s Suggested Bid Price and Our Model’s Predicted Bid Price for Buying Opportunities

The third hypothesis is to use the positive difference between the Google’s suggested bid price and the model’s predicted bid price to signal an opportunity to purchase a keyword at below market value, i.e., those keywords with higher than average search volume for lower than average competition. Because of the low accuracy of the model a difference between suggested bid price and the predicted bid price is meaningless. Given that there is little correlation between the variables, it is understandable why the predicted values are poor. If the model’s prediction accuracy had been better, the third hypothesis could have been favored in addition of buying keywords that appeared undervalued. However, the evidence does not support the third hypothesis. Instead, online advertisers should be wary of heuristics suggesting a correlation between the three variables provided by Google: average monthly search volume, competition, and suggested bid price.

DISCUSSION

Online advertisers compete in a generalized-secondary auction to place ads on search engine result pages. Each keyword phrase becomes a market for an auction. Online advertisers may participate in thousands of keyword phrase auctions at one time. In a generalized-secondary auction the winning bid receives the top ad spot but only pays the second place bid price. The ad for the second place winner is displayed in the second ad spot and only pays the price of the third place bid price, etc.

One strategy available for advertisers is to bid on keyword phrases, measure the results (e.g., click through rate and actual bid price), exit keyword auctions that don't reach a desired click through rate, and enter new keyword phrase auctions. This is a reactive strategy where money must be spent apriori to gain the information needed to make auction decisions. It can result in a waste of money by bidding on overpriced or highly competitive keyword phrases.

This research suggests a predictive strategy based on supervised machine learning algorithms to give an online advertiser a temporary competitive advantage compared to someone simply using a reactive strategy. This study does not suggest that an online advertiser abandon monitoring click through rate and bid prices, but with an accurate model, a predictive strategy may result in lower advertising budgets or a faster targeting of valuable keyword phrases.

The data available to all online advertisers on Google AdWords includes related keyword phrases, average monthly search volume, competition score measuring how completive a keyword phrases is bid on, and a suggested bid price. The first hypothesis tested is that there is a relationship between search volume, competition, and suggested bid. Advertisers value keyword phrases with higher search volumes because the keywords provide more opportunity for prospects to view their ad. Advertisers are assumed to desire a lower competition score because less competition means fewer people bidding for the limited 11 ad spots on a search results page. Advertisers also desire a lower bid price so that their advertising costs are lower. Putting these desires together, online advertisers desire high search volume, lower competition, and low bid price. A keyword phrase that is more desirable with higher competition should suggest a higher bid price and this study attempts to verify this relationship. These possible relationships are evaluated using data on 5,269 keyword phrases collected from Google AdWords.

No correlation is the consistent finding obtained between suggested bid price and average monthly search volume and competition score. Some strange possibilities emerge from the findings. For instance, online advertisers may not value higher search volumes or they are not willing to participate in auctions with perceived high bid prices. Additionally, either the auctions are erratic or the calculated competition score is more complicated than the definition provided by Google. The findings of no correlations leads this study to the following recommendations.

Recommendation #1: Avoid Bidding Heuristics Reliant on Relationships Between Search Volume, Competition, and Suggested Bid.

The ability to make sense of data and find patterns in numbers is an important part of human intelligence (Kotovsky & Simon, 1973; Simon & Kotovsky, 1963). Humans have a propensity to find patterns in data, even when one does not exist (Whitson & Galinsky, 2008). This ability is built into the human mind. The majority of advertisers are not professionally managed, i.e., lower budgets, fewer resources, and potentially less experience advertisers (Richardson, Dominowska, & Ragno, 2007). Earlier in this paper an example is presented of an online advertiser attempting to describe a bidding heuristic that implies a correlation between the Google AdWords data. Internet forums are replete with such examples of speculative bidding advice. This study supports ignoring the Google AdWords data because of the lack of correlation between the variables we consider. Trying to force correlations where they don't exist can lead to poor decision making.

Recommendation #2: Set Maximum Bid at Near Zero and Slowly Raise the Bid Over Time Until Desired Ad Position is Reached

Another major finding of this research is that a linear regression model using average monthly search volume, competition, and the interaction of the two does not predict suggested bid price. A machine learning algorithm based on linear regression is trained and tested using 5,269 data points. Despite this training, the model performs no better than simply taking the average of all the suggested bid prices. This result is understandable given the lack of correlation described earlier. If the model had performed effectively, then the online advertiser may have a tool to more accurately identify undervalued bid prices (i.e., higher than average search volume with lower than average competition) and focus auction efforts on those opportunities. Because the model in this study performs poorly, it suggests that a predictive strategy remains unattainable, at least with the publically available data from Google AdWords. Despite the results, a reactive strategy still remains an option.

Practitioners can at least avoid using flawed heuristics based on correlations between the Google AdWords data. Edelman et al. (2007) provides a reactive strategy that is lent support by the findings of this paper given no correlation and no reliable predictive modeling of suggested bid price. Edelman et al. (2007) recommend starting with a near zero maximum bid for a desired keyword phrase and increasing the bid over time until the desired ad position on the search results page is achieved. At first, the low maximum bid will fail to win any ad spot on the search engine results page, but over time, and as the maximum bid price is slowly increased, the ad will appear in the last ad spot (11^{th} on Google's search engine). This results in the lowest true bid price that achieves the aim of being on the search results page. Once an ad is observable on the search results page for a desired keyword phrase, the click through rate can be monitored. This process can be repeated with all the keywords in an ad group. While this is a reactive strategy and contrary to the original goal of this paper, Edelman el al.'s (2007) strategy prevents making unfounded decisions using heuristics reliant on non-existing relationships among the available data.

This paper concludes that it is best for the advertiser to ignore the data provided by Google AdWords when selecting keyword phrases for advertising. Attempting to mentally force a correlation between the data may lead to incorrect decisions and a waste of money or time. The empirical data show there is no reliable relationship. Instead, it is recommended that Google's Keyword Planner be used to find related keyword phrases in an ad group, then follow Edelman et al.'s (2007) strategy of starting with one cent maximum bid price. Each day raise the maximum bid price while monitoring for the appearance of and rank placement of the ad. While Google does not permit robot bidding, monitoring for the appearance and placement of an ad on a search engine results page can easily be automated with inexpensive software.

Recommendation #3: When to Stop Raising a Bid for a Keyword Phrase

The Edelman et al. (2007) strategy prompts the question of when to stop raising the maximum bid. For example, should the online advertiser stop when the ad appears in the 11th spot or keep raising the bid until the ad appears in the 1st spot? This paper's findings show that the variance within a keyword ad groups for suggested bid price is very large. There may be a diminishing return on ad position, where the cost of increasing the ad spot by one spot is greater than the return on investment. Choosing the wrong stopping point could minimize profit. While this is a great future avenue for research, existing research suggests some possibilities regarding when to stop raising a bid.

Eye tracking research evaluates search engine results pages to discover where a user's eye is directed. Richardson et al. (2007) report that ad position can result in a 90% decrease in click through rate from the top position to the bottom position with the top ad position receiving the most visual attention. The findings indicate a mental concentration on the top of the page with reduced attention on the lower part of the page. Furthermore, the results display a concentration of eye pauses in the left of the screen, where organic searches are found, and a brief, limited viewing of the right side advertisements. Within the right ad panel, the top ads are viewed more than the bottom ads. These observations suggest that the top three advertisements above the organic searches will be viewed longer by the user than the other ad spots. The second desirable spot would be the top ad in the right ad panel, as the ads at the top right panel appear to be viewed more by users than the ads at the bottom. One limitation of the Richardson et al.'s (2007) research is that it is a simulation asking users to pretend to search for keywords and may not correspond to an actual user with intent to purchase something online.

Instead of focusing only on click through rate, Agarwal et al. (2011) use profitability as a measure to determine where to place the ad. Similar to others, they note a decrease in the click through rate as the ad position decreases on the page (top to bottom). But they also observe a non-linear change in bid price with higher ad spots, which decreases profit at a greater rate than the decrease in click through rate. They demonstrate that ads in the top position do not always maximize revenue. For firms that seek to maximize profits, they recommend a middle ad position as the *profit maximizing position*. One limitation of the study is that it only looks at one industry. Meanwhile, this paper supports Agarwal et al.'s (2011) recommendation by observing the large variability in suggested bid prices across 5,269 keyword phrases in multiple industries. Agarwal et al. (2011) also observe that longer keyword phrases have a higher conversion rate than shorter keyword phrases. The authors conclude that the intent to purchase is greater for users searching with longer keyword phrases.

When combining the recommendations from Edelman et al. (2007) and Agarwal et al. (2011), advertising firms have a strategy to set bid price in an auction and know when to stop raising a bid. The recommendation is as follows:

1. Choose an ad group of related keyword phrases to target relevant to the product or service the firm provides.

2. Set a very low maximum bid price for each keyword phrase in the ad group

3. Regularly increase the maximum bid price. Use automated software to monitor for whether the ad appears on the search engine results page and in what position the ad appears.

4. Record the cost per click and conversion-to-customer rate. Use a landing page (i.e., a web page with a single offer and clear call to action) to not conflate results with multiple products. Both Google Analytics and Google AdWords have tools to monitor these data by placing JavaScript code on the advertiser's landing page.

5. Increase the maximum bid price until the profit per keyword phrase is maximized. According to Agarwal et al. (2011) the middle ad position is the expected profit-maximizing position and the top positions are not.

6. Repeat this process with other keyword phrases.

Software can be created to aid the human in managing thousands of keyword auctions using this combined strategy. Once a revenue maximizing position is found, the text in the ad can be adjusted to increase click through rate and conversion rate. The ads on search engines are text based and the word choice in the ad can be tested with an AB testing strategy. Graepel et al. (2010) present a software model that can evaluate the text in an ad to increase click through rate and empirically demonstrate that ads with longer titles and those that include a company phone number increase click through rate. Future research can evaluate if the increase in click through rate from word choices in the ad also relate to increase customer conversion rates.

Economic Implications

From an economic perspective, the findings from Edelman et al. (2007) and Agarwal et al.'s study (2011) suggest that online auction for ad placement is an unusual market because the middle position, not the most desirable top position, is the revenue maximizing position. The inability to predict bid price using a competition score and search volume adds additional support that online auctions of keywords may not work efficiently. Online auctioneers may be fighting for top ad positions by focusing on higher click through rates instead of higher profitability. Future research could look at alternative auction structures that better align bid prices with expected profitability. One possibility is the *click per action* model where search engines charge for a result, such as a customer conversion, filling out an interest form, or signing up for a free product trial (Nazerzadeh, Saberi, & Vohra, 2008).

In the interim, new software tools could be created to help monitor ad placement and conversation rates per keyword. Such software may help bidders set maximum bid prices and set goals for desired ad placement that maximized profits rather than maximize click through rates. Setting maximum bid prices in alignment with profit has the potential to make the keyword auction markets more efficient overall and raise revenue for the search engine.

LIMITATION

One limitation of this study is that the data comes only from Google. Yahoo and Bing provide similar data related to their search engines. It is common for research to be limited to one search engine (e.g., Agarwal et al., 2011; Graepel et al., 2010; Richardson et al., 2007; Yan et al., 2009) to investigate the unique auction characteristics. Future research can replicate these findings using other sources of publically available data. The other limitation is that suggested bid prices are investigated instead of actual bid prices. Google does not release actual bid prices. Partnering with an online advertiser could yield actual bid prices, but would be limited to a smaller dataset from only one industry for one type of product. This research collects 5,269 data points from multiple ad groups corresponding to multiple industries. Collecting actual bid prices for 5,000 data points could potentially cost hundreds of thousands of dollars in cost per clicks. The advantage of the dataset in this study is that the data mirrors what is publically available to all online bidders before they participate in the auction. Therefore, it can investigate what might or might not influence an auction. Actual bid prices are proprietary to one firm and not available to all bidders. There are tradeoffs in every methodology decision.

CONCLUSION

Online advertisers bid for ad placement on major search engines. Advertisers bid for limited ad spots in a generalized-secondary auction employed by search engines. Advertisers can use data provided by the search engines (e.g., keyword phrases, search volume, competition, and suggested bid) to determine which keyword phrases to bid on and how much to bid. Contrary to expectations, this study does not find any correlation between search volume, competition, and suggested bid. Nor could a machine learning algorithm create an accurate model predicting suggested bid price, which could have been useful to online advertisers in identifying undervalued keywords. The recommendation from this study is that online advertisers should not rely on heuristics to depend on correlations between these variables. Instead, online advertisers should use a strategy that set a low maximum bid price and slowly raises the maximum bid price until the desired ad spot is reached (Edelman et al., 2007). Furthermore, bidders should not seek the top ad position, which is generally the position that maximizes click through rate and presumes to be the most desirable position based on eye tracking findings. Instead, research suggests that the ad spot that maximizes profit is somewhere in the middle of the ad lists (Agarwal et al., 2011). Combining these two strategies should give the online advertiser a better chance of achieving increased sales.

REFERENCES

Agarwal, A., Hosanagar, K., & Smith, M. (2011). Location, Location: An Analysis of Profitability of Position in Online Advertising Markets. *Journal of Marketing Research*: *December 2011, 48*(6), 1057–1073.

Akaike, H. (1973). Information theory and an extension of the maximum likelihood principle. 2nd International Symposium on Information Theory, 267–281.

Baeza-Yates, R., Hurtado, C., & Mendoza, M. (2007). Improving search engines by query clustering. Journal of the American Society for Information Science and Technology, 58(12). http://doi.org/10.1002/asi.20627

Berry, M. W., & Brown, M. (2005). Understanding Search Engines: Mathematical Modeling and Text Retrieval (Second Edition). 3600 University City Science Center, Philadelphia PA 19104: Society for Industrial and Applied Mathematics.

Caragiannis, I., Kaklamanis, C., & Kanellopoulos, K. (2015). Bounding the inefficiency of outcomes in generalized second price auctions. Journal of Economic Theory, 156, 343–388.

Denny. (2012, Originally posted). How is competition worked out in Google AdWords Tool [EP forum]. Retrieved from http://experienced-people.net/forums/showthread.php/5736-How-is-competition-worked-out-in-the-Google-Adwords-Tool

Doubleclick Ad Research. (2015, April 21). Display advertising clickthrough rates. Retrieved from http://www.smartinsights.com/internet-advertising/internet-advertising-analytics/display-advertising-clickthrough-rates/

Edelman, B., Ostrovsky, M., & Schwarz, M. (2007). Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords. The American Economic Review, 97.

Google AdWords. (2015). Google AdWords Keyword Planner [in Google]. Retrieved from https://adwords.google.com/KeywordPlanner#search

Google Investor Relations. (2015). 2015 Financial Tables [in Google]. Retrieved from https://investor.google.com/financial/tables.html

Google Support. (2015). Understanding your Keyword Planner statistics and traffic forecasts [in Google]. Retrieved from https://support.google.com/adwords/answer/3022575?ctx=tltp

Graepel, T., Borchert, J. Q., & Herbrich, R. (2010). Web-scale bayesian click-through rate prediction for sponsored search advertising in Microsoft's Bing search engine. Proceedings of the 27th International Conference on Machine Learning (ICML-10), 13–20.

Hall, M., Frank, E., Holmes, G., Pfahringer, B., Reutemann, P., & Witten, I. (2009). The WEKA data mining software: An update. ACM SIGKDD Explorations, 11.

Klemperer, P. (1999). Auction theory: A guide to the literature. Journal of Economic Surveys, 13(3), 227–286.

Kotovsky, K., & Simon, H. (1973). Empirical tests of a theory of human acquisition of concepts for sequential patterns. Cognitive Psychology, 4(3), 399–424.

Nazerzadeh, H., Saberi, A., & Vohra, R. (2008). Dynamic cost-per-action mechanisms and applications to online advertising. Proceedings of the 17th International

Conference on World Wide Web, 179–188.
Richardson, M., Dominowska, E., & Ragno, R. (2007). Predicting Clicks: Estimating the Click-Through Rate for New Ads. Proceeding WWW ’07 Proceedings of the 16th International Conference on World Wide Web, 521–530.
Simon, H., & Kotovsky, K. (1963). Human acquisition of concepts for sequential patterns. Psychological Review, 70(6), 534.
Varian, H. (2007). Position auctions. International Journal of Industrial Organization, 25(6), 1163–1178.
Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. The Journal of Finance, 16(1), 8–37.
Whitson, J. A., & Galinsky, A. D. (2008). Lacking Control Increases Illusory Pattern Perception. Science, 322(5898), 115–117. http://doi.org/10.1126/science.1159845
Xie, M., Wang, H., & Goh, T. N. (1998). Quality dimensions of Internet search engines. Journal of Information Science, 24(5), 365–372. http://doi.org/10.1177/016555159802400509
Yan, J., Liu, N., Wang, G., Zhang, W., Jiang, Y., & Chen, Z. (2009). How much can behavioral targeting help online advertising? Proceeding WWW ’09 Proceedings of the 18th International Conference on World Wide Web, 261–270.

FIGURE 1. SEARCH ENGINE'S AD PLACEMENT

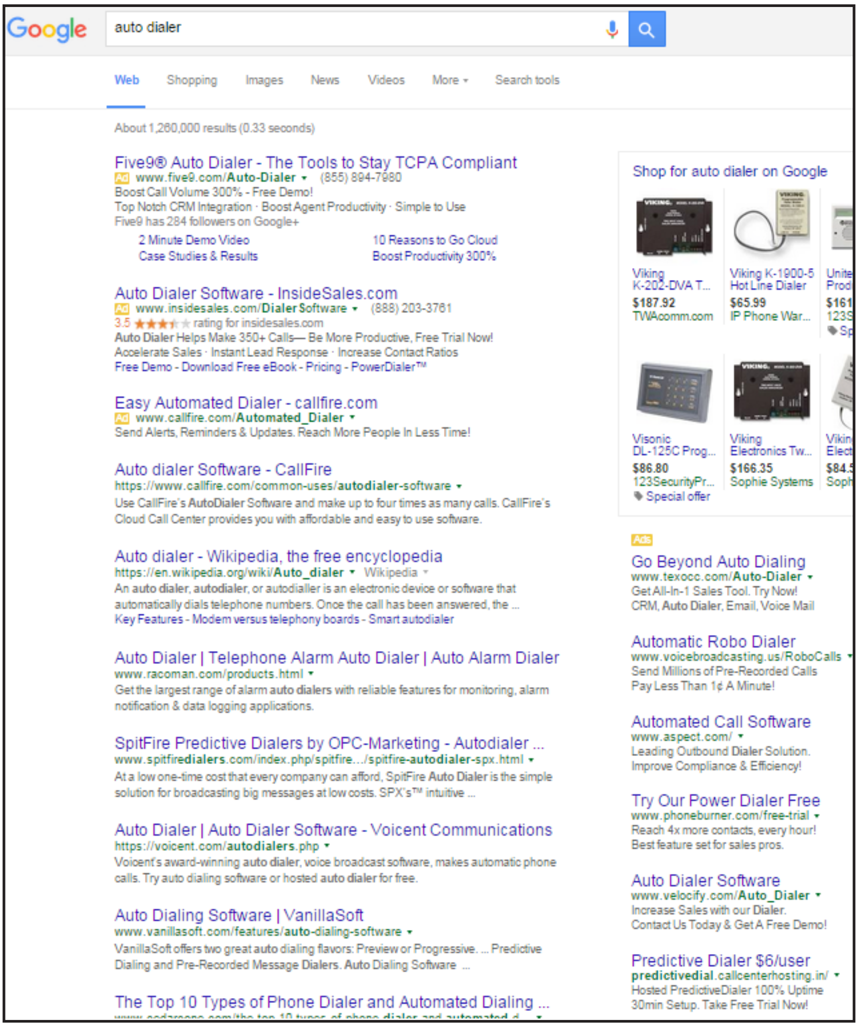

Note. Search engines place paid ads above organic search results. For Google, there are 11 ad slots available for paid advertising (3 at the top left and 8 in the right panel).

TABLE 1. DATA FOR THE KEYWORD PHRASE "AUTO DIALER"

Keyword Phrase	Avg Monthly Searches	Competition (scale 0.0 - 1.0*)	Suggested Bid (USD)
auto dialer	3,600	0.99	$29.46
auto dialer software	880	0.93	$54.71
auto dialer system	50	0.86	$9.22
auto dialer systems	40	1.00	$74.53
auto dialers	260	0.92	$48.11
call center software	4,400	0.94	$66.33
free auto dialer	170	0.84	$12.85
predictive dialer	2,400	0.82	$68.80
voip auto dialer	70	0.98	$34.97

Note. Google AdWords returned 401 related keyword phrases to "auto dialer". * The higher the number the more competitive the auction.

TABLE 2. DESCRIPTIVE STATISTICS ON KEYWORD PHRASES CATEGORIZED BY AD GROUP

Ad Group	# of Keyword Phrases	Suggested Bid (USD) Mean	Suggested Bid (USD) Max	Suggested Bid (USD) SD	Avg Monthly Searches Mean	Competition Mean
accounting software	782	11.06	97.4	7.98	853	0.83
auto dialer	343	18.36	119.66	22.98	618	0.74
corn	755	1.85	21.07	2.76	3,141	0.40
iPhone	792	1.58	59.69	3.09	22,459	0.63
shoes	631	0.92	3.19	0.41	15,467	0.87
smart phone	587	2.00	42.95	3.13	1,2072	0.72
telephony	583	18.06	157.11	24.42	2,460	0.78
women clothes	796	1.01	4.15	0.45	14,847	0.89

Note. N = 5,269 keywords. Ad group is a set of related keyword phrases. Each keyword phrase is a separate auction. Minimum suggested bid within an ad group was one cent. Data came from Google AdWords.

TABLE 3. PEARSON CORRELATION BETWEEN VARIABLES

Variables	Avg Monthly Searches	Competition	Suggested Bid
Avg Monthly Searches	1.00	-	-
Competition	-.022	1.00	-
Suggested Bid	-.039**	.156**	1.00

Note. N = 5,269 keyword phrases. ** Significant at 0.01 (2-tailed).

FIGURE 2. RELATIONSHIP BETWEEN AVERAGE MONTHLY SEARCH VOLUME AND SUGGESTED BID (N = 5,269; R^2 = .001)

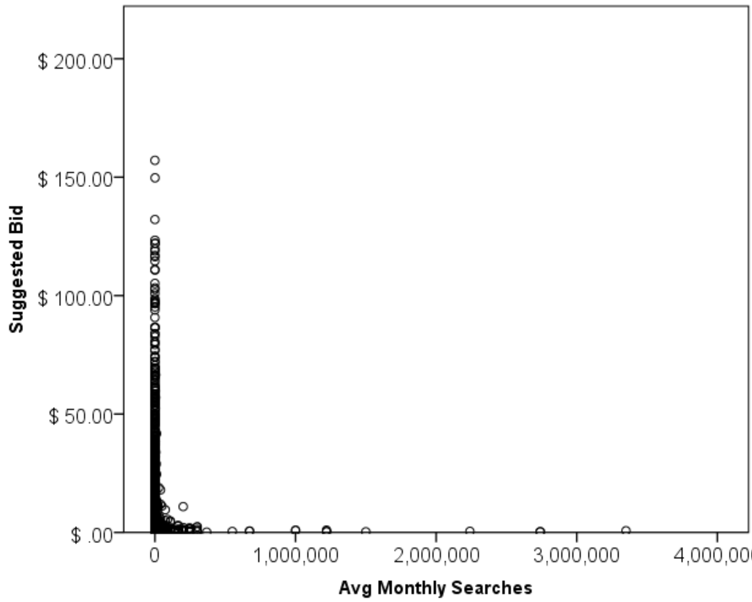

FIGURE 3. RELATIONSHIP BETWEEN AVERAGE MONTHLY SEARCH VOLUME AND SUGGESTED BID, FILTERED FOR THOSE KEYWORDS WITH LESS THAN 2,000 SEARCHERS PER MONTH (N = 3,539; R^2 = .04).

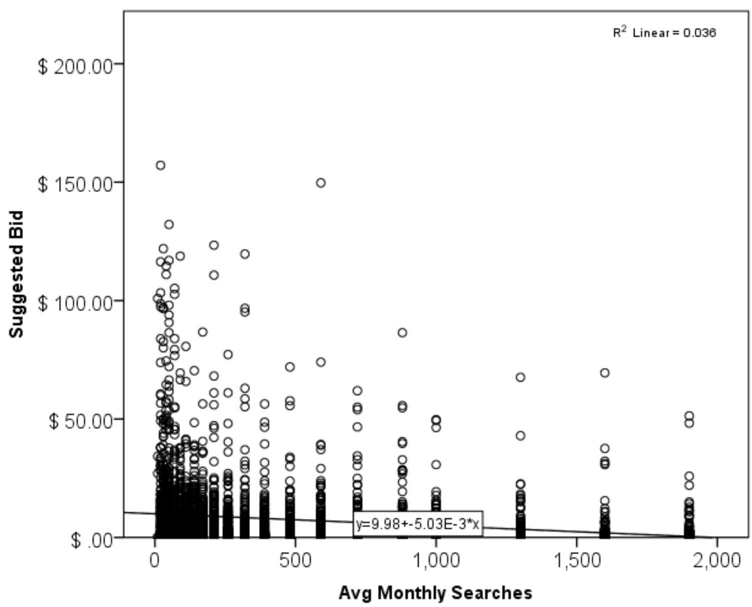

FIGURE 4. RELATIONSHIP BETWEEN COMPETITION AND SUGGESTED BID (N = 5,269).

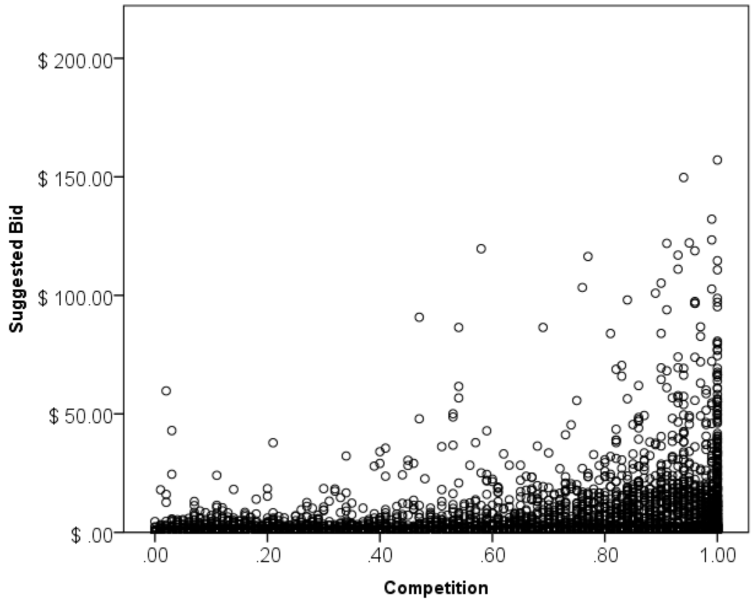

TABLE 4. LINEAR REGRESSION MODELS TO PREDICT SUGGESTED BID PRICE (USD)

Ad Group	Intercept	B for Avg Search Volume	B for Competition	B for Interaction
Accounting Software	0.27	0.000	13.03	0.00
Auto Dialer	2.37	-0.002	21.05	0.01
Corn	1.60	0.000	0.64	0.00
IPhone	2.17	0.000	-0.93	0.00
Shoes	0.46	0.000	0.52	0.00
Smart Phone	2.00	0.000	0.00	0.00
Telephony	-5.72	0.000	30.64	0.00
Women Clothes	0.64	0.000	0.42	0.00

Note. Unstandardized *b*-values reported. Average monthly search volume and the interaction term were automatically removed by the machine learning algorithm's feature selection because of weak or no contribution to accuracy.

TABLE 5. ACCURACY MEASURES OF LINEAR REGRESSION MODELS USED TO PREDICT SUGGESTED BID

Ad Group	Sample Size	Mean Absolute Error (USD)	Relative Absolute Error
Accounting Software	782	$5.23	92.7%
Auto Dialer	343	S16.44	98.0%
Corn	755	$1.74	101.4%
IPhone	792	$1.25	101.8%
Shoes	631	$0.29	93.8%
Smart Phone	587	$1.61	101.0%
Telephony	583	$15.49	92.2%
Women Clothes	796	$0.32	98.4%

Note. A lower value of relative absolute error indicates a more accurate model. A value equal to or greater than 100% indicates a worse model than simply using the mean to predict the dependent variable.