ABSTRACT

Title of Document: THREE EMPIRICAL STUDIES IN MARKET DESIGN

Andrew James Stocking
Doctor of Philosophy, Resource Economics
2009

Directed By: Peter Cramton
Professor
Economics

Andreas Lange
Assistant Professor
Agricultural and Resource Economics

Market design is the development of mechanisms that improve market efficiency and build on an understanding of the interaction between human behavior and market rules. The first chapter considers the sale of a charitable membership where the charity poses the market design question of how to price these memberships to capture the maximum value from donors' altruism. Using an online natural field experiment with over 700,000 subjects, this chapter tests theory on price discounts and shows large differences in donation behavior between donors who have previously given money and/or volunteered. For example, framing the charity's membership price as a discount increases response rates and decreases conditional contributions from former volunteers, but not from past money donors. This chapter

thereby demonstrates the importance of conditioning fundraising strategies on the specifics of past donation dimensions.

The second chapter examines an auction used to solve the assignment and price determination problems where price depends on the propensity to own or farm the land, a non-market good. This chapter studies bidder behavior in a reverse auction where landowners compete to sell and retire the right to develop their farmland. A reduced form bidding model is used to estimate the role of bidder competition, winner's curse correction, and the underlying distribution of private values. The chapter concludes that the auction enrolled as much as 3,000 acres (12 percent) more than a take-it-or-leave-it offer (i.e., non-auction program) would have enrolled for the same budgetary cost.

Finally, the third chapter considers the online advertising word auction. The pricing determination and assignment problem must occur for over 2,000 consumer searches each second. Theory is developed where asymmetric advertisers compete and an advertiser-optimal equilibrium bidding strategy is presented that is robust to this asymmetry. Within this rich strategy space, it is shown that advertiser subsidization can be revenue increasing for the search engine. Using a novel dataset of more than 4,500 keyword bids by three firms on four search engines, a simulation of the auction environment illustrates that bidder subsidization is indeed revenue positive and can be improved upon by imposing bid caps or fixed bids on the subsidized bidder.

THREE EMPIRICAL STUDIES IN MARKET DESIGN

By

Andrew James Stocking

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park, in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy,
Resource Economics
2009

Advisory Committee:
Professor Peter Cramton, Co-Chair
Andreas Lange, Co-Chair
Professor Lawrence M. Ausubel
Professor John K. Horowitz
Professor Loretta M. Lynch

Dedication

To Jack Hubbell, Michael Kavanaugh, Jerry Marxman, and Kerry Smith – four great mentors

And To mom, dad, Peter and Lisa – one great family

Acknowledgements

There are many who have helped to produce this dissertation and deserve a call out of appreciation. First, I would like to thank the two chairs of my committee: Peter Cramton for introducing me in concept and practice to market design principles and Andreas Lange for encouraging me to work on theory and fundraising research. Thanks also to the rest of my advising team and co-authors: Lori Lynch, John Horowitz, and Larry Ausubel. All five have contributed greatly to my understanding of economic research methodology. Thanks also to a host of other professors and co-authors who have helped along the way. Thanks to Ramon Lopez for teaching me how to write proofs, John List for first exposing me to field experiments, and Bill Evans, Richard Just, Marc Nerlove, and Anna Alberini for econometrics.

I am also grateful for the students at Maryland who were often helpful and willing to take time to talk about questions. Particular thanks go to my economics partner of the last two years, Nate Higgins, with whom I wrestled endlessly regarding various auction and game theory concepts. In addition, I will always be grateful to Katya Vasilaky for her study partnering before the qualifying exam.

Next, I turn to friends and colleagues outside of economics. Thanks go to M+R strategic services for their help running several field experiments, Jeff Regen and Cecilia Snyder for offering the opportunity to analyze large data sets, and all my friends at Care2.com who provided access to data and financial support. Thanks to Randy, Marlin, Clint, Joe, Eric, Eric, Justin, and Heather (plus others) for giving me something to do when I needed a break from economics over the last four years. And finally, thanks to my family (to whom this dissertation is dedicated) and friends.

Special thanks go to Jim Chapman for the generous sharing of his home when I needed a place to live and his pool when the other students and I needed a place to relax and escape the heat. Thanks also to my friends at the Perch and Emerson House for many evenings of laughing and food. And most importantly, thanks to Amy Myers for supporting me through to the end.

Table of Contents

List of Tables

List of Figures

Preface

Ask most undergraduate economic majors what a "market" is and they would respond: "a place where scarce goods and services are priced and sold." Expand your questioning to other economists, and you will hear more nuanced answers that relate market participant behavioral assumptions to the current market systems and their combined ability to achieve desired market outcomes. For example, Adam Smith, assuming a competitive market with many consumers, would describe how an invisible hand leads the market to socially efficient outcomes despite the economically selfish behavior of market participants.[1] Frederich Hayek, leader of the Austrian School, would probably make the point that markets price goods and services valued by society and efficiently ignore goods and services not valued by society. In response, Herman Daly might respond with behavioral economic results to show that socially valuable goods and services can be ignored when market participants are myopic and financial markets are incomplete with respect to future generations. Then there are many economists who would argue about the role of government interference and market rules for ensuring efficient outcomes. And the discussion could go on and on – each offering a slightly different take on how markets achieve desired outcomes based on some assumed behavior of market participants interacting with market rules. No one would disagree with the

[1] For purposes of this dissertation, efficient outcomes are defined with respect to assignment and price. The efficient assignment is one where the market participants with the highest values can buy the goods (or those with the lowest values sell the goods). The efficient price is that price that clears the market, i.e., the price that produces zero excess supply or demand.

undergraduate. But many economists would argue about how well markets work today given the range of consumer attitudes and markets in operation.

Enter the market designer. Market design is the development of mechanisms that improve market efficiency by building on an understanding of the interaction between human behavior and market rules. In many markets, market design might be limited to determination of the price or the unit for sale. For example, in a competitive market such as food prices at a grocery store, the market design question may involve how to optimally bundle products or incorporate discounts to encourage consumers to purchase new products. Alternatively, in the design of a highway congestion market where the government could be considered a monopoly provider of highway services, the design questions might include the toll fees, the interval of collection, and toll exemptions (e.g., high occupancy lanes) to achieve a socially efficient outcome. These are one class of market design problems.

A second and significantly more complicated class of market design problems involves situations where neither the buyer nor seller is a price taker, and the market is responsible for determining the socially efficient price and assignment of goods and services to market participants. In these cases, the market designer establishes rules for operation of the market such that the naturally occurring market forces, or invisible hand, can arrive at an optimal assignment and price. This area of economic research has recently started to receive significant attention (Roth 2002; Horowitz and Gayer, 2007; Cramton 2008; Ausubel et al., 2009). The fundamental question throughout this line of research is how different market rules influence participant behavior toward the achievement of efficient market outcomes.

The two standard paradigms for solving the efficient assignment and pricing problem in economics are the matching market and the auction. Typically, matching markets are used in situations where price is difficult to determine or deemed by society to be repugnant to determine. Typical matching markets involve the matching of medical interns to residency programs or prospective undergraduate students to colleges or children to elementary schools or college students to dormitory rooms (Roth and Sotomayor, 1990; Roth and Peranson, 1999; Roth 2007; Ausubel and Morrill, 2008). In all cases, there are multiple differentiated goods with participants who are able to express an ordinal preference ranking over those goods. Determining a cardinal (e.g., price) ranking may be awkward and, more importantly, not binding.[2] For example, even if parents could determine how much they would pay to have their daughter attend public school X, these parents would never be charged to attend the public school. Thus, ordinal rankings are all that is necessary for the operation of these matching markets.[3]

Alternatively, auction markets are used when market participants are accustomed to describing their demand or supply in terms of price. The canonical example of a one-sided auction market is the estate sale where an auctioneer accepts bids for estate assets. When neither side of the market is a price taker, these markets are modeled as a two sided auction where buyers and sellers simultaneously submit their demand and supply schedules. A good example of a two sided auction is the determination of the opening morning price for any of the thousands of stocks bought and sold on the NYSE or NASDAQ. Each morning before the opening of the market,

[2] Non-binding cardinal ranks might result in participants misrepresenting their cardinal values.
[3] Matching markets are robust to cardinal ordering of preferences or where prices are involved. See Chapter 3 of this dissertation or Cramton, 2009.

the NYSE or NASDAQ market managers accept sell and buy prices from multiple bidders for all stocks. And the opening price is the equilibrium price necessary to clear the market (i.e., eliminate excess demand/supply). These two specific auction examples are fairly straightforward and are widely believed to achieve socially efficient outcomes with fairly simple market rules.

There are many situations, however, where market efficiency is threatened by complications in the market layout or participant attributes. Several examples of complications that may result in market failures include:

- **Asymmetric Information.** There are times when one party in a market has significantly more information about a good being auctioned than other parties. The achievement of efficient outcomes depends on equal access to information by all market participants. However, if one party misrepresents their knowledge such that they can manipulate the assignment or price, the outcome may be sub-optimal.
- **Asymmetric Demand.** At times, market participants will have different demand schedules that are not substitutable with a competitor's demand schedule or that involve complementarities.[4] For example, in a telecommunication spectrum auction, some participants desire licenses across the entire United States; whereas, others only need licenses in a particular metro or regional area. Similarly, in an auction for airport landing slots, some airlines only want slots during peak times with other airlines needing a large number of slots during off-peak times.

[4] For a good summary of auctions that incorporate this type of complementary demand, see Cramton et al., 2006.

- **Market Power.** In highly concentrated markets there may be insufficient competition to determine an efficient market price or assignment. For example, a single very large buyer may misrepresent their demand to win the goods at a very low price. Or the large player may use pressure to force the smaller players into tacit or outright collusion, resulting in an inefficient price and assignment.
- **Externalities.** The market may fail to price either positive or negative externalities associated with a given product. For example, the price for electricity did not originally include the price associated with carbon emissions during the production of electricity. This market failure could have resulted from ignorance about the negative environmental repercussions of carbon emission or the myopia of market participants with respect to the impact of global warming on future generations.

There are many other situations where poorly designed markets might result in sub-optimal outcomes. Instead of describing all the situations where markets fail, it is easier to point to seven attributes that a well-running market should possess. These are adapted from Cramton (2007b).

1. *Efficient price formation*. The market should produce reliable and stable price signals based on market fundamentals.
2. *Transparency*. The market should be highly transparent. It should be clear how the pricing and assignment is done, and why one offer to buy is accepted and another rejected. The market should facilitate easy regulatory oversight and encourage regulatory certainty.

3. *Neutrality*. All participants (large or small) should be treated equally.
4. *Risk management*. The market should minimize risks for market participants, yet be responsive to long-run market fundamentals. The market should shield participants from short-term transient events and address counterparty risk.
5. *Liquidity*. The market should promote a secondary market, such that goods or services traded in the market are liquid for owners.
6. *Simplicity*. The market should be simple for participants, for the market operator, and for the regulator.
7. *Consistency*. The market should be consistent with, or improve upon, best-practices in other relevant markets. This might mean that the units being traded in the market are standard or easily interchangeable with other markets and that strategies can be implemented that incorporate multiple markets.

With this background, I turn to the three empirical market design examples considered in this dissertation. In each of the three markets, the market manager had a particular objective in mind that required innovative market design elements to influence buyer/seller behavior. Each chapter evaluates how the market designer changes behavior and the implications of this changed behavior on market efficiency. The first example conforms to the case above where the firm is free to set the price but wants to do so in such a way as to effect a particular type of consumer behavior. In the second two examples, the market operators do not know the optimal price or allocation of goods and thus must use an auction to solve the price determination and assignment problem. In both of these auctions, the established market rules are important for achieving the optimal outcome.

The first chapter explores the funding of a public good within the context of charitable fundraising. The specific market considered is the sale of a charitable membership where the charity poses the market design question of how to price these memberships such that they can capture the maximum value from donors' altruism. Typically, charities assume that altruism with respect to a charity increases with the quality of the charity. However, communicating charitable quality to the public is difficult and is thus primarily accomplished through the minimum membership level established by the charity. This chapter contains all of the interesting components of a market design problem: the established price conveys information about quality and, as a result, drives consumer behavior in the market. The chapter disentangles price and quality using price discounts from a standard level to fund the public good.

Specifically, co-author Professor Andreas Lange and I tested three donation requests across 720,890 potential donors: a request to become a member of the organization when the minimum membership level is $35, $25, and $25 represented as a special $10 discount from the standard $35 level. We concluded that the use of a discount provides potential donors with both a higher quality signal (the undiscounted level) and a lower price and this generates the most donations from the most donors. Our analysis of the market design implications of this mechanism goes further. The fundraising literature commonly assumes that potential donors differ along various dimensions, specifically their donation history, demographics, or affinity toward the organization. However, all of the literature to date analyzes fundraising mechanisms as if all potential donors interact similarly with the mechanism in question, i.e., use of a discount. In this analysis, we identify two distinct types of donors – those who have donated time in the past and those who have only donated money in the past. We find that these two pools of donors interact differently with the discount mechanism. For example, framing the charity's membership price as a discount increases response rates and decreases conditional contributions from former volunteers, but not from past money donors. This chapter thus demonstrates the importance of conditioning fundraising strategies on the specifics of past donation behavior.

The second chapter, co-authored with Professor John Horowitz and Professor Lori Lynch, moves to a market design question where assignment and price determination problems are best answered using an auction. This chapter explores the market design elements associated with pricing a non-market good, similar to altruism discussed in the first chapter. Here, though, the state of Maryland is interested in quantifying the propensity to own or farm land for each landowner or farmer. The market design question is how to design an auction where the state purchases landowner's development rights for the least cost. Development rights vary across landowners based on the fundamentals of the quality of the land, the proximity to development, and, most important, on each landowner's individual propensity to own (or farm) their land. The market design solution uses a reverse auction where landowners compete to sell their development rights to the state using a normalized price that allows the state to select those landowners willing to take the largest discount from the base price.

In this auction, run by the Maryland Agricultural Land Preservation Foundation (MALPF), the base price is unknown at the time of bidding and determined by an appraisal based only the fundamentals of land quality. Those land owners willing to accept the largest percent reduction from the base price will win the auction and receive payment in exchange for their property's development rights. We collected a unique dataset from 22 auction rounds with over 300 unique parcels to analyze the auction results. To model the bidding behavior in this common value auction, we develop a reduced-form model that incorporates the potential for differing private signals about the true development rights value. We also, however, allow bidders to differ in their affinity for farming and thus their willingness to accept a reduction on the sale of the development rights of their land. Our model disentangles these two factors and demonstrates that by capitalizing on different affinities for farming between landowners, the MALPF market performs better than competing land conservation systems.

Finally, in the third chapter, I – working alone this time – consider a market that is significantly more complex than the two analyzed in the previous chapters: the online-advertising word auction. In this market, advertisers bid for position in a list of advertisers that appears following any consumer search on a search engine.[5] The pricing determination and assignment problem must occur in real time for over 61 billion keyword searches annually or approximately 2,000 searches each second. This large volume of transactions requires a set of market rules that transparently assigns advertisers to positions and quickly determines a price that these advertisers should pay to remain in that position. The auction must be easy to understand by advertisers and, if designed well, should have a socially efficient equilibrium.

This chapter of the dissertation involves the evaluation of a number of market design solutions to address this problem. First, I evaluate the robustness of the current market to the presence of asymmetric advertisers – those advertisers who want to sell something, brand themselves, or generate more visitors – and their differing advertising objectives. Traffic bidders rank the top position as their best choice, branding bidders rank the bottom as the best choice, and transaction bidders may rank any position in between as their best choice. I conclude that an equilibrium bidding strategy exists that is optimal for the advertisers and is socially efficient.

Within this rich bidding environment, I then evaluate an innovative program whereby certain advertisers are subsidized to participate in the keyword auctions. I theoretically show that it is least costly for search engines to subsidize transaction bidders who preferentially bid on high quality words and may optimally appear in the middle of the list of advertisers. I prove that in some cases this subsidization strategy will increase search engine revenue.

[5] As described in detail in the chapter, when a consumer performs a search for any keyword or combination of keywords on a search engine (e.g., Google, Microsoft Bing, Yahoo!, Ask), they are taken to a page that displays organic results and paid advertiser results. This chapter describes the allocation of paid advertiser listings.

Finally, using a novel dataset of more than 4,500 keyword bids over two years by three firms on four different search engines, I simulate the auction environment to show that capping the bids of these subsidized bidders can also improve the likelihood of increasing search engine revenue.

Common to all three chapters described above is the use of market rules to influence participant behavior and ultimately market outcomes. The intent of this dissertation is to demonstrate both the breadth of markets where market designers can play a role and the efficiency benefits achievable when market designers are involved.

References

Ausubel, Larry and Thayer Morrill. 2008. "The Hungarian Auction." Working paper. University of Maryland.

Ausubel, Lawrence, Peter Cramton, Nathaniel Higgins, Emel Filiz-Ozbay, Erkut Ozbay, and Andrew J Stocking. 2009. "Common-Value Auctions with Liquidity needs: An Experimental Test of a Troubled Assets Reverse Auctions." Draft Paper, AREC University of Maryland College Park (May).

Cramton, P., Shoham, Y., and Steinbert, R. 2006. Combinatorial Auctions. MIT Press: Boston, MA.

Cramton, Peter. 2007a. Columbia's Forward Energy Market. Working Paper, University of Maryland, August.

Cramton, Peter. 2007b. "Comments on the RGGI Market Design." Submitted to RGGI, Inc. by ISO New England and NYISO, November.

Cramton, Peter. 2008. "A Review of the L-Band Auction." Office of Communications, United Kingdom.

Cramton, Peter. 2009. "Market Design: Auctions and Matching." In John Siegfried (ed.), *Better Living Through Economics,* Harvard University Press, forthcoming.

Horowitz, J.K. and Gayer, T. 2007. Market-based Approaches to Environmental Regulation. From Foundations and Trends in Microeconomics. In press.

Krishna, Vijay. 2002. "Action Theory." Academic Press: New York.

Roth, Alvin and Marilda A Oliveira Sotomayor. 1990. “Two-sided Matching: A study in Game-Theoretic Modeling and Analysis.” Econometric Society Monographs. Cambridge Press.

Roth, A.E. and Peranson, E. 1999. The Redesign of the Matching Market for American Physicians : Some engineering Aspects of Economic Design. American Economic Review. 89(4) : 748-780.

Roth, A.E. 2002. The Economist as Engineer: Game Theory, Experimentation, and Computation as Tools for Design Economics. Econometrica. 70(4): 1341-1378.

Roth, A.E. 2007 "Efficient Kidney Exchange: Coincidence of Wants in Markets with Compatibility-Based Preferences," American Economic Review. 97(3):828-851

Chapter 1: Charitable Memberships, Volunteering, and Discounts: Evidence from a Large-Scale Online Field Experiment [6]

By: Andreas Lange and Andrew Stocking

I. Introduction

The year 2007 marked the first time that more than $300 billion were donated to charity with individual donations accounting for more than 80% of this total (Giving USA, 2008). To solicit these donations, charities are aggressively turning to the internet due to its low marginal cost of use and rapid implementation time[7] (COP, 2008). However, there remains significant uncertainty within the fundraising community as to which rules from offline fundraising translate online and how to incorporate new features enabled by the internet into the charity's optimal fundraising strategy. Recent empirical publications on the determinants of giving have studied mail, phone, and door-to-door solicitation.[8] There are almost no studies, however, that consider the internet as a means of solicitation.[9] For economists, the internet provides a unique opportunity to actively contribute to a rapidly growing industry using field experiments that can be easily, inexpensively, and identically implemented

---

[6] We are grateful to Mike Price and John List for valuable comments on the paper. The authors would also like to thank M+R Strategic Services for their help and suggestions in implementing the field experiment.

[7] Notably, Barack Obama raised roughly $450 million dollars during his 2008 Presidential campaign with a large fraction of that coming online. In January 2008, the Obama campaign raised $32 million total with $28 million coming from online sources.

[8] For example, Falk (2007) uses postal mailings, Shang and Croson (2006) use phone solicitation, while Landry et al. (2006) employ a door-to-door strategy.

[9] Chen et al. (2006) is a notable exception. They conduct an online experiment soliciting visitors to a charity's website using a pop-up. This experiment yielded 24 donors who contributed a total of $1,128. Here we actively solicit potential donors via email instead of passively waiting for them to visit a particular website.

across large populations. As such, in this paper we conduct a natural field experiment with a sample of over 700,000 subjects, including donors and volunteers, to better understand the role of discounts and memberships in charitable fundraising.

Fundraisers generally believe that an individual who has donated money in the past, regardless the amount, is much more likely to donate money in the future. This effect, commonly called the "warm list" effect, has been recently discussed in the empirical literature with respect to offline donations (Landry et al., 2009). There is much less known, however, about anticipated money donations from those who have volunteered in the past, i.e., donated time. One unique feature of the internet is the ease with which charities can solicit and use individual time donations to accomplish their legislative, policy or other communication objectives. While it might be expected that past volunteers are also more likely to donate money in the future, to the best of the authors' knowledge, this topic is unstudied in the empirical or experimental economics literature.

Understanding how donation dimension affects the propensity to give, charities can optimize their solicitation appeals by targeting different donor types with specific mechanisms. We therefore define an "augmented warm list" which includes both the magnitude and the dimension of past donations. We demonstrate its beneficial use for targeted fundraising using the example of different pricing schemes for membership offers. In particular, we study the role of price discounts.

The analysis of linking donations with charitable memberships represents a unique contribution to the economic literature and complements a rich literature that links public good contributions with private benefits (see e.g., Cornes and Sandler, 1984). Andreoni (1989, 1990) introduced a model of impure altruism to the public goods model where donors receive some extra utility (i.e., warm glow) from their donations. Building on this, the recent economic literature considers a variety of specific private incentives to motivate contributions to public goods.[10] For example, gifts may lead to increased donations if they trigger reciprocity from the potential donors, or if their receipt is conditioned on donations exceeding a certain threshold, gifts may induce agents to contribute in order to obtain (or buy) the offered good. The ideal gift or conditional good would be one that (i) has no or low costs to the charity,[11] (ii) is difficult to obtain without making a contribution to the charity; (iii) has a high consumption value for the agent such that, if charged, agents are willing to pay a high price, (iv) triggers a large increase in warm-glow or reciprocity. Charitable membership and its associated benefits, potentially meet all of these criteria.

The use of price discounts has been well-studied in the marketing literature with respect to for-profit applications (e.g., Thaler 1985; Blattberg and Neslin 1989; Folkes and Wheat 1995, Gupta and Cooper 1992, Crewal et al. 1998). Discounts are observed to work via different mechanisms: (i) announcing the undiscounted higher price may serve as a quality signal and thus may enhance demand;[12] (ii) offering the

[10] Morgan (2000), Morgan and Sefton (2000), Landry et al. (2006), Lange et al. (2007) show the potential welfare-enhancing effects from introducing fixed-price lotteries to charitable fundraising. Falk (2007) and Landry et al. (2008) consider the role of gifts in triggering increased donations.

[11] This includes opportunity costs: even if goods are donated to the charity, the nonprofit might have alternative ways to sell the good and thereby to generate income to the charity. In order to be beneficial to the charity, the good must induce increases in giving in excess of those opportunity costs. Vesterlund (2003) discusses an additional mechanism by which this may occur: seed money or goods and prices which are donated to the charity may serve as a signal to subsequent donors about the quality of the charity.

[12] Milgrom and Roberts (1986) and Bagwell and Riordan (1991) discuss prices as quality signals.

discounts may create some perception of savings; (iii) lowering the price may increase demand due to a downward sloping demand curve. The first two mechanisms are primarily short run in nature because consumers are expected to update their beliefs and reference-points.[13] The authors were unable to identify, however, any peer-reviewed literature that connects discounts with charitable giving or charitable memberships despite the increased use of these two mechanisms by charities.

Returning to our initial premise, we combine these two mechanisms to determine the effect of offering the charitable membership with and without a price discount on 1) response rate (i.e., growth of the warm list) and 2) profitability (i.e., total dollars raised). We do this by making identical membership offers to two treatment groups except that one group must make a donation of at least $35 to receive the membership and the second need only make a donation of at least $25 which was said to represent "a special $10 discount" from the standard membership level. The literature suggests that there may be price effects related to the private benefits from donations. Harbaugh (1998) shows the potential benefits from selling reputation signals at different price points and Landry et al. (2009) consider gifts that are only given if donations exceed a minimum donation threshold. To disentangle the effect of the discount from possible unrelated price effects, we include a third treatment group that is offered an identical membership for a minimum donation of

[13] That is, in the short run unexpected price discounts may provide demand boosts due to reference-dependent preferences (e.g., Heidhus and Köszegi 2008), while in the long-run these effects might be smaller as the quality signal may be diluted by continuous price discounts (Folkes and Wheat 1995, Gupta and Cooper 1992,), and references may shift (Köszegi and Rabin 2006).

$25 without the mention of the special discount. We also develop a theoretical model in Section II that allows us to derive hypotheses as to the results.

Our online field experiment was conducted on an unprecedented scale and raised a total of $77,026 from 1,691 donors. Framing a $25 threshold price to become a member as a *special discount* from a standard $35 level induces a significantly larger proportion of people to donate and become members without reducing the total dollars raised. Thus the discount increased the size of the charity's warm list without compromising profitability. Without the discount framing, we find that $25 and $35 minimum donation levels for a membership induce a similar proportion of people to donate; however, the overall donation amount was 17% less with the lower membership threshold. Interestingly, this price effect suggests that a charity could therefore exploit the relatively inelastic demand for membership by requiring a larger minimum contribution to receive membership benefits. These results conform to our theory.

This is not, however, the entire story. We believe that our study is the first to explicitly consider the interaction between fundraising mechanism and mode of previous contributions to the charity. Volunteering or the donation of time accounts for roughly 55% of total giving in the United States (Salamon et al. 2007) to combine with financial contributions to represent 5.1% of US GDP. Using a complete history of the treatment pool's online contributions of time and money to the charity, we analyze the heterogeneity of response with respect to the above treatments. We show that previous time and money donors respond to the discount treatment quite differently. Past financial donors contribute at a conditional level consistent with the undiscounted \$35 minimum membership threshold; whereas, past time donors contribute at a conditional level consistent with the undiscounted \$25 minimum membership threshold. Those who have only given money in the past donate at an unconditional level that is higher than the undiscounted \$35 threshold; those who have given both money and time donate at an unconditional level equal to the \$25 threshold. These differences within the discount treatment do not appear when presented with the \$25 minimum threshold without the discount. This demonstrates differences in the perception of charitable discounts by past time or money donors. Furthermore, it demonstrates the importance of studying the effectiveness of charity fundraising mechanisms along different modes of giving.

The remainder of our study proceeds as follows. The next section provides the theoretical framework, which illustrates behavioral motivations for giving. Section III introduces our field experimental design. Section IV describes our findings and we conclude in Section V.

II. Theoretical Model

We provide a simple model to illustrate the most important determinants of giving behavior. We use a variant of Andreoni's (1989, 1990) impure altruism model and include consumption utility from a private good which in this case is the membership in the organization.

Agent $i \in \Omega$ receives utility from consuming a numeraire good, y_i, a private good (membership) at expected value m_i, a public good provided at level G, and (possibly) some extra utility (*warm-glow*) $f_i(g_i)$ from her own contribution to the public good g_i. Agents might also perceive the membership as a gift from the organization, depending on the way the membership is offered (e.g., with a discount). We model the resulting reciprocity component of utility as a function of both the contribution level and the perceived consumption value of the membership: $r_i(g_i, m_i)$. Importantly, the extent to which membership offers trigger reciprocal actions might depend on the historical interaction between donors and organization as will be described below.

Assuming additive separability of these different utility components, agent i's utility facing a budget constraint $y_i + g_i \le w_i$ and receiving expected membership benefits m_i is defined as:

$$U_i(g_i, m_i) = w_i - g_i + m_i + h_i(g_i + G_{-i}) + f_i(g_i) + r_i(g_i, m_i) \tag{1}$$

where $h_i(\bullet)$ and $f_i(\bullet)$ are twice continuously differentiable, non-decreasing and concave, and $\partial^2 r_i / \partial g^2 \le 0$, $\partial^2 r_i / \partial g \partial m \ge 0$, and $G_{-i} = \sum_{j \ne i} g_j$. The optimal solution to maximizing (1) for a given m_i is denoted by $\hat{g}_i(m_i)$. The assumptions ensure that $\partial \hat{g}_i(m_i) / \partial m_i \ge 0$.

We assume that the membership is awarded if contributions exceed a threshold T. Taking this threshold into account, the agent's optimal contribution level is denoted by $g_i^*(m_i, T)$. This optimal contribution may fall into three regions: (I) the threshold is not binding and the agent chooses to contribute at a higher level ($\hat{g}_i(m_i)$), (II) the minimum donation threshold is binding and the agent contributes at the threshold, (III) the agent is not willing to contribute at the threshold level and therefore does not receive membership benefits. The agent's willingness-to-pay for the membership is then given by:

$$WTP_i(m_i) = \max\left\{g_i \mid U^i(g_i, m_i) \geq U^i(\hat{g}_i(0), 0)\right\} \tag{2}$$

Note that $WTP_i(m_i) \geq \hat{g}_i(m_i)$. We can now formally state the optimal donation when offering a membership of quality m_i at a threshold of T:

$$g_i^*(m_i, T) = \begin{cases} \hat{g}_i(m_i) & \text{if} \quad T \leq \hat{g}_i(m_i) \\ T & \text{if} \quad \hat{g}_i(m_i) < T \leq WTP_i(m_i) \\ \hat{g}_i(0) & \text{if} \quad WTP(m_i) < T \end{cases} \tag{3}$$

We illustrate the effect of the threshold level T on contributions $g_i^*(m_i, T)$ in Figure 1 for a representative individual. If the threshold level is below the contribution level $\hat{g}_i(m_i)$, a marginal increase in T naturally has no effect (Region I). In Region II, the threshold exceeds the unconditional contribution level, but falls short of the willingness-to-pay such that the agent decides to contribute exactly the minimum donation level. If T exceeds the willingness-to-pay (Region III), contributions fall to the voluntary contribution level without being a member ($\hat{g}_i(0)$) which may be zero.

For a given perceived membership value m_i, we would therefore expect the number of donors (i.e., response rate) to be downward sloping in the required minimum donation level (*price effect*). The impact of this price effect on average conditional contributions, however, is less clear: (i) donors whose WTP now is less than the threshold T may still contribute (if $\hat{g}_i(0) > 0$), but at a lower level; (ii) other subjects whose WTP still is larger than the required threshold increase their contributions in order to obtain the membership.

In addition to this price effect, the threshold price for membership may serve as a quality signal. This link between prices and perceived product quality has been established in several economic studies (e.g., Milgrom and Roberts 1986; Bagwell and Riordan 1991; Gerstner 1985). We therefore assume that the expected value of membership benefits to the individual will depend on the announcement of the threshold (T): $m_i = M_i\left(T\right) > 0$.[14]

A higher (expected) quality of membership may increase donations via two distinct channels: it may trigger increased donations via the reciprocity term of equation (1) as $\hat{g}_i(m_i)$ is non-decreasing in m_i (*quality-reciprocity effect*).[15] Furthermore, a higher (expected) quality increases the consumption value m_i of the membership and therefore the willingness-to-pay to consume the product (*quality-consumption effect*).

With our experimental design we attempt to provide insights into the price and quality effects. For this we will compare donations at a high threshold level ($g_i^*(M_i(T^{high}), T^{high})$) with those at a reduced threshold level ($g_i^*(M_i(T^{low}), T^{low})$), as

[14] Similarly, treatment relevant information like the announcement of a threshold for the membership might influence the perceived quality of the charity and thereby the utility from the public good or from the own contributions. To simplify the presentation, we abstain from modeling those effects explicitly.
[15] Falk (2007) shows a positive correlation between the size of a gift and the response rate.

well as consider a discount treatment where a threshold of T^{low} is announced as a price discount from T^{high} ($g_i^*(M_i(T^{high}), T^{low})$).

When comparing individual contributions at a high vs. low threshold level, i.e.

$$g_i^*(M_i(T^{high}), T^{high}) \quad \text{vs.} \qquad g_i^*(M_i(T^{low}), T^{low}).$$

the quality effect and the price effect interact (see Figure 1). An increased threshold ($m_i' \equiv M_i(T^{high}) > m_i$) increases $\hat{g}_i(m_i')$ and may also increase the willingness-to-pay (T_2'). As a result, the higher threshold generates larger contributions for subjects contributing above the threshold (quality-reciprocity effect), or at the threshold (price effect or quality-consumption effect). A negative net effect might prevail if the threshold T^{high} exceeds the corresponding willingness-to-pay for membership. (i.e., if T becomes so high that the charity moves into Region III of Figure 1 for a large fraction of the potential donors despite the possible shifting of Region III to the right). These predictions are summarized in Hypothesis 1:

Hypothesis 1: ***Due to the interaction of quality and price effect, an increase in the minimum donation level (threshold) required for membership has an ambiguous effect on the number of donors and conditional contributions. Conditional on contributing above the threshold, conditional contributions increase in the minimum donation level if this minimum level provides a quality signal.***

We now turn to the role of discounts and compare the contribution level where the lower threshold (T^{low}) is framed as a discount $g_i^*(M_i(T^{high}), T^{low})$ to the contribution level without the discount $g_i^*(M_i(T^{high}), T^{high})$. The former should result in a higher participation rate as more people are acquiring the membership (price effect). The effects on conditional contributions are ambiguous as some agents may move down along the threshold if a discount is offered (stay within Region II in Figure 1)), while others start contributing at the (lower) threshold level (move from Region III into Region II).

A similar effect occurs when a price T^{low} is announced as a price discount from T^{high}: that is $g_i^*(M_i(T^{high}), T^{low})$ vs. $g_i^*(M_i(T^{low}), T^{low})$. The increased quality may induce more people to be willing to pay the minimum contribution level such that participation should increase. The effect on individual contributions is predicted to be positive, while the conditional contributions are ambiguous as newly entering agents may drive down the average. We therefore can formulate the following hypothesis:

Hypothesis 2: ***Formulating a minimum contribution level for membership as a price discount from a standard higher level increases participation compared to the higher threshold (due to price effect), but also compared to a lower threshold not announced as a discount (due to quality effect). The effect on conditional contribution levels is ambiguous.***

Complementary to these main treatment effects which are the focus of the current study, we are interested in the heterogeneous treatment affects based on certain donor characteristics. We use the remainder of this section to discuss additional predictions based on donor specifics.

Heterogeneity of the donor pool – Money vs. time donors

While much of the fundraising literature focuses on financial contributions, donors are also often requested to contribute time. In the following, we discuss how previous money or time donors may differ in their reaction to a (money) fundraising attempt.

Subjects who make money or time contributions in the past should be more likely to give than non-donors as their valuation of the public good ($h_i'(\bullet)$) and/or their warm glow from giving is larger ($f_i'(\bullet)$). This effect is well-known in the fundraising literature as a reason to value warm-lists.[16] Subjects who previously donated money might also be anchored by their previous donation such that their reaction to price changes is less elastic. For example, a donor who always contributed \$40 in response to donation requests might not change her contributions if membership benefits are given at \$25 or \$35; whereas, a donor who has never donated money in the past is not anchored by their first gift. As such, past money and time donors may differ in their reaction to manipulations in the price required for membership benefits. In terms of our model, previous money donors may be more likely to have $\hat{g}_i(m_i) > T^{high}$. Similar differences between subjects may occur with respect to the price effect when framing a lower threshold as a discount.

Price discounts might, however, also work through the reciprocity channel. That is, the discount might be interpreted as a nice offer from the charity such that an

[16] Landry et al. (2009) identify warm-list benefits in a field experimental setting. They show that the warm-list value depends on how donors are solicited: warm list value is significant when charities ask for voluntary contributions, but is reduced when additional gifts/incentives are offered to induce contributions.

individual agent might reciprocate gifts by increasing donations. In aggregate, this could result in a higher response rate and in larger conditional donations. This type of reciprocal gift exchange has been demonstrated in the field (e.g. giving post cards in exchange for money donations as in Falk (2007)) as well as in the lab (gift exchange game).

If the interaction between donors and the charity is repeated, however, actions by the charity may also be interpreted as a reaction to past donor behavior. That is, a price discount could potentially also be seen as a "thank you" by past donors. As such, a discount may trigger less reciprocal action among previous donors than non-donors. We argue that this "thank you" effect is also more likely to occur for previous time donors than for those who previously gave money. As argued above, donations from money donors may react less elastically to price changes. Furthermore, giving a *money* discount as a "thank-you" for a *money* donation may seem illogical to money donors who view the donor/charity relationship as one of raising money. Conversely, past time donors might be expected to interpret the discount as an indication of value placed on their time-donations by the charity and thus a reasonable "thank-you" for volunteering. Consequently, price discounts may trigger a stronger reciprocity reaction from past money donors and "thank-you" reaction from past time donors. We therefore formulate the following hypothesis:

Hypothesis 3: ***Former time donors are more likely to reduce their conditional contributions as a reaction to price discounts than past non-donors and former money donors are.***

With the establishment of these hypotheses, we now turn to the experiment design and results.

III Experimental Design

Following our theory, we designed a large scale natural field experiment with the goal of understanding how discounting a charitable membership affects the response rate, the unconditional donation level and the conditional donation level. The study was conducted online with 702,890 individuals. In addition to the experimental data, we obtained information on previous time and money donations from the individuals to the charitable organization.

The charitable organization is a large left-leaning advocacy organization that fights for civil liberties in the United States. The organization qualifies as a 501c(4) under IRS tax code which means donations to the organization cannot be deducted from income for federal tax purposes. It possesses an online email list with more than 1 million subscribers. This list serves as the basis for all of the organization's online fundraising efforts as well as other communications between the charitable organization and the individual.

Our experimental pool consisted of the 702,890 individuals who have not been a member of the charitable organization in the past and are thus referred to as Prospects by the organization. Some of these Prospects (10,077) have made financial donations in the past to the organization, but these individuals were not considered members because their donation was either a) to an affiliated entity such as the organization's PAC or 527[17]; or b) too small to be considered a member. Membership to the organization in 2008 usually requires a donation of at least $35. Donors making gifts below this level are not considered members for fundraising purposes. During typical membership drives, these non-donors and sub-$35 donors are requested to become a "member" of the organization by giving a gift of at least $35. The only additional tangible benefit to "membership" for the individual is the receipt of a quarterly magazine which is not emphasized during the membership process. The organization mainly promotes the psychic or altruistic benefits of membership.

The usual communications between the organization and their email list consist of three primary types of activities: 1) fundraising appeals which request financial contributions and emanate from the fundraising department; 2) action alerts, surveys, and event invitations which request non-financial contributions of time and emanate from the policy department; and 3) education or other general notifications which are informational in purpose and come from the policy or communications department. Requests for contributions of time to support the lobbying efforts of the organization represent a large majority of the communications. These emails notify the list subscribers of some situation that the charitable organization believes is important to their constituency and ask the subscribers to take "action." Typically "action" involves a 10-15 minute process whereby the member visits a page on the charitable organization's website and personalizes a letter for emailing or faxing to

[17] A PAC (political action committee) or 527 represent two different legal representations of the nonprofit that allow for different activities, mostly related to electioneering, for the nonprofit. Each legal entity is considered separate and must perform their own fundraising and email list building activities.

one or more of their elected officials. Approximately 268,504 list subscribers had donated time in the 12 months prior to the experiment.

In our experiment, we therefore carefully take the donation history of individuals into account and differentiate between time-donors (alternatively referred to as activists in the language of the organization) and money-donors.

Our experiment involved three treatments for which we divided the pool of Prospects into three equally-sized groups. Table 1 describes the demographic breakdown of the three groups based on information collected from a subset of the sample during previous surveys.[18] Table 2 summarizes the past behavior of the three prospect groups. All characteristics balance across the three treatments. For example, past money donors are evenly divided across the three groups in terms of number (1.35%), average conditional donations ($88.30), and number of past donations (1.37). The "Messaging and Time-Donation (Action) Behavior" rows describe the communication relationship between the charitable organization and the individual. On average the organizations sends 62 messages per year to list subscribers and members make on average 1.6 time donations in the form of actions. Finally, the table describes the activist (time donor) breakdown of the three groups. Activists are described within the organization as being 1) Super Activist if they made more than 4 time donations in the last 12 months; 2) Active if they made 1-3 time donations in the last 12 months, joined the email list within the past 6 months, or attended an offline event for the organization in the past year; or 3) Inactive if they do

[18] While the demographic makeup of the list based on survey participation may be biased due to survey selection issues, it can be used to demonstrate the identical nature of the three groups since selection into the three groups was done orthogonally to any survey participation variable.

not meet either the criteria for Super Activist or Active. As Table 2 shows, the three activist groups are evenly divided across the three treatment groups. For the remainder of this analysis we will combine the Super Activist and Active groups into a single group of time donors.

The experiment consisted of sending three sequential fundraising emails to the experimental pool. Each email consisted of identical language with the exception of the treatment language, which appeared in three different places in the email. If an email recipient clicked on any of the links in the email, they were taken to an online donation page that reinforced the treatment language. At this point in the process, the individuals decided whether to enter their donation amount, credit card information, and submit a donation or abandon the process and not make a financial contribution. The general theme of the appeal was to ask Prospects to become a member within the next few days to support the organizations general outreach and education work in the 2008 general election (approximately 90 days in the future).

The first email was sent on a Sunday (day 1), the second email was sent on the following Thursday (day 4), and the third email a week later on the next Thursday (day 11). The first two emails urged members to become a member by day 4 midnight. The third email extended the deadline to midnight on day 12. Prospects who donated any positive amount were removed from subsequent mailings. Treatment language did not vary across the three messages within each treatment or control group. Results are aggregated across the three messages for the treatment and control groups.

The three groups of Prospects were approached with the following donation requests:

- Control: “Become an [*organization*] member with a gift of $35 or more.”

- Treatment 1: “Become an [*organization*] member with a gift of $25 or more.”

- Treatment 2: “Become an [*organization*] member with a gift of $25 or more – that's a $10 discount off our normal membership.”

Given that none of these Prospects had been members before, it is reasonable to assume that those who were offered the $25 membership were not aware that the price was normally $35. The donation landing page reinforced this by stating “Minimum membership” next to the $35 or $25 donation level for the Control and Treatment 1 group, respectively. In Treatment 2, the phrase “Special $10 Discount on Minimum Membership” appeared next to the $25 donation level.

In all treatments, the webpage included radio buttons from among which donors could select when making a donation. The ask strings for the treatments were as follows:

- Control: $35, $50, $75, $100, $250, $500, Other

- Treatment 1 and 2: $25, $50, $75, $100, $250, $500, Other

The “Other” radio button had a text box into which donors could write any number they wanted above $10. These ask strings mimic the strings typically used by the organization in fundraising.[19]

[19] Ideally, we would have liked to offer identical ask strings for all treatments. However, the charity did not want to have a $25 ask in the $35 membership treatment (Control), as this was thought to confuse potential donors. Donors who wanted to make a contribution less than indicated by the minimum radio button, therefore had to choose the “other” option. This option to make a donation less than the minimum threshold was used by 3.8% of donors in the control (where the minimum button was $35), while 1.1% (1.7%) used this option in Treatment 1 (Treatment 2) where the minimum button

Finally, all donors in any group who made a gift of $20 or more were given the option to receive a branded picture frame. The language for the second sequential email stated "If you donate before midnight, we'll send you a magnetic picture frame as our gift." The language for the other two messages was similar. Donors who did not want the frame could check a box at the bottom of the donation page stating: "Please don't send a frame: use my full donation to fight for [cause]."

IV. Experimental Results

Table 3 provides summary statistics for our experimental treatments. Throughout this section, we discuss four different indicators: 1) "Response Rate" is the total number of donors conditional on being solicited (in %); 2) "Dollars given, unconditional" are the average dollars of all gifts for all solicited Prospects; 3) "Dollars given, conditional on giving" are the average dollars of all contributions conditional on giving; 4) "% who said "no frame" is the percent of donors who checked the box at the bottom of the donation page to not receive the picture frame. For example, Table 3 shows that we contacted 231,183 subjects in the control treatment, of which 521 decided to give which corresponds to a response rate of 0.23%. Those donors contributed an average of $45.21 and 62.2% (324 donors) requested not to receive the picture frame. In total, our study raised $77,026 from 1,691 individuals across the three treatments.[20]

---

was $25. Overall, the "Other" box was used 5%, 4.5% and 3.9% in the Control, Treatment 1, and Treatment 2, respectively.

[20] In the empirical analysis, we exclude one donor in Treatment 1 who made a donation of $10,000. Given that it is unlikely that this gift resulted from the specifics of Treatment 1, the gift was

Decreasing Minimum Donation Level

As stated in the introduction, we establish a baseline for the effect on membership demand from dropping the membership minimum threshold without announcing a discount. This will allow us later to disentangle the price and quality effects in the discount treatment. Comparing the Control and Treatment 1 in Table 3 shows that a reduction of the minimum level from $35 to $25 leaves the response rate unchanged at 0.23%, while the conditional contributions decrease from $45.21 to $36.32. This immediately leads to the following result:

Result 1: A change in the donation threshold for obtaining membership decreases average unconditional and conditional contributions, while the participation rate remains unaffected.

Further evidence for this result can be seen in regression 1 from Tables 4, 5, and 6. Table 4 (regression 1) shows that the difference in the participation rate for donors in Treatment 1 relative to the Control is statistically insignificant. However, the conditional donation amount was approximately $9 lower (Table 5 regression 1) and thus the unconditional contribution amount was statistically significantly lower by $0.02/Prospect (Table 6).

These results provide an empirical answer to the theoretically ambiguous effects described in Hypothesis 1: conditional contributions decrease with a decrease in the threshold. A $10 decrease in the threshold produces a $9 drop in the

eliminated. The largest gift after this was a $1000 gift in Treatment 1; 3 $500 gifts in Treatment 2; and a $365 gift in the Control. All other gifts were at or below $250.

conditional donation. Mapped to the theory, this suggests that many donors are operating in Region II of their demand curve (as illustrated in Figure 1), i.e. contribute at the threshold level. However, given that the conditional donation averages approximately $10 above the threshold, there are also donors in both treatment groups donating above this threshold and thus operating in Region I of their demand curve.[21]

These results already suggest an interaction between price and perceived quality of the charity. Changing the donation threshold does not change the response rate between the two treatments. As described by the theory, this could be because (i) the willingness-to-pay for membership is greater than either of the thresholds used in this study for most of the treatment group; or (ii) lowering the minimum donation threshold does not induce more subjects to give because it simultaneously reduces the willingness-to-pay. This later effect would indicate an interaction between price and perceived quality (i.e., willingness-to-pay). A second observation that supports the presence of a price-quality interaction as described in Hypothesis 1 is shown by the fact that a significantly larger fraction of donations are above the threshold in the Control treatment (36.7% above $35) than in Treatment 1 (26.5% above $25). If there were no interaction effect, this relationship should be reversed.

[21] In the Control (Treatment1, Treatment2), 36.7% (26.5%, 25.4%) of donors give above the threshold; the differences between the Control and Treatment 1/Treatment 2 are statistically significant at the 0.5% level. Unconditionally, 0.083% (0.060%, 0.068%) gave above the threshold in the Control (Treatment 1, Treatment 2). The difference between the Control and Treatment 1 (Treatment 2) is statistically significant at the 0.5% (10%) level; whereas, the difference between Treatment 1 and Treatment 2 is not statistically different.

In terms of the theoretical model, this suggests that – for some subjects – the higher threshold resulted in an increase in their utility-maximizing donation amount, or $g_i^*(M_i(T^{high}),T^{high}) > g_i^*(M_i(T^{low}),T^{low})$.

Effect of a Special Discount

With this impact of a price change now characterized, we can consider the role of a price discount as implemented in Treatment 2. A comparison with the Control corresponds to the pure price effect of our theory, a comparison with Treatment 1 to the quality effect. Table 3 supports Hypothesis 2 by showing that under the discount treatment we observed an increased response rate compared to both Control and Treatment 1 (0.27% in Treatment 2 vs. 0.23% in Control and Treatment 1). The conditional donations in Treatment 2, however, are smaller than in the Control treatment ($37.94 vs. $45.21), but almost identical to those in Treatment 1 ($36.32). This leads us to formulate the following result on the pure price effect:

Result 2: Framing a lower membership threshold as a discount from a given standard level decreases conditional donations, but increases the participation rate such that average unconditional donations are stable.

Further evidence for Result 2 can be found in the first regression in Tables 4, 5 and 6. The response rate for Treatment 2 is statistically significantly greater than for either the Control or Treatment 1 (1% level of significance). Specifically, introducing the discount causes the response rate to increase by 0.04 percentage points above a baseline of 0.24% (see baseline observed probability in Table 4). That is, discounting the charitable membership price by 29% triggered an increase in participation of 18%. However, the reduction in conditional donations (significant at 1%), leads to an insignificant change in average unconditional contributions (Table 6).

This result is consistent with our theoretical model: the response rate is predicted to increase if the utility-maximizing gift for many Prospects at the high threshold is zero (i.e., $g_i^*(M_i(T^{high}), T^{high}) = 0$ at $T^{high} = \$35$), but positive with a $10 discount (i.e., $g_i^*(M_i(T^{high}), T^{low}) > 0$ at $T^{low} = \$25$). That is, the result indicates that the quality signal in the discount treatment may be determined by $T^{high} = \$35$.

Comparing the discount Treatment 2 to Treatment 1 (i.e. considering the quality effect), we see that the unconditional amount raised per Prospect in Treatment 2 is statistically larger as more donors give (significant at 5% in Table 4 and 6, respectively) while conditional contributions are equivalent between these two treatments.

The results, therefore, suggest that it is beneficial to frame a given minimum donation level as a discount from a higher level. Going to the lower minimum level without framing it as a discount may decrease conditional donations, but does not increase the participation rate. Alternatively, a price discount appears to increase the participation rate while showing no effect on the unconditional donation rate, thus making their use beneficial. In the long-term, the relative benefits (Control vs. Treatment 2) depend on the future donation behavior of the additionally attracted donors which is beyond the scope of the current paper.

Heterogeneous Treatment Effects

We now explore the potential benefits from using discounts in charitable fundraising at a deeper level, that is, if this mechanism has particular benefits for specific *subsets* of prospects. For this, we differentiate our subject pool by means of their previous interaction with the charity; in particular, if they have donated time or money in the past. We denote past time donors as Active/Activists with Inactive being those who did not donate time in the past. Past money donors are denoted by M-Donor while NoM-Donor describes those who did not donate money in the past.[22]

Our analysis thus far has not shown a significant change in gross income (unconditional contributions) to the charity from using a price discount (Treatment 2 vs. control). Splitting the sample based on previous interactions, Table 6, regression (4) shows important differences for the treatment effect: while unconditional donations do not change for NoM-Donors, they do decrease for donors who gave time *and* money in the past ($0.39 lower), but increase for money donors who were inactive ($0.30 higher).

A charity could therefore profit by differentiating its fundraising strategy based on knowledge of previous Prospect interactions. That is, while money donors who did not give time should be contacted via the discount treatment, the charity could lose money by contacting money donors who had given time with the discount treatment.

[22] As a sensitivity analysis, we also considered using actual past donation amount instead of the binary M-Donor/NoM-Donor and found the results robust to either specification.

Result 3: An optimal fundraising strategy differentiates fundraising mechanisms by donor types. That is, the charity should exploit information on if and how (money or time) subjects have contributed in the past.

Result 3 establishes the benefits from targeted fundraising. We now study the causes of the differences across donor types in a more detailed way. For this, we again consider the effects on participation (Table 4) and conditional contributions (Table 5). We first consider former time donors (activists) regardless their previous money donations. Regression 2 from Tables 4 and 5 show that former time donors contribute under the discount treatment at a higher rate (0.035 percentage points higher), but with a lower conditional donation amount ($9.77 lower). This effect is primarily driven by those time donors who did not give money in the past (Active*NoM-Donor). This can be seen from regression 4 in Tables 4 and 5. While the conditional donation results for activists appear invariant to subdividing them by past money donations ($8.04 and $9.78 lower for Active*M-Donor and Active*NoM-Donor, respectively) only activists who have not donated money exhibit a statistically significant increase in response rate for the discount offer (0.039 percentage point increase). As a result, the unconditional donation amount (regression 4 from Table 6) for the discount offer is statistically indistinguishable for activists who are not money donors and statistically lower ($0.39 lower) for activists who are money donors.

Different results prevail for subjects who did not donate time in the past (Inactive). These Prospects do not exhibit an elevated participation rate for Treatment 2 like their active counterparts. Also unlike their active counterparts, the conditional contribution of inactives is higher in Treatment 2 than Treatment 1 and statistically indistinguishable from the Control (regression 2, Table 5). This is consistent with the discussion above regarding repeated interaction between charity and donor. Whereas activists may consider the discount to be a "thank you" for previous action, inactives have done nothing to warrant a "thank you" from the charity and thus may view the discount as a gift that should be reciprocated. As a result, inactives donate at a higher conditional contribution level than actives when presented with the discount. This inactive effect with respect to the response rate and conditional donation is invariant across both subgroups of inactive past money donors and inactives who have not donated in the past.

Turning to past money donors, we focus on those who have only interacted with the charity along the single dimension of money (i.e., no time donations). Relative to the Control treatment, we observe an increase in the response rate (0.021 percentage point) and conditional contribution amount ($12.20 increase) for these inactive M-donors in Treatment 2, though neither are statistically significant. As a consequence, inactive M-Donors were the only group to donate unconditionally more at a significant level in the discount treatment relative to the control ($0.30/prospect more).

We summarize these results as follows:

Result 4: The effect of framing a lower threshold level on participation and donation levels differs depending on the dimension of past donations (time or money). Specifically, *(i) Past time-donors who had not given money before respond to the discount at a higher rate and lower conditional gift leaving the unconditional donation level unchanged.*

(ii) Time-donors who also gave money in the past react to a discount on membership with a marginally lower conditional gift, but no change in response rate relative to the control, such that unconditional contributions decrease.

(iii) Past money-donors who did not give time respond to a discount with positive (but insignificant) increases in both participation and conditional contributions such that unconditional gift are significantly increased.

(iv) Prospects who had interacted with the charity along neither a time nor money dimension are unaffected by the discount treatment relative to the control.

Taken together, Result 3 and 4 suggests that a charity can benefit from offering the discount treatment to those prospects who have interacted with the charity in only one dimension, i.e., past inactive money donors and past active non-money donors. For the former an immediate increase in contributions results while for the latter the benefit potentially is given by an enlarged donor pool. The charity would do no harm to offer the discount to those prospects who had never interacted with the charity.

These differences in behavior are consistent with Hypothesis 3: time-donors may see the discount as a “thank you” for previous actions and therefore decrease their (conditional) contributions, while this “thank you” interpretation does not apply to non-donors. Furthermore, past money-donors react differently to the discount which uses the same (monetary) dimension as the initial donation.

These findings indicate that for the analysis of repeated reciprocal relationships it is necessary to comprehensively analyze exchanges in the different dimensions of the commodity space (money, time, consumption goods). It is therefore important for charities to keep track of all former interactions with potential donors: the value of "warm-lists" to a charity stems from the potential to discriminate future solicitation attempts based on the whole history of interactions. Traditionally, a "warm list" is defined as a group of people who have made at least one past contribution to the charity, agnostic to the specifics of the past contribution. Given the results described above, we therefore formulate the following definition of an "augmented warm list":

Definition: An "augmented warm list" is a list of the contact information for those who have agreed to accept communications from the charity along with their historical time and/or money contributions to the charity.

The benefits of an augmented warm list to a nonprofit are clearly delineated in the above discussion. Table 7 summarizes these effects for all treatments. From Table 4, we see that on average past time-donors contribute at a rate that is 2.6 times (0.41 percentage points above) the contribution rate of inactive donors. Past money-donors contribute at a rate that is 14.5 times (3.05 percentage points above) the contribution rate of non-money donors. The average unconditional donation of past time-donors is $0.18 greater than that of inactive donors (Table 6). Similarly, the average unconditional donation of past donors was $1.45 greater than that of non-donors. Overall, we obtain the following result:

Result 5: Subjects on a warm list are more likely to contribute and give – on average – larger donations per contact than those not on the warm list. Past money donors are more likely to donate money relative to past time donors.

The analyses conducted here do not consider any temporal dimension to the augmented "warm list". Intuition suggests that a donor whose last money or time donation lies more in the past is not as "warm" as a more current donor. Exploring the lapse rate of a "warm list" is a subject of future research.

Donor Quality

The previous results provide strong evidence for the value that past donors to the charity have for future fundraising drives. For the long-run analysis, it is therefore important to see if the marginal donors attracted by the price discount will make future contributions at a frequency and magnitude relative to donors attracted in the Control treatment. Intuitively, enlarging the donor pool must come at a cost for the charity as individual motivations to give must be lower for the marginal donor.

While a full analysis of donor quality can only be done through the long term analysis of donor behavior, we can gain insights into the "cost" and possible motivation of a donor based on their acceptance or rejection of the donation-conditional picture frame. Table 3 indicates that the fraction of donors who reject the gift is smaller in the discount treatment (56.2% compared to 62.2% in the Control and

59.3% in Treatment 1). Table 7 further indicates a larger rejection rate among past money and time donors in all treatments but one. We formulate the following result:

Result 6: Warm list subjects do not only generate larger revenues to the charity by contributing more, but also generate less fundraising costs: the additional gift is the more likely to be turned down the more past donations (money/time) a subject has made. The rejection rate of the gift is smaller if membership is framed as a discount.

For all donors, the rejection rate of the frame averaged 59.1%. Table 8 displays the marginal effects from a probit regression on the binary decision to accept the picture frame for various groups and treatments. When the $25 minimum membership is framed as a discount, there is a statistically significant increase in the frame acceptance rate by 6.0 percentage points relative to the 41% acceptance level for the control. Past activists were 7.1 percentage points more likely to accept the frame under the discount treatment relative to the control. Similarly, those who had not donated money in the past were 6.5 percentage points more likely to accept the frame under the discount treatment. Both of these results are statistically significant. These results suggest that former time-donors joining the charity under the discount treatment might be of lower value to the charity. In the short run, sending more gifts is costly to the charity. In the long run, their future contributions will be decisive.

Past money-donors are most beneficial to the charity. Not only do they contribute at a higher rate, they are also less likely to accept the additional gift. Consistent with intuition, they are high quality donors as their motivation to help the charity (e.g., their public good utility or warm-glow) is manifested in the same dimension (money) as is requested by the charity. As the charity attracts an increasing number of donors, the marginal donor has an increasingly lower intrinsic motivation to give and needs increasingly greater extrinsic incentives to give.

The Effect of Political Environment

Having established that fundraising mechanisms work different depending on time and money donation history, it is natural to check differences due to political environments since time donations are largely linked to actions addressing elected officials. A link between political voting patterns and the effectiveness of fundraising mechanisms was suggested by Karlan and List (2007). They find that a matching grant treatment was ineffective in Blue states, but quite effective in Red states.[23]

We therefore finally study the links between donation behavior and voting outcomes in the 2008 Presidential Elections. That is, Blue refers to states won by Barack Obama, while Red states are those won by John McCain.[24] The results of this analysis are presented in Table 9-11 and show no significant difference in donation behavior linked to the political outcome if controlling for the specific donor types. That is, the donor types react similarly to the respective treatments in Blue and in Red states.

These results indicate that observed differences in reactions to specific fundraising mechanisms at the state level could be driven by different compositions

[23] They defined Blue and Red states relative to the 2004 Presidential Election outcome. Following popular terminology, Blue states are those won by the Democrat (here, John Kerry) and Red states are those won by the Republican (George W. Bush).

[24] Our results are invariant to defining Blue and Red according to the 2004 Presidential Election outcome. Similarly our results do not change when we use actual vote percentages for the two candidates in lieu of blue/red dummy variables.

of donor types at the state level. While donor type characteristics are balanced for our study as shown in Table 12, any analysis of fundraising mechanisms with other charities might involve a different composition of donor types across states.[25] Our results call for a careful interpretation of aggregated effects in the literature, in particular when relying on smaller samples.

V. Conclusion

We conducted an online natural field experiment with 702,890 subjects designed to analyze the charity membership as a fundraising instrument. Due to the low cost and ease of use which allow for a much larger pool of experimental subjects than other experimental settings, we thereby propose using online fundraising platforms as a new and extremely beneficial way to conduct natural field experiments.

We found that donation behavior is affected by the minimum donation level required for membership. Reducing the minimum donation threshold did not lead to more subjects donating, but to lower average donations. While such a reduction is thereby extremely costly to the charity, we showed that by framing the reduction as a special discount the reduction in conditional contributions can be offset by attracting more donors.

These general findings suggest that the use of discounts as a charitable organization marketing instrument can be beneficial. Our results are consistent with

[25] For example, a charity may do political organizing in particular states to combat ballot measures which might result in a higher percentage of Actives in some subset of states.

the economic literature relating prices to quality: a charity that requires a larger donation to become a member appears to be signaling that it is a higher quality charity and thus membership has a higher value to the individual. In addition, our results suggest that there is a range of thresholds such that the charity can increase the threshold without reducing the response rate. Identifying the optimal threshold and whether that threshold is specific to a particular charity is beyond the scope of this experiment but appears to be warranted given our results.

This analysis also highlights for the first time the important distinction between donor types (financial or volunteer) and their differing reaction to fundraising mechanisms. This suggests that not only is there value in having a warm-list of previous donors, but also in denoting the nature of past donations as either money and/or time donors. Our results show that charities could benefit from differentiating fundraising mechanism across donor types. The differing effects of discounts on past time vs. money donors furthermore indicate that it is important to understand fundraising as a multi-dimensional activity. A full understanding of the economics of charities can only be achieved if different modes of giving as well as their interaction with fundraising mechanisms are studied. This paper provides a first step in this direction.

VI. References

Andreoni, James. 1989. "Giving with Impure Altruism: Applications to Charity and Ricardian Equivalence," *Journal of Political Economy* 97: 1147-1458.

Andreoni, James. 1990. "Impure Altruism, and Donations to Public Goods: A Theory of Warm-Glow Giving," *The Economic Journal* 100: 464-477.

Bagwell, Kyle & Riordan, Michael H, 1991. “High and Declining Prices Signal Product Quality,” *American Economic Review* 81(1), 224-39.

Blattberg, Robert C. and Scott A. Neslin. 1989, “Sales Promotion: The Long and the Short of it,” *Marketing Letters* 1(1), 81-97.

Chen, Yan, Xin Li and Jeffrey MacKie-Mason. 2006. Online Fund-raising Mechanisms: A Field Experiment. *Contributions to Economic Analysis & Policy.* 5(2): Article 4.

COP (Chronicle of Philanthropy). 2008. After the Flood: Online gifts still flow to charities, but the stream slowed in 2007. http://philanthropy.com/premium/articles/v20/i17/17000701.htm

Cornes, Richard and Todd Sandler. 1984. “Easy Riders, Joint Production, and Public Goods,” *Economic Journal* 94: 580-598.

Crewal, Dhruv, R. Krishnan, Julie Baker, and Norm Borin. 1998.. “The Effect of Store Name, Brand Name and Price Discounts on Consumers’ Evaluations and Purchase Intentions”, *Journal of Retailing* 74(3): 331-352.

Falk, Armin. 2007. “Gift Exchange in the Field,” *Econometrica* 75(5): 1501-1511.

Folkes, Valerie and Rita D. Wheat. 1995. “Consumers’ Price Perceptions of Promoted Products,” *Journal of Retailing* 71(3): 317-328.

Gerstner, Eitan. 1985. “Do Higher Prices Signal Higher Quality?” *Journal of Marketing Research* 22(2): 209-215.

Giving USA. 2008: http://nonprofit.about.com/od/trendsissuesstatistics/a/giving2008.htm

Gupta, Sunil and Lee G. Cooper. 1992. “The Discounting of Discounts and Promotion Thresholds”, *Journal of Consumer Research* 19: 410-411.

Harbaugh, W. 1998. “What do Donations buy? A model of philanthropy based on prestige and warm-glow,” *Journal of Public Economics*: 269-284.

Heidhus, Paul and Botond Köszegi. 2008. “Competition and Price Variation when Consumers Are Loss Averse,” *American Economic Review* 98(4): 1245-1268.

Karlan, Dean and John A. List. 2008. “Does Price matter in Charitable Giving? Evidence from a Large-Scale Natural Field Experiment.” *American Economic Review* 97(5): 1774-1793.

Köszegi, Botond and Matthew Rabin. 2006. “A Model of Reference-Dependent Preferences.” *Quarterly Journal of Economics*, 121(4): 1133-1166.

Landry, Craig, Andreas Lange, John List, Michael Price, and Nicholas Rupp. 2006. "Toward an Understanding of the Economics of Charity: Evidence from a Field Experiment", *Quarterly Journal of Economics* 121: 747-782.

Landry, Craig, Andreas Lange, John List, Michael Price, and Nicholas Rupp. 2008, "Is a Donor in Hand Better Than Two in the Bush? Evidence from a Natural Field Experiment," NBER Working Paper 14319.

Lange, Andreas, John A. List, and Michael Price. 2007. "Using Lotteries to Finance Public Goods: Theory and Experimental Evidence," *International Economic Review*, 48(3): 901-927.

Milgrom,. Paul and John Roberts. 1986. "Price and Advertising Signals of Product Quality," *Journal of Political Economy* 94(4): 796-821.

Morgan, John. (2000). "Financing Public Goods by Means of Lotteries". *Review of Economic Studies,* 67: 761-784.

Morgan, John and Martin Sefton. (2000). "Funding Public Goods with Lotteries: Experimental Evidence". *Review of Economic Studies*, 67: 785-810.

Salamon, Lester M., Megan A. Haddock, S. Wojciech Sokolowski, and Helen S. Tice. 2007. "Measuring Civil Society and Volunteering: Initial Findings from Implementation of the UN Handbook on Nonprofit Institutions." Working Paper No. 23. Baltimore: Johns Hopkins Center for Civil Society Studies.

Shang, Jen and Rachel Croson. 2006. "The Impact of Social Comparisons on Nonprofit Fundraising," Research in Experimental Economics, Vol 11: 143-156.

Thaler, Richard. 1985, "Mental Accounting and Consumer Choice," Marketing Science 4(3): 199- 214.

Vesterlund, Lise. 2003. "The Informational Value of Sequential Fundraising." *Journal of Public Economics*. 87(3-4): 627-57.

VII. Tables and Figures

Table 1 Demographic Comparison of Control and Treatment Groups

	Demographics	All	Control ($35)	Treatment 1 ($25)	Treatment 2 ($25 w/SD)
	All	702,890	231,183	236,234	235,473
1	Gender (% Female)	60.9%	61.1%	60.9%	60.8%
2	Age	42.7 (14.6)	42.7 (14.6)	42.7 (14.5)	42.8 (14.6)
	>25	10.3%	10.4%	10.3%	10.3%
	25-35	35.4%	35.5%	35.3%	35.5%
	35-50	21.3%	21.3%	21.4%	21.3%
	50+	32.9%	32.8%	32.9%	32.9%
3	Party Affiliation				
	Democrat	73.5%	73.5%	73.2%	73.8%
	Republican	3.6%	3.3%	4.0%	3.4%
	Other	22.9%	23.2%	22.8%	22.7%
4	Ethnicity				
	White	84.7%	84.1%	84.9%	84.9%
	African American	2.6%	2.6%	2.4%	2.6%
	Hispanic	5.1%	5.1%	5.1%	5.2%
	Other	7.7%	8.2%	7.6%	7.2%
5	Relationship Status				
	Single	37.6%	38.2%	37.0%	37.7%
	Married	11.5%	11.3%	11.4%	11.7%
	Partnered / Civil Unions	48.2%	47.9%	48.7%	48.1%
	Other	2.7%	2.6%	2.9%	2.5%
6	Region				
	Midwest	21.4%	21.3%	21.5%	21.5%
	Northeast	21.4%	21.5%	21.5%	21.2%
	South	29.2%	29.1%	29.1%	29.3%
	West	28.0%	28.2%	27.9%	28.0%
7	2004 Election (% Living in Blue State)	57.4%	57.6%	57.4%	57.3%
	2008 Election (% Living in Blue State)	78.2%	78.3%	78.1%	78.1%

Note: not all information reported above is available for every prospect. Specifically Gender is available for 54% of file; Age 28.2%; Party Affiliation 1.5%; Ethnicity 1%; Relationship Status 1.5%; Region 77%. Numbers in parentheses are standard deviation.

Table 2 Behavioral Comparison of Control and Treatment Groups

	Psychographics	All	Control ($35)	Treatment 1 ($25)	Treatment 2 ($25 w/SD)
1	Time on File (yrs)	2.38 (1.73)	2.37 (1.73)	2.38 (1.73)	2.38 (1.73)
Money Donor Behavior					
2	% Previous Money Donors	1.35%	1.35%	1.36%	1.34%
3	Average M-Donation, cond on giving	88.3 (248.5)	91.2 (287.5)	85.9 (210.7)	87.8 (242.4)
4	No. of M-Donations, cond on giving	1.37 (0.954)	1.36 (0.873)	1.37 (1.01)	1.37 (0.976)
5	Months Since Last M-Donation, cond on giving	18.8 (17.9)	18.3 (17.5)	18.9 (17.9)	19.2 (18.2)
6	No. of M-Donations, unconditional	0.0184 (0.192)	0.0182 (0.186)	0.0186 (0.197)	0.0183 (0.193)
7	Average M-Donation, unconditional	1.19 (30.6)	1.23 (35.0)	1.16 (26.5)	1.17 (29.8)
Messaging and Time-Donation (Action) Behavior					
8	No. of Msgs/yr	62.2 (32.5)	62.2 (32.4)	62.3 (32.7)	62.1 (32.3)
9	No. of Time-Donations (Actions)/yr	1.64 (2.90)	1.64 (2.88)	1.64 (2.92)	1.65 (2.90)
10	No. Msgs/yr, cond on 1+ Action	66.06 (27.1)	66.0 (26.8)	66.1 (27.4)	66.0 (27.1)
11	No. T-Donations/yr, cond on 1+ Action	2.36 (3.22)	2.35 (3.20)	2.36 (3.24)	2.37 (3.22)
12	No. Msgs to M-Donors/yr	81.8 (44.0)	82.0 (45.4)	82.1 (44.0)	81.2 (42.8)
13	No. Actions from M-Donors/yr	2.89 (3.73)	2.85 (3.68)	3.00 (3.84)	2.83 (3.66)
14	Activist (T-Donor) Category				
	Super Active (4+ time donations/yr)	4.0%	4.0%	4.0%	4.0%
	Active (1-3 time donations/yr)	34.2%	34.3%	34.3%	34.1%
	Inactive (0 time donations/yr)	61.8%	61.7%	61.7%	61.9%

Note: Activist Category is available for 100% of prospect file.

Table 3 Summary of results for the three treatments (means shown)

	Control ($35)	Treatment 1 ($25)	Treatment 2 ($25 w/SD)
Combined 3 Messages			
Response rate	0.23%	0.23%	0.27%
Dollars given, unconditional	0.10	0.08	0.10
Dollars given, conditional on giving	45.21	36.30	37.94
% who said "no frame"	62.2%	59.3%	56.2%
observations	231,183	236,234	235,473
donors	521	536	633

Table 4 Probit, marginal effects (dependent variable=donated (binary))

		(1)	(2)	(3)	(4)	Add'l Tests
1	T1 (d)	1.616e-05				(1) ≠ (2)***
		[1.468e-04]				
2	T2 (d)	4.355e-04***				
		[1.497e-04]				
3	T1*Active (d)		7.417e-05			(3) ≠ (4)**
			[1.250e-04]			
4	T2*Active (d)		3.514e-04***			
			[1.363e-04]			
5	T1*Inactive (d)		-2.205e-04			(5) ≠ (6)**
			[1.980e-04]			
6	T2*Inactive (d)		2.684e-04			
			[2.176e-04]			
7	Active (d)		4.111e-03***			
			[2.986e-04]			
8	T1*M-Donor (d)			6.640e-06		(8) = (9)
				[4.146e-04]		
9	T2*M-Donor (d)			-3.252e-05		
				[4.103e-04]		
10	T1*NoM-Donor (d)			1.313e-05		(10) ≠ (11)***
				[1.414e-04]		
11	T2*NoM-Donor (d)			4.669e-04***		
				[1.451e-04]		
12	M-Donor (d)			3.049e-02***		
				[3.329e-03]		
13	Active*M-Donor (d)				6.576e-02***	
					[7.596e-03]	
14	Active*NoM-Donor (d)				3.697e-03***	
					[2.973e-04]	
15	Inactive*M-Donor (d)				2.742e-02***	
					[6.065e-03]	
16	T1*Active*M-Donor (d)				2.898e-04	(16) = (17)
					[3.934e-04]	
17	T2*Active*M-Donor (d)				-7.706e-05	
					[3.256e-04]	
18	T1*Inactive*M-Donor (d)				-7.868e-04**	(18) ≠ (19)**
					[3.082e-04]	
19	T2*Inactive*M-Donor (d)				2.124e-04	
					[6.369e-04]	
20	T1*Active*NoM-Donor (d)				3.411e-05	(20) ≠ (21)***
					[1.191e-04]	
21	T2*Active*NoM-Donor (d)				3.899e-04***	
					[1.350e-04]	
22	T1*Inactive*NoM-Donor (d)				-8.828e-05	(22) ≠ (23)*
					[1.970e-04]	
23	T2*Inactive*NoM-Donor (d)				2.520e-04	
					[2.123e-04]	
	Baseline Obs Prob	2.40E-03	1.60E-03	2.10E-03	1.40E-03	
	Observations	702,890	702,890	702,890	702,890	
	Psuedo R-squared	4.922e-04	5.542e-02	4.272e-02	9.165e-02	
	Marginal effects; (d) for discrete change of dummy variable from 0 to 1					
	Standard errors in brackets; * significant at 10%; ** significant at 5%; *** significant at 1%					

Table 5 OLS (dependent variable=conditional contribution)

		(1)	(2)	(3)	(4)	Add'l Tests
1	T1 (d)	-8.914***				(1) = (2)
		[2.455]				
2	T2 (d)	-7.271***				
		[2.361]				
3	T1*Active (d)		-9.295***			(3) = (4)
			[2.716]			
4	T2*Active (d)		-9.765***			
			[2.634]			
5	T1*Inactive (d)		-7.225			(5) ≠ (6)*
			[5.719]			
6	T2*Inactive (d)		2.759			
			[5.279]			
7	Active (d)		0.985			
			[4.367]			
8	T1*M-Donor (d)			-13.01**		(8) ≠ (9)*
				[5.711]		
9	T2*M-Donor (d)			-3.312		
				[5.755]		
10	T1*NoM-Donor (d)			-7.990***		(10) = (11)
				[2.716]		
11	T2*NoM-Donor (d)			-7.773***		
				[2.589]		
12	M-Donor (d)			5.347		
				[4.505]		
13	Active*M-Donor (d)				6.453	
					[6.307]	
14	Active*NoM-Donor (d)				0.371	
					[4.853]	
15	Inactive*M-Donor (d)				2.571	
					[9.913]	
16	T1*Active*M-Donor (d)				-13.90**	(16) = (17)
					[6.239]	
17	T2*Active*M-Donor (d)				-8.039	
					[6.529]	
18	T1*Inactive*M-Donor (d)				-9.227	(18) = (19)
					[14.96]	
19	T2*Inactive*M-Donor (d)				12.20	
					[12.18]	
20	T1*Active*NoM-Donor (d)				-8.277***	(20) = (21)
					[3.017]	
21	T2*Active*NoM-Donor (d)				-9.783***	
					[2.883]	
22	T1*Inactive*NoM-Donor (d)				-6.741	(22) = (23)
					[6.224]	
23	T2*Inactive*NoM-Donor (d)				0.683	
					[5.857]	
24	Constant	45.21***	44.42***	44.23***	43.93***	
		[1.748]	[3.907]	[1.934]	[4.347]	
	Observations	1690	1690	1690	1690	
	R-squared	0.00767	0.0109	0.00964	0.0111	

(d) for discrete change of dummy variable from 0 to 1

Standard errors in brackets; * significant at 10%; ** significant at 5%; *** significant at 1%

Table 6 OLS (dependent variable=unconditional contribution amount)

		(1)	(2)	(3)	(4)	Add'l Tests
1	T1 (d)	-0.0195**				(1) ≠ (2)**
		[0.00808]				
2	T2 (d)	0.000101				
		[0.00809]				
3	T1*Active (d)		-0.0367***			(3) ≠ (4)*
			[0.0131]			
4	T2*Active (d)		-0.0124			
			[0.0131]			
5	T1*Inactive (d)		-0.00912			(5) ≠ (6)*
			[0.0103]			
6	T2*Inactive (d)		0.00839			
			[0.0103]			
7	Active (d)		0.182***			
			[0.0118]			
8	T1*M-Donor (d)			-0.399***		(8) ≠ (9)***
				[0.0694]		
9	T2*M-Donor (d)			-0.118*		
				[0.0697]		
10	T1*NoM-Donor (d)			-0.0145*		(10) ≠ (11)**
				[0.00812]		
11	T2*NoM-Donor (d)			0.00184		
				[0.00813]		
12	M-Donor (d)			1.447***		
				[0.0498]		
13	Active*M-Donor (d)				2.038***	
					[0.0644]	
14	Active*NoM-Donor (d)				0.148***	
					[0.0119]	
15	Inactive*M-Donor (d)				0.714***	
					[0.0781]	
16	T1*Active*M-Donor (d)				-0.375***	(16) = (17)
					[0.0900]	
17	T2*Active*M-Donor (d)				-0.392***	
					[0.0906]	
18	T1*Inactive*M-Donor (d)				-0.425***	(18) ≠ (19)***
					[0.109]	
19	T2*Inactive*M-Donor (d)				0.300***	
					[0.109]	
20	T1*Active*NoM-Donor (d)				-0.0295**	(20) ≠ (21)*
					[0.0132]	
21	T2*Active*NoM-Donor (d)				-0.00377	
					[0.0132]	
22	T1*Inactive*NoM-Donor (d)				-0.00547	(22) = (23)
					[0.0103]	
23	T2*Inactive*NoM-Donor (d)				0.00570	
					[0.0103]	
24	Constant	0.102***	0.0324***	0.0824***	0.0261***	
		[0.00574]	[0.00731]	[0.00578]	[0.00733]	
	Observations	702890	702890	702890	702890	
	R-squared	8.3E-06	8.5E-04	0.00289	0.00403	

(d) for discrete change of dummy variable from 0 to 1

Standard errors in brackets; * significant at 10%; ** significant at 5%; *** significant at 1%

Table 7 Results by Activist Type (means shown)

	Control ($35)	Treatment 1 ($25)	Treatment 2 ($25 w/SD)
Cross Section A: Past Time-Donors (Activists)			
Response rate	0.47%	0.49%	0.57%
Dollars given, unconditional	0.21	0.18	0.20
Dollars given, conditional on giving	45.41	36.11	35.64
% who said "no frame"	62.8%	59.8%	55.8%
observations	88,444	90,595	89,614
donors	417	445	507
Cross Section B: No past time donation (Inactives)			
Response rate	0.07%	0.06%	0.09%
Dollars given, unconditional	0.03	0.02	0.04
Dollars given, conditional on giving	44.42	37.20	47.18
% who said "no frame"	59.6%	57.1%	57.9%
observations	142,739	145,639	145,859
donors	104	91	126
Cross Section C: Past Money-Donors			
Response rate	3.08%	3.09%	3.05%
Dollars given, unconditional	1.53	1.13	1.41
Dollars given, conditional on giving	49.57	36.57	46.26
% who said "no frame"	65.6%	65.7%	63.5%
observations	3,112	3,202	3,147
donors	96	99	96
Cross Section D: Not Past Money-Donors			
Response rate	0.19%	0.19%	0.23%
Dollars given, unconditional	0.08	0.07	0.08
Dollars given, conditional on giving	44.23	36.24	36.45
% who said "no frame"	61.4%	57.9%	54.9%
observations	228,071	233,032	232,326
donors	425	437	537

Table 8 Probit, mfx (dependent variable=accept conditional gift (picture frame))

		(1)	(2)	(3)	Add'l Tests
1	T1 (d)	2.900e-02			(1) = (2)
		[3.054e-02]			
2	T2 (d)	5.979e-02**			
		[2.931e-02]			
3	T1*Active (d)		3.113e-02		(3) = (4)
			[3.396e-02]		
4	T2*Active (d)		7.077e-02**		
			[3.295e-02]		
5	T1*Inactive (d)		2.478e-02		(5) = (6)
			[7.128e-02]		
6	T2*Inactive (d)		1.683e-02		
			[6.562e-02]		
7	Active		-3.279e-02		
			[5.460e-02]		
8	T1*M-Donor (d)			-3.334e-04	(8) = (9)
				[7.183e-02]	
9	T2*M-Donor (d)			2.188e-02	
				[7.285e-02]	
10	T1*NoDonor (d)			3.546e-02	(10) = (11)
				[3.383e-02]	
11	T2*NoDonor (d)			6.487e-02**	
				[3.222e-02]	
12	PastM-Donor (d)			-4.317e-02	
				[5.538e-02]	
	Baseline Obs Prob	4.09E-01	4.09E-01	4.09E-01	
	Observations	1690	1690	1690	
	Pseudo R-squared	1.841e-03	2.175e-03	4.083e-03	
	Marginal effects; (d) for discrete change of dummy variable from 0 to 1				
	Standard errors in brackets; * sig at 10%; ** sig at 5%; *** sig at 1%				

Table 9 Probit, marginal effects (dependent variable=donated (binary))

	(1) Active M-Donor	(2) Active*NoM-Donor	(3) InActive*M-Donor	(4) InActive*NonM-Donor
T1 (d)	7.193e-03	-7.528e-05	-8.529e-03*	-4.323e-05
	[7.549e-03]	[3.891e-04]	[4.686e-03]	[1.469e-04]
T2 (d)	-1.138e-03	1.007e-03**	1.509e-03	1.394e-04
	[7.436e-03]	[4.002e-04]	[4.840e-03]	[1.503e-04]
C*Red (d)	4.808e-04	-8.837e-04*	-2.738e-04	1.079e-04
	[1.196e-02]	[5.303e-04]	[7.544e-03]	[2.428e-04]
T1*Red (d)	-8.655e-03	-1.339e-04	3.638e-03	2.807e-05
	[9.927e-03]	[5.698e-04]	[1.039e-02]	[2.350e-04]
T2*Red (d)	-1.695e-03	-6.231e-04	-4.315e-04	1.198e-04
	[1.176e-02]	[4.960e-04]	[7.452e-03]	[2.243e-04]
Baseline Obs Prob	4.22E-02	4.58E-03	1.35E-02	7.89E-04
Observations	5595	240655	3826	290946
R-squared	9.479e-04	1.120e-03	8.906e-03	8.458e-04
Test T1=T2	T1 = T2	T1 ≠ T2***	T1 ≠ T2**	T1 = T2

Marginal effects; (d) for discrete change of dummy variable from 0 to 1
Standard errors in brackets; * significant at 10%; ** significant at 5%; *** significant at 1%

Table 10 OLS (dependent variable=conditional contribution amount)

	(1) Active M-Donor	(2) Active*NoM-Donor	(3) InActive*M-Donor	(4) InActive*NonM-Donor
T1 (d)	-14.31***	-8.021**	-6.562	-8.767
	[4.320]	[3.280]	[30.03]	[10.02]
T2 (d)	-7.408	-10.02***	-7.648	-3.083
	[4.547]	[3.115]	[23.54]	[9.477]
C*Red (d)	-0.475	-4.496	-9.687	-4.921
	[7.177]	[5.309]	[38.77]	[14.54]
T1*Red (d)	1.836	-5.411	-16.88	2.963
	[7.060]	[4.904]	[46.96]	[15.33]
T2*Red (d)	-3.782	-3.730	103.0***	7.994
	[7.403]	[4.589]	[38.16]	[13.11]
Constant	50.48***	45.36***	48.44***	45.80***
	[3.189]	[2.324]	[17.34]	[7.018]
Observations	237	1110	54	231
R-squared	0.0331	0.0103	0.0579	-0.0160
Test T1=T2	T1 = T2	T1 = T2	T1 = T2	T1 = T2

Marginal effects; (d) for discrete change of dummy variable from 0 to 1
Standard errors in brackets; * significant at 10%; ** significant at 5%; *** significant at 1%

Table 11 OLS (dependent variable=unconditional contribution amount)

	(1) Active M-Donor	(2) Active*NoM-Donor	(3) InActive*M-Donor	(4) InActive*NonM-Donor
T1 (d)	-0.321	-0.0381*	-0.447	-0.00799
	[0.367]	[0.0208]	[0.459]	[0.00970]
T2 (d)	-0.351	-0.00757	-0.0494	0.00365
	[0.369]	[0.0209]	[0.453]	[0.00969]
C*Red (d)	0.00399	-0.0543*	-0.169	0.000554
	[0.590]	[0.0311]	[0.722]	[0.0150]
T1*Red (d)	-0.274	-0.0276	-0.0757	0.00309
	[0.573]	[0.0307]	[0.684]	[0.0148]
T2*Red (d)	-0.214	-0.0432	1.768**	0.0138
	[0.587]	[0.0308]	[0.746]	[0.0148]
Constant	2.068***	0.200***	0.781**	0.0340***
	[0.261]	[0.0148]	[0.326]	[0.00689]
Observations	5595	240655	3826	290946
R-squared	-0.000504	0.0000182	0.00104	-0.00000495
Test T1=T2	T1 = T2	T1 = T2	T1 = T2	T1 = T2

Marginal effects; (d) for discrete change of dummy variable from 0 to 1
Standard errors in brackets; * significant at 10%; ** significant at 5%; *** significant at 1%

Table 12 Composition of Blue and Red State List Members

Blue States	Active	Inactive	Totals
M-Donor	1.1%	0.7%	1.8%
Non M-Donor	43.9%	54.3%	98.2%
Totals	45.0%	55.0%	

Red States	Active	Inactive	Totals
M-Donor	0.9%	0.7%	1.6%
Non M-Donor	46.4%	52.0%	98.4%
Totals	47.3%	52.7%	

Figure 1 Discount Theory.

Note: The solid line represents an individual's demand curves based on one specific membership signal. As the charity announces a higher threshold for receiving the membership benefit (T), the individual adjusts their beliefs about the expected value of the membership and thus operates along the demand curve depicted by the dotted line (i.e., $m_i' \geq m_i$).

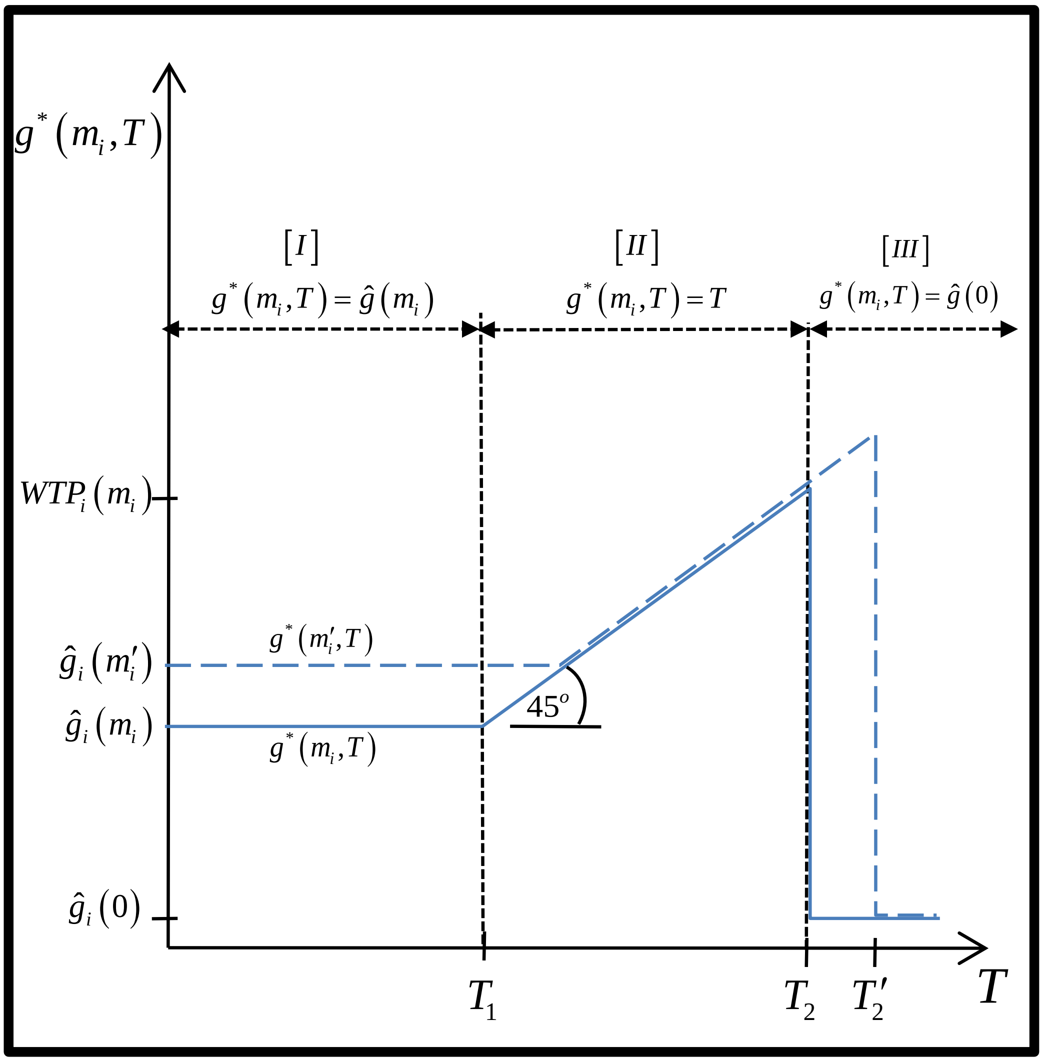

Chapter 2: Competition-based Environmental Policy: An Analysis of Farmland Preservation in Maryland[26]

By: John K. Horowitz, Lori Lynch, and Andrew Stocking

I. Introduction

A large proportion of environmental problems stem from private land use decisions that do not incorporate negative externalities. Such problems include biodiversity loss and threat of species extinction due to land conversion, carbon release due to deforestation and land cultivation practices, and nonpoint water pollution from agriculture, flood control, and loss of amenities such as scenery, local climate, and wildlife viewing. Indeed, it seems that a large portion of U.S. environmental problems would be solved if land use decisions could be optimized to include environmental externalities.

This prominent role for land use is important because the legal means used to regulate land use in the U.S. are very different from those used to regulate pollution. In the case of pollution, federal regulations directly limit what polluters can do. In contrast, land use in the U.S. outside of urban areas is much more loosely regulated. Instead, the primary policy tool has been to *pay* landowners to undertake the actions that we as a society would like them to undertake. The general principle is that for pollution, the environmental rights are held by the public, in the hands of the government, whereas for land use the rights are inherent in the property. This difference affects the policies and policy implementation issues that must be

[26]We thank Gary Biglaiser, Peter Cramton, Bill Neilson, Harry Paarsch, and Dan Vincent for helpful discussions.

considered in addressing land use issues.

Due in part to the property rights movement and the shift away from regulations since the 1970's (Echevarria 2005), policymakers have sought voluntary-enrollment land use policies with low budgetary costs but that are nonetheless cost-effective. The most promising of these are *competition-based policies* in which landowners compete for shares of a fixed budget in return for implementing environmentally desirable land uses. A prominent example is the Conservation Reserve Program (CRP), a Federal program under which landowners submit bids to take environmentally sensitive agricultural land out of production for a 10 or 15 year period. Lowest bids (per environmental benefit) are enrolled first, with enrollment continuing until the budget is exhausted. Non-U.S. examples include Australia's Auction for Landscape Recovery, the U.K.'s Challenge Funds and initial auctions for greenhouse gas reductions, and Germany's Grassland Pilot. Competition-based policies are likely to be important components of future U.S. policies for two major environmental issues: nonpoint water pollution and carbon sequestration.

Although competition-based schemes have many desirable properties there is much that remains unknown about the bidders' behavior and its implications for policy design. Auctions modeled on environmental policies have been studied extensively in laboratory and field experiments. Schilizzi and Latacz-Lohmann (2007) conducted induced-value, private-value experiments to compare two bid acceptance rules for conservation auctions, and Cummings, Holt and Laury (2004) conducted induced-value, private-value experiments to study selection rules for the Georgia irrigation auctions. Induced value experiments, however, dictate the bidders' information structure rather than work with the complex real-world information structure that characterizes most auctions. Such information structures may distort bidding and may therefore cause competition-based programs to be potentially more costly than other voluntary-enrollment programs with more certain and straightforward payment structures. This paper attempts to shed light on this question using data from an actual competition-based policy.

Competition-based real-world land use programs such as the Conservation Reserve Program (CRP) have been the subject of some theoretical work (Smith 1995; Latacz-Lohmann and Van der Hamsvoort 1997) but only a small number of empirical articles have made use of the auction paradigm. Kirwan, Lubowski, and Roberts (2005) estimate a reduced form model of the CRP and infer reservation values from assumptions based on the portion of the reservation value that is observable. Vukina *et al*. (forthcoming) also use a reduced form model to examine how plot-specific environmental scores, which increase the probability of winning, ceteris paribus, affect bids. Our auction set-up and rich data allows us to estimate a broader set of relationships related to bidder behavior including the impact of competition, potential winner's curse as well as infer the underlying reservation values.

This paper studies bids in the Maryland Agricultural Land Preservation Foundation (MALPF) program, an innovative program in which Maryland farmers compete to sell their right to develop their land to the State while retaining ownership of their land. The State then retires these development rights, ensuring that the property remains in agricultural or related use.

The MALPF program, like other competition-based policies, relies on the principles that underlie auction theory which posits that participant behavior depends on the underlying information structure. Several unique features of this auction allow us to identify the informational components. Using data from 24 auction rounds over 19 years, we examine the effects of competition, information, and bidder entry and selection on bidding behavior in the MALPF auction. We find that on average bids are 5 to 15 percent above the underlying reservation value and show that increased competition (in the form of lower budgets or more bidders) reduces this mark-up. We also find evidence that bidders adjust for a possible winner's curse by increasing their bids by 8 to 14 percent.

This framework then allows us to back out the underlying distribution of reservation values. This distribution of values constitutes the supply curve for MALPF enrollment. We use the inferred reservation values to compare the performance of the MALPF auction to an alternative policy under which administrators make a take it or leave it offer to all potential enrollees. We find that the MALPF auction enrolled 5 to 12 percent more acres than this alternative for the given budget.

The remainder of the paper proceeds as follows. Section II describes the details of the MALPF auction. Section III describes the empirical model. Section IV describes the data. Section V discusses the results and Section VI uses the results in a comparison to an alternative farmland preservation policy. Concluding remarks are in Section VII.

II. MALPF Auction Design

Maryland established the MALPF program in 1977. It was one of the nation's first statewide programs to restrict development on agricultural land by purchasing from landowners their rights to develop their property. Preservation is achieved through the attachment of a "conservation easement," a restriction on the deed that proscribes most forms of development, essentially in perpetuity. The enrolling owner of a parcel is free to sell the land but future owners continue to be bound by the development restriction. By 2003 the MALPF had preserved 228,854 acres, which is 4 percent of State land and 10 percent of its agricultural acres. Acquisition expenditures for these development rights for the fiscal year 2002 were $37.6 million, with a statewide average per-acre cost of $1,960.

We study bids from Carroll County, an urbanizing county just west of Baltimore. We chose this county because: (i) it experienced substantial development pressure during the period of study yet also had 178,000 acres of agricultural land at the start of our study period, (ii) it actively promoted the MALPF auctions, and (iii) there were no other competing preservation programs during the study period that paid compensation.[27]

In each round of bidding, eligible landowners submit offers to sell their parcels' development rights. After the offers are submitted, the program pays for two professional appraisals of the market value of the unconstrained property. The state

[27]The Conservation Reserve Program operates in Carroll County but the maximum rental rates over this time period were not high and enrollment was quite low. Therefore, the CRP's impact on agricultural value calculations and enrollment decisions is likely to have been essentially non-existent. The Maryland Environmental Trust accepts donations of development rights. Landowners who donate their development rights could apply for tax deductions.

selects the appraisal "which in their judgment reflects the most accurate value" (MALPF, 1984). A new appraisal is conducted each time a parcel is bid. Parcel values for parcels re-bidding a second or third time can change substantially from round to round.

The program then computes the agricultural value of the property (that is, its value if it were constrained to remain in agriculture) based on specified rules. The state calculates the "market easement value" as the difference between the property's unconstrained market value and its agricultural value. The landowner's submitted offer is then converted into a ratio by dividing it by the market easement value. These ratios are ranked and the program purchases development rights starting with the lowest ratio offer. The landowner is paid the amount of his offer, with exceptions described below; thus, this program is a type of first-price auction. The program works its way up this line-up of ratios until the annual budget is exhausted. In this way, the program buys development rights that are the least expensive relative to the assigned market easement value. Because of the auction-like approach we refer to the landowners' offers as bids.

The ratio approach represents an innovative design feature for farmland preservation. If the state were to purchase development rights from the lowest bidders, it would preserve those properties that were least likely to be developed, although it would be able to enroll the most acres. By comparing bids to the market's assessment, the state acquires easements that are cheapest relative to the market price; this adjustment is presumably no longer biased toward low development probability parcels. Our paper does not address the optimality of this feature nor MALPF's implicit objective function. The selection rule does not account for the effect of preservation of a given parcel on the development rates of other parcels, nor for other public goods provided by the parcels.

Like all government programs, the MALPF program has twists that complicate the analysis. For ratios greater than one, the program may offer to purchase the development rights at the market easement value, assuming the budget has not been exhausted by parcels with ratios less than one. For these parcels, this payment would be lower than the landowner's requested payment. Landowners can either accept this offer or decline it. The administrator selects these parcels for auxiliary offers based on unspecified criteria and not necessarily on the lowest-ratio-first rule.

Participants who are not accepted can re-bid in any future round. Multiple rebidding is allowed and indeed is common.

The program had an escape clause under which, if after 25 years a landowner can demonstrate there exists no profitable agricultural use and is willing to refund the MALPF the current value of the development rights, the landowner may petition to remove the restrictions. This clause has never been utilized and considerable opposition exists toward allowing any parcel to exit.

III. Estimation Model and Equations

This section provides a conceptual model that defines the informational structure of the auction, lays out our research questions, and forms the basis for our econometric model.

Land market

We start with a standard model of land conversion. Let *V(t)* represent the discounted value of services on a given parcel were it to be developed at time *t*. This value is assumed to follow a random walk process. Development results in foregone profits from agriculture. Let *x* represent annual net returns from agriculture on this parcel; for simplicity these returns are assumed non-stochastic. Let ρ be the risk-free discount rate. Consequently, the market value of the land absent any opportunity for development is $X = x/\rho$.

The landowner's decision is the date τ at which to develop the parcel. When the parcel is converted the landowner gives up the remaining stream of *x*, valued at $Xe^{-\rho\tau}$. The parcel is converted when *V(t)* exceeds *X* plus the option of waiting for a higher offer. Thus the value of an undeveloped parcel given current realization *V* is:

$$F^m(V) = \max_\tau E\left\{X\left(1 - e^{-\rho\tau}\right) + e^{-\rho\tau}V(\tau)\right\} \tag{1}$$

The cost of land conversion is assumed to be zero in this model but such a cost would not change the underlying structure of the problem. Making agricultural returns stochastic would complicate the presentation but again would not change the result.

This is a standard investment under uncertainty problem, solvable by methods described in Dixit and Pindyck (1994) and regularly applied in the land conservation context (Plantinga, Lubowski, and Stavins 2002). The solution, described by Dixit and Pyndick (1994) is $F^m(V) = X + aV^\beta$ where $a = \beta^{-\beta}\left(X/(\beta - 1)\right)^{1-\beta}$ and $\beta > 1$ is a parameter that depends on ρ and the process governing *V*(*t*). The market easement value, M, is the difference between the value of an undeveloped parcel, $F^m(V)$, and the value of the land restricted to remain undeveloped, X:

$$\textit{Market easement value} = F^m(V) - X = aV^{\beta} = M(V) \qquad (2)$$

where M depends on the observation of V.

Since the MALPF program is based on both the landowner's and "the market's" valuation of a parcel's development right, it is necessary to consider how these values may differ. There are two possible sources.

First, landowners may derive utility from owning and operating the farm beyond the agricultural income it provides. This utility would make them less likely to convert the land at the "cash flow optimal" date to convert. We model this effect by multiplying agricultural income by an individual-specific parameter ψ_i in the landowner utility function. Thus landowner i's value for his land is given by:

$$F^f(V) = \max_{\tau} E\left\{\psi_i X\left(1 - e^{-\rho\tau}\right) + e^{-\rho\tau}V(\tau)\right\} \qquad (3)$$

again with the expectation conditional on the current observation of V. We use superscript f to denote the farmer's valuations; the m subscript denotes the market's valuation. We expect $\psi_i > 1$, since this corresponds to landowners placing additional value on farming utility.[28] ψ_i may also be interpreted as representing the individual's private observation of agricultural income, with $\psi_i > 1$ indicating that agricultural income is higher than the market belief. A higher ψ represents a greater weight on farming utility as an income-generating occupation or a higher assessment of

[28]Lynch and Lovell (2003) find both agricultural income indicators and non-consumptive values affect the likelihood of enrollment. These variables include farm size, cropland use, a child planning to take over the farm, and the share of family income from the farm.

agricultural potential. In other contexts, ψ also forms the basis of the unobserved likelihood that the parcel will be developed. This likelihood is a common concern in the design of land use policies (e.g., Sánchez et al. 2006)

We assume ψ_i is known to the individual bidder but is unobserved by the administrator or by other bidders and is distributed independently of ψ_{-i} and of $V(t)$. It thus operates as an *independent private value* which forms the basis of the landowner bid, as in Milgrom and Weber (1982). For ease of notation, we drop the *i* subscript.

Equation (3) has solution $F^f(V) = \psi X + a_f V^{\beta}$ where $a_f = \beta^{-\beta}\left(\psi X/(\beta - 1)\right)^{1-\beta}$. When the farmer observes the same V as the market his valuation of the parcel's development right is:

$$F^f(V) - \psi X = \psi^{1-\beta} a V^{\beta} \equiv \theta M \tag{4}$$

with $\theta \equiv \psi^{1-\beta}$. When $\psi > 1$ we have $\theta < 1$, which implies that for any observation of M the landowner requires less than the market easement value for his development rights.

Second, the market and the individual landowner may have different observations of $V(t)$ and thus of $aV(t)^{\beta}$. Such disagreement about land values is a natural element of any market although it is missing from the standard asset pricing model.

One way to incorporate this difference is to have both parties draw separate observations of $V(t)$, which then correspond to different assessments of $aV(t)^{\beta}$. Let the landowner's assessment be denoted $D = aV^f(t)^{\beta}$, where $V^f(t)$ is his belief about

$V(t)$, analogous to the (un-superscripted) $V(t)$ in equation (2). The landowner's reservation value for his development right, from equation (4), is then:

$$\textit{Landowner easement value} \equiv \theta D \tag{5}$$

This formulation makes clear the two elements that constitute the landowner's reservation value: the extra utility he derives from owning and operating the farm, θ, and his observation of the price his unconstrained property would receive on the market, D.

Bidding

To participate in the MALPF auction an eligible landowner submits a bid, b_i, that represents the one-time payment he would accept in return for his parcel's development rights. After the bid, appraisals are conducted to determine M. For each bid the administrator constructs the ratio of the bid to the market easement value:

$$R_i = \frac{b_i}{M_i} \tag{6}$$

This ratio R_i forms the basis for selecting the winning bids. Winning bidders are paid the amount of their bid, with the exceptions described in Section 2. Based on (6), the probability of acceptance at R □ 1 is $\text{Prob}\left[b \leq \tilde{R}^{*} \cdot \tilde{M}\right]$ where R^* is the (random) cut-off ratio, M is the market appraisal (which is unknown when the bid is submitted) and a tilde denotes a random variable.

To construct the bid, a landowner should first construct the value of his development right conditional on bid acceptance and therefore conditional on the *ex post* report, M. There are many possible informational assumptions. Suppose the landowner adopts the market easement value as his ex post valuation. The landowner's expected value of his development right conditional on winning when submitting bid b would then be:

$$\tilde{v}(b) = E_M(\theta\tilde{M} \mid b < \tilde{R}^* \cdot \tilde{M}) \tag{7}$$

Under this assumption, there is a *common value* between the landowner and the program administrator (e.g., Haile, Hong, and Shum, 2003). Our common value is different from standard models in which different bidders share a common value but its effect on bidding is comparable. Since b must be above θD, equation (7) implies that the expected value of the easement conditional on winning is above the unconditional expectation, $E(\theta M) = \theta D$. [29]

The informational assumption underlying (7) is not the only possible assumption. Rather than treat the market observation M as "correct" in forming his *ex post* valuation, the landowner may treat (5) as a true reservation value that need not be updated based on appraisers' assessments of M. In this case, there would be no common value element to the auction and no winner's curse. An intermediate case would be that landowners use M as informative but not definitive and therefore use Bayesian updating on D. In each case, (7) would be modified accordingly.

Because of these possibilities we allow for a winner's curse correction in our econometric model and estimate its magnitude. Note that we do not assume that the winner's curse exists; we merely accommodate the possibility in our estimation and

[29]We do not consider the case where other bidders' information affects the ex post valuation other than indirectly through the realization of the cut-off ratio, R^*. An alternative assumption would be that M is less informative if all bidders have $D < M$ than if the D's are distributed around M. Likewise, we do not consider the case where D's are affiliated across bidders. These assumptions are obviously open to investigation.

interpretation.

Expression (7) is analogous to the winner's curse correction in standard common value auctions (Milgrom and Weber 1982; Athey and Haile 2002). Landowners whose bids are selected will be those who have most underestimated the market value, ceteris paribus, yielding a winner's curse (under the assumption that M is at least partially informative for the landowner's *ex post* valuation). They therefore have an incentive to correct for this risk by bidding as if they have values above θD. The more that the landowner relies on M in formulating his ex post valuation, and assuming that she incorporates this anticipated information in her bid, the greater is the bid above θD.

Entry and Selection

To illustrate bidder entry and selection we write out the bidder's objective function and use its notation to derive the final component of our model. We do not derive the optimal bidding strategy because of both the complex information environment and the dynamic nature of the auction. Expression (7) allows us to write the objective function for a risk neutral bidder facing a one-time auction as:

$$\pi(\theta, D) = \max_b E_{R,M}\left[(b - \tilde{v})(1 - \Phi(b)) + \left((1-\theta)\cdot\tilde{M}\right)\Phi(b)Prob[offer]\right] - k \qquad (8)$$

The bidder selects his bid to maximize the expectation of profits from winning outright (the first term) or being made an offer less than his bid ex post (second term). Φ represents the cumulative distribution of the product $R^* \cdot M$, Prob[*offer*] is the probability that an ex post offer of $M < b$ is made, and k is a bid preparation cost. In a static framework this ex post offer would always be accepted, regardless of the bid.

In the MALPF program, bidders are faced with the additional decisions of when to enter the auction and when to rebid. Computational derivation of this result is complex and beyond the scope of this paper and thus we describe the results only heuristically.

Consider bidder entry. We assume that landowner surplus falls as θ increases (□□/□□ < 0) which corresponds to the assumption that surplus falls with decreasing private utility from farming or owning land. Although this is a standard assumption it is difficult to establish definitively given the complexity of the winner's curse. Under this assumption we argue, again heuristically, that lower θ's lose more surplus, □□□, from waiting for the next round. This framework further yields □b/□□$_i$ > 0, which implies that lower □'s are more likely to win a given auction.[30]

These conditions together suggest that the lowest reservation price bidders enter the auction earlier and win more frequently when they do enter. As these bidders leave the bidding pool, the remaining pool of potential and actual bidders contains a larger fraction of higher reservation price bidders. This relationship forms the basis for our identification of reservation values since the pattern of bids and ratios over time allows us to trace out the underlying distribution of θ's. This distribution in turn forms the supply curve for development rights.

[30] Formally, this condition relies on D being sufficiently independent of □$_i$. When entry is costless, D will indeed be independent of □$_i$ even under more general assumptions about the winner's curse, because D contains all landowner information about M. Therefore, no landowner believes he has a "high" D relative to M. When entry is costly, D will be negatively correlated with □$_i$.

Estimation Setup

We specify a reduced form bidding function based on the above model. Bids will be above the reservation value, θD, due to the information rent that bidders accrue from their private values, represented by γ, and to the winner's curse correction captured by (7), represented by ω. Our breakdown of bids into a winner's curse correction and a mark-up over the conditional reservation value due to competition is useful for intuition, although it is not strictly consistent with the theoretical model in (8), which does not yield a clear distinction between the two components.

We specify a functional form in which these mark-ups enter multiplicatively above the reservation value. The bidding function is thus:

$$b_{it} = \gamma_{(t)} \omega \theta_i D_i \tag{9}$$

Equation (9) includes bidder and auction-round subscripts, *i* and *t*, to make clear how components vary across observations.

We cannot directly observe the elements of (9) but we do observe related variables that allow us to infer their values:

Private Values (θ). Although we do not observe individual θ_i's, we can infer the average θ among first-bidders in each bidding-round, a sequence we denote as $\{ \bar{\theta}_{(t)} \}$ where t indexes the bidding round. Following the discussion above, we note that as cumulative acceptances increase, there are fewer bidders left in the pool and these are bidders with higher values of θ. Bidders with high θ's, relative to the remaining pool, should sit out auction rounds until they are competitive; that is, until their θ's are in the low range of remaining □'s. Therefore, the range of □'s in a given auction round should be relatively narrowly distributed around "local mean" $\bar{\theta}_{(t)}$, which will be an increasing function of the number of previously accepted parcels, denoted CA_t for cumulative acceptances. This component of our reservation price estimation strategy relies on the assumption that bidders are drawn from an otherwise invariant underlying distribution. We therefore restrict our main regressions to first-time bidders. We adopt the general specification:

$$\ln(\bar{\theta}_{(t)}) = \square - \exp(\phi_1 \square (CA_t+10)) \tag{10}$$

We expect bids to be increasing in cumulative acceptances, $\phi_1 < 0$.[31] The parameter □ measures the lower bound of □'s since $\square = 1 + \ln(\bar{\theta}_{(0)})$.

Information Rent Mark-up (□). Following the standard independent private values paradigm, we allow information rent mark-up, □, to depend on the competitiveness of a given auction round. There are multiple possible measures of competition, such as the number of bidders or the (negative of the) available budget per bidder. We label these generically as $COMP_t$. Let $\ln(\square_{(t)}) = \square_0 + \square_1 COMP_t$. Greater competition implies that bids will be lower, yielding the prediction $\square_1 < 0$. We assume that $\square_{(t)}$ is the same for all bidders in a given round.

---

[31]We were able to identify one Carroll county parcel whose development rights were donated before the MALPF program began. Other unidentified parcels may also have donated their rights. Thus, we add 10 parcels to cumulative acceptances to account for an estimated 0.5 percent of land that was preserved prior to the MALPF program.

Winner's Curse Correction (ω). The winner's curse correction, ω, sets the bid above the unconditional estimate of the easement value to compensate for the fact that the winners will be those who have most under-estimated their market easement value. Theoretically consistent specifications of the winner's curse correction are difficult to derive. Gordy (1998) notes that closed-form equilibrium bidding functions are quite rare except for the simplest of common value assumptions. Paarsch (1992) describes the assumptions that yield a multiplicative or additive winner's curse correction but these derivations assume a different information structure from the MALPF setting.

We use a multiplicative specification in (9) and define $\ln(\omega) = □_0$.[32] In general, we expect ω to be (i) decreasing in θ, (ii) increasing in competition, and (iii) decreasing in bidder experience. We briefly consider each of these issues. The intuition for each of these claims can be gleaned directly from (7), although the full winner's curse correction comes from (8). Interactions among competition, experience and θ greatly complicate formal statements about the winner's curse.

<u>Relationship with θ</u>. The winner's curse correction ω should be higher for bidders with a low θ_i because individuals with low θ_i's are more likely to have their bid accepted for a given signal F^f. Therefore, if their bid is accepted, they must have a lower *ex ante* estimate of the market easement value relative to its draw. They should shade their bids upward further to account for this greater winner's curse. This effect applies to bidders within a round, however, and not across rounds. The

[32]An alternative specification would use an additive mark-up □(□D+□). An additive mark-up has the valuable property that it is a greater percentage mark-up for low reservation values, which are the ones most susceptible to the winner's curse. Unfortunately, estimation did not converge with an additive mark-up.

range of □'s within a round is much smaller than the range of □'s over the entire set of rounds. Therefore, this effect is likely to be small.

<u>Relationship with competition</u>. The winner's curse correction should be higher as competition increases because greater competition means that the cut-off R^* will tend to be lower and therefore winning bidders will have underestimated M to a greater degree. For example, we might write ln(□) = $□_a$ + $□_b COMP_t$, with $□_b > 0$. In this case, the estimated parameter $□_1$ in (11) might be interpreted as measuring the *net* effect of the competition and winner's curse effects, as Hong and Shum (2002) discuss, and may be positive or negative.

<u>Relationship with bidding experience</u>. Second-bidders should be more informed about their parcel's market easement value than first-bidders due to the appraisals that were conducted following their first bid. As a result, second-bidders should have smaller winner's curse corrections. To examine this assumption we estimated a version of (9) with only second bidders.

Estimation Equations

These substitutions now allow us to specify the estimated regression, assuming that M is an unbiased estimate of the landowner's ex ante beliefs, yielding $E(D_i) = M_i$. Taking the log of equation (9) and adding an error term produces:

$$\ln(b_{it}) = \alpha_0 + \gamma_1 \cdot COMP_t + \ln(M_i) - \exp(\phi_1 \cdot (CA_t + 10)) + \varepsilon_i \tag{11}$$

where $□_0 = □_0 + □_0 + □$.

It is straightforward to interpret the slope coefficients in this model, $□_1$ and ϕ_1. The challenge is to identify the constant terms, $□_0$, □, and $□_0$, from $□_0$. An extensive set of identification procedures is necessary because each of the relevant relationships has a constant term that cannot be distinguished without such structure.

When the auction is at its most competitive (represented by $COMP_{max}$) we should have □ □ 1, which implies □$_0$+□$_1 COMP_{max}$ □ 0. Thus, we specify □$_0$ = – □$_1 COMP_{max}$. We discuss the specification of $COMP_{max}$ in Section IV.

In the calculations below we first assume $\omega_0 = 0$, which then yields $\zeta = \alpha_0 - \gamma_0$. The assumption of □$_0$ = 0 captures two possible scenarios. First, the assumption may be valid if there is no winner's curse. This would occur if the value the landowner places on his development right is unaffected by the state's declaration of the land's value; in other words, a pure independent private values setting. Using the notation of (7), let $\tilde{v} = E_M(\theta D \mid b < \tilde{R}^* \cdot \tilde{M})$. Because D is nonrandom at the time of bidding, the conditional expectation is exactly equal to the unconditional expectation, hence no winner's curse. Second, equation (7) could be correct but individuals fail to condition their expectation in forming their bids. In other words, they suffer the winner's curse. This is not an uncommon finding in the experimental literature (Kagel and Levin 2002), although empirical analysis of high-stakes auctions usually finds at least some correction for the winner's curse (Hendricks, Pinkse, and Porter 2003). (The implications of this second explanation are quite different from the first. Since we find at least some winner's curse correction, we leave further discussion of this issue to a separate paper.)

An alternative approach is to use bidder experience to estimate □$_0$. We estimate the analog to (11) using second-bidders, who are more experienced than first-bidders. Note that this estimation must account for the selection effect, recognizing both that second bidders were rejected in the first round and have chosen to rebid and that the timing of the second bid is endogenous. Our heuristic model of entry provides a straightforward approach. Second-bidders should re-enter precisely when they are "competitive," namely at round *t* such that their $□_i$ is in the range of $\overline{\theta}_{(t)}$.[33] It is therefore sufficient to treat second-bidders as equivalent to first-bidders but with a different winner's curse correction. Let $\ln(\hat{\overline{\theta}}_{(t)}) = \hat{\zeta} - \exp(\hat{\phi}_1 \cdot (CA_t + 10))$ with $\ln(\hat{\overline{\theta}}_t)$ derived from a separate first-bidder regression. We then estimate:

$$\ln(b_{it}) = \alpha_1 + \gamma_1 COMP_t + \ln(M_{it}) - \ln(\hat{\overline{\theta}}_{(t)}) + \varepsilon_{it} \tag{12}$$

Let $□_1$ = $□_1$ + $□_0$ + □ where $□_1$ is the winner's curse correction for second-bidders. Suppose that $□_1$ reflects only information-rent effects and is unaffected by the winner's curse. Since our econometric results suggest this to be the case, we feel comfortable in imposing it here. Then we can identify $□_0$ as described above. If we assume that second-bidders are perfectly informed then $□_1 = 0$ and we can identify □ using the same procedure as for first-bidders. Using this □ estimate we can then infer $□_0$, the winner's curse correction for first-time bidders. These assumptions yield the prediction $□_1 < □_0$ which serves as an additional test to their validity.

[33]Our arguments further imply that second-bidders should have □'s in the upper range of the $□_t$'s at the time of first bid and that the greater the lag after which a second bidder enters, the higher his bidder was likely to have been in the set of □'s at the time of first bid. We do not model these effects in the current paper.

Alternative specifications. The estimated □ distribution depends on the functional form. We estimated two alternative specifications to (11). The first uses logged cumulative acceptances:

$$\ln(b_{it}) = \alpha_0 + \gamma_1 \cdot COMP_t + \ln(M_i) + \phi_2 \cdot \ln(CA_t + 10) + \varepsilon_i \tag{13}$$

We then calculate $\bar{\theta}_{(t)} = (CA_t + 10)^{\beta_2} \exp(\zeta)$. We expect $\phi_2 > 0$.

We also estimated a specification in which we replaced □*M* with *F*-□*X*, from (4). This specification applies when the landowner, in forming the expected value of his development right conditional on winning, takes the market observation of $\tilde{F}^m$ as correct but does not update any market assessment of *X*. That is, the reservation value conditional on winning is $\tilde{v} = E(\tilde{F}^m - \psi X \mid b < \tilde{R}^*(\tilde{F}^m - X))$ with the expectation taken over $\tilde{F}^m$. This expression is identical to (7) whenever the landowner and the market share identical assessments of *X*.

Since □ □ 1 we use the functional form □ = 1+exp(ϕ_3+ ϕ_4(CA_t+10)). We converted the results to □ using □ = □$^{1-□}$ where □ comes from the underlying stochastic process, as in (2). The estimated equation is:

$$\ln(b_{it}) = \alpha_0 + \gamma_1 COMP_t + \ln(F_i^m - (1 + \exp(\phi_3 + \phi_4(CA_t + 10))) \cdot X_i) + \varepsilon_i \tag{14}$$

We expect bids to be increasing in cumulative acceptances, which requires $\phi_4 < 0$. There is no prediction for ϕ_3.

IV. Data

We collected a unique data set that includes parcels that satisfied the eligibility requirements and have enrolled in an agricultural district in Carroll County starting with the program's inception in 1977 to 1999. Because of the timeframe, we returned to the original parcel-level files via microfiche to track each parcel through each round of bidding. Changes of ownership presented the largest data integrity challenge. We constructed a panel dataset with an observation for each agent-parcel in each bidding round. The final data set is carefully constructed to ensure that we correctly followed each agent-parcel combination. By checking the MALPF's annual published reports, which publish the number of bidders and total acquisition expenditures, against the sum of accepted bids based on the parcel histories, we were able to ensure that we had as complete a record as possible of all bids. Final data include bids, parcel sizes, appraisals, agricultural values, and outcomes, as well as assorted parcel characteristics. We convert all dollars to $2002 using the Northeast Housing CPI.

Several events in the early 1990's affected the operation of the MALPF auction. Between 1990 and 1995, to ease the administrative pressure of one bidding round per year, the state decided to accept bids twice during each fiscal year. Thus, we have two round of bidding for some of these years. A state budget crisis led the state to rescind the MALPF budget after bids had been submitted in the first round of 1991 and thus no bids were accepted. No MALPF auctions were held for two years, although a few bids were submitted in the second round of 1991 and the first round of 1992. Funding was restored in 1993. In the first round of 1993, offers were made to 35 bidders on hold from the first round of 1991 bidding (those who would have been accepted had funding been available). And regular bidding began again in the second round of 1993.

In total, we analyze the bids in 22 auction rounds between 1980 and 1999. A second, competing preservation program became active in the County around 2000 which did not use an auction format. Therefore, we chose to use the data prior to this event. Data summaries are provided in Tables 1 and 2.

There are 306 unique agent-parcels with submitted first bids (a total of 574 bids are included in the dataset when multiple bids are counted). We lose first bid observations for three reasons. (1) The state did not conduct appraisals for all bids in all years, primarily because the bids were high relative to the budget, (3 first bids lost). (2) In rounds with no budget (1991 and 1992), no appraisals were conducted (21 first bids lost). (3) Some bids are clearly wild guesses or pure gambles. We drop bids with ratios greater than 3.31, the highest ratio ever accepted (5 first bids lost). The remaining usable first bids equaled 277.

The MALPF changed the formula by which agricultural values were calculated twice, once in 1990 and again in 1996. To adjust agricultural values to a common formula in this model, the post-1996 formula is assumed to be correct, since the changes were adopted by program administrators solely to make the values closer to the true value of agricultural production on the land. We estimated a model of average per-acre agricultural values on dummy variables for pre-1991 and for 1991-1996 and used the coefficients to calculate a consistent agricultural formula. The formula we use is $X = \{0.263X_t$, t□ 1991; $0.638X_t$, 1991□ t□ 1996; X_t, 1996□ t□ 1999$\}$ where X_t is a given parcel's agricultural value per-acre as reported by the program. We subtract this adjusted measure of X to obtain $M = F - X$.

There are several reasons why we think this solution adequately corrects for measurement error in *X*. First, the regression coefficients to adjust *X* were essentially invariant to the sample we used to estimate them. In other words, these changes in *X* were "across the board" and not restricted to any class of bidders, such as those with large parcels or rebidders. Second, other adjustments to *X* (such as no adjustment or half-adjustment) yielded results that were nonsensical, such as bids being far below the estimated reservation values. This suggests that landowners' assessment of agricultural value were closest to those computed under the post-1996 formula. Third, our estimates of equations (11) and (14) yielded quite similar results. These expressions are identical if and only if the *X*'s are measured correctly.

In the analysis of competition, we consider two possible budgets: (i) a statewide program budget for the given bidding round that was publically available as bids were being prepared. Since counties received more-or-less constant shares of this budget, this statewide figure is a good measure of budget availability in a given round; and (ii) acquisition expenditures at the county level for the previous bidding round, inflated for the previous year. Acquisition expenditures in 1985, for example, were inflated using the July 1985 CPI.

V. Results

Competition Effects

Program administrators are often eager to know the role of competition in reducing bids. Some situations might naturally have few eligible bidders, so policy-makers would like to know how successful an auction might be in driving down procurement costs. In other situations, administrators may be able to increase the number of eligible bidders either by relaxing the eligibility criteria or publicizing the auction more widely. They would then like to know whether the costs of these actions would likely be covered by the savings generated from lower bids. Finally, administrators may want to use the effects of competition to argue to policymakers that a competition-based design is worth the increased complexity.

We examined several measures of competition: (i) the total number of bidders in each round, (ii) the State's announced budget for each round divided by the number of bidders, and (iii) the County's expenditure on parcels in the previous round divided by the number of bidders. Each competitiveness measure is also associated with a value for $COMP_{max}$, needed to identify $\square_0$. When competitiveness is measured by the budget or expenditure per bidder, we set $COMP_{max} = 0$, since the most competitive auction has a zero budget; this then yields $\square_0 = 0$.[34] When competitiveness is measured by the number of bidders, we set $COMP_{max} = 60$, which is above the highest number of bidders we observed, 53. This latter assumption is somewhat ad hoc and for this reason, among others, we rely mostly on the budget and expenditure measures.

Endogenous competition is a problem in many empirical auction analyses. The largest potential problem in our context arises when competition is measured by the number of bidders. A large budget may attract more bidders but leave the budget-per-bidder unchanged, therefore yielding no change in the likelihood of a given bid being accepted and thus no change in the auction's competitiveness. Endogeneity problems are probably less severe when competition is measured in terms of budget-

---

[34]These values for $COMP_{max}$ are out-of-sample. Kirwan, Lubowski, and Roberts (2005) note that predictions based on an out-of-sample regressor are typically undesirable. Such prediction is unavoidable here, however, since out-of-sample prediction is key to inferring true reservation values.

per-bidder and therefore we focus our attention on these regressions. Year-to-year variation in the funds allocated to the MALPF reflects both variation in state revenues and in the conversion of agricultural land. Although these phenomena may be correlated with bidder valuations, this correlation should be captured by the market easement value. Therefore, these budget measures are likely to be exogenous regressors.

Still, to account for potentially endogenous numbers of bidders, we developed regressions to predict both number of bidders and budget-per-bidder. (See Haile, Hong, and Shum (2003) for a discussion of instruments used to address endogeneity.) Much of the variation in the number of bidders depends on the State and County's promotion of the MALPF program and on publicized problems or successes. Both elements have distinct time components which we capture through period dummies. Intuition also suggests that successful bidders will beget more bidders in subsequent rounds and therefore we include as an instrument the number of bidders accepted in the previous round. Results are shown in Table 4 for the full 24 bidding rounds.

We use the predicted values from Table 4 either directly, as a predicted number of bidders, or indirectly to construct the State budget per predicted bidder, predicted State-budget-per-bidder, County expenditure (at t-1) per predicted bidder, and predicted County expenditures at t-1 per bidder at t. We refer to these latter four measures generically as budget-per-bidder.

Results are shown in Tables 5 and 6. All of the predicted competition variables show that greater competition leads to lower bids. The possible permutations to instrument for competition are large; Table 5 represents a range of possible choices. Table 6 shows alternative specifications with two measures of competition: state budget per predicted bidder and predicted County expenditures (at t-1) per bidder (at t). These results are therefore most directly comparable to regressions 2 and 4.

To judge the magnitude of the effects and demonstrate the range of estimates, we calculated for each regression the implied mark-up above the reservation value at the median level of competitiveness, denoted $\square_{median}$. We find $\square$, the bid multiplier, ranges from 1.05 to 1.39, depending on our measure of competition. For the more appropriate budget-per-bidder measures of competition, we find $\square$'s ranging from 1.05 to 1.15. In other words, MALPF bids are roughly 5 to 15 percent above bidders' conditional reservation values.

We also calculated the predicted percentage change in the median bid due to one additional bidder, holding the budgets fixed. This number represents, roughly speaking, the value to the auctioneer of attracting one more bidder. We found that one additional bidder on top of the mean 23 bidders reduced bids by 0.1 percent to 1.4 percent, with a mode of 0.2 percent.[35] This calculation has the advantage that it does not depend on identification of $\square_0$. Its interpretation does, however, rely on the additional bidder being the same as existing bidders. In dollar terms, for example, in 1999, the average bid per acre of $2,839 would have been lower by $2.84 to $39.75 per acre.

[35]It is difficult to know how these findings compare to other auctions since this calculation is rarely reported in the literature. Bajari and Hortacsu (2003) found that adding one more bidder resulted in bids that were 3.2 percent closer to the reservation value. This number may be higher than ours because the number of bidders in their auction was much smaller than in the MALPF.

Private Values Distribution

One of our main objectives is to estimate the distribution of □, the bidder-specific parameter that captures the landowner's taste for owning and operating a farm and thus forms the basis for his reservation value. This private-value is the motivation for competition-based policies, which are designed to drive bids as close to the unobserved □ as possible. The distribution of □ is a key element of participation in land use policies generally, not only those that are competition-based, and thus affects enrollment in all voluntary enrollment land programs. This parameter is a crucial factor determining the probability of land conversion, which in turn is the main variable of interest in many environmental land use policies (Sánchez et al. 2006). Finally, □ can also indicate the compensation that should be awarded for eminent domain takings, since it can be used as a measure of how much a landowner values owning her property above the market price.

We estimate this parameter using (11), (13), and (14). Results from these regressions are presented in Tables 5 and 6. In all cases, the coefficients are consistent with higher cumulative acceptances leading to higher bids: $\phi_1 < 0$ (regressions 1-5 using equation (11)), $\phi_2 > 0$ (regressions 6 and 7 using equation (13)), and $\phi_4 < 0$ (regressions 8 and 9 using equation (14)).

With these results we can infer the distribution of □'s, using the group-means implied by equation (10) and first making the assumption that the winner's curse correction is zero. Graphs of the implied distribution functions are shown in Figure 1 for regressions 1 through 5 and Figure 2 for regressions 4 and 6 through 9; regression 4 is included in both graphs for comparison. We use □ = 1.54 to convert the estimated □'s from regression 12 to □'s, using □ = $□^{1-□}$, where □/(□-1) is the trigger between investment value and investment cost from (1).[36]

Recall that □ < 1 and that lower values of □ represent higher values for farm ownership and thus lower reservation easement values. Figures 1 and 2 show distributions of □'s that are far below one. Among regressions 2-7, which use the most appropriate measures of competition and do not require the additional parameter β, the highest □'s at the end of the period we analyze are below 0.80. These relatively low values are perhaps not too surprising since the program is never valuable in expectation to a risk-neutral bidder with □ = 1; the program is designed to attract and enroll low □ farmers. Our findings are a bit surprising, however, since they suggest that all of the parcels could have been obtained with a take-it-or-leave-it offer of 0.80, even under a conservative assumption of no winner's curse correction. In contrast, the median ratio among accepted parcels over this period was 0.89.

Our estimates of the □ distribution are robust to functional form and competition measure. They are, however, sensitive to auxiliary assumptions about the level of competition that would drive bids to the reservation value and, in regressions 8 and 9, to the value of β, sometimes called the volatility parameter. For example, the highest value of $\hat{\hat{\theta}}_{(0)}$ comes from regression 8. If we instead take □=2 for this regression then $□_{(0)}$ = 0.37, which is quite similar to the regressions based on (9).

[36]We apply Dixit and Pindyck's notation (1994, p. 142) with parameters from Quigg's study of land development: □=0.08, □=0.03, and □=0.20 (Quigg 1993).

Winner's Curse

MALPF auction bidders would appear potentially vulnerable to a winner's curse since landowners bid without knowing the state's appraisal of the easement value. This market easement value should be informative to bidders and bids that are selected are those that are lowest relative to this value. Bidders who recognize the winner's curse should raise their bid above their unconditional reservation value.

More informed bidders should be less vulnerable to a winner's curse. An obvious instrument for bidder information is the comparison between first and second bidders. Therefore, to get a rough estimate of the winner's curse correction, we estimated (12) using only second-bidders, using results from regressions 2, 3, 4, and 5 for $\hat{\hat{\theta}}_{(t)}$, and compare the results to the same regressions for first-bidders There are 109 second-bidders with full data. Estimates are shown in Table 7.

We find substantially smaller intercepts for second-bidders in all four cases, as expected. If we assume that second-bidders are perfectly informed and that the winner's curse correction for first-time bidders is invariant to all other factors, then the difference $\hat{\alpha}_0 - \hat{\alpha}_1$ is a measure of the winner's curse correction. Based on Tables 5 and 7 we find winner's curse corrections of 8-14 percent of the reservation value.

We have some evidence that the winner's curse correction is indeed invariant to other factors. Initially, we suspected that greater competition would increase the winner's curse correction and therefore lead to higher bids (e.g., Pinkse and Tan 2005). If these conditions held then γ_1 should be larger in magnitude for second-bidders than first-bidders. We find the opposite result, however, since the $\hat{\gamma}_1$s in Table 7 are uniformly below their counterparts in Table 5. This result leads us to conclude that the effect of competition on the winner's curse correction is non-existent and allows us to focus on models with a winner's curse correction that is invariant to all factors except bidder information.

A more sophisticated analysis of the role of information and experience would be valuable but lies beyond the scope of this paper. Joint estimation of first- and second-bidders would appear to be the most valuable approach but multiple sample selection issues would arise under this treatment, including whether offers were made to bidders at $b > M$ and whether they were accepted. We therefore leave this topic for future research.

VI. Comparison to Alternative Policy

This section describes alternatives to competition-based environmental policy and presents our assessment of one prominent alternative to the MALPF auction, the take-it-or-leave-it offer, based on the results of Section V.

It is useful first to define categories of non-competition-based voluntary-enrollment programs. One such category is *formula-based* payments or point systems, in which the state offers a payment to enrollees that is based solely on observable parcel characteristics. Formula-based payments are used for the Conservation Reserve Enhancement Program, and have begun to achieve popularity in farmland preservation programs (Maryland Rural Legacy Program, Ohio's Agricultural Easement Purchase Program) and are also the main format of programs internationally (see Sánchez-Azofeifa *et al*. 2007). The take-it-or-leave-it offer is a form of formula-based payment. The third category is *negotiation-based*, which is the way most land ownership transactions are conducted. This is also the method by which non-governmental organizations such as The Nature Conservancy acquire property or development restrictions.

Our analysis of the MALPF auction leads naturally to the question of whether auctions might be expected on theoretical grounds to be generally superior to a formula-based or negotiation-based approach. The MALPF auction is similar to Bulow and Klemperer (1996) in which an English auction with N+1 bidders is compared to an auction with N bidders in which the auctioneer makes a take-it-or-leave-it offer to the winner after observing all submitted bids. They argue that regardless of the information structure, the pure auction always yields higher expected revenue than any take-it-or-leave-it offer with fewer bidders. This finding suggests that the MALPF set-up might be superior to any take-it-or-leave-it approach, although there remain many differences between their model and the real world MALPF auction.

A few empirical papers have addressed environmental policy questions similar to ours. Stoneham et al. (2003) analyzed auctions for conservation contracts in Australia and compared the existing auction with an alternative take it or leave it offer. They assumed that bids were equal to reservation values rather than inferring them. In this case, an auction always does better than a take-it-or-leave-it offer. Connor, Ward, and Bryan (2007) examined a second Australian program and compared an auction format to various uniform payment schemes and negotiated payments, but did not analyze bids econometrically and did not attempt to infer reservation values. Cummings, Holt, and Laury (2004) examined bids to sell irrigation rights in the Georgia Irrigation Auction and suggested that the rights could have been obtained more cheaply if landowners had been allowed to revise their bids after an initial bidding round. They did not analyze bids econometrically and their recommendation is based on their induced-value auctions. Messer and Allen (2008), examining a land preservation auction in Delaware that is similar to the MALPF auction, use existing bids and benefit measures to show the consequences of alternative parcel selection procedures. They also did not analyze bids econometrically and did not consider how bidding behavior might change if an alternative selection rule were used. Reichelderfer and Boggess (1988) addressed a similar question for the CRP.

We compare the MALPF program to a take-it-or-leave-it (TIOLI) offer under which the state would offer to purchase development rights from all interested landowners at a set percentage of their market easement value.[37] We selected this policy for comparison because it contains the same elements as the current MALPF program yet does not rely on competition among enrollees.[38] The TIOLI would be less expensive than the current set-up if the MALP auction's bids were greatly above landowners' reservation values (due either to high information rents or a large

[37]In the real world, policymakers would have to worry about running into the budget constraint if too many landowners were willing to accept the take-it-or-leave-it offer. The TIOLI policy would then have to include some sort of rationing rule, such as first-come-first-served. In our simulations we can choose the TIOLI offer that would exactly meet the budget, so this concern does not arise. We thank a referee for pointing out this important distinction.

[38]The Delaware Agricultural Land Preservation Foundation (DALPF) program is a competition-based program under which appraisals occur before the bidding thus eliminating the largest potential opportunity for a winner's curse (DALPF 2007). Landowners submit ratio bids in a sealed-bid first-price auction. The DALPF purchases the easements with the lowest ratios until the budget is exhausted.

winners' curse correction) and if the distribution of reservation values were relatively flat. These conditions can be verified only through empirical analysis.

Let t be the announced ratio offer. A bidder is predicted to accept the offer if his θ is below t. He would be paid tM and his surplus from the program would be $(t-\theta)M$. To assess the TIOLI we then find the lowest t that would result in acceptances totaling approximately \$53.3 million, the amount spent by the state in Carroll County over the period of study. Using regression 4 as an example, we find that $t = 0.602$ would cost \$53.3 million and would enroll 24,896 acres.

We next simulate MALPF auction bidding.[39] We multiply the reservation values θM by a uniform mark-up, γ. We use a uniform mark-up because otherwise we must simulate budgets and competition year-by-year. Accepted bidders are paid $\gamma\theta M$; we first assume no winner's curse correction. We rank bidders based on θ and sum the bids, starting with the lowest θ, until acceptances total \$53.3 million. Again using regression 4 as an example and a mark-up of 1.06, we find that all bidders with θ below 0.67 would be accepted. The program would cost \$52.6 million and would enroll 27,841 acres. (Because the θ's form a step function it is not possible to hit the budget-target dead on.)

In other words, the MALPF auction would enroll as much as 3,000 more acres

[39]For the comparison we simulate the bids rather than use actual bids and enrollments for multiple reasons: (1) Actual enrollments depend on yearly budgets and competition, which do not have direct counterparts in the TIOLI. (2) Actual enrollments include rebidders but the proper comparison is with first bidders. (3) The regressions are an unbiased representation of bidding behavior but not an unbiased representation of winning bids. This affects the simulation because actual bids include an error term; because accepted bids are more likely to have large negative errors they are not a representative sample for comparison with TIOLI. We can either simulate these errors in constructing a simulated TIOLI (for comparison with actual auction enrollment and costs) or can simulate the bids so that they are directly comparable to the TIOLI. We chose the latter approach. (4) Changes in the agricultural formula complicate comparison of raw bids across time.

for the given budget than the comparable TIOLI approach, or roughly 12 percent more acreage. Other regressions yield similar results, with the MALPF auction consistently enrolling more acres for the given budget. In addition, an added plus is that when using the MALPF auction approach, the administrator does not need to know the distribution of θ's. Under the TIOLI when the distribution of θ's is not known, the administrator must guess the right t to meet the budget.

The degree of superiority of the auction mechanism is sensitive to assumptions about the winner's curse. If we assume that bidders apply a winner's curse correction of, say, 12 percent then the true reservation values are lower than the values shown in Figure 2 by 12 points. Under this assumption about reservation values, the TIOLI would enroll 26,547 acres for \$52.5 million. The MALPF auction calculations remain unchanged. Thus, even when bids contain a winner's curse correction, the MALPF program remains superior to the TIOLI. The MALPF auction would enroll 1,300 more acres for the given budget with a 12 percent winner's curse correction, or a little over 5 percent more acreage. In other words, even though the TIOLI's relative performance is improved when we assume auction participants correct for a winner's curse, the MALPF continues to outperform the TIOLI.

VII. Concluding Comments

The unique set-up of the MALPF program has allowed us to estimate several components of landowner bidding behavior. Our model includes an independent private value component and a possible common value component between each bidder and the administrator.

We find that bids are 5-15 percent above the conditional reservation value. Our finding is similar to a field experiment on which the Georgia Irrigation Reduction Auction is based which found that hypothetical water rights sold for 7-12 percent above reservation values (Cummings, Holt, and Laury 2004). We further find consistent evidence that greater competition leads bidders to reduce their bids. Policy makers have been motivated to adopt competition-based policies based on this theoretical feature which we uncover empirically to be true. We find that each additional bidder reduces bids by 0.1-1.4 percent, or \$2.84 - \$39.75 per acre. This suggests efforts to encourage more competition among bidders are worthwhile.

Because of the prominent role for the market easement value in the MALPF auction, we considered the possibility that bidders adjust for a possible winners' curse. Through a comparison of first-time bidders with more informed repeat bidders, and correcting for endogeneity in the timing of a second bid, we estimate that first-time bidders adjust for a possible winner's curse by increasing their bids between 8 and 14 percent. Our finding suggests that a modification to the MALPF program to conduct appraisals before bidding, as under Delaware's DALPF auction, might lead to lower bids by reducing a key source of uncertainty.

We derive and demonstrate a unique approach to inferring reservation values based on bidder entry and selection. These reservation values allow us to compare the MALPF auction with an alternative policy that does not involve bidding, called a take-it-or-leave-it offer. We find that the MALPF program would enroll 5 to 12 percent more acreage for a given budget than an "ideal" take-it-or-leave offer. In short, the competition-based design of the MALPF program appears to have paid off in this instance. The MALPF approach also has the advantage that the administrator does not need to know the underlying distribution of landowners' valuations to reach the budget target.

Many empirical issues warrant further exploration. These include possible bidder collusion, asymmetric bidders, within-round distribution of ▯, bidding by re-bidders, and the decision by bidders with ratios above one to accept a payment of *R*=1. Each of these topics would shed light on bidder behavior and thus potentially lead to better design of the MALPF auction and competition-based policies more generally. We leave these issues for future research.

VIII. References

Athey, Susan, and Philip Haile. 2002. "Identification of Standard Auction Models." *Econometrica* 70 (6): 2107-40.

Bajari, Patrick, and Ali Hortacsu. 2003. "The Winner's Curse, Reserve Prices, and Endogenous Entry: Empirical Insights from eBay Auctions." *RAND Journal of Economics* 34 (2): 329-55.

Bulow, Jeremy, and Paul Klemperer. 1996. "Auctions Versus Negotiations." *American Economic Review* 86 (1): 180-94.

Connor, Jeffrey, John Ward, and Brett Bryan. 2007. "Exploring The Cost Effectiveness of Land Conservation Auctions and Payment Policies." American Agricultural Economics Association Annual Meeting. Portland, Oregon.

Cummings, Ronald, Charles Holt, and Susan Laury. 2004. "Using Laboratory Experiments for Policymaking: An Example from the Georgia Irrigation Reduction Auction." *Journal of Policy Analysis and Management* 23 (2): 341-63.

Delaware Agricultural Land Preservation Foundation (DALPF). 2007. Farmland Preservation in Delaware. http://dda.delaware.gov/aglands/lndpres.shtml November.

Dixit, Avinash K., and Robert S. Pindyck. 1994. *Investment under Uncertainty*. Princeton, New Jersey: Princeton University Press.

Echevarria, John. 2005. “Regulating Versus Paying Land Owners to Protect the Environment.” *Journal of Land, Resources, and Environmental Law* 26 (1): 1-46.

Goldberger, Arthur S. 1968. “The Interpretation and Estimation of Cobb-Douglas Functions.” *Econometrica* 35 (3-4): 464-472.

Gordy, Michael B. 1998. “Computationally Convenient Distributional Assumptions for Common Value Auctions.” *Computational Economics* 12 (1): 61-78.

Haile, Philip A., Han Hong, and Matthew Shum. 2003. “Nonparametric Tests for Common Values in First-Price Sealed Bid Auctions.” NBER Working Papers 10105. National Bureau of Economic Research, Inc.

Hendricks, Kenneth, Joris Pinkse, and Robert H. Porter. 2003. “Empirical Implications of Equilibrium Bidding in First-Price, Symmetric, Common Value Auctions.” *Review of Economic Studies* 70 (1): 115-45.

Hong, Han, and Matthew Shum. 2002. “Increasing Competition and the Winner’s Curse: Evidence from Procurement.” *Review of Economic Studies* 69: 871-98.

Kagel, John H., and Levin, Dan. 2002. “Common Value Auctions and the Winner’s Curse: Lessons from the Laboratory.” In *Common Value Auctions and the Winner’s Curse* ed. J. H. Kagel and D. Levin, Princeton, NJ: Princeton University Press.

Kirwan, Barrett, Ruben Lubowski, and Michael J. Roberts. 2005. “How Cost-Effective Are Land Retirement Auctions? Estimating the Difference between Payments and Willingness to Accept in the Conservation Reserve Program.” *American Journal of Agricultural Economics* 87 (5): 1239-47.

Latacz-Lohmann, Uwe, and Carel Van der Hamsvoort. 1997. “Auctioning Conservation Contracts: A Theoretical Analysis and an Application.” *American Journal of Agricultural Economics* 79 (2): 407-18.

Lynch, Lori, and Sabrina Lovell. 2003. “Combining Spatial and Survey Data to Explain Participation in Agricultural Land Preservation Programs.” *Land Economics* 79 (2): 259-76.

Messer, Kent D., and William Allen III. 2008. “Applying Optimization and the Analytic Hierarchy Process to Enhance Agricultural Preservation Strategies in the State of Delaware." Department of Economics: Working Paper No. 2008-18. University of Delaware, Newark, Delaware.

Milgrom, Paul R., and Robert J. Weber. 1982. "A Theory of Auctions and Competitive Bidding." *Econometrica* 50 (5): 1089-1122.

Paarsch, Harry. 1992. "Deciding Between the Common and Private Value Paradigms in Empirical Models of Auctions." *Journal of Econometrics* 51 (1-2): 191-215.

Pinkse, Joris, and Guofu Tan. 2005. "The Affiliation Effect in First-Price Auctions." *Econometrica* 73 (1): 263-277.

Plantinga, Andrew, Ruben Lubowski, and Robert Stavins. 2002. "The Effects of Potential Land Development on Agricultural Land Prices." *Journal of Urban Economics* 52 (3): 561-81.

Quigg, Laura. 1993. "Empirical Testing of Real Option-Pricing Models." *Journal of Finance* 48 (2): 621-40.

Reichelderfer, Katherine, and William G. Boggess. 1988. "Government Decision Making and Program Performance: The Case of the Conservation Reserve Program." *American Journal of Agricultural Economics* 70 (1): 1-11.

Sánchez-Azofeifa, Arturo, Alexander Pfaff, Juan Robalino, and Judson Boomhower. 2007. "Costa Rican Payments for Ecological Services Program: Intention, Implementation, and Impact." *Conservation Biology* 21 (5): 1165-1173.

Schilizzi, Steven, and Uwe Latacz-Lohman. 2007. "Assessing the Performance of Conservation Auctions: An Experimental Study." *Land Economics* 83 (4): 497-515.

Smith, Rodney. 1995. "The Conservation Reserve Program as a Least-Cost Land Retirement Mechanism." *American Journal of Agricultural Economics* **77** (1): 93-105.

Stoneham, Gary, Vivek Chaudhri, Arthur Ha, and Loris Strappazzon. 2003. "Auctions for conservation contracts: an empirical examination of Victoria's Bush Tender trial." *Australian Journal of Agricultural and Resource Economics* 47 (4): 477-500.

Vukina, Tomislav, X. Zheng, Michele Marra, and Armando Levy. Forthcoming. "Do Farmers Value the Environment? Evidence from the Conservation Reserve Program Auction." *International Journal of Industrial Organization* in press.

IX. *Tables and Figures*

Table 1 Summary statistics For MALPF Program, Carroll County 1980-1999: Budgets and bidders

Bidding year and round	Statewide budget, thousands ($2002)	Carroll County expenditures, thousands ($2002)	Number of bidders	Number of first-time bidders	First-time bids used in estimation
1980	$5,208	$3,528	11	11	11
1981	$10,221	$4,647	32	32	32
1982	$10,933	$2,977	35	22	21
1983	$12,174	$3,048	53	38	36
1984	$8,762	$1,953	42	24	24
1985	$11,666	$1,565	29	14	13
1986	$13,627	$1,911	17	7	7
1987	$13,539	$1,096	20	17	15
1988	$12,740	$1,931	12	3	3
1989	$17,064	$3,682	27	19	19
1990	$24,794	$6,006	25	24	23
1990-2*	$17,835	$1,821	6	2	2
1991	$0	$0	29	22	22
1991-2	$0	$0	1	1	0
1992	$0	$0	4	1	0
1993	$6,795	$558	0	0	0
1993-2	$8,008	$867	22	13	4
1994	$7,120	$812	25	4	4
1994-2	$6,605	$806	30	5	3
1995	$6,535	$1,342	30	4	2
1995-2	$6,720	$795	21	1	1
1996	$12,030	$3,074	26	7	7
1997	$18,872	$3,199	22	6	4
1998	$22,986	$5,734	34	25	21
1999	$25,466	$2,522	21	4	3
Mean	**$11,188**	**$2,155**	**23**	**12**	**11**
Total	**$279,701**	**$53,875**	**574**	**306**	**277**

*In 1991 the MALPF Program funding was cut to cover a statewide budget deficit; however, the 29 bidders from 1991 were ranked and the lowest bids were funded in 1993 Round 1. All bidders in 1991 Round 2 and 1992 were summarily rejected. No bids were accepted in 1993 Round 1. Normal bidding resumed in 1993 Round 2.

Table 2 Summary statistics for MALPF program, Carroll County, 1980-1999: Bids and ratios (N=545)

Bidding year and round	Mean ratio	Ave. bid per acre ($2002)	Highest accepted ratio	Accepted bids	Cumulative acceptances (CA)	Ave. payment per accepted acre ($2002)
1980	1.39	$3,407	2.10	10	10	$3,410
1981	1.15	$2,169	2.43	18	28	$1,963
1982	1.00	$1,613	0.95	16	44	$1,390
1983	1.26	$1,619	1.00	17	61	$1,332
1984	1.30	$1,528	1.00	11	72	$1,334
1985	1.20	$1,337	1.00	10	82	$1,175
1986	1.34	$1,305	3.31	12	94	$1,346
1987	1.52	$1,766	1.35	7	101	$1,591
1988	0.92	$1,662	1.20	7	108	$1,683
1989	1.29	$2,703	1.70	16	124	$2,520
1990	1.23	$3,827	1.79	20	144	$2,855
1990, 2nd round	1.16	$3,202	1.65	6	150	$3,202
1991	1.17	$3,139	1.31	6	156	$3,593
1991, 2nd round	.	$2,290	.	0	156	$0
1992	1.58	$2,676	0.00	0	156	$0
1993	.	.	.	0	156	$0
1993, 2nd round	0.86	$2,512	0.92	5	161	$2,053
1994	0.97	$2,428	0.59	1	162	$2,504
1994, 2nd round	0.83	$2,548	0.68	3	165	$1,906
1995	0.93	$2,679	0.74	3	168	$2,475
1995, 2nd round	0.92	$2,482	0.84	5	173	$1,481
1996	0.92	$2,287	0.87	10	183	$2,120
1997	0.92	$2,639	0.94	17	200	$2,197
1998	0.90	$2,980	0.87	16	216	$2,395
1999	0.87	$2,839	0.81	7	223	$2,567

Table 3 Summary Statistics for Variables Included in Regression Analysis

	Mean	Median	Std. Dev.	Min.	Max.	n
Tables 5, 6, & 7 (first bidders)						
Bid per acre	2282	2021	1120	563	11444	277
ln(Bid per acre)	7.64	7.61	0.41	6.33	9.34	277
Number of bidders	31.7	30	10.99	6	53	277
State budget/bidder	529	402	340	208	2972	277
County exp/bidder	91.6	73.6	84.3	25.3	1001	277
Market easement value	3556	3377	1123	1432	10153	277
Cumulative acceptances	90	82	63.08	0	222	277
ln(Cumulative acceptances+10)	4.33	4.52	0.83	2.3	5.45	277
Table 4 (rounds)						
Number of bidders	24	25	11.88	1	53	24
State budget/bidder	604	432	600	0	2972	24
County exp./bidder	117.0	73.1	196.4	0	1001	24
Accepted last round	9.25	9	5.35	0	18	24

Table 4 Estimated Coefficients for Predicting Three Measures of Competitiveness: Number of Round (n = 24)

	Dependent variable		
	Number of bidders in round *t* #A	**State budget at *t* per bidder at *t* #B**	**County expenditure at *t*-1 per bidder at *t* #C**
County expenditure at t-1	--	0.33 (5.07)	0.12 (8.01)
State budget announced for t	-7.82×10^{-5} (0.26)	--	--
# accepted parcels at t-1	1.11 (2.59)	-53.61 (2.61)	-19.46 (4.05)
1991-1992 = 1	-15.63 (3.02)	404 (1.77)	182 (3.41)
Post-1992 = 1	1.29 (0.31)	-206 (1.13)	-49.84 (1.17)
Constant	16.64 (3.08)	400 (1.77)	22.02 (0.42)
Prob. > F	0.0076	0.0008	0.000
R^2	0.50	0.62	0.80

Table 5 Estimated Coefficients for ln(Bid per acre), first-time bidders only (equation (11), n = 277)

	#1	#2	#3	#4	#5
$\square_0$	0.57 (6.47)	0.091 (2.46)	-0.10 (1.03)	0.033 (0.55)	-0.027 (0.35)
$\square_1$: Predicted Bidders (#A)	-0.014 (5.85)	--	--	--	--
$\square_1$: State Budget at *t* per Predicted Bidder (#A)	--	0.00031 (6.69)	--	--	--
$\square_1$: Predicted State-Budget-per-Bidder (#B)		--	0.00023 (3.27)	--	--
$\square_1$: County Exp. at *t*-1 per Predicted Bidder (#A)	--	--	--	0.00082 (2.79)	--
$\square_1$: Predicted County-Exp.-per-Bidder (#C)	--	--	--	--	0.00065 (3.46)
ϕ_1	-0.0046 (6.10)	-0.0028 (4.90)	-0.0073 (4.37)	-0.0057 (5.71)	-0.0071 (4.58)
$\square_{median}$[b]	1.39	1.13	1.10	1.06	1.05
$\bar{\theta}_{(0)}$[b]	0.31	0.42	0.35	0.40	0.38
R^2	0.56	0.58	0.53	0.52	0.53

[a]*t*-ratios in parentheses. [b]These estimates include Goldberger's correction (Goldberger, 1968).

Table 6 Estimated Coefficients for Alternative Specifications of ln(Bid per acre), first-time bidders only (equations (13) and (14), n= 277)

	#6 (Eqn. (13))	**#7 (Eqn. (13))**	**#8 (Eqn. (14))**	**#9 (Eqn. (14))**
$□_0$	-1.27 (13.89)	-1.55 (15.02)	-0.50 (8.53)	-0.31 (5.67)
$□_1$: State Budget at *t* per Predicted Bidder (#A)	0.00035 (7.59)	--	0.00032 (7.21)	--
$□_1$: County Exp. at *t*-1 per Predicted Bidder (#A)	--	0.00087 (2.83)	--	0.00087 (2.91)
ϕ_2	0.13 (5.70)	0.23 (10.34)	--	--
ϕ_3	--	--	0.65 (5.43)	0.97 (14.71)
ϕ_4	--	--	-0.011 (3.38)	-0.011 (4.75)
$□_{median}$ [b]	1.15	1.07	1.14	1.07
$\bar{\theta}_{(0)}$ [b]	0.38	0.36	0.59[c]	0.52[c]
R^2	0.55	0.47	0.57	0.50

[a] *t*-ratios in parentheses. [b] These estimates include Goldberger's correction (Goldberger, 1968).
[c] Assumes $□ = □^{-0.54}$

Table 7 Estimated Coefficients for ln(Bid per acre), second-time bidders only (equation (12), n = 109)

	#10	**#11**	**#12**	**#13**
β_0	-0.03 (0.56)	-0.18 (3.26)	-0.09 (2.06)	-0.17 (4.44)
β_1: Predicted Bidders (#A)	--	--	--	--
β_1: State Budget at *t* per Predicted Bidder (#A)	0.00026 (3.18)	--	--	--
β_1: Predicted State-Budget-per-Bidder (#B)	--	0.00004 (0.53)	--	--
β_1: County Exp. at *t*-1 per Predicted Bidder (#A)	--	--	0.0001 (0.37)	--
β_1: Predicted County-Exp.-per-Bidder (#C)	--	--	--	0.00017 (0.88)
$\ln(\hat{\hat{\theta}}_{(t)})$	from #2	from #3	from #4	from #5
$\hat{\alpha}_0 - \hat{\alpha}_1$	0.12	0.08	0.12	0.14
R^2	0.44	0.49	0.48	0.50

[a] *t*-ratios in parentheses.

Figure 1 Distribution of taste parameter, regressions 1-5

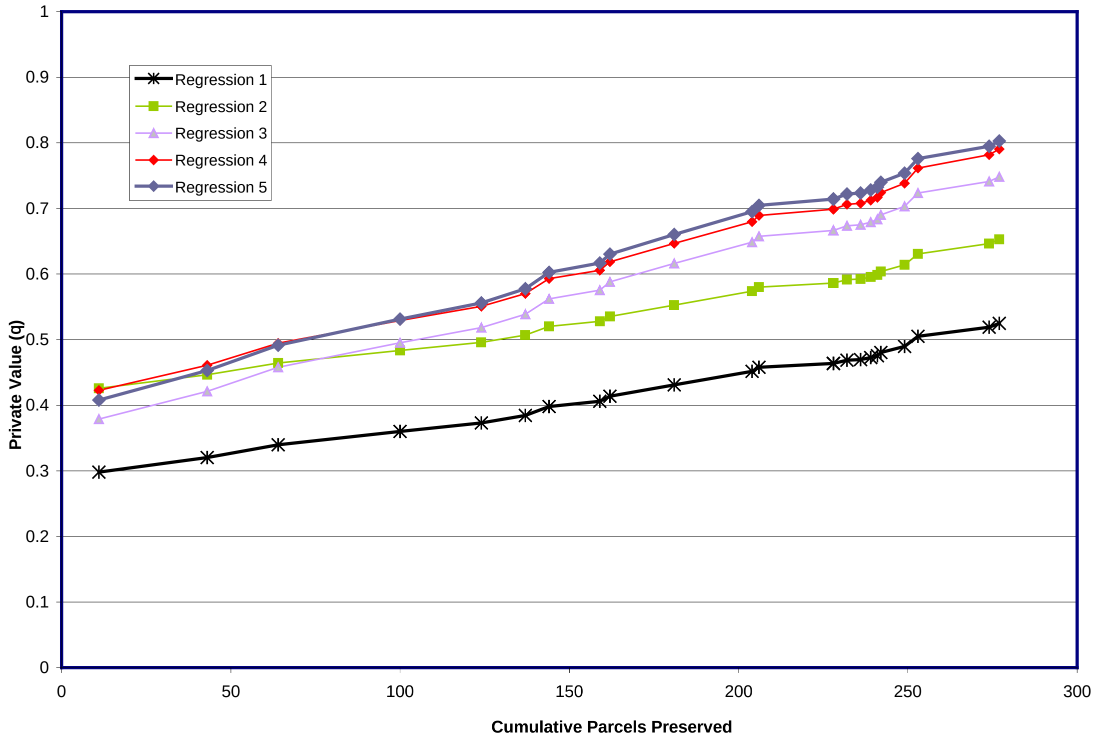

Figure 2 Distribution of taste parameter

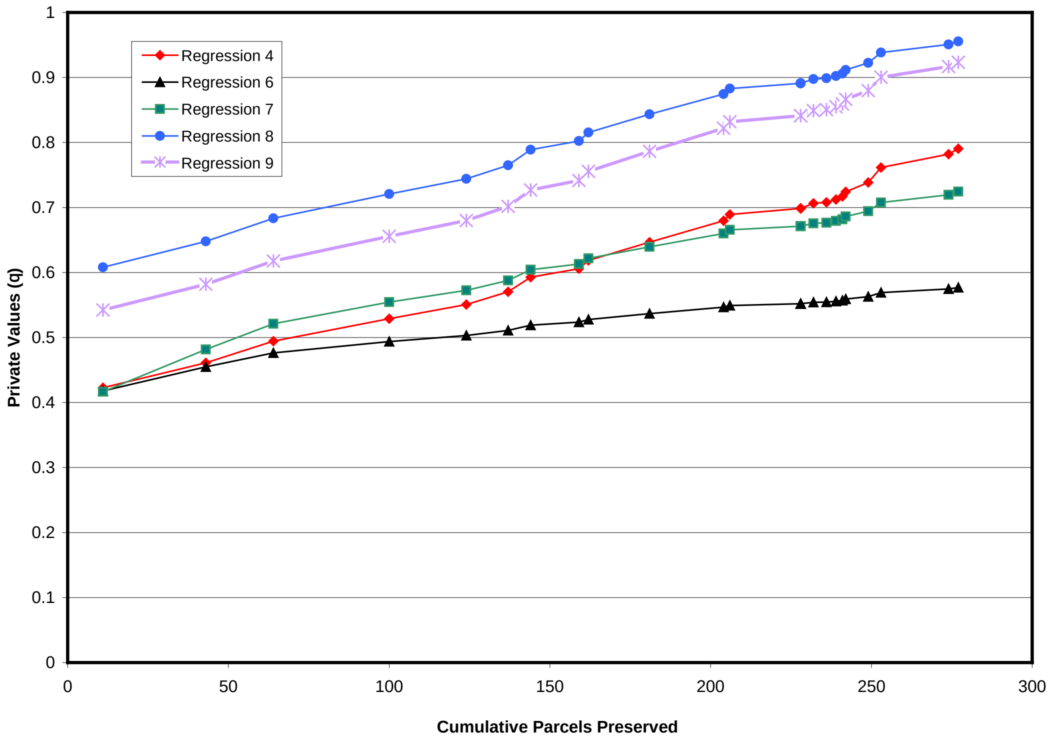

Chapter 3: Market Design with Asymmetric Bidders in Online Keyword Auctions[40]

By: Andrew Stocking

I. Introduction

Auctions are increasingly being used in both government and private contexts to exchange goods and services. Cramton (2007) estimates that over 450% of world GNP (~$30 trillion in 2000) is traded each year through auctions and the applications continue to grow[41]. Auctions have the benefit of a rich economic theoretical literature and significant flexibility such that they can handle many different types of products and significant variation in expressiveness by auction participants[42].

One specific implementation of an auction to buy goods is the online advertising auction, which is currently used by Google, Yahoo!, Microsoft, Ask, among others to allocate search engine advertising space to advertisers. Stated simply, the objective of the search engine is to solve a searching consumer's need. The more successful the search engine is in meeting consumer needs, the more the search engine will be used for consumer search in the future and the more revenue the search engine will earn. For each of the 9.82 billion searches completed annually on search engines in the US (or 61.04 billion searches worldwide), these search engines

[40] Special thanks to Andreas Lange, Peter Cramton, Lori Lynch, John Horowitz, Larry Ausubel, Susan Athey, and Michael Ostrovsky for comments and conversations. Thanks also to Rebecca Young, Marya Gomez, Jeff Regen, Cecilia Snyder for providing access to their data.

[41] Recently the Federal Aviation Administration has considered using an auction to allocate takeoff and landing slots in New York's airports. More relevant to this discussion, the last 5 years have witnessed the growth of online advertising (search engines) or goods (eBay, Amazon) auctions.

[42] Expressiveness is that attribute of demand that includes complementarities, substitutes, and other specifics.

attempt to solve the searching consumer's need by displaying both organic results according to their proprietary search algorithm and advertisements based on real-time continuous auctions for each search term. The keyword or "adword" auction determines the order to display advertisers and the associated cost charged to each advertiser. These auctions accounted for 40 percent or $8 billion of total online advertising spending in 2007. This figure is projected to increase to $25 billion by 2012 (Klaassen, 2007).

Given the relative newness of these markets, economic research is just starting to materialize and generally followed one of two bifurcations. First, computer scientists have analyzed these auctions from an algorithmic perspective. Due to the complexity of the strategy space they are forced to work under very confining assumptions toward the goal of estimating lower bounds on optimal revenue for the search engine (see for example, Roughgarden and Sundararajan, 2007; Goel, et al., 2008).

The second strain of research is the one I will extend with this analysis. In this area of the literature, economists have analyzed optimal search algorithms from the consumer perspective (Athey and Ellison, 2008) and optimal bidding behavior under a simplified version of the current adword auction setup (Edelman et al., 2007; Varian, 2007, Aggarwal et al., 2006).[43] Their research extends a rich literature on auction design (Ausubel and Cramton, 1999; Maskin, 1992; McAfee and McMillan, 1987; Milgrom and Weber, 1982; Myerson, 1981). The research presented here

[43] For purposes of this analysis, "consumers" will refer to the individuals performing the search and asking a question of the search engine; "advertisers" or occasionally "bidders" will refer to the firms engaged in the strategic bidding process; "search engines" or "platforms" will refer to the search engine hosting the market which caters to both consumers and advertisers.

further extends their analysis by relaxing two simplifying assumptions in an effort to make the adword auction models richer in inputs and conclusions. This more robust model, as presented below, represents the first contribution to the literature.

The first relaxed assumption relates to the homogeneity of keywords. In the past, specific keywords have been characterized as homogenous in their value to advertisers (Edelman et al., 2007; Varian, 2007). That is, advertisers have a marginal value per click that is independent of the keyword on which advertisers are bidding. I enlarge this equilibrium bidding environment such that advertisers must select both keywords on which to bid and their respective bids for those keywords. This is an important extension because keywords vary by their specificity and relevance to an advertiser's product. The existence of a continuum of keyword specificity and relevance with respect to any advertiser is then used by the search engines to determine the quality-weight of the advertiser's bid. This heightens the importance of understanding how advertisers select the "quality" of the keywords on which they are bidding. Thus, the enlarged equilibrium model allows one to model the advertiser's profit-maximizing selection of keywords when advertisers are asymmetric over their advertising strategies.

Asymmetric advertising strategies are the second extension to the current models. I begin by deconstructing the advertising firm's objective function to allow them to vary in the value they receive from advertising. Stated simply, the objectives for online advertisers can be aggregated into three categories: those advertisers who are most interested in branding, traffic, or transactions where advertiser value is measured in terms of impressions, clicks, and actions/purchases, respectively. Realistically, an advertiser's true objective function is likely to be a weighted combination of these three components; however, I will assume that advertisers fall distinctly into one of these three pure advertiser types.

Each of these advertiser-types is present on the internet, though in different proportions. Historically advertising in newspapers, the radio, and television has been thought of as primarily a mechanism for branding. This is due to the fact that there was no explicitly trackable action[44] following an advertisement and thus advertisers relied simply on the number of people reached as a proxy of the effectiveness of the advertisement. This is not true on the internet because there are a number of trackable actions, including visitors (traffic) and purchases (transactions). Thus, those advertisers only interested in branding represent a small fraction of the firms advertising on the internet. They are primarily large firms with recognizable brands.

Since the advent of the internet, traffic – or number of visitors – has been an important metric in determining the value of a website. Over the past decade, many of the large acquisitions of online entities have been based exclusively on volumes of website traffic, independent of the revenue of the firm.[45] In general there are two types of advertisers seeking traffic. The first is composed of community sites that rely on traffic to create a vibrant community; these advertisers often have answers to consumers' needs and are thus are complementary to the search engine's mission. The second and significantly larger group of traffic advertisers is composed of aggregator sites. These sites sell advertising to specific targeted groups of advertisers

[44] Coupons are a notable exception. Advertisers could publish coupons and track the redemption over a period of time to understand the efficacy of the advertising outlet.

[45] While difficult to disentangle the exact valuation method for most acquisitions, it's clear from press statements and other information that traffic is a big factor in the equation. iVillage.com was acquired for $600 million by NBC in 2006 because NBC wanted to "access a critical mass of Web users" (Davis, 2006). iVillage reportedly had 14.5 million visitors and $91 million in revenue. Similarly MySpace was purchased for $580 million by Fox in 2005. At the time MySpace had 90 million members and an annual revenue of ~$20million.

(e.g., Caribbean travel firms, machine parts, low price airline travel) and then advertise on search engines to attract people to their search engine-like site. As such, these traffic advertisers are substitutes for the search engine. Their goal is to attract advertisers and consumers away from the search engine on which they are advertising such that, in the future, consumers will come directly to their aggregator site. As might be expected, the nature of these advertisers as substitutes makes them unattractive to search engines as market participants.

The final type of advertiser is the transaction and conversion advertiser. These advertisers would appear to be most directly interested in solving a consumer's need and thus are the most complementary to any search engine's mission. That is, one can reasonably assume that a consumer who makes a purchase from an advertiser has accomplished what she set out to do when initiating the search. A transaction could include buying a book from Amazon.com, signing up for newsletter of NPR, or making a donation to World Wildlife Fund.

With these two model extensions in place, I make a second contribution to the literature by describing how the subsidization of advertisers affects search engine revenue. Outside of the search engine context, the subsidization of bidders and in particular, bidders who might otherwise not participate in the auction, has been studied from a number of different perspectives (Ayres and Cramton, 1996; Higgins and Roberts, 2009). Following Bulow and Klemperer (1996), it is well-established that adding an $n+1^{st}$ bidder to a second price auction is superior to allowing the auctioneer to negotiate the best price possible with the n original bidders. Furthermore, Ayres and Cramton (1996) demonstrated that this result holds in FCC auctions even if the $n+1^{st}$ bidder is subsidized by the auctioneer to participate. In both cases, this result holds because of the additional competition brought to the auction by the $n+1^{st}$ bidder. This mechanism also holds in the case of keyword auctions, but a second mechanism for increasing revenue may also be present: the addition of a high-quality advertiser.

The analysis of advertiser subsidization is inspired by an innovative program at Google that entices a relatively small player in terms of time and financial resources – the charitable organization – to become an active participant in the adword auction when they might otherwise have not participated. Between its inception in 2003 and 2008, this program provided subsidies valued at more than $100 million to over 5,000 nonprofits (Wynes, 2008). Based on optimal bidding strategies for advertisers, I will elucidate the conditions under which subsidization of advertisers is a revenue positive activity for search engines.

At the conclusion of the theoretical analysis, I use actual keyword auction data and the associated auction outcomes to verify the direction and determine the magnitude of the effects in the developed model. The novel data set comes from three different advertising firms across four different search engines. With this, I show that the model developed here is robust to different advertisers and different search engines. This is the third contribution to the literature. Finally, I will use this data and regression results to illustrate through simulation when advertiser subsidization is revenue enhancing for search engines. This is the fourth contribution to the literature. This analysis complements the existing literature and sets the stage for follow-up work on other equilibrium bidding strategies.

II. *Online Adword Auction Description*

The general format for online advertising auctions is well-characterized in the literature (Athey and Ellison, 2008; Edelman et al., 2007; Varian, 2007, 2009; Lahaie et al., 2007; Aggarwal et al, 2006). Given the large volume of transactions and the revenue implications for the search engines, there is significant incentive by the search engines to implement continuous enhancements or improvements to the current auction design. As a result, several of the above publications describe auction implementation features that are no longer in use. Here I will describe the current implementation, noting where it differs from those described in the above publications.

Following a consumer search on any of the popular search engines (Google, Yahoo!, MSN, Ask, among others), the consumer is shown a combination of organic search results and paid search results. Organic search results are results generated by the search engine's proprietary algorithm intended to deliver to the consumer the locations on the internet most relevant to the consumer's search terms. The search engines are not paid for displaying these organic results in contrast to the paid search advertisements which also appear. These two types of results are separated on the results page such that the consumer can easily differentiate advertisements from organic results (see Figure 1). As would be expected, the consumer is free to click on none of the advertising links or one or more of the links. The number of advertisements shown to the consumer varies, but is usually less than 20[46].

[46] The average advertisements per search for Google, Yahoo! and MSN in June 2008 were 4, 5, and 8 respectively. This represents a decrease compared to January, 2008 where the three search firms

Consumer welfare, as described in the literature (Athey and Ellison, 2008), is improved when the combination of organic and paid results answer whatever question the consumer had in mind when initiating the search. As a result, the search engine wants to present a combination of organic and paid results that have the highest likelihood of answering a consumer question. This motivation will become important later during the discussion of how different firm-types contribute to consumer welfare and why the search engine might choose to subsidize one particular type of firm.

The presentation of advertisements begins with advertisers selecting specific keywords that, if searched for, will result in the display of their advertisement and a link back to their website. When bidding on keywords, advertisers submit the following four items: 1) the keyword; 2) the specific text of an advertisement to be displayed in the paid results area following a consumer search using that keyword; 3) a link back to a page on the advertiser's website called the landing page; and 4) an amount the advertiser is willing to pay for each click on their link within the advertisement. In addition, advertisers can specify daily budgets for how much they want to spend. Search engines implement two strategies with respect to maximum budgets: Standard or Accelerated[47]. Under the Accelerated delivery method, the advertiser's ad is displayed until their budget is exhausted and then not displayed again until the next day when a new budget is in effect. Under the Standard delivery method, search engines "throttle" display of the advertisement such that the ad does not appear every time the associated keyword is searched, but appears frequently

---

showed 6.5, 6, and 8 paid advertisements, respectively (Stokes, 2008). However, in the data from Section V, the search engines report up to 170 positions for some keywords.

[47] These are Google's terms, but the other search engines have similar capabilities and use similar words.

enough to exhaust the daily budget by the end of the day. Designing an optimal allocation algorithm for search engines using the Standard strategy is a difficult problem that has been well-studied in the computer science literature (Goel et al., 2008).

Advertisers are charged by the search engine anytime their ad appears and the corresponding link receives a click. This type of payment scheme is referred to as pay-per-click (PPC) or cost-per-click (CPC) and differs from traditional advertising which charges advertisers based on the view or impression. Television, radio, or some online banner advertisements are sold using view- or impression-based advertising models which is often referred to as pay-per-view (PPV) or cost-per-1000 impressions (CPM) advertising[48]. There is a third advertising model that charges advertisers only after a transaction has been completed. For example, Amazon.com's affiliate program reimburses online websites some percentage (e.g., 5%) of the purchases made by customers driven to Amazon by that website. This is referred to as the pay-per-action (PPA) or cost-per-action (CPA) model[49]. Today, most search engines employ the CPC model of pricing for adword auctions.

From the search engine perspective, the most important auction design questions relate to the ordering of advertisers and the payment made by each advertiser following a click on their advertisement. Ordering is important because in the current auction implementation, the order is the only mechanism used by the search engine to signal the quality of advertisers with respect to consumer search. A

[48] In search engine advertising, an impression or view of the advertisement is created anytime a search is done which results in the display of the advertisement.

[49] see Agarwal, Athey, and Yang (2009) or Dellarocas and Viswanathan (2008) for more information on the benefits and weaknesses associated with CPA pricing relative to CPC or CPM pricing.

transparent ordering will induce consumers to click on the highest quality advertiser most often with decreasing likelihood for lower quality advertisers (Athey and Ellison, 2008). Given that increasing clicks translate directly to higher search engine revenue, a search engine will have more control over revenue if they provide a consistent and transparent ordering of advertisers. In practice, search engines have suggested that advertisers are ordered efficiently from the top of the page to the bottom of the page. From a consumer perspective, an efficient ordering places advertisers in order of likelihood of meeting their search need. From the search engine perspective, efficiency is defined as an ordering that places the advertisers in order based on their revenue contribution to the search engine. While these two efficiency definitions may seem at odds, recent search engine market design changes suggests they may be more similar than not.

Originally search engines ordered advertisers solely by their bid, which provides a first approximation of maximal search engine revenue when consumers expect the advertisers to be ordered by likelihood of meeting their need. However, search engines found that this approach was not an optimal long term strategy. The ordering largely ignored any propensity of the advertisement to meet the consumer's need, and included only the willingness for each advertiser to pay for a click. While willingness to pay for a click may be related to the ability to meet a consumer's need, it proved to be suboptimal for the search engines. Thus, after a few searches and clicks, consumers were found to be less likely to click on advertisements, ostensibly because their prior belief about the utility of a click (i.e., the advertiser meeting their need) was low. This hurt search engines because fewer clicks resulted in lower search engine revenue. To address this problem, search engines have recently modified their objective to be more aligned with the consumer's efficient ordering. Search engines now rank advertisers by quality-weighted bids (i.e., ordering based on the product of the quality-weight and the bid). This is an attempt by the search engine to balance short term gains against long run profits. In addition, it demonstrates the convergence of the consumer and search engine definitions of an efficient ordering.

Quality-Weight (q). Search engines initially defined the quality-weight almost exclusively based on the click through rate (CTR)[50] of the advertisement; however, since consumers expect an efficient ordering, advertisements at the top of the page are likely to receive more clicks than those at the bottom of the page. Consequently, CTR is an endogenous (and unstable) measure of quality from the consumer perspective. Ordering advertisers by CTR-weighted bids provides a very clear ordering according to the advertiser's revenue contribution to the search engine in the short term. But if this CTR-weighting does not correspond to true quality in terms of meeting the consumer's need, over time the CTR will fall as consumers update their beliefs downward about the ability of the advertisement to meet their needs. To address this problem, various search engines implemented a detailed proprietary algorithm for defining quality that is based on a variety of factors that attempt to determine a more accurate measure of quality with respect to meeting consumer's needs.[51] Search engines assign a unique quality-weight to every

[50] Click through rate (CTR) is defined as the probability of an advertisement receiving a click contingent on it being shown to consumers.
[51] One search engine, Google, claims that they are continually refining their Quality Score formulas; however, today the quality score is a function of the advertisement's CTR, the usefulness and

advertiser-keyword pair. The product of this quality weight and the advertiser's willingness to pay for a click is believed to result in an ordering that is more consistent with the consumer's definition of efficient ordering. It should be noted that an ordering based exclusively on the quality-weight ignores the advertiser's willingness to pay, which in a competitive marketplace should also be correlated with their ability to meet a consumer's need. That is, the advertisers are ordered by the search engine's best estimate of their ability to meet the consumer's need. And as a result, this approach to ordering should not cause a degradation of consumer clicks over time and is thus consistent with the effective long term revenue contribution to the search engine.

A second mechanism implemented by search engines to improve revenue is the use of bidder- and word-specific reserve prices assigned by the search engine. These reserve prices serve two purposes: 1) they guarantees revenue to the search engine from auctions with little competition or from the lowest ranked bidder in each auction, assuming the number of bidders is less than the number of slots; 2) they serve as a second mechanism through which the search engine can enforce minimum quality guarantees to protect the consumer. I will assume that these bidder- and word-specific reserve prices are only used by the search engine to ensure every advertiser pays a positive amount per click and advertisers are ordered by decreasing quality-weighted bid.

---

relevance of the landing page, the ease of navigation of the landing page, the load time for the landing page (i.e., time it takes a user to view the landing page after clicking on an advertising link), the number of links on the landing page, the relevance of the keyword to the ads used by the advertiser, among other things.

Following a click on an advertisement, the associated advertiser is charged the least she would have to pay to remain at that position given the quality of her advertisement. More generally, each advertiser pays the expected revenue of the advertiser below them at the time of the click divided by their own quality (i.e., advertiser *i* pays $b_{i+1}q_{i+1}/q_i$ where position $i+1$ is below position i and b_{i+1} is the bid of the advertisers in position $i+1$). A higher quality weight allows the advertiser to pay less while maintaining a high position. This auction environment means that an advertiser can bring two currencies to the auction with which she can increase her position: the bid (b) and the keyword quality-weight (q). These can be increased together to achieve a top position, or an advertiser can use high word quality-weight to offset lower bids. The obvious result of this approach is that advertisers are not necessarily ranked by their bids alone, unless all advertisers have the same quality ranking. I now turn to two components of quality for the advertisers: keyword specificity and keyword relevance.

Keyword Specificity (θ). As the name suggests, keyword specificity relates to the specificness of the keyword[52]. For example, the keywords "*clothing*", "*book*",

[52] When an advertiser places a bid for a keyword, the advertiser can specify whether they are bidding for the keyword using a Broad, Phrase, or Exact match (Google terminology but other search engines use similar language). The Broad match is the default option and will display the advertisement when similar keywords are searched, using search engine algorithms to determine similarity. For example, if one bids for the keyword *auction design* using a Broad match, the associated advertisement will appear following searches for *design auction*, *designing auctions*, *how to design an auction*, and *building auctions*. The phrase match is more targeted in that the advertisement only appears when searches are conducted for the keyword(s) in the order specified by the advertiser with or without the addition of lead or lag words. For example, bidding for the keyword *auction design* would display the associated advertisement following searches for *good auction design*, *auction design examples*, or *what is auction design* but not *design auctions*. Finally, the exact match only shows an advertisement when a consumer enters a search query that includes exactly the keyword(s) bid on without additional words. This is an area of the auction design that is under constant flux, as search firms are constantly optimizing the best approach to displaying advertisements to maximize consumer welfare.

and “*car*” are less specific than “*University of Maryland t-shirt*”, “*Combinatorial Auctions book*”, and “*Toyota Prius.*” As the keyword specificity increases, advertisers earn higher quality-weight from the search engine, ceteris paribus.

Keyword Relevance (R). This component of the quality-weight describes how relevant the purchased search term is to the advertisement. For example, “*Toyota Prius*” is highly relevant to a Toyota dealership’s website, less relevant for a bookstore’s homepage even if they have books on the Prius, and not at all relevant for a cooking school. More relevant keywords earn advertisers higher quality weights, ceteris paribus.

As stated earlier, the exact formulation for calculating the quality-weight is proprietary and a function of many components; however, keyword specificity and keyword relevance are known to play both a direct and indirect role in the calculation. Directly, these attributes measure the relationship of the search term to the advertiser’s landing page and the likelihood that the advertiser is answering the question posed by the search. Indirectly, search engines use the average quality-weight for all of the keywords within an advertiser’s portfolio of keywords to adjust the quality-weight for each specific keyword. That is, as the advertiser adds lower quality keywords (lower relevance or lower specificity words) to their portfolio, this lowers the quality-weight for every keyword in the portfolio. The implications of this will be discussed in more detail below.

III. Model

I present a simple model for evaluating advertiser equilibrium strategies within an equilibrium framework. This equilibrium is an enlarged concept relative to the existing literature which primarily considers bidding behavior given a particular keyword. In practice advertisers are jointly selecting both keywords and a bidding strategy. The model below is unique to the literature though some parts draw from the presentations in Edelman et al., (2007), Varian (2007) and Athey and Ellison (2008).

Firm $j \in J$ generates profit by selecting a series of keywords $i \in I$ and a bidding strategy B^* which places those keywords in position $m \in M$. The selection of a particular keyword influences profits through the corresponding quality of that keyword and other unknown firm and keyword characteristics that without loss of generality are assumed to implicitly depend on the keyword quality. Thus, the overall keyword quality-weight for firm j is modeled as a function of keyword specificity (θ_i) and relevance between the keyword and the firm (R_{ij}).

$$q_{ij} = g_j\left(\theta_i, R_{ij}\right) \tag{1}$$

Where $g_j\left(\bullet\right)$ is twice continuously differentiable and $g_\theta \geq 0$ and $g_R \geq 0$. For ease of exposition I will assume there are two word relevancies – high relevance and low relevance keywords – and that advertisers are only bidding on the highly relevant words (i.e., a car dealership is not bidding on keywords related to candy). This facilitates further simplification of (1) such that $q_{ij} = g_j\left(\theta_i, R_H\right) \equiv g_j\left(\theta_i\right)$.

In practice, keywords span a complete surface of relevancy and specificity values and advertisers are constantly attempting to determine the most relevant keywords for a given specificity (or vice versa). Selection of the optimal quality keywords could be modeled as a probabilistic process by the advertiser, but this complication is unnecessary for this analysis beyond a brief discussion below with respect to search engine marketing (SEM) consulting firms.

Advertising firm profit is modeled as the benefit derived from consumer clicks on a portfolio of keywords net of the cost per click charged by the search engine. In addition, advertisers incur an administrative fee ($\alpha_j(n)$) associated with the cost of placing bids, tracking performance, and reporting and analyzing results which increases with each incremental keyword. This administrative fee also captures the penalty levied by the search engine on the quality-weight for each of the keywords in the portfolio as the average quality of the marginal selected keyword falls. Thus, the general profit function for firm j is:

$$\pi_j\left(\Theta, B^*\right) = \sum_i (\Pi_i) - \alpha_j\left(n_j\right) \tag{2}$$

Where Π_i is the profit derived from the i keywords bid on by the advertiser given the bidding strategy (B^*) and a vector of keyword specificities (Θ); $n_j = \#\{i | b_i > 0\}$ is the number of keywords purchased by firm j, and $\alpha_j(\bullet)$ is twice continuously differentiable, strictly increasing, convex ($\alpha_j' \geq 0$ and $\alpha_j'' \geq 0$) and satisfies standard Inada conditions ($\lim_{n \to 0} \alpha_j'(n) = 0$ and $\lim_{n \to \infty} \alpha_j'(n) = \infty$). In addition, equation (2) assumes additive separability of the profit from each keyword. This realistic assumption stems from the fact that every search reaches a new person and thus the benefits are additive to the advertisers, i.e., reaching 20 consumers is twice as valuable as reaching 10 consumers. This same logic does not apply on the television or radio where each incremental advertisement during a particular time slot reaches the same group of people and consequently, garners diminishing marginal returns.

Given this setup, each advertising firm will optimally purchase N_j search terms, where $0 \le N_j < \infty$. In practice, firms that use an SEM consultant can often manage a much larger portfolio of words[53]. This is because the consultant is better at selection and organization of keyword portfolios that contain the maximum word quality. Consider an example. From the perspective of the search engine quality algorithm, there exists a complete cardinal ranking of quality-weights for every search term in the English language with respect to any particular advertiser. Thus if the advertiser selects the highest quality words, she can at most select $\hat{N}_j$ words. The selection of words with the maximum specificity and highest relevance, however, is probabilistic, and as a result each advertiser must employ some technology to select N_j words with the highest quality-weights where $N_j \le \hat{N}_j$. The success of this technology to select the highest quality-weighted words is observed in the administrative fee function ($\alpha_j(\bullet)$) since lower quality words lower the quality-weight for all words.

[53] Anecdotal evidence suggests that an advertiser working alone can typically manage a few thousand words; a search engine consulting firm (or advertiser with the technological capacity of a search engine consulting firm) can manage tens of thousands of words or more.

Figure 2 illustrates this example. In this case, $\alpha_x'(n_x)$ is the administrative marginal fee schedule representative of Advertising Firm X's ability to choose and manage the optimal quality words. Without loss of generality, we assume that the marginal benefit from each keyword is fixed at $\overline{\Pi}$. Under scenario 1, Firm X purchases N_1 words and generates a profit equal to the size of region A. Under scenario 2, Firm X decides to retain a search consulting firm to purchase and manage the keywords. If this consultant has better technology for choosing and managing the words such that their administrative marginal fee schedule is $\hat{\alpha}_x'(n_x)$, Firm X can instead purchase N_2 keywords and generate profits equal to the combined size of region A and B. However, Firm X must pay the consultant their retainer fee, causing gains to be lowered. This discussion provides a basis for bounding a firm's optimization problem; i.e., firms will never choose to bid the reservation price for an infinite number of keywords.

Asymmetric Advertisers

The advertising firm profit function (denoted as Π_i in equation (2)) for keyword i must be deconstructed in order to consider the specific bidding strategies for each advertiser type. I discuss each advertiser type below with their respective profit function, starting with those advertisers primarily interested in branding.

Branding (B-type advertisers). Translated to online search engine marketing, branding is measured by reach or number of impressions delivered which is equivalent to the number of times a particular keyword is searched. Thus, advertisers interested in branding are maximizing the following profit function for each of *i* auctions:

$$\Pi_i = \max_{b_i,\theta_i} I(\theta_i) \cdot \left[\beta - CTR(r,\theta_i,\overline{q}_i) \cdot CPC(r)_i \right] \tag{3}$$

Where β is the value of an impression to the advertiser measured in terms of \$/impression; $CPC(r)_i$ is the cost per click to the advertiser and is simplified notation for $(b_i q_i)_{m+1} / q_{i,m}$ or the score of the advertiser in position $m+1$ divided by own quality-weight; $M+1$ denotes the first rejected advertiser; $I(\theta_i)$ is the number of advertising impressions as a function of the keyword specificity; and $CTR(r,\theta_i,\overline{q}_i)$ is the click through rate which is a function of keyword position on the page, the keyword specificity, and the average quality of all of the other advertisements on the page ($\overline{q}_i$). These functions will be described in more detail below. The value of appearing as a paid advertiser following a consumer search on keyword i is: $v_i = I(\theta_i) \cdot \beta$, assuming keyword i is in the set of N_j optimally chosen keywords.

Traffic (T-type advertisers). Those advertisers who are interested in traffic are selecting keywords to maximize the number of clicks they receive from the search engine. In other words, T-type advertisers maximize the following profit function for auction *i*:

$$\Pi_i = \max_{b_i,\theta_i} I(\theta_i) CTR(r,\theta_i,\overline{q}_i) \left[\tau - CPC(r)_i \right] \tag{4}$$

Where τ is the value of a visitor to the advertiser measured in terms of \$/click; the other terms are as described above. The only difference between equation (3) and (4) is the inclusion of the click through rate in both the benefit and cost side of the objective function. This reflects the fact that traffic-focused advertisers achieve their value following receipt of the traffic, or a click, and thus have a value for keyword $i \in N_j$: $v_i = I(\theta_i) \cdot CTR(r, \theta_i, \overline{q}) \cdot \tau$.

Conversion (K-type advertisers). These advertisers want to convert a searcher into a consumer on their site or encourage some type of online transaction (e.g., account registration). Examples may include cars.com wanting to induce a car purchase, Greenpeace wanting a donation, or Care2.com wanting to consumers to subscribe to their newsletter. Each of these advertiser's objective is represented by the following profit function for auction *i*:

$$\Pi_i = \max_{b_i, \theta_i} I(\theta_i) CTR(r, \theta_i, \overline{q}_i) \left[CR(r, \theta_i) \kappa - CPC(r)_i \right] \tag{5}$$

Where κ is the value of a transaction to the advertiser measured in terms of \$/transaction; $CR(r, \theta_i)$ is the conversion rate which is a function of the position of the advertiser on the page and keyword specificity; and the other terms are as described above. The value to the conversion advertisers from appearing in auction $i \in N_j$ is: $v_i = I(\theta_i) \cdot CTR(r, \theta_i, \overline{q}) \cdot CR(r, \theta_i) \cdot \kappa$

Functional Relationships

Position (r). As described above, ceteris paribus, position is a weakly decreasing function of the advertiser score, defined as the product of the bid and the search engine-assigned quality-weight for that word/advertiser pair:

$$r = f_i\left(b_i q_i\right) \qquad (6)$$

Here $f_i\left(\bullet\right)$ is a piecewise constant function mapping real-valued bids to integer-valued positions and monotonically decreasing (see Figure 3). This function takes as inputs the bids[54] of all advertisers for a given keyword i and returns the position for each advertiser. Given that position depends on the score, advertisers can compensate for low quality-weights with a higher bid and low bids with higher quality-weights. Thus, as described above advertisers are ranked by the product of their search-engine defined quality-weight and their own bid. By convention, as the advertiser score for a particular word increases, the position (r) moves up the page since the top position is defined as rank or position one. Thus, the highest score earns the top position and low scores result in positions further down the page.

Impressions ($I\left(\theta_i\right)$). The number of impressions a keyword elicits is generally a function of the specificity of that keyword (θ). Consider for example a nonprofit dedicated to wolf preservation that is selecting keywords to advertise a request for donations. Clearly the words *donate to [nonprofit]* or *save wolves* are more specific than *gift* or *donate*. However, *gift* or *donate* are much more likely to be searched for

[54] Thus for an advertiser to know her position, she must assume a bidding strategy for the other advertisers.

because their intended meaning is much broader than *donate to [nonprofit]*. Exceptions to this rule exist, particularly following a natural disaster when *donate to Red Cross* might be both highly specific and generate many impressions for the Red Cross, but these exceptions occur only infrequently. The impression function is assumed to be twice continuously differentiable, decreasing and the second derivative is positive for low specificity words (i.e., there is a long tail of low specificity keywords) $I'(\theta) \leq 0$. I use this heuristic model in Assumption 1 to define keyword specificity.

Assumption 1. Specificity. For any two keyword specificity types θ^H (high specificity) and θ^L (low specificity), specificity is defined such that $I(\theta_i^H) < I(\theta_i^L)$ and thus when there exists a continuous distribution of specificities, $I'(\theta) \leq 0$.

Click-through-Rate ($CTR(r, \theta_i, \bar{q}_i)$ ***).*** The CTR is defined as a function of the position or rank, the keyword specificity, and the average advertiser quality on the page. Inherent in this definition is the assumption that consumers do not glean any additional information beyond ordering from the advertisements (i.e., advertiser brand, advertiser text) as to the advertiser's ability to satisfy a consumer's need. This is consistent with the assumption made by Edelman et al (2007)[55].

[55] An analysis of the validity of this assumption requires a dataset of advertisers by keyword which can only be obtained from the search engines. The data described below and used to validate the theoretical setup described here is a dataset of keywords by advertisers which will not support the requisite analysis.

To motivate an understanding of both the CTR and the conversion rate (CR), I will present a model of consumer behavior that extends the work of Athey and Ellison (2008) to consider two general groups of consumers instead of just one. These two groups are composed of: 1) those consumers (c_A) who generally know what they want and have entered the appropriate search term; and 2) those consumers (c_B) who are unsure what they want and/or have entered a non-optimal search term such that the search results page may or may not contain the information they want. Similar to Athey and Ellison's setup, consumers receive a benefit of 1 if their need is met and incur a cost for clicking on an advertisement ($s_z, z \in \{A, B\}$). For the first group of consumers this cost (s_A) is drawn from a cumulative probability distribution $G_A(s_A), s_A \in [0,1]$ which is first order stochastically dominated by the search cost for the second group of consumers (s_B) that is drawn from a cumulative probability distribution $G_B(s_B), s_B \in [0,1]$. The ability of an advertiser to meet a consumer's need (μ) is modeled as the product of the quality-weight (q) and the value derived by the advertiser (v) from appearing in the auction and distributed on the CDF F with support [0, $\hat{\mu}$] where $\hat{\mu}$ is some upper limit on ability of an advertiser to meet a consumer's need. I assume that advertisers are sorted from top position to bottom position by this indicator ($\mu = q \cdot v$) where the top advertiser is most likely to meet the consumer's need. This is consistent with the interpretation that the current advertiser ordering is the search engine's best estimate of the consumer efficient order. Inherent in this setup is the assumption that a Bayesian Nash equilibrium bidding strategy exists that monotonically maps value into bid such that the bidding strategy does not upset the ordering. This generic probability of meeting a consumer's need ($F(\mu)$) determines the probability that the advertiser meets the need of a consumer for the respective groups defined by $\eta_z = \phi_z(F(\mu))$ where $\phi_A(F(\mu)) \geq \phi_B(F(\mu))$.

Intuitively this setup maps quite clearly to the behavior of the search engines with respect to the two groups of consumers. The consumers in group B do not know what they want or if their search query is accurate. Consequently, they also know that they are further from satisfying their need and only use clicks on advertisers to narrow their search. Given that they may have to search again with a new and better defined search query, they are less disposed to clicking on advertisers. This behavior is modeled by giving consumers in group B higher search costs than consumers in group A (G_B first order stochastically dominates G_A) (see Figure 4). Once a consumer from either group clicks on the advertiser, η_z is the probability that the advertiser meets the consumer's need. For the A group of consumers, "meeting their need" is obvious since they know what they want and have searched for that need appropriately. For the B group of consumers, their "need" is less well defined. They must determine if they want what their search query has produced since they may not know exactly what they want. Alternatively, if they knew what they wanted but did not know the appropriate search term, they must determine if their search query produced what they truly wanted. This explains why for any given advertiser, the probability that that advertiser meets the need of a consumer from group A is greater than the probability that the same advertiser meets the need of a consumer from group B (see Figure 4).

The behavior of these two consumer groups is assumed to be independent. Given the decreasing likelihood that an advertiser will meet a consumer's need for advertisers lower in the list, intuition suggests and Athey and Ellison show that the optimal consumer strategy is to follow a cascade model of search (2008). That is, consumers will start clicking on the top advertiser and continue down the list of advertisers until their need is met (η_z) or the expected utility from the next advertiser is less than their search costs. In addition to this, I make the following assumption:

Assumption 2. Retention. A much larger fraction of consumers leave the search process because of negative expected utility rather than having their need met. Stated mathematically:

$$\left(1-\eta_{A,m}\right)G_A\left(\overline{\eta}_{A,m+1}\right) \geq \left(1-\eta_{B,m}\right)G_B\left(\overline{\eta}_{B,m+1}\right) \tag{7}$$

Where $\left(1-\eta_{z,m}\right)$ is the fraction of group-z consumers who have not met their need in the m^{th} position and $G_z\left(\overline{\eta}_{z,m+1}\right)$ is the fraction of group-z consumers who will click on the $m+1^{st}$ advertiser.

The retention rate is the percentage of consumers who have not stopped clicking on advertisers, either because their need has not yet been met or because they still estimate a positive expected profit from clicking. The retention assumption states that the retention rate for group A consumers is greater than the retention rate for group B consumers. This is not a particularly demanding assumption. Consider, for example, the following setup: $F(\mu) \sim U[0,1]$, $\eta_A = \frac{1}{10}\mu + \xi$, $\eta_B = \frac{1}{10}\mu$ where both groups of advertisers have identically distributed search costs ($s_A, s_B \sim U[0,1]$). As long as the maximum probability of an advertising meeting the need of a consumer is less than 0.9, condition (7) holds. The maximum probability occurs for an advertiser where $\mu = 1$ and thus condition (7) holds when or $\xi \in [0, 0.8]$ using the above setup. If search costs are not equal – group B consumers have higher search costs than group A consumers – this increases the likelihood of condition (7) holding. That is, the search costs of group B consumers need only be shifted higher than those of group A by a small margin to make the retention assumption hold even when the maximum probability of an advertising meeting the need of a consumer is 0.99.

Finally, by nature of the group descriptions, it is natural to assume that consumer group B is more likely to search lower specificity keywords than consumer group A. This setup produces the following results, summarized in Lemma 1 and 2 below.

Lemma 1. *Differences between the two groups of consumers result in the following CTR conclusions:*

1) *The CTR increases with keyword specificity for a given position, i.e.,* $\partial CTR/\partial\theta \geq 0$.
2) *The CTR is falling in position for a given keyword auction, i.e.,* $\partial CTR/\partial r \leq 0$.
3) *The CTR increases with the average quality for all advertisements on the search results page, i.e.,* $\partial CTR/\partial\overline{q} \geq 0$.

Proof: *See Appendix.* ⊗

The three conclusions presented in Lemma 1 provide a theoretical basis for understanding the relationship between the click through rate and its arguments. The literature on click through rate has remained agnostic to keyword specificity and thus the first relationship is new to the literature. The second relationship between CTR and position has been shown in the empirical literature (Goldfarb and Tucker, 2008). Similarly, Kempe and Mahdian (2008) find that high quality advertisements on a search results page impose a positive externality on the CTR of other advertisements on the page, regardless of their specific location. Intuitively this is because consumers see a high quality advertiser on the page while maintaining the assumption that advertisers are sorted in order of decreasing ability to meet their need. Thus, a high quality advertiser in the middle of the list of advertisers suggests to the consumers that the top advertisers are even more likely to meet their need.

Conversion Rate ($CR(r,\theta_i)$). The conversion rate is defined as the probability of a consumer making a purchase, subscribing to a newsletter, answering a questionnaire or completing any type of transaction requested from a given advertiser conditional on clicking on that advertiser in the paid search results. The conversion rate is a function of the position and keyword specificity and for purposes of this analysis does not have an advertiser-specific argument. That is, I assume that any two advertisers will have the exact same conversion rate for the m^{th} position. Contrary to the CTR, there is no indication or intuition that a higher quality advertisement on a page imposes a positive externality on the conversation rate and thus average quality is not an argument in the conversion rate. Using the presentation of the two groups of consumers presented above, I present the following conclusions in Lemma 2:

Lemma 2. *Differences between the two groups of consumers result in the following CR conclusions:*

1) *The CR increases with keyword specificity, i.e.,* $\partial CR/\partial\theta \geq 0$.
2) *The CR could be weakly increasing in position, i.e.,* $\partial CR/\partial r \geq 0$.

Proof: *See Appendix.* ⊗

The intuition behind the first result in Lemma 2 follows from the assumption that consumers in group A are more likely to search higher specificity keywords. The second result is possibly less intuitive. Figure 5 illustrates an example for a given keyword and six associated advertisers. Each advertiser has an intrinsic ability to meet a consumer's need (μ) that falls for advertisers lower in the list. This ability maps to a probability that the advertiser will meet the need of a generic consumer ($F(\mu)$). . The corresponding order statistics for the six advertisers are shown as $F(\mu)^{(m)}$. The expected overall probability of an advertiser in the m^{th} position to meet a consumer's need (η_m) is a weighted combination of meeting the need of group A's consumers and meeting the need of group B's consumers. For each lower position starting from the top, a higher percentage of total clicks are coming from consumers in group A and the conversion rate is increasing to reflect this. After the fourth position, there are no consumers from group B left clicking on advertisers and thus the conversion rate starts to decrease. This is only to show that it may be possible for the conversion rate to increase with position, especially near the top of the list of advertisers. There is some empirical evidence and significant anecdotal evidence corroborating this conclusion (Agarwal et al., 2009; Lahaie et al., 2007).

Cost-per-Click ($CPC(r)$ ***).*** The final functional relationship we consider is the cost per click ($CPC(r)$). Holding keyword specificity and search engine-assigned quality constant, it follows from equation (6), that the CPC is weakly increasing with the score (see Figure 3). That is, as the advertisement reaches the top of the list, the CPC will be higher than at the bottom of the list of advertisements. Determination of the exact $CPC(r)$ depends on the equilibrium bidding strategy employed by the bidders and will be discussed in extensive detail below.

Competition Between Advertiser-Types

With this understanding of the various functional relationships, I return to equations (3)-(5) to draw some behavioral conclusions about the best choices for three types of advertisers. Prior to considering best choices in an equilibrium setting, I consider just the revenue for each advertiser type by word type and position. These results offer insight into the Bayesian Nash Equilibrium bidding strategy discussed below. I determine each advertiser-type's best choice of word specificity (and thus assigned word quality) from their marginal revenue per click, given that word specificity affects all positions equally. And second, I can draw conclusions about best positional choices using each advertiser's total revenue by position.

Proposition 1 (Best Choice Behavior by Advertiser-Type): *Using the above model of the functional relationships and the profit functions for each advertiser, the following represent the respective bidder-type's best choice for keyword specificity and ordinal ranking over position:*

- *T-type advertisers (Traffic): The marginal revenue (* $MR_T = \tau$ *) is independent of keyword specificity and thus they are indifferent to keyword specificity. However, given that their portfolio of words is capped at a finite N, they may select keywords with more impressions despite the lower CTR and lower assigned quality-weights for these words, ceteris paribus. Their ordinal ranking over positions starts with the top position as their best choice and continues in order with lower positions (e.g., position 2 = 2nd choice).*
- *B-type advertisers (Branding): The marginal revenue (* $MR_B = \beta / CTR(r, \theta_i, \overline{q})$ *) is decreasing in keyword specificity and thus they choose low specificity keywords which earns them lower quality-weights. Given that all positions generate the same revenue benefits, B-type advertisers are indifferent to position.*
- *K-type advertisers (Conversion): The marginal revenue (* $MR_K = \kappa \cdot CV(r, \theta_i)$ *) is increasing in keyword specificity and thus they choose high specificity keywords which earns them higher quality-weights. Their best choice of position is the maximum of the product of the* $CR(r, \theta_i)$ *and* $CTR(r, \theta_i, \overline{q}_i)$

function which could result in a best choice of position in the middle of the list of advertisers.

Proof: *Follows from Lemma 1 and Lemma 2.* ⊗

The implications of Proposition 1 are that the various advertiser types are more likely to bid on different word specificities and exhibit different best positional choices. B-type advertisers migrate toward low specificity words and are indifferent to position, since all positions generate the same benefit. However, in an equilibrium setting where costs for any particular advertisers are monotonically increasing for higher positions, B-type advertisers will always prefer the bottom position. These low specificity words will result in B-type advertisers earning low quality-weights for the keywords in their portfolio. T-type advertisers are indifferent to word specificity but always select the top position as their best choice ranking lower positions in order by position. They will bid on low and high quality keywords, but by Assumption 1, will migrate toward low specificity, high volume keywords. Thus, T-type advertisers are also likely to earn lower quality-weights for their portfolio of words. K-type advertisers like high specificity words. Their ordinal ranking of position depends on the specific functional relationships for each word. As a result of their choice for high specificity keywords, they will earn higher quality-weights for their portfolio of words.

Despite the differences in advertiser keyword best choices, it is possible and perhaps likely that they could all appear in the same keyword auction. For example, the search term "Toyota Prius Washington DC" might be purchased by a Toyota-

Prius dealership wanting to generate sales leads (K-type), the Toyota corporate website wanting to brand Toyota (B-type), and cars.com wanting traffic so that people will search for Toyota Prius cars on their site (T-type). Similarly, the terms "doctor", "breast cancer" and "medical tests" could attract all three advertiser types. Each advertiser type brings a different bid and quality-weight. The next question from an auction design perspective relates to how the various advertiser types arrange themselves optimally on the search results page. This represents the first equilibrium behavioral result analyzed.

The optimal arrangement of various advertiser types by position has obvious parallels to the matching literature (see e.g., Roth and Sotomayor, 1990). Without quality-weights the optimal ordering can be solved using the canonical one-sided preference school-choice matching environment (Edelman et al., 2007; Varian et al., 2007). That is, advertisers can ordinally rank positions but the search engine (representing the positions) does not have ordinal preferences over advertisers. Once quality weights are added, the problem becomes a two-sided matching problem more similar to the college admissions environment. Now the search engine can ordinally rank advertisers reflected in the quality-weight assignments. These rankings are translated to the positions such that all positions have identical rankings over advertisers and advertisers maintain their rankings over positions, as determined in Proposition 1. Using a simple Gale-Shapley algorithm where the advertisers choose first results in the bidder-optimal ordering.

Definition 1. Bidder Optimal Ordering (BOO) is determined using a two-sided matching algorithm (e.g., Gale-Shapley, 1962) where the search engine assigns rankings over advertisers to positions equal to the quality-weights and advertisers express rankings over positions equal to the cardinal values of total revenue per position. The product of the quality weights and the cardinal values determines which advertiser wins each position. The BOO is the equilibrium outcome when advertisers choose first.

The BOO describes an equilibrium ranking of advertisers that should be supported by any equilibrium bidding strategy. Initially, the BOO will be determined by each advertiser's cardinal values for total revenue per positions based on Proposition 1. Due to the structure of the auction where each advertiser must submit only one bid value per click, however, the ultimate arrangement may not be BOO. That is, there may be cases where advertisers submit identical scores and the search engine must arrange advertisers in some alternative way e.g. randomly. Imagine for example two B-type advertisers vying for the bottom position in a two position auction. They will both submit bids such that their score is slightly greater than that of the first rejected advertiser who is bidding her full value per click. The search engine will be forced to choose one advertiser, which may not be the advertiser with the highest value for the bottom position. Despite these tie-breaking scenarios, determination of the BOO is the first step to determining the equilibrium bidding strategy.

Using the results from Proposition 1, one can quickly see that any equilibrium auction outcome will result involve several distinct advertiser blocks, each of which may contain one or more advertisers. Figure 6 illustrates how Proposition 1 maps into various advertiser block arrangements. The exact arrangement of blocks depends on the product of the conversion rate and click through rate for the K-type advertisers. That is, if this relationship is strictly decreasing from the top such that their ordinal rankings are identical to T-type advertisers, then the top frame represents the block structure of the auction result. Alternatively, if this relationship is strictly increasing from the top such that K-type advertisers rank the bottom position as their top choice, then the middle frame of Figure 6 represents the block structure of the auction. The third frame of Figure 6 occurs when K-type advertisers rank a middle position as their best choice.

Defining the various blocks shown in Figure 6 is important for being able to succinctly describe the equilibrium bidding strategies of the three advertiser types. I will do that now. At the very top will be a block of T-type advertisers (M_1) who rank the top position as their best choice, followed by positions further down the list of advertisers. The advertiser in the top position will be satiated, in that she occupies her best choice position. Below M_1 may be a block of K-type advertisers who rank their best position in the middle (M_2) where at least one advertiser will occupy their best choice and thus be satiated. As described above, if K-types have a total revenue per position as describe in the top or bottom frame of Figure 6, this M_2 block will not exist. Next is another group of K-type or T-type advertisers who are strictly below their best choice of position (M_3) where no advertisers are satiated (unless

$M_1 = M_2 = \{\varnothing\}$ in which case the top advertiser is satiated). And finally the group of B-type advertisers who rank the bottom position as a best choice are denoted as M_4 with the lowest advertiser satiated. In a given auction, each block may or may not be present.

When K-type advertisers rank their best choice of position in the middle, their second choice of position is not necessarily above or below their first choice (see third pane of Figure 6). That is, the ordinal rankings could be bi-nodal such that their best choice is position 2, followed by position 5 as their second choice. To address this possibility and simplify the exposition of bidding strategies, I impose the following Assumption that removes the possibility of a bi-nodal optimal position.

Assumption 3. Bi-Nodal. Given the complete information auction environment, any search engine that determines K-type advertisers have a bi-nodal positional ranking structure will set reserve prices for those K-type advertisers to ensure that they are competing only for the highest optimal position node.

This assumption imposes a restriction on all K-type advertisers that they have a single best choice with ordinal rankings strictly decreasing above and below this best choice of position. Inherent in Assumption 3, as well as determination of the BOO and throughout the remaining discussion is an assumption about the information setting of the auction. The keyword auction is set-up such that advertisers can make continuous modifications to their bid and receive feedback on the effect of this change with respect to position, CTR, CR, and CPC. As a result of these near-infinite opportunities to learn how competitors are behaving, the keyword auction can be modeled as a simultaneous-move, one-shot game of complete information (Edelman et al, 2007; Varian, 2007).

Given this complete information setting, Assumption 3 and an advertiser arrangement according to the BOO, Lemma 3 describes the equilibrium nature of M_2:

Lemma 3. *If $M_2 \neq \{\varnothing\}$, it must contain only K-type advertisers who rank a position other than the top position as their best choice and all K-type advertisers above their best positional choice.*

Proof: *M_2 is defined from below (i.e., further down positions) to include only K-type advertisers (i.e., the lower boundary of M_2 is defined by a T-type or B-type advertiser) thus the only way to violate Lemma 3 is the existence of an unsatiated K-type advertiser above her best choice position and above a T-type advertiser. By nature of the model described above, K-type advertisers experience identical CTR and CR functional relationships and thus only the \$ per conversion valuations (κ) are unique to each K-type advertiser. Thus, all K-type advertisers have the same ordinal ranking over positions. Assume that Lemma 3 is not true and it is necessary to include one T-type advertiser in M_2 to contain all K-type advertisers above their best choice position. But given that T-type's rank lower positions lower and K-type's above their best choice position rank higher positions lower, then any K-type above a*

T-type will switch positions to the benefit of both advertisers. Thus, M_2 must include all K-type advertisers above their best choice position. ⊗

The above advertiser setup now lays the foundation for describing the advertiser optimal bidding strategy. First I consider a scenario where only one advertiser is satiated (the top most advertiser) and all other advertisers rank the position above their BOO assignment above their assigned position. This does not require that all advertisers rank the top position as their best choice, only that no advertiser, except the top advertiser is satiated. With the exception of the top advertiser, this is the definition of group M_3. This could include K-type advertisers competing alongside T-type advertisers for the top position or T-types competing for the top position and K-types competing for (and not achieving) some lower position. It cannot include, however, B-type advertisers who always rank the bottom position as their best choice or satiated K-type advertisers.

In a related setting, Edelman et al., (2007) prove that a Vickery-Clark-Groves (VCG) strategy is an Bayesian Nash equilibrium that is optimal for bidders and achieves the BOO. The VCG strategy can be characterized as truth-dominant which means that advertisers reveal their true per click valuations ($\kappa \cdot CR(r,\theta_i)$, τ, or $\beta / CTR(r,\theta_i,\bar{q})$) for use in: 1) determining the BOO and 2) calculating each advertiser's equilibrium bid. In other words, advertisers cannot do better by misrepresenting their true value from a click in each position. The VCG strategy presented by Edelman et al. must be modified, however, to account for the two new features of the model described above: 1) advertisers are ranked by their score (or product of the quality-weights and bid) and 2) K-type advertiser's per click valuations change as a function of position. A modified-VCG strategy that incorporates these two features is presented in Proposition 2. Valuation v_1 is the top advertiser's per click valuation.

Proposition 2 (Limited-Advertiser Bidding Strategy). *When all advertisers fall into group M_3 with the top advertiser being the only satiated advertiser, the following is a truth-dominant strategy that produces a locally envy-free equilibrium among bidders. This optimal strategy (B^*) is defined as:*

$$\begin{aligned} b_1^* &= v_1 \\ b_m^* &= \frac{1}{q_m \cdot CTR_{m-1}}\left[\left(v_m^{m-1} \cdot CTR_{m-1} - v_m^m \cdot CTR_m\right) q_m + CTR_m \cdot b_{m+1}^* \cdot q_{m+1}\right] \qquad (8) \\ b_{M+1}^* &= v_{M+1}^M \end{aligned}$$

Where b_j^* is the equilibrium bid for the j^{th} advertiser, v_j^m is the per click valuation of the j^{th} advertiser for the m^{th} position, there are M positions (the $M+1$ position is the first rejected advertiser), and the advertiser in the top position bids their marginal value per click (or anything greater than the advertiser below her).

Proof: *See Appendix.* $\otimes$

Equation (8) is a recursive formula for stating that each advertiser bids the quality- and click through rate corrected negative externality imposed by the next higher adjacent advertiser. It is mathematically equivalent to the algebraic formula:

$$
\begin{aligned}
&b_1^* = v_1 \\
&b_m^* = \frac{1}{q_m \cdot CTR_{m-1}} \sum_{k=m}^{M+1} \left(v_k^{k-1} \cdot CTR_{k-1} - v_k^k \cdot CTR_k \right) q_k
\end{aligned}
\tag{9}
$$

In other words, the m^{th} advertiser determines the maximum amount she would pay to move into the $m-1^{st}$ position and knows that the $m-1^{st}$ advertiser must pay that plus the third advertiser's payment to remain in that position. The m^{th} advertiser then converts this into a per click bid by dividing the amount by the click through rate for the $m-1^{st}$ spot and by her own quality.

Fundamentally, this strategy requires advertisers to pay the negative externality they exert on advertisers below them or equivalently, the lowest amount necessary to make the next lower advertiser indifferent between their current position and the position above them. The implementation, however, is different from that proposed in Edelman et al. (2007) when quality-weights were not included. That is, the negative externality is not defined based on the benefit to the advertiser from a higher position, but by the benefit to the consumers from having that advertiser in a higher position. In other words, the negative externality is measured in terms of the score (value times quality-weight) and not just the advertiser's value. In equation (8), the first term for the m^{th} advertiser is: $\left(S_m^{m-1} \cdot CTR_{m-1} - S_m^m \cdot CTR_k \right)$ where $S_j^m = q_j \cdot v_j^m$, i.e., the score of the j^{th} advertiser for the m^{th} position. The strategy requires that each advertiser bid such that they are indifferent to their position and the position above them. Operationally this means that they bid the full negative externality imposed on them by the advertiser above them (i.e., the full incremental amount they would pay to move up a position) on top of their current payment (i.e., their own negative externality on advertisers below them). Once the full externality is determined, an advertiser has two currencies with which to pay it: 1) dollars or 2) quality-weight. That is, an advertiser with a high quality-weight need pay less for the same position than an advertiser with a low quality-weight. This multiple currency auction has the obvious consequence that advertisers may not be listed in order of decreasing financial expenditure or decreasing bid.

The B^* equilibrium strategy is locally envy free, defined to mean that advertisers are indifferent to exchanging *scores* with the advertiser above her. This differs from the Edelman et al. (2007) definition where locally envy free occurs when advertisers are indifferent to exchanging *bids*. This indifference only applies to exchanging scores with the next higher advertiser; moving to a lower position generates strictly less revenue to the advertiser, as demonstrated in the proof to Proposition 2.

Similar to the strategy proposed by Edelman et al. (2007), the above strategy is the advertiser-optimal strategy such that any other equilibrium bidding strategy that is locally envy-free will produce more revenue for the search engine and lower profits for the advertisers. Intuitively this is because the strategy shown in Proposition 2 requires that advertisers bid the lowest amount necessary to maintain their current position. If they bid any less, they would move to a lower position. Bidding more, but less than the advertiser above them would not change their surplus for the keyword, but would generate more revenue for the search engine. Thus, the bidding strategy shown in Proposition 2, generates the lowest amount of revenue for the search engine of any equilibrium bidding strategy. This is shown formally in the Corollary to Proposition 2.

Corollary to Proposition 2 (Search Engine Revenue). *The strategy B^* described in Proposition 2 represents the lowest possible revenue to the search engine and the highest possible profit to the advertisers in the class of equilibrium strategies that are locally envy free.*

Proof: *See Appendix.* ⊗

Regardless of whether all advertisers in the Proposition 2 auction rank the top position as their best choice, Proposition 2 requires that advertisers are currently forced to a position below their best choice (except for the top advertiser who is assumed to be satiated). By definition, this requires that only T-type and K-type advertisers are present, since B-type advertisers rank the lowest position as their best choice and thus do not receive a negative externality from those advertisers above them.

When satiated K-type advertisers (group M_2) or B-type advertisers (group M_4) exist in the auction, the bidding strategy must be revised. To gain some intuition into the bidding strategy, consider the following set of bidders:

Advertiser Scenario 1					
Position	Clicks	Value (T1) ($10/click)	Value (T2) ($5/click)	Value (K1)	Value (K2)
Top: 1	6	60	30	0	0
2	3	30	15	90	81
3	1	10	5	45	41
4	0	0	0	0	0

With the following BOO arrangement and Proposition 2 bidding outcome as follows:

Position	Clicks	Advertiser	Negative Externality*	Prop 2 Bid*	Revised Bid*
Top: 1	6	T1	30	10	5+2ε
2	3	K1	45	5	5+ε
3	1	K2	5	15	5+ε
4	0	T2	0	5	5

*Note that the negative externality and bid are measured in units of score or assumes that all bidders have equal quality-weights.

If the bids shown in the "Prop 2 Bid" column above are submitted, it will clearly change the ordering of advertisers to an order that is not consistent with BOO. The resulting order will be K2, T1, K1, and T2 which is suboptimal for the bidders. Consequently, K1 and K2 will both move their bids to 10-ε to return to the BOO. However, this causes T1 to pay almost her full value for the top spot. In response, T1, will lower her bid knowing that K1 and K2 will never choose the top position over position 2. T1 will lower her bid to just enough to prevent T2 from gaining the top position (i.e., 5+2ε). K2 and K1 will follow and bid slightly less than T1 (i.e., 5+ε). Given that they submit equal scores, the search engine will randomly select the ordering (assumed here to be K1, K2). Thus, all advertisers will bid approximately 5 with payments = (30, 15, 5, 0) for (T1, K1, K2, and T2), respectively.

Next consider a slightly revised valuation scenario such that K-Type advertisers rank the top position strictly above the first rejected advertiser:

Advertiser Scenario 2					
Position	Clicks	Value (T1) ($10/click)	Value (T2) ($5/click)	Value (K1)	Value (K2)
Top: 1	6	60	30	50	46
2	3	30	15	90	81
3	1	10	5	1	6
4	0	0	0	0	0

Again, the advertisers cannot bid a Proposition 2 strategy or they will be placed in an order inconsistent with the BOO. In this example, it is still optimal for T1 to retain the top position; however, the next highest value for that position is K2 and thus T1 will bid to prevent this advertiser (i.e., the advertiser experiencing the largest negative externality from the presence of T1) from taking the top position. If K2 had a value for the first position of 46 in Advertiser Scenario1, K2 would only be willing to pay 5 extra units to achieve the top spot (46-41=5) resulting in a total negative externality of 10 (5 from K2 plus 5 from T2). Thus, the largest negative externality would still be 30 from T2 moving to the top position. In Advertiser Scenario 2, however, the largest negative externality from the presence of T1 is experienced by K2. As a result, T1 must bid an amount such that K2 is indifferent between her BOO position (position 3) and the top position. The resulting bids are:

Position	Clicks	Advertiser	Negative Externality*	Prop 2 Bid*	Revised Bid*
Top: 1	6	T1	45	10	7.5+ ε
2	3	K1	80	7.5	7.5
3	1	K2	5	26.7	7.5
4	0	T2	0	5	5

As in the previous example, the bid intended to drive the payment of the satiated advertiser (K1) is replaced with a bid that ensures BOO or at least the possibility of the BOO (the exact ordering of the two tying bids will be determined through a randomization by the search engine). As in the previous example, the top advertiser bids such that she provides a ceiling on the two K-type advertisers who will bid up to a level just below the bid of T1. These examples provide insight into the behavior of advertisers when there exists satiated or near satiated K-type advertisers who do not rank the top position as their best choice. Proposition 3 develops an equilibrium strategy when all three advertiser types are present in the same auction.

Proposition 3 (Tri-Advertiser Equilibrium Bidding Strategy): *Order M advertisers according to the BOO and designate as belonging to one of the following groups (starting from the top of the advertiser list (M_1) to the bottom (R_6):* $M = \{M_1, M_2, M_3, M_4, R_5, R_6\}$ *where* $M_1 \in \{1,...K_1\}$, $M_2 \in \{K_1+1,...,K_2\}$, $M_3 \in \{K_2+1,...,K_3\}$, $M_4 \in \{K_3+1,...,M\}$, R_5 *is the first rejected advertiser who would like to move into* M_4*, and* R_6 *is the first rejected advertiser who would like to* *move into* M_3 *and is assumed to be T-type. The following bidding strategy represents the advertiser-optimal strategy when all three advertising types are present in the same auction. The optimal score (* $\overline{S}^*$ *) and bids (* $\overline{B}^*$ *) are as follows (Note: the notation below always shows the* k^{th} *advertiser in the* k^{th} *position):*

- M_4 *:* $\overline{S}_m^* = \left(\max\left\{S_{R5}, S_{R6}\right\} + \varepsilon\right)\ \forall m \in \left\{K_3+1,\ldots,M\right\}$

- M_3 *:* $\overline{S}_m^* = \dfrac{1}{CTR_{m-1}}\left[\left(v_m^{m-1}\cdot CTR_{m-1} - v_m^m \cdot CTR_m\right)q_m + CTR_m \cdot \overline{S}_{m+1}^*\right]$

  $\forall m \in \left\{K_2+2,\ldots,K_3-1\right\}$ *;*

  $\overline{S}_{K_3}^* = \dfrac{1}{CTR_{K_3-1}}\left[\left(v_{K_3}^{K_3-1}\cdot CTR_{K_3-1} - v_{K_3}^{K_3}\cdot CTR_{K_3}\right)q_{K_3} + \left(v_{R_6}^{K_3}\cdot CTR_{K_3}\right)q_{R_6}\right]$ *and*

  $\overline{S}_{K_2+1}^* = \dfrac{1}{CTR_{K_1}}\left[\left(v_{K_2+1}^{K_1}\cdot CTR_{K_1} - v_{K_2+1}^{K_2+1}\cdot CTR_{K_2+1}\right)q_{K_2+1} + CTR_{K_2+1}\cdot \overline{S}_{K_2+2}^*\right]$

- M_2 *:* $\overline{S}_m^* = \begin{cases} \min\limits_S\left\{M_1\right\} - \varepsilon, \text{ if max value closer to } M_1 \\ \max\limits_S\left\{M_3\right\} + \varepsilon, \text{ if max value closer to } M_3 \end{cases},\ \forall m \in \left\{K_1+1,\ldots,K_2\right\}$

- M_1 *:* $\overline{S}_1^* = q_1 \cdot v_1^1$, $\overline{S}_m^* = \dfrac{1}{CTR_{m-1}}\left[\left(v_m^{m-1}\cdot CTR_{m-1} - v_m^m\cdot CTR_m\right)q_m + CTR_m\cdot \overline{S}_{m+1}^*\right]$

  $\forall m \in \left\{2,\ldots,K_1-1\right\}$ *;*

  $\overline{S}_{K_1}^* = \dfrac{1}{CTR_{K_1-1}}\left[\left(v_{K_1}^{K_1-1}\cdot CTR_{K_1-1} - v_{K_1}^{K_1}\cdot CTR_{K_1}\right)q_{K_1} + CTR_{K_1}\cdot\Omega\right]$

where $\Omega = \dfrac{1}{CTR_{K_1}}\left[\max\limits_{m,k}\left\{\left(v_m^{K_1}\cdot CTR_{K_1} - v_m^k\cdot CTR_k\right)q_m\right\} + CTR_{K_2+1}\cdot\overline{S}_{K_2+2}^*\right]$,

$\forall m,k \in \left\{K_1+1,\ldots,K_2,K_2+1\right\}$. *Each advertiser's bid is equal to the optimal score (* $\overline{S}^*$ *) divided by their quality weight (* $\overline{b}_j^* = \overline{S}_j^*/q_j$ *).*

Proof: *See Appendix.* ⊗

The intuition underlying Proposition 3 is fairly straightforward. The BOO places advertisers in an order such that any advertiser's quality-weighted value for their assigned position is higher than any other advertiser's quality-weighted value for that same position. Thus, any advertiser should be able to bid an amount sufficient to maintain their optimal position, as defined by the BOO. The one situation where this does not occur is when an advertiser's best ranked position is not the top position. In these cases, the advertiser in her best ranked position is prevented from bidding her full value or a value large enough to guarantee placement in that position because of the auction rule that ranks advertisers by score. As a result, those advertisers who rank a position other than the top position as their best choice position will bid an amount equal to other neighboring advertisers.

This is what happens to the branding advertisers in the lowest group (M_4). Given that they rank the lowest position as their best choice, this group of advertisers will all compete for the lowest position. The advertiser who values the lowest position the most is prevented from bidding a lot for that position, as that will move her up to a higher position and she will be undercut by other advertisers in the group. Thus, all the advertisers in this group bid slightly more than is necessary to keep the first rejected advertiser out of the auction.

The second group of advertisers from the bottom (M_3) all want to move up a position and could be either T-type or K-type advertisers. Thus, they bid very similar to a Proposition 2 strategy; they bid an amount that makes them indifferent between their assigned position and the next position above them. The lowest advertiser in this group bids with respect to the negative externality she imposes on R_6, ignoring group M_4. The top advertiser in this group bids the largest amount she would pay to achieve the lowest position in M_1.

The group of advertisers in M_2, like the M_4 group, is near their best choice position and one of these K-type advertisers is satiated at her best choice position. Also similar to the setting in M_4, the one satiated advertisers in M_2 cannot bid enough to guarantee her best choice position due to the auction rules. As a result, she and the other advertisers in M_2 all bid such that they are either ε above the highest score in M_3 or ε below the highest score in M_1. The selection of one option will be based on which of the two positions has the highest rank for the group of K-type advertisers in M_2.

Finally, the M_1 group all bid according to the Proposition 2 strategy with the exception of the lowest advertiser in this group. This advertiser must bid sufficiently high to make the advertiser with the highest value from the M_2 or M_3 group indifferent to their current position or the lowest advertiser in M_1's position. Thus, the lowest advertiser bids the most she would pay to move up one position plus the largest negative externality imposed by her on those advertisers below her.

As described in Proposition 3, the presence of multiple advertiser types causes satiated, near-satiated, and bottom-desiring advertisers to misrepresent their scores. The effect of this on search engine revenue is presented in the Corollary to Proposition 3:

Corollary to Proposition 3 (Search Engine Revenue): *When T-type advertisers or unsatiated K-type advertisers (i.e., advertisers in M_1 and M_3) are replaced with satiated and bottom-desiring advertisers, respectively, and all advertisers have the same quality-weights, search engine revenue declines.*

Proof: *See Appendix.* $\otimes$

This concludes the analysis of the equilibrium bidding strategy. I now turn to a theoretical exploration of bidder subsidization.

IV. Subsidization of Bidders

With the above background, all of the tools are in place to evaluate bidder subsidization schemes. This represents the second contribution of this paper to the literature. This analysis is framed by the market design elements of Google's innovative program to subsidize charities. There are several important market design elements encapsulated in a search engine subsidization program of charitable organizations:

- **Participant Identification.** The program is limited to charitable organizations recognized under section 501c3 of the US tax code. Given the separate government entities tasked with ensuring that organization which receive this 501c3 classification are legitimate charitable organizations, search engines can capitalize on this institution to very easily identify eligible participants.

- **Historical Participation.** While same charities may participate to a limited extent without the program, the vast majority of charities will never expend the time and financial resources to understand optimal bidding behavior and benefits from participating in adword auctions.
- **Participant Valuations.** While charities occasionally do some B-type advertising, their primary objective is best characterized as generating donations of time or money (i.e., volunteering). To verify this, one is encouraged to visit the website of nearly any charitable organization. It is a rare charity that uses up a large percentage of the homepage to build a brand; instead the charity is asking for money donations, requesting time donations, or seeking some other transaction (e.g., newsletter signups). Thus, charities are readily classifiable as K-type advertisers.
- **Program Rules.** The Google program initially involved an application by organizations to prove they were indeed eligible and then the charity was provided with a $10,000/month credit by Google in their adword auction. This credit lasted indefinitely and did not require any financial contribution by the nonprofit. The credit had to be evenly spaced over the month at approximately $300-$350/day. Recently the Google program has been expanded to offer $40,000/month grants which require a $1500/month contribution by the charity to participate. Both grant programs renew indefinitely for eligible charities.
- **Deviation from Normal Rules.** Charities in the program are limited to a maximum bid of $1 for each keyword on which they bid.

Considered simply, the charity subsidization program described above brings an otherwise non-existent market participant to the adword auction. This participant can be classified as having a homogenous bidding strategy (K-type) and is easily identifiable. Requiring that the subsidy be equally extended throughout the month removes any incentive for participants to become T-type or B-type bidders. That is, if the subsidy could be spent in a single day, it may be in the charity's best interest to avoid spending energy to learn optimal K-type behavior and instead bid on expensive, high volume keywords that will bring small short term gains to the charity. This is not possible given the program rules. Participants must understand the rules and their own per click valuations enough to evenly spend their subsidy throughout the month.

In addition, by selectively targeting K-type bidders, the search engine is likely to see the subsidized advertiser bid on high specificity words with higher quality-weights. This lowers the total cost of the subsidy for the search engine and increases the likelihood that charities will have high enough scores with the bid cap to appear in the paid search results. Had the subsidization program been open to any advertisers, two challenges may have arisen: 1) the program may have attracted aggregator T-type advertisers which, as described earlier, are substitutes to any search engine and thus not preferred recipients for a subsidy; 2) the program may have attracted B-type advertisers who will aggregate toward the bottom of the search listings and not increase the quality of the advertiser offering to the consumers. Conversely, the objectives of K-type advertisers are consistent with those of the search engine in that both want to satisfy a consumer's need. Thus, subsidizing K-type advertisers is a best choice strategy for the search engine.

The question raised initially, however, is whether this best choice extends not only to long term revenue benefit but may also increase the short term revenue of the search engine platforms. As an example, consider the following five advertisers competing for the same four positions. The base case does not include the subsidized charity (S1).

Advertiser Scenario 3						
Position	Clicks	Value (T1) ($10/click)	Value (T2) ($5/click)	Value (T3) ($3/click)	Value (T4) ($2/click)	Subsidized Charity (S1) ($6/click)
Top: 1	6	60	30	18	12	36
2	3	30	15	9	6	18
3	1	10	5	3	2	6
Reject: 4	0	0	0	0	0	0

Clearly the other four advertisers (T1, T2, T3, and T4) will be ranked according to their τ valuations and will bid and pay according to Proposition 2. This is noted as the Original Payment in the table below. Then when the new subsidized charity (S1) enters the bidding, she is placed according to the BOO. The New Payment column describes the payments to the search engine with the additional bidder. In this example, it is the case that the addition of a new bidder in position 2 results in more profit (34 vs. 33) for the search engine, even when the entire payment of the incremental advertiser (13) is subsidized.

Position	Clicks	Advertiser	Original Payment	New Payment
Top: 1	6	T1	5*3 + 8 = 23	6*3 + 13 = 31
2	3	S1	--	5*2 + 3 = 13
3	1	T2	3*2 + 2 = 8	3
Reject: 4	0	T3	2	0
5	0	T4	0	0
Revenue			**33**	**47 – 13 = 34**

While the payments from the T2 advertiser and the T3 advertiser are reduced, the payment from the T1 advertiser is raised enough to offset the subsidy. One can observe from this example the benefits of a subsidization location right below a large drop in the click through rate. This causes the subsidized advertiser to bid a lot to make her indifferent between her current position and the position above her. And this higher bid causes the advertiser above the subsidized advertiser to pay significantly more. For simplicity, this example assumes that all advertisers have the same quality-weight. If this were not the case and the subsidized charity had a higher quality weight, the revenue difference would not change, but the amount paid to the charity would fall. Thus the subsidy from the search engine to the subsidized bidder falls as the quality-weight increases.

With this intuition, I present a more formal analysis of the search engine revenue effects of advertiser subsidization. Consider an auction that does not contain any satiated advertisers. That is, the auction is either only T-type in group M_1 or T-type and K-type in group M_3. The search engine that subsidizes a charity such that the charity enters a given keyword auction at position K, moving the original advertiser (K) down one position and the M^{th} advertiser to the $M+1^{st}$ position to become the first rejected advertiser. The score of the new subsidized advertiser is then denoted as $S_{K'}$. I assume there are M positions and that all scores maintain their pre-subsidization positional label. That is, before the subsidization, the last advertiser has score S_M and click through rate CTR_M; after subsidization this advertiser maintains score S_M but moves to click through rate $CTR_{M+1}=0$. With this notation, the change in revenue to the search engine following subsidization of an advertiser is defined as:

$$\begin{aligned}\Delta\text{Rev} &= \left(S_{K'}-S_K\right)\left(CTR_{K-1}-CTR_K\right)\cdot\left(\sum_{m=2}^{K}\frac{1}{q_{m-1}}\right)\\ &\quad+\sum_{j=K}^{M-1}\left[\left(\left(S_j-S_{j+1}\right)\cdot\left(\sum_{m=2}^{j}\frac{1}{q_{m-1}}\right)-\frac{S_{j+1}}{q_j}\right)\cdot\left(CTR_j-CTR_{j+1}\right)\right]\\ &\quad+\left(\left(S_M-S_{M+1}\right)\cdot\left(\sum_{m=2}^{M}\frac{1}{q_{m-1}}\right)-\frac{S_{M+1}}{q_M}\right)\cdot\left(CTR_M\right)\end{aligned} \tag{10}$$

Where ΔRev =(Rev$_{\text{Subsidy}}$ – Rev$_{\text{NoSubsidy}}$ – Subsidy), $S_K=q_K v_K$. For notational ease, equation (10) also assumes that scores are not position-dependent. This is not expected to be the case for K-type advertisers and thus the analog to (10) with position-dependent scores would be required. Derivation is trivial; multiple out the above score and CTR differences such that each term includes the product of a click through rate and score, where the value of the score is always taken at the position of the multiplied click through rate (e.g., $S_m^{K*}\cdot CTR_{K*}$).

An advertiser subsidization program involves three revenue adjustments which can be mapped to the three terms presented in equation (10). First, assuming the subsidized advertiser is not the top advertiser the first term represents the additional revenue earned by the search engine from the advertiser above the subsidized advertiser. The magnitude of this term increases with the size of the difference in scores between the new advertiser (K') and the displaced advertiser (K) and the size of the CTR between the new advertiser (CTR_K) and the position above her (CTR_{K-1}). This confirms the intuition generated by Advertiser Scenario 3. Thus, a subsidization scheme that brings in subsidized advertisers at the position below the largest drop-off in CTR will lead to the largest possible first term. As long as the subsidized advertiser does not win the top position, the first term will always be positive; otherwise, this first term is zero.

Second, search engine revenue will likely be decreased from each of the bidders below the subsidized advertiser who is pushed down a position. This effect is represented by the second term. These terms are not necessarily negative as a very large drop off in advertiser scores can cause these terms to be positive, i.e., for equal quality, the second term reduces to (for the m^{th} position): $m \cdot S_m - (m+1) \cdot S_{m+1} > 0$. Under competitive conditions (i.e., $S_m = S_{m+1} + \varepsilon$), however, these terms will most likely be negative.

Third, the search engine will receive no revenue from the last bidder who was knocked out of the bidding by the subsidized advertiser, as represented by the last term in equation (10). As in the discussion of the second term, this term may be positive resulting in an actual increase in revenue to the search engine from rejecting the original last bidder. Consider for example a situation where all advertisers have quality-weights equal to one and the pre-subsidization values for the $M-1^{st}$ and M^{th} advertiser are \$3/click and \$2.5/click, respectively. Assuming 10 advertisers and insertion of the subsidized bidder in the 4th position (i.e., $K'=4$), the third term above is $\left((3-2.5)\cdot 9-2.5\right)\cdot CTR_M = 2\cdot CTR_M$. If the click through rate is large for the last position relative to zero (i.e., the last accepted bidder exerts a large negative externality on the first rejected bidder) this term could be significantly larger than any of the components of the second term of (10), especially if the click through rate rises gradually over position. Continuing the example, consider an extreme case where the CTR is a constant 1% in every position. Thus, the first and second terms in (10) are zero and the third term is positive. In this case, even if the subsidized bidder earns the top position, the revenue to the search engine rises with subsidization.

In addition to these three mechanisms for changing revenue, a fourth change in revenue from subsidization could occur as a result of the third conclusion in Lemma 1: there is a CTR benefit on all positions to the search engine from attracting an advertiser who brings a higher quality-weight to the auction. Thus, bringing a new K-type advertiser to the auction – particularly one who bids on high quality-weight keywords –would be expected to increase the revenue to the search engine through this mechanism; whereas, bringing a new B-type advertiser would be likely to decrease search engine revenue due to their choice of lower quality keywords. Moreover, by capping the bid of the subsidized advertiser, the search engine is effectively forcing the charity to bid on only high quality keywords which further reinforces this effect. The magnitude of this average quality effect is not studied here but causes any estimates of revenue increase from subsidization to be conservative.

The value of the subsidy to the m^{th} advertiser in the K' position is defined as:

$$Sub_m^{K'} = \frac{1}{q_m} \cdot \sum_{j=K'}^{M+1} q_j \cdot \left(v_j^{j-1} CTR_{j-1} - v_j^{j} CTR_j \right) \tag{11}$$

Proposition 4 summarizes the revenue implications of advertiser subsidization, as shown by equations (10) and (11):

Proposition 4 (On Bidder Subsidization). *Search engine subsidization of a new advertiser by paying the entire cost to the advertiser such that the advertiser joins a particular keyword auction at position K' causing the lowest bidder to drop out leads to the following conclusions:*

1) *The change in search engine revenue as a result of the subsidizing event (and incorporating the cost of subsidization) is increasing in the degree of click-through-rate curvature in the position directly above K' and increasing in the level of click-through-rate flatness in the positions below K'.*
2) *Subsidizing a B-type advertiser (M_4) if the number of advertisers exceeds the number of positions decreases revenue.*
3) *Subsidizing a high quality advertiser has the effect of allowing them to participate in more auctions relative to subsidizing a low quality advertiser at an equal rate.*

Proof: *1) Follows from Equation (10). 2) Follows from the fact that B-type advertisers bid only* ε *higher than the first rejected advertiser. Thus adding a B-type does not increase the value of the first term and only reduces the revenue from the rejected bidder. 3) The revenue value any advertiser brings to an auction is equal to their score which is the product of their bid and quality-weight. Advertisers with higher quality weights will be required to bring lower bids to achieve the same score. Thus for equal total subsidies, an advertiser with a higher quality-weight will be able to participate in more auctions.* ⊗

Proposition 4 does not make a statement delineating when or if subsidizing a bidder is revenue improving. From inspection of equation (10), one can see that it is possible for subsidization to be revenue increasing but it is surely not necessary.

A final component deserving of analysis in the subsidy program described above is the inclusion of bid caps. The bid cap can serve multiple purposes, one of which is to respond to the possibility raised in Proposition 1 that K-type advertisers may become positionally satiated near the top, but not at the top of the list of advertisers. As shown in Proposition 3, multiple satiated and near-satiated K-type advertisers uniquely determine at most one advertiser's bid (i.e., the bottom advertiser in M_1). Thus, if any of the K-type advertisers who are not uniquely determining another advertiser's bid are subsidized, they are less likely to result in an increase in search engine revenue. With respect to equation (10), the first term is zero in this case and so the only effect of introducing a subsidized bidder is to move all of the advertisers below the subsidized bidder down a position. Under certain conditions described above, revenue may still increase. However, if this subsidized bidder were forced to move to a lower position by a bid cap (i.e., out of M_2 and into M_3), she would now be uniquely determining the bid of the advertiser above her. Intuition suggests that this condition is more likely to result in a revenue increase for the search engine; however, the model above does not support a Proposition on when bid caps are revenue improving for the search engine.

A second effect of the bid cap is that it will force advertisers to bid on only the highest quality-weight words or auctions with insufficient competition in order to win a position in the auction. Through the mechanisms discussed above – positive externality on CTR from higher quality advertisers, lower cost of subsidization – this will benefit the search engine. Finally, a bid cap may serve as an anchor for bids and thus, may influence the actual bid value for participating charities. Remember that the strategy described in Proposition 2 and 3 generates the lowest level of revenue for the search engine. If the anchor causes these advertisers to deviate from this strategy by bidding more[56], this will result in more revenue for the search engine. This is the third mechanism whereby bid caps may increase the revenue to the search engine.

This concludes the theoretical analysis of search engine keyword bidding. In the next two sections I bring data, econometrics, and simulations to the problem in an effort to better characterize the magnitude of the effects described in the theory and better understand the revenue implications of bidder subsidization.

[56] In practice, it appears that subsidized charities often bid the cap making this mechanism for increasing search engine revenue a very real possibility.

V. Functional Relationship Calibration

As stated in the introduction, a third contribution to the literature is to examine the effects described above using a novel dataset of keywords from three different firms across four different search engines (i.e., Google, Yahoo, MSN, and Ask). The objectives of the firms in placing advertisements varied as follows:

- Firm 1 is a for-profit firm wanting consumers to visit their site and sign up for their regular newsletter. This data was not generated under a bidder subsidization program.
- Firm 2 is a charitable organization wanting consumers to make a symbolic adoption of a wild animal as part of a charitable fundraising campaign. This data was generated under a bidder subsidization program.
- Firm 3 is also a charitable organization wanting consumers to donate directly to their organization. This data was also generated under a bidder subsidization program.

Table 1 summarizes the available data. Data for firms 1 and 2 are available on a daily basis; whereas, data for firm 3 is only available on a monthly basis. Table 1 reflects the adjustment of this monthly data from firm 3 into approximate daily statistics assuming 30 days/month. As can be seen from Table 1, the three firms vary by the quantity of keywords purchased on each search engine, the price they paid and the results they observed. It should not be surprising that Firm 1 has a higher conversion rate than the other two firms (3.8% - 11.5% compared to 0.6% - 1.3%) given the lower transaction cost of their advertisement's request to the consumer (free newsletter signup compared to financial donation).

My strategy in analyzing the data is as follows. Firm 2 has the richest collection of data on a single search engine and thus I will start by analyzing Firm 2's data to support the functional relationship conclusions drawn in Lemma 1 and Lemma 2. In addition and where possible, I will estimate the magnitude of these relationships. Then as a robustness check on this analysis, I will compare the three firms across the Google search engine and similarly I will compare Firm 1 across the four platforms (Ask, MSN, Yahoo, and Google).

Firm 2 purchased 2,300 keywords over 108 day period with 2,188 keywords purchased on at least two days. Using this data I will estimate the CTR and CR as a function of their arguments. Search engines do not currently provide information on the specific quality-weight or word specificity for each keyword (θ) and thus I cannot directly determine any of the relationships with respect to keyword specificity. To establish the relationship with specificity, I employ Assumption 1; namely, that keyword specificity is defined such that the number of impressions for low specificity keywords are greater than the number for high specificity keywords: $I'(\theta) \le 0$.

Using Assumption 1, I can solve for the effect of keyword specificity on the click-through-rate and conversion rate using the following equalities:

$$\frac{\partial CTR / \partial \theta}{\partial I / \partial \theta} = \frac{\partial CTR}{\partial I}; \qquad \frac{\partial CR / \partial \theta}{\partial I / \partial \theta} = \frac{\partial CR}{\partial I} \tag{12}$$

If the empirical analysis finds that $\partial CTR/\partial I \le 0$ controlling for position, then it must be the case that $\partial CTR/\partial \theta \ge 0$ because by Assumption 1 $\partial I/\partial \theta \le 0$. This confirms the first conclusion of Lemma 1. Similarly if $\partial CR/\partial I \le 0$, it must be the case that $\partial CR/\partial \theta \ge 0$ given that $\partial I/\partial \theta \le 0$ which is consistent with the first conclusion of

Lemma 2. The use of Assumption 1 precludes finding a magnitude of the effect of keyword specificity on CTR and CR. Figure 7, however, illustrates the distribution of daily impressions across the 1,256 keywords on Google.

I should also note that the available data does not allow me to evaluate the third conclusion of Lemma 1 regarding the increase in the click-through-rate as the average advertiser quality ($\bar{q}$) increases. This would require a dataset of advertisers by keyword so I can observe the affect on CTR of adding an advertiser with a higher quality weight or lower quality weight.

For purposes of the regressions below, I ignore keywords that do not receive any clicks over the term of the analysis and I ignore any days where the keyword is in a position greater than 20[57]. Both of these events describe low probability anomalies that fall outside the model described above.

Click-through-rate (CTR). I verify the first conclusion of Lemma 1 ($\partial CTR / \partial \theta \geq 0$) using the following regression:

$$CTR_i = \beta_0 + \beta_1 I_i + \beta_2 r_i + \beta_3 r_i^2 + \varepsilon_i \tag{13}$$

Given that I am ultimately interested in knowing the effect of keyword specificity on CTR, I run this regression across words and not within words. Specifically, I consider the average CTR for each word and regress that on the average impressions for each word, controlling for position. This regression, including various

[57] An advertisement in a position greater than 20 does not appear on the search results page but on a deeper page of advertisers. Thus, for a consumer to see this advertiser, the consumer must go to the bottom of the list of advertisers and click the link to see more advertisers. While infrequent, this behavior falls outside the activity modeled above and thus any days where the advertisement is in a position greater than 20 are removed from the dataset.

transformations of the dependent and independent variables is shown in Table 2. Given that I am only using keywords that receive at least one click over the period of analysis, there are no censored data in this analysis. The row labeled "Word Dummy Included" is used to describe the case when a dummy variable for each word is included. When marked "yes", the regression describes a linear regression with fixed effects for each keyword. In each of the regressions from Table 2, the coefficient on impressions is negative and it is significant in regressions 1 and 3 at the 5% and 0.1% levels respectively. The results suggest that lower impression words have a higher CTR than higher impression words, controlling for position. Using the first equality in equation (12), one can conclude that the click through rate is rising for higher specificity keywords ($\partial CTR / \partial \theta \geq 0$), as predicted in the first conclusion from Lemma 1.

The second conclusion from Lemma 1 regarding the effect of position on click through rate ($\partial CTR / \partial r \leq 0$) requires a more detailed discussion of the data and data generating process. Given that over 80% of the word*day observations have zero clicks (96% of the words with positive clicks have zero conversions), prediction of the CTR or CR requires either a two-staged hurdle model or alternatively a tobit or Heckman model. Determination of the appropriate model requires an assumption of the underlying mechanism for generating the zero click or zero conversion days. If the mechanism that generates zero clicks (conversions) is independent of the mechanism that generates positive clicks (conversions), then a two-staged hurdle model is appropriate. Alternatively, if the same mechanism generates zeros and positive clicks (conversions) then a tobit model or two-stage Heckman model is appropriate. An example of this later case could be that in the consumer clicking model, consumers do not follow a cascade model as predicted by Athey and Ellison (2007) but instead choose never to click on the top two advertisements unless the likelihood of those advertisements solving their need exceeds some very high threshold. Instead, consumers jump to the third advertisement and start clicking there. There is no theory in the literature to support such a model, nor are there any empirical analyses of search engine data that reinforce this as a possible model. However, the possible existence of such a model warrants analysis.

Including this model in our range of data generating processes, there are three mechanisms through which zero clicks could be generated: 1) the associated keyword receives very few impressions and thus no advertiser receives any clicks; 2) the advertiser is in a very low position on the page and receives no clicks; or 3) there is some nonlinear consumer rule about clicking on a particular advertisement (as described above). To disentangle these three causes of zero clicks, one would need to know clicks by keyword for many advertisers on the same keyword. The dataset available here does not include this information and thus arriving at a definitive conclusion may not be possible. However, one can gain insight into the determinants of receiving zero clicks through the following regression:

$$\Pr\left(CTR_i > 0\right) = \beta_0 + \beta_1 I_i + \beta_2 r_i + \beta_3 r_i^2 + \varepsilon_i \tag{14}$$

The results are shown in Table 3 and confirm that a low number of impressions is a strong determinant of zero clicks, as is a low position. Reducing the number of impressions from 54 per day to 20 per day reduces the probability of a click by 7 percentage points. The effect of position is less: dropping from the third to the fourth position reduces the probability of a click by 1.5 percentage points. This provides support for the cascade model of consumer clicking described in Section III, but does not exclude the possibility of a non-monotonic clicking strategy of consumers.

As a further test of the model, I run both a random effect panel tobit model and the second stage of a two-stage panel hurdle model. By using only keywords with at least one click over the period of analysis, I exclude a large number of the low impression words which receive zero clicks; however, over 70,000 keyword*day combinations remain with zero clicks. The random effects tobit model assumption would say that every data point with zero clicks provides information on the propensity of that position to elicit clicks. As an example of where this might break down, consider two words. Word 1 receives on average 5 impressions a day and only has 2 advertisers competing for advertising slots. Word 2 receives on average 500 impressions a day and has 10 advertisers competing for advertising slots. Firm 2 is present in the top position 1 of Word 1 but never receives any clicks. Due to the more competitive nature of Word 2, Firm 2 is present in position 6 and receives many clicks in this position. The assumption underlying the tobit model would use this data to support a model of increasing click through rates as an advertiser moves further down the page (i.e., zero clicks in the first position and many clicks in the 6th position supports an increasing CTR further down the page).

The alternative assumption necessary for the two-stage hurdle model to be the appropriate model is that when you exclude words that receive zero clicks for low impression reasons, the zero click events are caused because the advertiser is too far down the page relative to the number of consumers seeing the page and not for any other consumer behavioral clicking reason. If this were the case, then the second stage of the two-stage model could be a regression of CTR on position conditional on a positive number of clicks.

Unfortunately, the dataset available here does not provide sufficient data to completely disentangle these two effects and as a result the presence of zero click days confounds the results. However, by running both the random effect panel tobit and the fixed effect OLS model one can gain yet further insight as to the appropriate model. The results are presented in Table 4 and are consistent in their analysis of the CTR – position relationship. Both show that as the advertiser falls to a lower position, the CTR also falls. This confirms the second conclusion of Lemma 1.

This is further evidence supporting the hypothesis that there does not exist some non monotonic consumer behavior that causes consumers to be more likely to click further down the list of advertisers. Thus, I use the following fixed effect OLS regression as my base regression in the robustness analysis:

$$\ln\left(CTR_{it}\right) = \beta_0 + \beta_1 \ln\left(I_{it}\right) + \beta_2 r_{it} + \beta_3 r_{it}^{\,2} + \mu_i + \varepsilon_{it} \qquad (15)$$

This is the second stage of a two-stage hurdle model and tests the assumption that if an advertisement is going to get any clicks, it will generate more clicks if higher in the list of advertisers. While this still may not be the case, the data available does not support further disentangling the mechanism causing zero click days.

Conversion rate (CR). Following a methodology similar to the click through rate, I verify the first conclusion of Lemma 2 ($\partial CR / \partial\theta \geq 0$) using the following regression:

$$CR_i = \beta_0 + \beta_1 I_i + \beta_2 r_i + \beta_3 r_i^2 + \varepsilon_i \qquad (16)$$

As in the CTR regression, I consider the average conversion rate and average daily impressions for each keyword. Like the regressions in Table 2, this regression does not have a keyword fixed effect which allows the interpretation of β_1 to include between-word differences. The results for this regression and various log transformations of the variables are shown in Table 5. As predicted in Lemma 2, the coefficient on the impressions term is negative in every specification regardless of whether the regression is conditioned on a positive conversion rate. Using the second equality in equation (12) and Assumption 1, this demonstrates that the conversion rate is rising for higher specificity keywords.

To illustrate the results from the second conclusion of Lemma 2, I use both a random effect panel tobit model and a fixed effect OLS regression to show that $\partial CR/\partial r \geq 0$. While the results are not always significant, Table 6 does show that as the position increases (i.e., advertisers further down the page), the conversion rate weakly increases. Selection of the correct model for the conversion rate follows a similar course of reasoning as discussed above with respect to the click through rate. Given the difficulty in disentangling the data generating process for the zeros, but the consistency of results across several model specifications, I choose to use the following fixed effect OLS specification as my base regression in the robustness analysis:

$$\ln\left(CR_{it}\right) = \beta_0 + \beta_1 \ln\left(I_{it}\right) + \beta_2 r_i + \beta_3 r_i^2 + \mu_i + \varepsilon_{it} \tag{17}$$

Robustness Checks

This richness of the dataset allows for a robustness analysis of Lemma 1 and Lemma 2 across different advertisers and different search engine platforms. First I consider the effect of different platforms. Figure 8 illustrates the average daily keyword impressions across the four search engine platforms. They generally reflect similar keyword behavior, i.e., most keywords on which Firm 1 bid are low impression keywords. Tables 7 and 8 summarize the results from all three firms on the same search engine platform (Google) using regression 2 from Table 4 as the base regression for CTR and regression 3 from Table 6 as the base regression for CR. In Table 7, the signs for the three firms are the same with respect to the effect of position on CTR; however, the magnitude vary. The CTR on keywords bid for by Firms 1 and 3 drops faster for positions further down the page than the CTR on keywords bid for by Firm 2. This is evidenced by the larger regression coefficient on position for Firm 1 and 3 (-0.420 and -0.307, respectively) compared to the coefficient on position for Firm 2 (-0.104). The implications of this on platform revenue and bidding strategy will be explored in the next section during the simulations.

Turning to Table 8 one can see that the conversion rate relations using regression 3 from Table 6 is also similar in sign and magnitude across advertisers. Thus, one can conclude that the results from Lemma 1 and 2 are robust to different advertisers.

A second check to robustness is to compare the same advertiser across search engines. As alluded to above, different search engines employ slightly different rules with respect to assignment of advertiser quality-weights, however, all search engines examined here employ a generalized second price algorithm for determining the cost per click. Table 9 and 10 demonstrate that with respect to click through rate, MSN, Yahoo, and Google are comparable. Ask is the smallest of the four search engines and thus has significantly less traffic than the other three. The effect of this is that, as shown in Table 1, Firm 1 was almost always in the top position on Ask for keywords. As a result of this lack of variation in position, one might expect the statistically insignificant coefficients observed in Regression 1 of Table 9. With respect to the conversion rate, Table 10 demonstrates that MSN, Yahoo, and Google again are similar while Ask appears to be an outlier. In all three search engines the conversion rate is rising for positions further down the list of advertisers, though this effect is only statistically significant on Yahoo. These results demonstrate that the conclusions from Lemma 1 and 2 are robust to different search engines.

With this setup, I now turn to the final section involving simulating the effect of bidder subsidization.

VI. Simulation

The subsidization of weak bidders is discussed in theory in Section III where I conclude that it may be revenue improving to the search engine to entice an additional bidder to the keyword auction even if all of the costs to that bidder are covered by the search engine. The model, however, states this only as a possibility and not a guarantee. It is clearly possible that there are cases where subsidization may cost the search engine more than it gains. To evaluate the conditions under which subsidization is beneficial to the search engine I have built a simulation of the auction environment using actual keyword data and click through rate results from the regressions. I then compare search engine revenue with and without a subsidized bidder. The results are presented in this section.

The model assumes that under the base case, there are 11 T-type or K-type advertisers competing for 10 positions. All of the advertisers rank the top position as their best choice with the second highest position ranked second and so on until the 10^{th} position. That is, ordinal rankings are strictly decreasing for position further down the page. To normalize competitive effects between advertisers, I model all advertisers as having a marginal value per click independent of position and drawn from a fixed uniform distribution. This is consistent with the model for T-type advertisers. With respect to K-type advertisers, this accounts for the fact that while the conversion rate conditional on positive conversions may be increasing in position, the vast majority of clicks result in zero conversions. Thus the unconditional conversion rate is much flatter with respect to position and for purposes of these simulations, I have assumed the unconditional conversion rate to be constant across positions. The quality-weights for each of these advertisers are also drawn from a uniform distribution. In each base case model the 11 advertisers are bidding on 10 keywords, each with a distinct click through rate, as determined by the data in Section V for keywords with a positive CTR[58].

Specifically I run regression 2 from Table 4 and extract the keyword fixed effect for each word. Then for 10 randomly chosen keywords, I use the regression coefficients, this keyword fixed effect, and the average number of daily impressions

[58] I do not attempt to model zero click events. Instead I only include in the subsidization simulations keywords where there are a positive number of clicks each day.

to estimate the click through rate for each position. As discussed above, for keywords with a low probability of receiving any clicks, these predicted CTRs may be slightly higher than in actuality. However, the occasionally elevated CTR should have no affect on the simulation results.

Once each advertiser's value by position is determined, I run the Gale-Shapley algorithm to determine the Bidder Optimal Ordering for the group of advertisers. Each advertiser's bidding strategy is based on Proposition 2 which is optimal in this setting because all advertisers rank the top position as their best choice. That is, every advertiser experiences a negative externality imposed by the advertiser above them, except for the top advertiser who is satiated. Equilibrium bids are determined and then based on the specifics of the selected keyword, the total cost to each advertiser and thus revenue to the search engine is determined. This process is repeated for 10 keywords and the sum of the revenue to the search engine becomes the search engine revenue for one simulation. Each model consists of 1,000 simulations such that the expected search engine revenue is the average revenue from each simulation:

$$\text{Rev} = \tfrac{1}{T}\sum_{sim}\sum_{w}\sum_{i}\frac{b_{i+q}q_{q+1}}{q_i} \tag{18}$$

Where i is the index for the 11 advertisers, w is the index for the 10 keywords and sim is the index for the 1,000 simulations ($T = 1,000$).

Following each model (of 1,000 simulations with 11 advertisers competing for 10 positions over 10 keywords), I introduce a 12th advertiser as the subsidized advertiser. The quality-weight for each model is drawn from a uniform distribution starting at $q \sim U[0,1]$. For each subsequent model, I increase the upper and lower bound on the quality-weight distribution for the subsidized bidder by 0.5. That is, in the second model, the quality-weight for the 12th subsidized bidder is drawn from the uniform distribution $q \sim U[0.5,1.5]$. This continues up to the 90th model where the quality-weight for the subsidized bidder is drawn from the uniform distribution $q \sim U[9,10]$. I assume that all advertisers have the same quality weight for each of the 10 keywords, but that the quality weight differs across days.

It should be noted that while the selection of the specific quality weights are arbitrary, all that matters for determination of the auction outcome is the relative quality weights. Thus, by varying the quality weight of the subsidized advertiser from 0-10 relative to a group of base case advertisers with quality weights between 3.5 and 4.5, I am effectively comparing a situation where the subsidized advertiser has a lower quality weight than the base advertisers, same quality weight as the base advertisers, and higher quality weight than the base advertisers. This encapsulates the three possibilities for the subsidized advertiser and is independent of the exact quality weights assigned by the search engine.

To summarize, a single session consists of the following:

1) Model 1: *a=0; b=1*

    a. Base Case Simulation (repeated 1000 times)

        i. 11 Advertisers compete over 10 position for each of 10 keywords

        ii. Per click value for these advertisers: $v \sim U[0.5, 2.5]$

        iii. Quality-weight for these advertisers: $q \sim U[3.5, 4.5]$

b. Subsidized Simulation (repeated 1000 times)

  i. Identical 11 Advertisers plus a 12th Advertiser compete over 10 positions for each of the same 10 keywords as in base case

  ii. Per click value for base advertisers (not re-sampled); Per click value for subsidized advertiser $v_{sub} \sim U[0,1]$.

  iii. Quality-weight for base advertisers (not re-sampled); Quality-weight for subsidized advertiser is drawn $U \sim [a,b]$

2) Model 2: *a=0.5; b=1.5*

  a. Repeat above

⋮

90) Model 90: a=9; b=10

Simulation results

Quality-Weight and Bid Caps. The purpose of the simulations is to determine the effect of advertiser subsidization on search engine revenue under several settings. First, I will explore the interplay between the quality-weight of the subsidized advertiser and bid caps for the subsidized bidder relative to the base case bidders. That is, I will run the above session under three bidder valuation scenarios: 1) $v_{sub} \sim U[0,1]$; 2) $v_{sub} \sim U[1,2]$; and 3) $v_{sub} \sim U[3,4]$. The first scenario mimics a tight bid cap and forces the expected value of the subsidized advertiser to be at the lower end of the possible values for the base case advertisers. The second scenario ($v_{sub} \sim U[1,2]$) has a relaxed bid cap which mimics the case where the expected value for the subsidized advertiser is equal to that of the base case advertisers, though drawn from a tighter interval. The tighter interval mimics the fact that the bid cap may result in anchoring for the subsidized advertiser which will decrease the variance of their values. And finally, I model a loose bid cap which results in the subsidized advertiser having a per click valuation that is in expectation greater than the base case advertisers.

The setup is described in Table 11 and the results are presented in Figure 9. The top three graphs in Figure 9 illustrate the revenue difference between the base case and the subsidized case under the three scenarios. The thick solid black line is the expected revenue difference and the thinner dotted lines are the standard errors on the revenue difference. In each case, there is a portion of the black line that is positive which indicates that in expectation, subsidizing a bidder can result in incremental revenue under the assumed model specification. The bottom three graphs indicate the position (smooth dark line) and the value of the subsidy (ridged lighter line) under each of the three scenarios.

There are several conclusions one can draw from these graphs. First, as the quality of the subsidized advertiser increases (represented by the X-axis in all graphs), their winning position increases up to the point of winning the top position. When the subsidized advertiser wins the top position, revenue for the search engine appears to decline. This is particularly evident in the third graph when the subsidized advertiser starts winning the top position around a quality weight of 4. At this point, the change in revenue for the search engine falls below zero and subsidization is no longer profitable.

A second observation is that as the bid cap tightens, there is an increased likelihood – under the assumptions made in this model – that advertiser subsidization is profitable for the search engine. And the profitability increases with increasing quality of the subsidized bidder. This is evident in the first graph. As the subsidized advertiser starts winning, the profitability of subsidization increases. The reason for this is that the size of the subsidy is very small (<4) but the large quality-weight of the advertiser causes them to have a large score capable of winning the 7^{th} or 8^{th} position. Thus, the subsidized advertiser increases the payment of the advertisers above her without requiring much actual subsidization from the search engine.

Finally, as the bid cap weakens, the total cost of the subsidy to the search engine increases. With a bid cap of 1, the largest subsidy is 4; at a bid cap of 2, the subsidy increases to 11-12; and a bid cap of 3 causes the subsidy to increase in places to over 22. As predicted, the bid cap serves the purpose of requiring the advertiser to bring quality to the auction if they want to win because they will not be winning based only on the size of their bid. As discussed above, this implied motivation for advertisers is a natural strategy for the K-type advertisers whose best choice words are high specificity, earning them a high quality score.

Role of CTR Curvature. The second point of exploration relates to the curvature of the click through rate equation. As described in Table 7 and 9, there is some indication that while the click through rate is declining as the position falls, it may be declining at differing rates. To test the implications of various curvature assumptions on subsidization I use three sets of coefficients on position and position-squared that represents an even more extreme curvature. The exact coefficient differences are described in Table 11 and the results are illustrated in Figure 10.

The top three graphs again illustrate the difference in revenue to the search engine between the subsidized case and the base case. And the bottom three graphs illustrate the average position of the subsidized advertiser and subsidy paid to the subsidized advertiser. The subsidized advertiser in all three cases is assumed to have a bid cap of 2 (i.e., the middle case in Figure 9).

The results are striking. As the click through rate curvature increases (faster drop off for positions further down the list of advertisers), the optimal target position for the subsidized advertiser, from the perspective of the search engine, narrows. In the first graph, the search engine earns incremental profits from subsidization as long as the subsidized advertiser wins a position lower than the 2^{nd} position from the top. By the third case, the search engine experiences lower profits if the subsidized advertiser wins anything higher than the 6^{th} position.

Intuitively this stems from the calculation of the negative externality for each advertiser in each position. If, for example, the click through rate drop off were exactly linear in position, each advertiser would pay to move from position 4 to position 3 as much as they would from position 3 to position 2. However, if there is a substantial drop in the click through rate in the first few positions, advertisers experience an increasing negative externality as they get closer to the top. That is, they would much rather move from position 3 to position 2 relative to moving from position 4 to position 3. In the third graph, the click through rate is much steeper in the first few positions and thus a subsidized advertiser moving into one of those positions exerts a large negative externality on those advertisers who she displaced. This results in a large negative magnitude of term 2 in equation (10) and thus a negative profit resulting from subsidization.

This result confirms the conclusions presented in Proposition 4 and illustrates that advertiser subsidization with certain bid caps can be profitable for search engines.

Potential Benefits of Anchoring. The final analysis reflects the possibility discussed earlier that the bid cap may serve as an anchor for the bids of subsidized advertisers. That is, given the potential lack of bidding sophistication of subsidized advertisers, particularly when they first start bidding, they may bid the bid cap instead of following the optimal strategy described in Proposition 2 or Proposition 3. As discussed in the Corollary to Proposition 2, any deviation from the Proposition 2 bidding strategy will result in strictly higher revenue for the search engine. To consider this through simulation I consider two sessions: 1) the subsidized bidder is known to have a value of one but follows the Proposition 2 bidding strategy; and 2) the subsidized bidder is known to have a value of one and deviates from the Proposition 2 strategy by bidding their value regardless of their position. In this second session, all other advertisers adjust their bid based on the new advertiser ordering with the subsidized advertiser deviating from the Proposition 2 strategy. The sessions and resulting figure are described in Table 11.

Figure 11 describes the results of these simulations. As is quickly evident from the figure on the right, this small deviation from the Proposition 2 bidding strategy results in more revenue to the search engine. The expected incremental revenue when the Proposition 2 bidding strategy is followed by the subsidized bidder is \$3.47 (i.e., revenue with subsidization minus revenue without subsidization is \$3.47 as shown from the figure on the left). When the subsidized advertiser deviates from Proposition 2 and bids their full value, this incremental revenue to the search engine jumps to \$7.06. From the lower two graphs in Figure 11 it is clear that the deviation moves the subsidized advertiser to a higher position and costs the search engine more money in terms of the actual subsidy. This additional expense to the search engine, however, is more than compensated for by the incremental fees paid by the advertisers ranked above the subsidized advertiser. Note that the subsidized advertiser is most often in position 2 or 3 when their expected quality exceeds 6 and thus, the advertiser in the first or second position is being charged significantly more with the subsidized advertiser present than without her present. This confirms the theoretical result presented in the Corollary to Proposition 2 and provides evidence of yet another benefit to the bid cap.

VII. Conclusion

Search engine marketing represents one of the largest and fastest growing means for generating revenue on the internet with over \$8 billion raised in 2007. The sheer magnitude of this market and the number of transactions estimated at over 2,000 each second provides economists with a tremendously rich environment for understanding the auction dynamics.

In this paper I extend the current literature in several dimensions. First, I incorporate search engine assigned quality-weights which are used by search engines to ensure that advertisers are sorted in order of their benefit to the consumer rather than just their bid. This consumer benefit measure is also consistent with the long run revenue benefit to the search engine. In addition, I present a model where keywords can be differentiated by their impressions, click through rate and conversion rate. With this framework, I present asymmetric advertisers who bring dramatically different strategies to the keyword auction. I first assume only a subgroup of these advertisers is present; the same subgroup assumed to be present in the current literature. With this subgroup, I derive an equilibrium bidding function that parallels the equilibrium bidding function currently presented in the literature. The introduction of quality-weights, however, causes the equilibrium bidding function to consider the overall value the advertiser brings to the auction and not just their bid. Then, I present an equilibrium bidding strategy that is robust to different advertiser types.

With this complex set of asymmetric advertisers, I explore the role of advertiser subsidization to determine if it can be successfully used to increase search engine revenue. This represents my second contribution to the literature. I find that while not guaranteed, there are conditions under which advertiser subsidization can be revenue increasing for search engines.

Following a theoretical evaluation, I present an extremely rich dataset which I use to fully characterize the relevant functional relationships (click through rate and conversion rate) important to keyword advertisers. The analysis and robustness checks using this dataset represent the third contribution to the literature. This extension and validation of predictions with respect to the click through rate and conversion rate provides a foundation on which to test alternative market design initiatives. This testing of market design initiatives represents the fourth and final contribution to the literature. Using actual search engine data and simulations, I demonstrate that under certain conditions, advertiser subsidization can be used to increase profitability for search engines.

VIII. References

Aggarwal, Gagan, Jon Feldman, S. Muthukrishnan, and Martin Pal. 2008. "Sponsored Search Auctions with Markovian Users." In ACM Conference on Electronic Commerce (EC), 2008.

Aggarwal, Gagan, Ashish Goel, and Rajeev Motwani. 2006. "Truthful auctions for pricing search keywords." In ACM Conference on Electronic Commerce (EC), 2006.

Agarwal, Nikhil, Susan Athey, and David Yang. 2009. "Skewed Bidding in Pay Per Action Auctions for Online Advertising." American Economic Review Papers and Proceedings, May 2009.

Anderson, Chris. 2006. "The Long Tail: Why the future of Business is Selling Less to more." Hyperion: New York.

Athey, Susan and Glenn Ellison. 2008. "Position Auctions with Consumer Search." Working paper Harvard and MIT.

Ausubel, Lawrence M. and Peter Cramton. 1999. "The Optimality of Being Efficient." Working paper University of Maryland.

Ayres, Ian and Peter Cramton. 1996. "Deficit reduction through diversity: How affirmative action at the FCC increased auction competition." Stanford Law Review. 48:761–815.

Bulow, J., and P. Klemperer. 1996. "Auctions Versus Negotiations." American Economic Review. 86(1): 180-194.

Caillaud, Bernard and Bruno Julien. 2003. "Chicken & egg: competition among intermediation service providers." RAND Journal of Economics. 34(2): 309-328.

Davis, Wendy. 2006. "NBC Expands Web Foothold with iVillage Acquisition." Mediapost archive. Available at: http://www.mediapost.com/publications/index.cfm?fuseaction=Articles.san&s=40634&Nid=18894&p=347224

Dellarocas, Chrysanthos and Siva Viswanathan. 2008. "The Holy Grail of Advertising? Allocative Efficiency and Revenue Implications of "Pay-per-Action" Advertising in Environments with Quality Uncertainty." In ACM Conference on Electronic Commerce (EC), 2008.

Edelman, B., M. Ostrovsky and M. Schwarz (2007), "Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords," *American Economic Review*, 97(1): 242-259

Gale, David, and Lloyd Shapley. 1962. "College admissions and the stability of marriage. American Mathematical Monthly." 69: 9-15.

Goel, Ashish, Mohammad Mahdian, Hamid Nazerzadeh, and Amin Saberi. 2008. "Advertisement Allocation for Generalized Second Pricing Schemes." Presented at the 2008 ACM Conference on Electronic Commerce, Chicago, IL.

Goldfarb , Avi and Catherine Tucker. 2008. Search Engine Advertising: Pricing Ads to Context. In ACM Conference on Electronic Commerce (EC), 2008.

Higgins, Nathaniel and Michael Roberts. 2008. "Increasing Competition through Quotas." Working paper. University of Maryland.

Krishna, Vijay. 2002. "Action Theory." Academic Press: New York.

Kempe, David and Mohammad Mahdian. 2008. "A Cascade Model for Externalities in Sponsored Search." In ACM Conference on Electronic Commerce (EC), 2008.

Klaassen, Abbey. 2007. "Search Marketing Fact Pack 2007: a Supplement to Advertising Age." Published by Advertising Age. Available at: http://adage.com/images/random/datacenter/2007/searchfactpack2007.pdf

Krishna, Vijay. 2002. "Action Theory." Academic Press: New York.

Lahaie, Sebastien, David M Pennock, Amin Saberi and Rakesh V. Vohra. 2007. "Sponsored Search Auctions," in Noam Nisan, Tim Roughgarden, Eva Tardos and Vijay V. Vazirani, eds., Algorithmic Game Theory. Cambridge: Cambridge University Press. Pp 699-716.

Maskin, E. (1992), "Auctions and Privatization," in Horst Siebert, ed., *Privatization: Symposium in Honor of Herbert Giersch*. Tubingen: Mohr (Siebeck), 115–136

McAfee, P. and J. McMillan (1987), "Auctions and Bidding," *Journal of Economic Literature*, 25: 699–738

Milgrom, P. and R. Weber (1982), "A Theory of Auctions and Competitive Bidding," *Econometrica*, 50: 1089-1122

Myerson, Roger B. (1981), "Optimal Auction Design," Mathematics of Operations Research, 6, 58-73.

Rochet, Jean-Charles and Jean Tirole. 2006. "Two-sided markets: a progress report." RAND Journal of Economics. 37(3): 645-667.

Roth, Alvin and Marilda A Oliveira Sotomayor. 1990. "Two-sided Matching: A study in Game-Theoretic Modeling and Analysis." Econometric Society Monographs. Cambridge Press.

Roughgarden, Tim and Mukund Sundararajan. 2007. "Is Efficiency Expensive?" Working paper Presented at the 2007 ACM Conference on Electronic Commerce. Banff, Canada.

Stokes, Richard. 2008. "Search Engine Advertiser Update – Q208." AdGooroo Advertising insight.

Varian, Hal. 2007. "Position Auctions." International Journal of Industrial Organization. 25: 1163-1178.

Varian, Hal. 2009. "Online Ad Auctions." American Economic Review. 99(2): 430-434.

IX. Tables and Figures

Figure 1 Example of Search Engine Search Results page

The assignment and pricing of advertisers to paid advertising slots are the focus of this paper. The organic results are determined through a proprietary algorithm unique to each search engine.

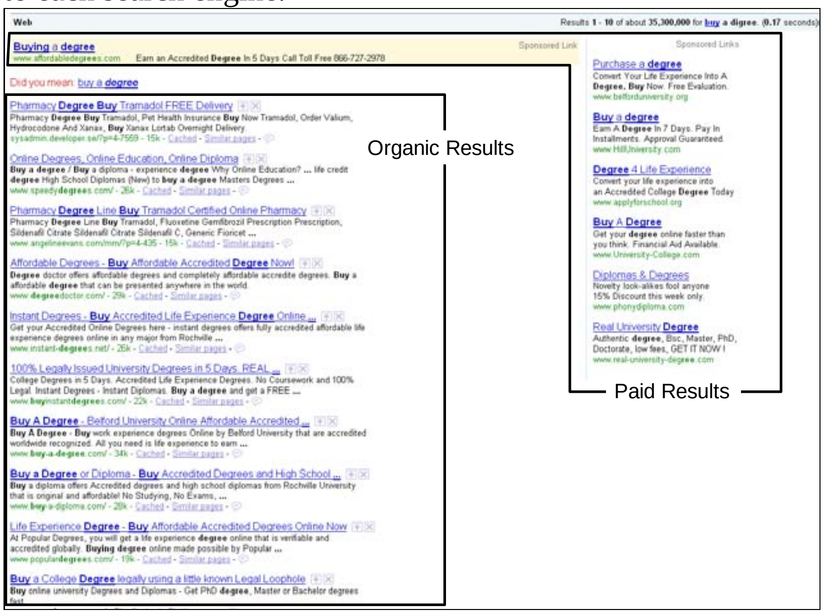

Figure 2 Determining the optimal number of words to purchase.

Here $\overline{\Pi}$ is the fixed surplus per word and $\alpha'_x(n_x)$ is the marginal cost of managing n_x words under one management system and $\hat{\alpha}'_x(n_x)$ is a more efficient system with a lower marginal cost per word. The equilibrium number of words under the two word-management systems is N_1 and N_2, resulting in surplus A and B, respectively.

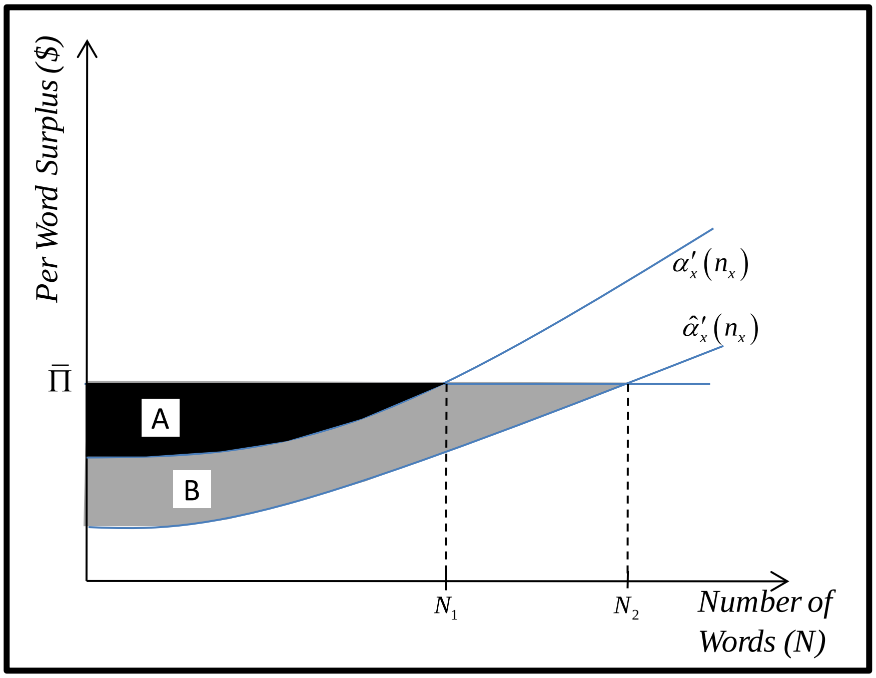

Figure 3 The discontinuous relationship between score (b_iq_i) and position

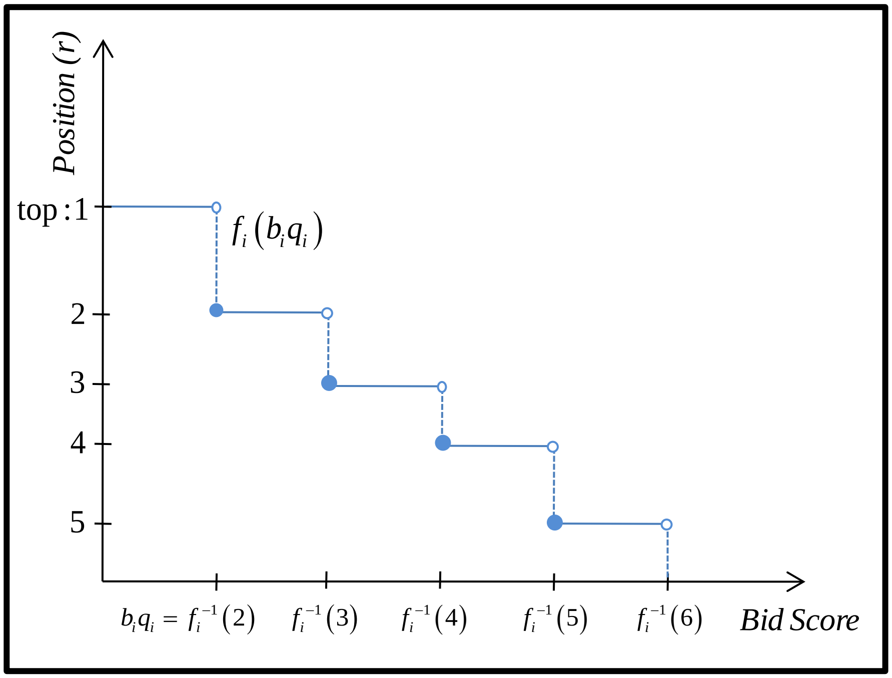

Figure 4 The distribution of search costs and likelihood of a search need being met for consumers in group A and B

Consumer group A are those consumers who know what they want and how to search for it, thus they have lower search costs and a higher likelihood of having their need met. Consumer group B either do not know what they want or do not know how to search for it giving them higher search costs and a lower likelihood of having their need met.

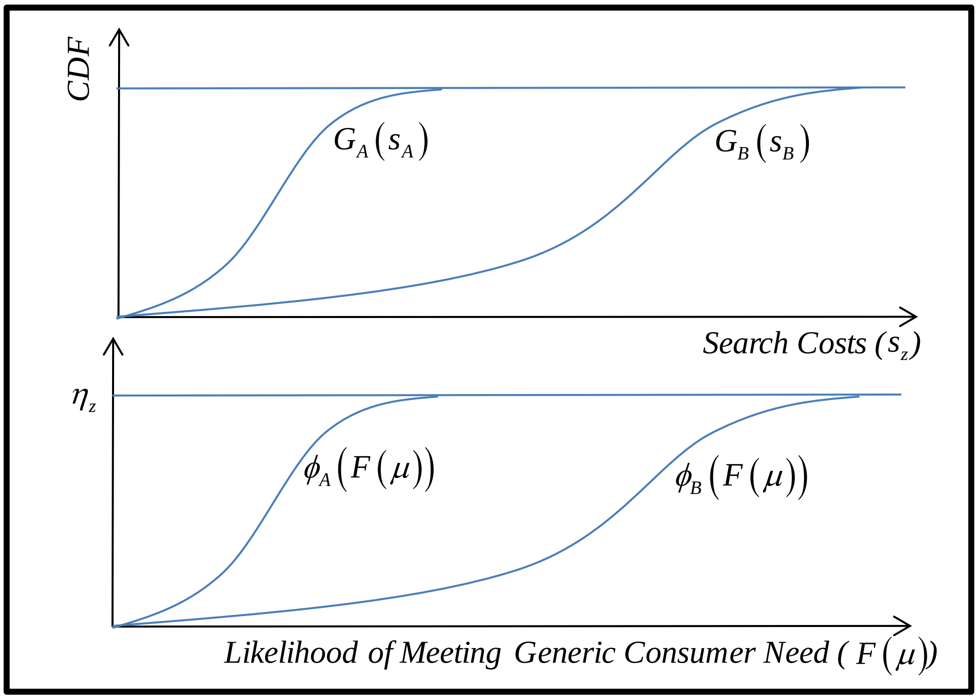

Figure 5 Example of increasing conversion rates

The overall conversion rate of an advertiser is a weighted average of the likelihood of that advertisers meeting the need of consumer group A (those who know what they want) and consumer group B (those who do not know what they want). $F(\mu)^{(1)}$ is the first order statistic of the generic probability of an advertiser to meet a consumer's need. This is transformed into the probability of meeting the need of a consumer in group A and B through the functions $\phi_A(F(\mu))$ and $\phi_B(F(\mu))$, respectively.

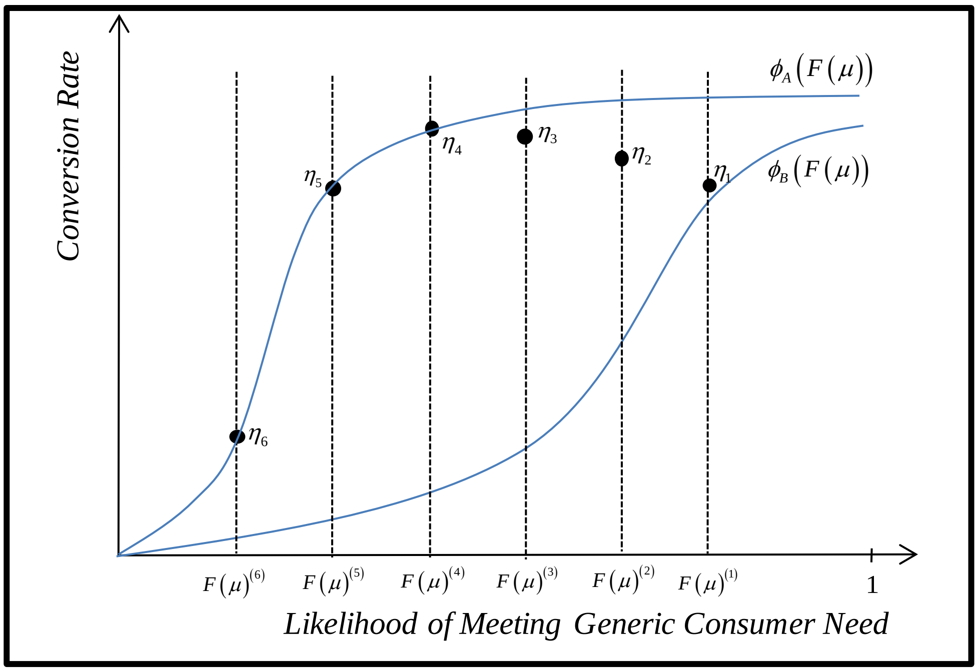

Figure 6 Possible Bidder Arrangement Blocks

T-type (traffic advertisers), B-type (branding advertisers), and K-type (conversion advertisers) in equilibrium will arrange themselves into one of four specific blocks depending on the revenue per position for the K-type advertisers. The revenue per position for B-type and T-type advertisers is fixed for any particular keyword but the revenue per position for the K-type will be proportional to the product of the conversion rate and click through rate.

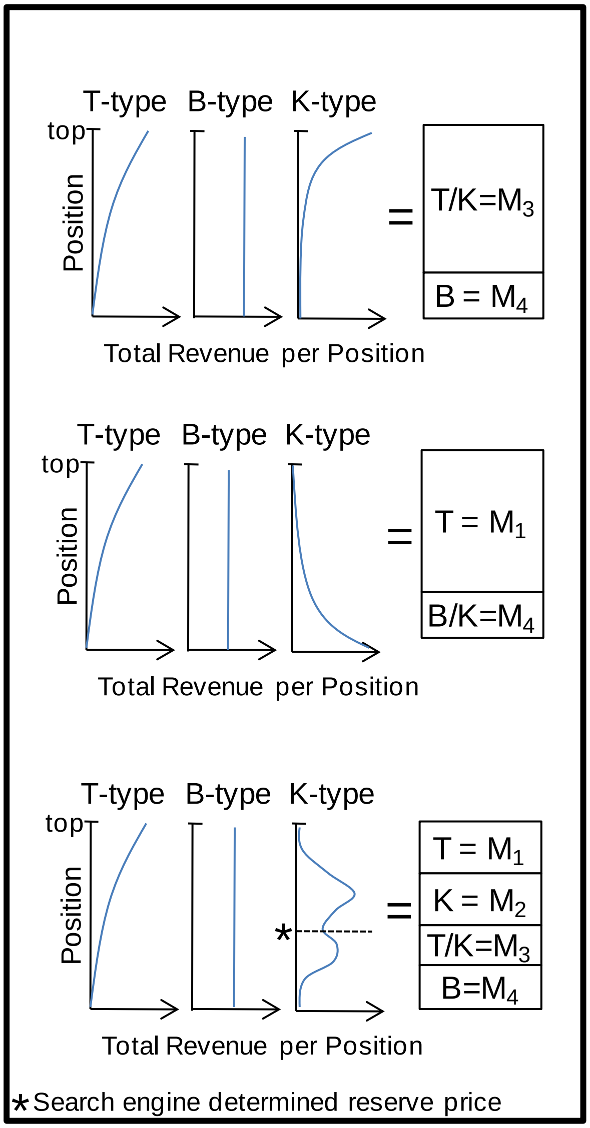

Figure 7 Distribution of Keyword Daily Impressions for Firm 2 on Google (1,120 words)

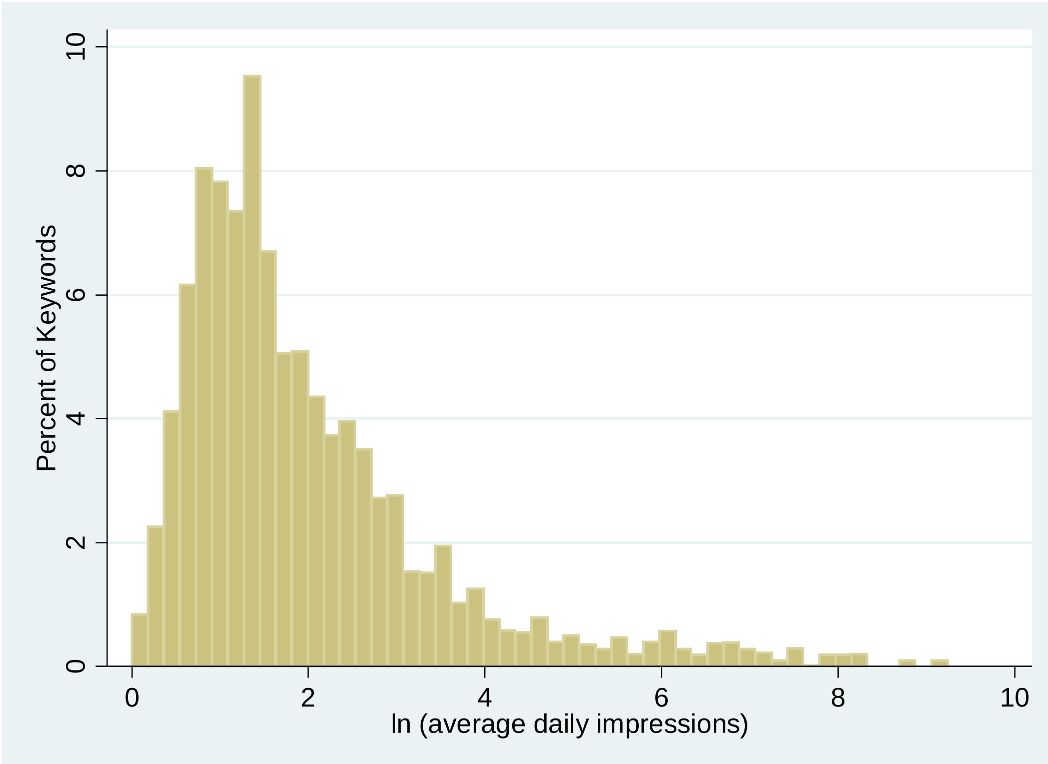

Figure 8 Distribution of Keyword Daily Impressions for Firm 1 across four Platforms

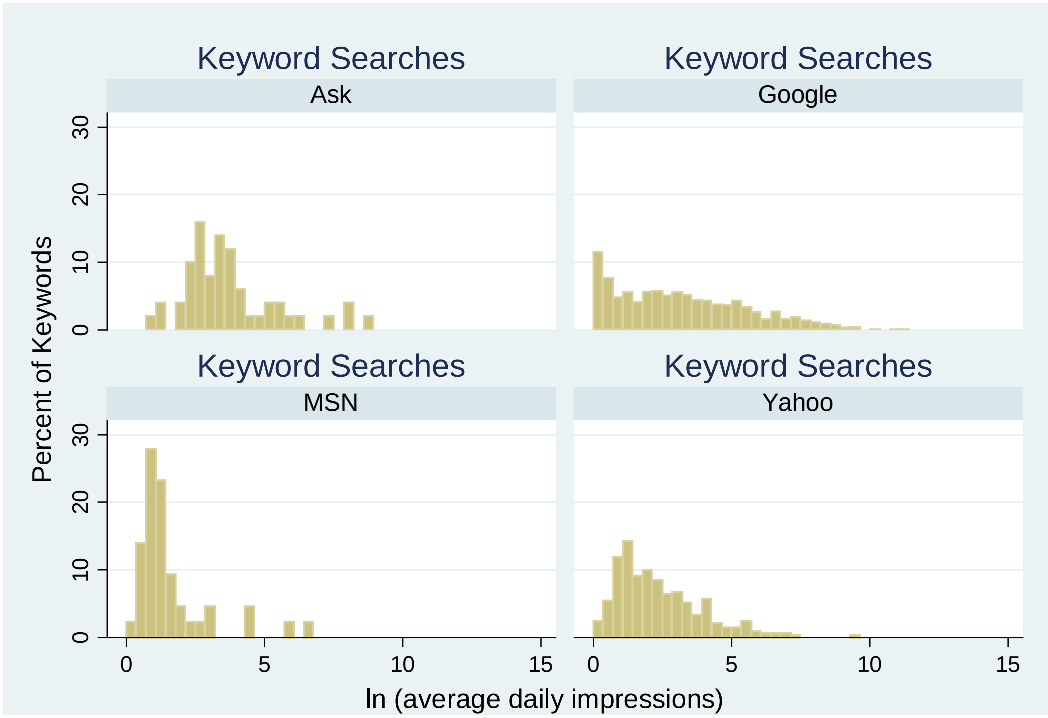

Figure 9 Simulation results for varying Subsidized Bidder Values

In the upper three figures, the middle dark line represents the mean additional search engine revenue under the subsidization program and the outer two lines are one standard deviation from the mean. The vertical line through four represents the expected quality of the eleven base advertisers. In the bottom three graphs, the smooth line with a maximum value of eleven is the position of the subsidized advertiser. The jagged line is the value of the subsidy. The underlying value of the subsidized advertiser (or the bid cap for the subsidized advertiser) differs across the three scenarios getting progressively higher in each case moving right.

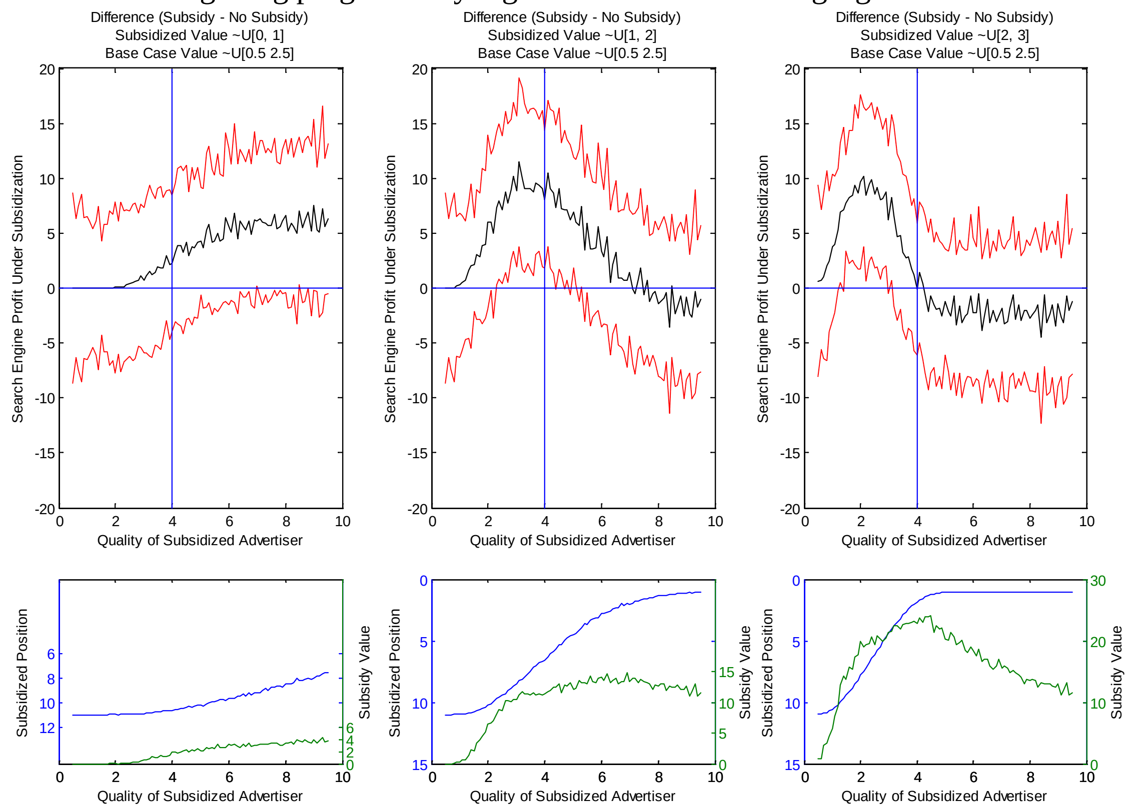

Figure 10 Simulation results for varying CTR curvatures

In the upper three figures, the middle dark line represents the mean additional search engine revenue under the subsidization program and the outer two lines are one standard deviation from the mean. The vertical line through four represents the expected quality of the eleven base advertisers. In the bottom three graphs, the smooth line with a maximum value of eleven is the position of the subsidized advertiser. The jagged line is the value of the subsidy. The slope of the CTR changes over the eleven positions becoming progressively steeper for the three scenarios below.

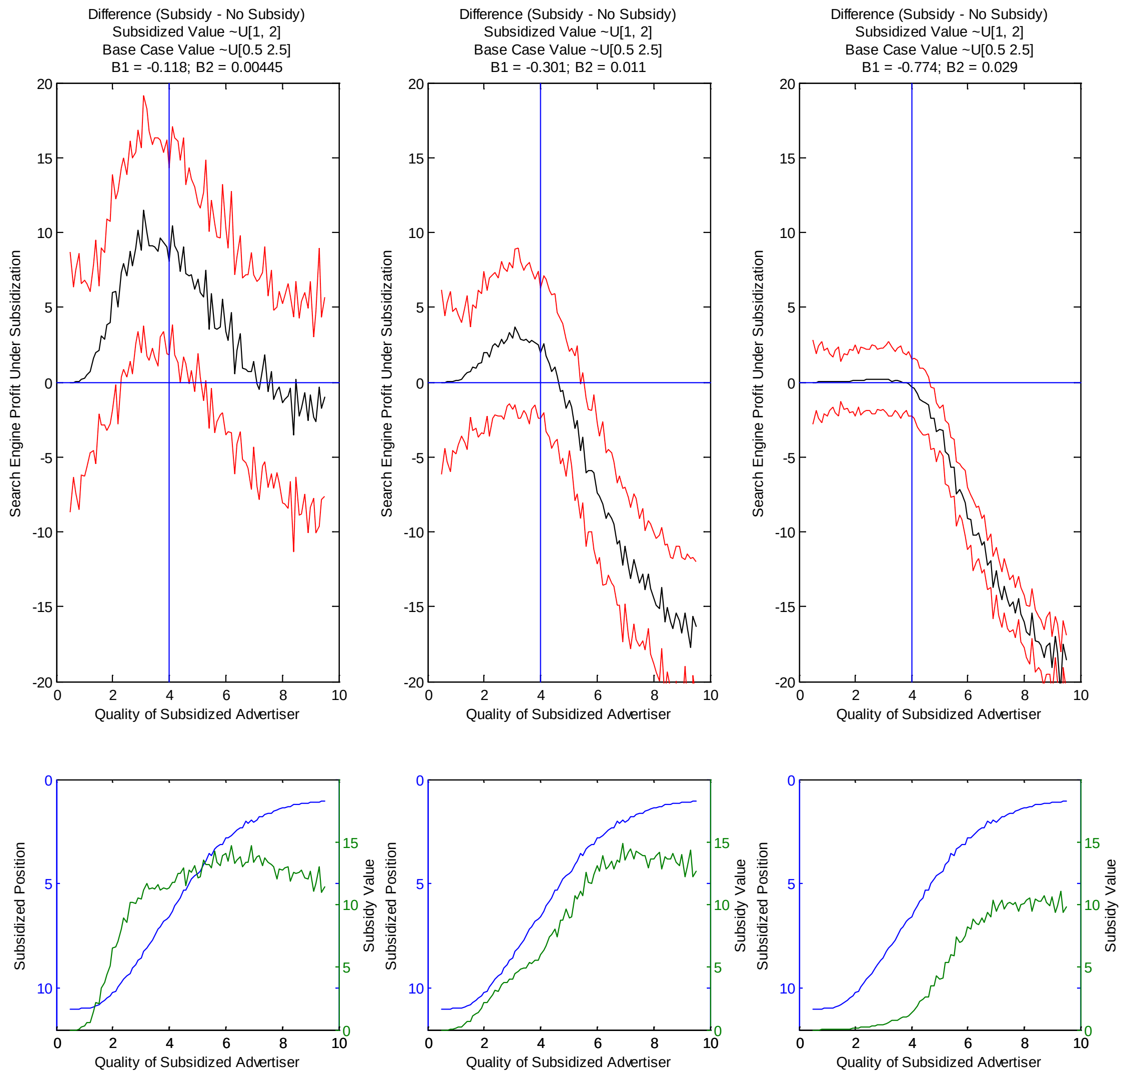

Figure 11 Simulation results comparing Proposition 2 to Value Bidding for Subsidized Advertiser

In the upper three figures, the middle dark line represents the mean additional search engine revenue under the subsidization program and the outer two lines are one standard deviation from the mean. The vertical line through four represents the expected quality of the eleven base advertisers. In the bottom three graphs, the smooth line with a maximum value of eleven is the position of the subsidized advertiser. The jagged line is the value of the subsidy. The two cases compare a case where the subsidized advertiser follows a Proposition 2 bidding strategy (left scenario) against a case where the subsidized advertisers bids their value (right scenario)

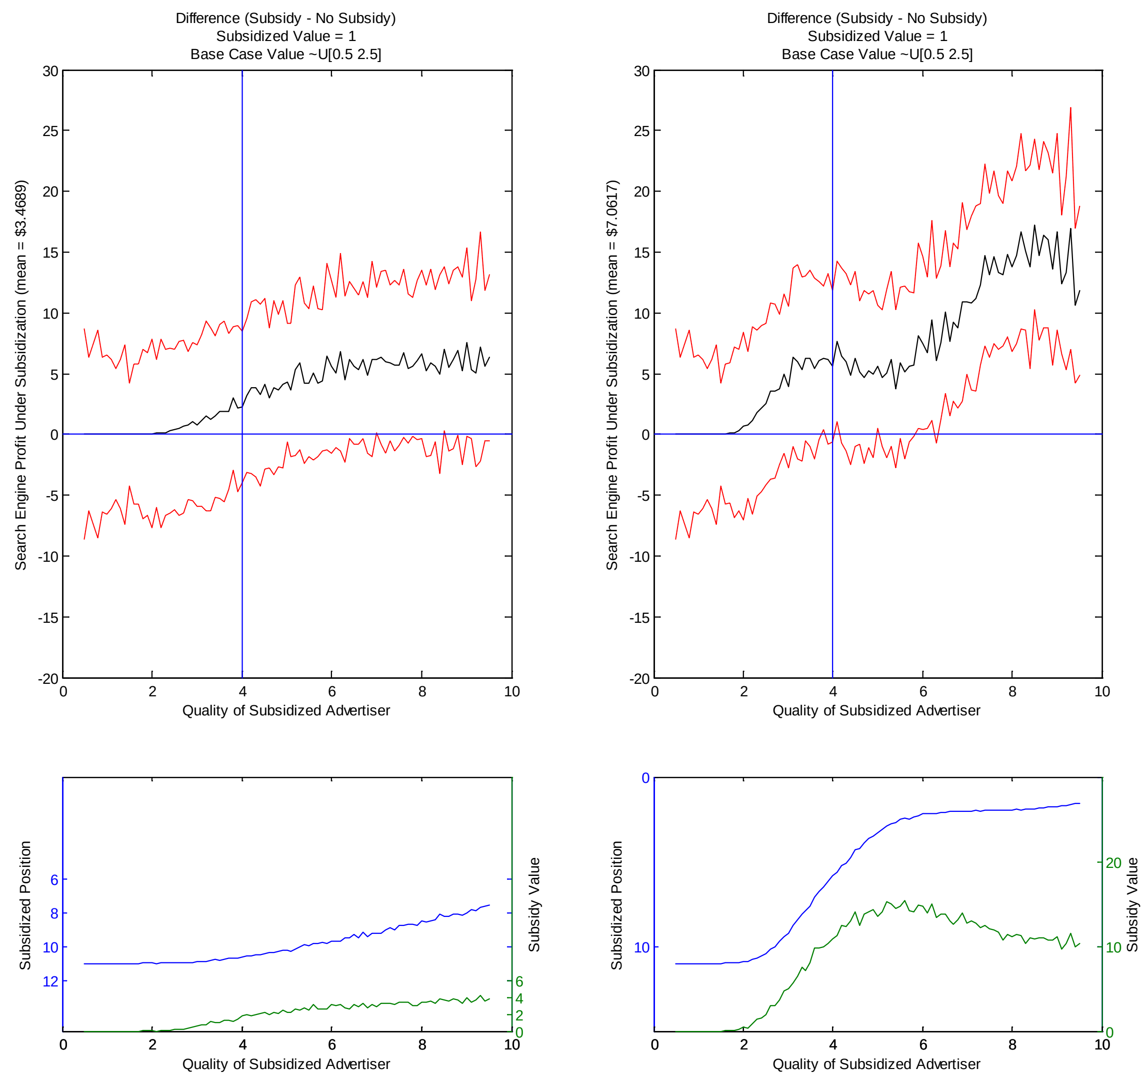

Table 1 Summary Statistics for 3 Firms across Four Platforms

		Firm 1 Ask	Firm 1 MSN	Firm 1 Yahoo	Firm 1 Google	Firm 2 Google	Firm 3 Google
# Days of Data		346	388	708	710	244	450
# Keywords Purchased / Day	Range	1 - 22	5 - 38	21 - 272	33 - 147	818 - 1,256	102 - 809
	Mean	13	28	130	81	1,052	406
	SD	(4)	(5)	(98)	(20)	(85)	(275)
Total Impressions/day	Range	92 - 28,434	25 - 29,418	145 - 57,605	1,579 - 189,206	37,790 - 117,025	6,180 - 133,906
	Mean	6,043	1,472	9,060	19,943	77,611	60,166
	SD	(5,457)	(2,197)	(11,333)	(25,578)	(15,837)	(28,740)
Average Daily Impressions / Keyword	Range	2 - 6,736	1 - 722	1 - 11,959	1 - 64,774	1 - 10,323	0 - 37,880
	Mean	437	51	58	224	74	256
	SD	(1,237)	(149)	(385)	(1,990)	(476)	(2,390)
Avg Daily Position	Range	1 - 3	1 - 24	1 - 67	1 - 170	1 - 88	1 - 82
	Mean	1.0	1.5	3.0	2.5	3.1	4.0
	SD	(0.1)	0.8	(0.1)	(4.4)	(3.1)	(3.7)
Clicks/Day	Range	1 - 480	1 - 78	3 - 654	53 - 2,483	306 - 1,294	46 - 1,708
	Mean	98	32	120	466	931	1,171
	SD	(106)	(18)	(148)	(358)	(223)	(540)
CTR	Mean	6.1%	8.7%	2.1%	6.2%	1.7%	1.7%
	SD	(17.8%)	(19.8%)	(7.4%)	(10.7%)	(8.5%)	(4.6%)
Conversions/day	Range	0 - 108	0 - 134	0 - 237	0 - 951	0 - 19	0 - 4
	Mean	5	4	28	99	6	1
	SD	(12)	(10)	(44)	(90)	(5)	(1)
CR	Mean	5.2%	3.8%	7.6%	11.5%	1.3%	0.6%
	SD	(19.3%)	(16.8%)	(21.7%)	(24.6%)	(9.3%)	(5.7%)
Average Daily Cost / Click	Range	\$0.03 - \$0.50	\$0.05 - \$0.65	\$0.05 - \$0.50	\$0.01 - \$10.00	\$0.02 - \$1.00	\$0.02 - \$1.00
	Mean	\$0.17	\$0.24	\$0.22	\$0.33	\$0.81	\$0.80
	SD	(\$0.10)	(\$0.11)	(\$0.11)	(\$0.17)	(\$0.19)	(\$0.21)

Table 2 Analysis of CTR across keywords testing the role of keyword specificity on CTR (standard errors are in parentheses)

Independent Variable	(1) OLS Avg CTR	(2) OLS ln (Avg CTR)	(3) OLS ln (Avg CTR)
Impressions	-5.765e-06*	-5.804e-05	
	[2.326e-06]	*[5.722e-05]*	
ln (Impressions)			-2.086e-01***
			[2.536e-02]
position	-3.634e-03***	-9.191e-02***	-9.083e-02***
	[1.017e-03]	*[2.192e-02]*	*[2.088e-02]*
position^2	5.218e-05**	7.551e-04*	6.924e-04
	[1.838e-05]	*[3.840e-04]*	*[3.691e-04]*
Const	4.306e-02***	-4.312e+00***	-3.957e+00***
	[3.858e-03]	*[8.209e-02]*	*[9.357e-02]*
Word Dummy Included	no	no	no
n	1,117	1,117	1,117
F	11.10	33.34	51.35
Adj. R-squared	0.0713	0.0230	0.0778

Robust standard errors in brackets (clustered on word); * significant at 5%; ** significant at 1%; *** significant at 0.1%

Table 3 Analysis of the determinants of zero click events

Independent Variable	(1) Probit, mfx Pr (lnCTR > 0)
ln (Impressions)	0.0745***
	[8.84e-4]
position	-0.0075
	[1.27e-3]
position^2	-9.33e-4***
	[1.35e-4]
Base probability	0.1393
n	128640
Likelihood Ratio Chi2	36636
Pseudo R2	0.353

Marginal effects for Probit; * significant at 5%; ** significant at 1%; *** significant at 0.1%

Table 4 Analysis of CTR within keywords testing the role of position on CTR (standard errors are in parentheses)

Independent Variable	(1) Random Effect Tobit ln (CTR)	(2) Fixed Effect, OLS ln (CTR)	(3) Fixed Effect, OLS CTR
Impressions			-9.794e-07*
			[4.719e-07]
ln (Impressions)	1.238***	-0.705***	
	[0.0162]	*[0.0210]*	
position	-0.504***	-0.104***	-9.268e-03***
	[0.0473]	*[0.0175]*	*[1.197e-03]*
position^2	0.0145**	0.00295	5.996e-04***
	[0.00517]	*[0.00235]*	*[8.993e-05]*
Const	-11.84***	-0.0630	4.349e-02***
	[0.101]	*[0.0747]*	*[2.524e-03]*
sigma_u	3.197***	--	--
	[0.0741]		
sigma_e	4.603***	--	--
	[0.0271]		
Word Dummy Included	--	yes	yes
No. Groups	1,117	1,083	1,083
n	90,719	17,799	90,719
F (OLS) / Wald Chi2 (Tobit)	6127	845.8	21.53
Adj. R-squared	--	0.9168	0.1938

Standard errors in brackets for Tobit, Robust standard errors in brackets for OLS (clustered on word); * significant at 5%; ** significant at 1%; *** significant at 0.1%

Table 5 Analysis of CR across keywords testing the role of keyword specificity on CR (standard errors are in parentheses)

Independent Variable	(1) OLS Avg CR	(2) OLS ln (Avg CR)	(3) OLS ln (Avg CR)
Impressions	-2.758e-06**		-5.752e-04***
	[1.051e-06]		*[1.239e-04]*
ln (Impressions)		-4.664e-01***	
		[3.558e-02]	
position	6.703e-04	9.354e-02	-9.521e-02
	[1.547e-03]	*[2.133e-01]*	*[2.656e-01]*
position^2	-1.723e-05	-7.198e-03	1.409e-02
	[2.746e-05]	*[2.790e-02]*	*[3.532e-02]*
Const	1.191e-02**	-2.415e+00***	-3.380e+00***
	[4.596e-03]	*[3.504e-01]*	*[4.290e-01]*
Word Dummy Included	no	no	no
n	1,117	181	181
F	3.282	58.08	7.811
Adj. R-squared	5.692e-04	0.4942	0.1873
Robust standard errors in brackets (clustered on word); * significant at 5%; ** significant at 1%; *** significant at 0.1%			

Table 6 Analysis of CTR within keywords testing the role of position on CTR (t statistics are in parentheses)

Independent Variable	(1) Random Effect Tobit ln (CR)	(2) Random Effect Tobit ln (CR)	(3) Fixed Effect, OLS ln (CR)
ln (Impressions)	0.372***		-0.447***
	[0.0726]		*[0.0565]*
position	0.978*	1.589***	0.156
	[0.454]	*[0.440]*	*[0.103]*
position^2	-0.208**	-0.274***	-0.00144
	[0.0641]	*[0.0633]*	*[0.0154]*
Const	-23.16***	-22.58***	-0.0502
	[0.946]	*[0.922]*	*[0.255]*
sigma_u	0.265***	0.261***	--
	[0.0253]	*[0.0249]*	
sigma_e	8.644***	8.617***	--
	[0.309]	*[0.307]*	
Word Dummy Included	--	--	yes
No. Groups	1,117	1,117	181
n	17,917	17,917	681
F (OLS) / Wald Chi2 (Tobit)	48.26	23.41	24.31
Adj. R-squared	--	--	0.9162
Standard errors in brackets for Tobit, Robust standard errors in brackets for OLS (clustered on word); * significant at 5%; ** significant at 1%; *** significant at 0.1%			

Table 7 Robustness Check on CTR Across Advertisers on Platform 4 (t statistics are in parentheses)

Independent Variable	(1) Firm 1 ln (CTR)	(2) Firm 2 ln (CTR)	(3) Firm 3 ln (CTR)
ln (Impressions)	-0.350***	-0.705***	-0.157***
	(0.032)	*[0.0210]*	*(0.027)*
position	-0.420***	-0.104***	-0.307***
	(0.061)	*[0.0175]*	*(0.034)*
position^2	0.031***	0.00295	0.010***
	(0.006)	*[0.00235]*	*(0.001)*
Const	-0.687***	-0.0630	-1.896***
	(0.126)	*[0.0747]*	*(0.191)*
Word Dummy Included	yes	yes	yes
No. Clusters	584	1,083	566
n	29,337	17,799	2,979
F	83.967	845.8	61.87
Adj. R-squared	0.736	0.9168	0.829

Robust standard errors in brackets (clustered on word); * significant at 5%; ** significant at 1%; *** significant at 0.1%

Table 8 Robustness Check on CR Across Advertisers on Platform 4 (t statistics are in parentheses)

Independent Variable	(1) Firm 1 ln (CR)	(2) Firm 2 ln (CR)	(3) Firm 3 ln (CR)
ln (Impressions)	-0.525***	-0.447***	-0.755***
	(0.033)	*[0.0565]*	*(0.077)*
position	0.105	0.156	0.027
	(0.122)	*[0.103]*	*(0.227)*
position^2	0.005	-0.00144	0.001
	(0.015)	*[0.0154]*	*(0.014)*
Const	1.542***	-0.0502	1.943*
	(0.234)	*[0.255]*	*(0.794)*
Word Dummy Included	yes	yes	yes
No. Clusters	284	181	95
n	9,429	681	235
F	85.698	24.31	36.731
Adj. R-squared	0.487	0.9162	0.935

Robust standard errors in brackets (clustered on word); * significant at 5%; ** significant at 1%; *** significant at 0.1%

Table 9 Robustness Check on CTR Across Platforms (t statistics are in parentheses)

Search Engine Independent Variable	(1) Ask ln (CTR)	(2) MSN ln (CTR)	(3) Yahoo ln (CTR)	(4) Google ln (CTR)
ln (Impressions)	-0.695***	-0.625***	-0.587***	-0.350***
	(0.049)	*(0.071)*	*(0.043)*	*(0.032)*
position	0.635	-0.191*	-0.226***	-0.420***
	(1.227)	*(0.075)*	*(0.063)*	*(0.061)*
position^2	-0.236	0.009	0.015**	0.031***
	(0.350)	*(0.011)*	*(0.005)*	*(0.006)*
Const	-0.520	-0.152	-0.328**	-0.687***
	(0.791)	*(0.226)*	*(0.117)*	*(0.126)*
Word Dummy Included	yes	yes	yes	yes
No. Clusters	48.00	41.00	249.00	584.00
n	1,951	3,713	10,280	29,337
F	91.373	47.682	236.316	83.967
Adj. R-squared	0.789	0.873	0.796	0.736

Robust standard errors in brackets (clustered on word); * significant at 5%; ** significant at 1%; *** significant at 0.1%

Table 10 Robustness Check on CR Across Platforms (t statistics are in parentheses)

Search Engine Independent Variable	(1) Ask ln (CR)	(2) MSN ln (CR)	(3) Yahoo ln (CR)	(4) Google ln (CR)
ln (Impressions)	-0.374**	-0.315***	-0.475***	-0.525***
	(0.113)	*(0.084)*	*(0.052)*	*(0.033)*
position	-18.774	0.335	0.250*	0.105
	(10.251)	*(0.454)*	*(0.098)*	*(0.122)*
position^2	6.858	-0.037	-0.015	0.005
	(3.403)	*(0.091)*	*(0.009)*	*(0.015)*
Const	12.819	0.228	1.220***	1.542***
	(6.721)	*(0.413)*	*(0.185)*	*(0.234)*
Word Dummy Included	yes	yes	yes	yes
No. Clusters	23	28	96	284
n	341	359	2,281	9,429
F	192.076	5.668	32.439	85.698
Adj. R-squared	0.576	0.579	0.424	0.487

Robust standard errors in brackets (clustered on word); * significant at 5%; ** significant at 1%; *** significant at 0.1%

Table 11 Simulation parameters

Fig. #	Session #	11 Base Case Advertisers		Subsidized Advertiser
		Quality	Value [$/click]	Value [$/click]
8	1	~U[3.5 , 4.5]	~U[0.5, 2.5]	~U[0 , 1]
	2	~U[3.5 , 4.5]	~U[0.5, 2.5]	~U[1, 2]
	3	~U[3.5 , 4.5]	~U[0.5, 2.5]	~U[2 , 3]
Fig. #	**Session #**	**Click-through-Rate Slope Coefficients**		**Subsidized Bidder**
		Eq. (13) B_1	**Eq. (13) B_2**	**Value [$/click]**
9	2	-0.118	0.00445	~U[1 , 2]
	4	-0.301	0.011	~U[1 , 2]
	5	-0.774	0.029	~U[1 , 2]
Fig. #	**Session #**	**11 Base Case Advertisers**		**Subsidized Advertiser**
		Quality	**Value [$/click]**	**Actual Bid [$/click]**
10	6	~U[3.5 , 4.5]	~U[0.5, 2.5]	v=1; b=Prop 2
	7	~U[3.5 , 4.5]	~U[0.5, 2.5]	v=1; b=1

X. Appendix

Proof of Lemma 1. 1) The probability of consumer k in group $z \in \{A, B\}$ clicking on the m^{th} advertisement ($CTR_{z,m}$), conditional on seeing the advertisement is: $CTR_{z,m} = \Pr\left(\overline{\eta}_{z,m} > s_{z,k}\right)$ where $\overline{\eta}_{z,m} = E\left(\phi_z\left(F\left(\mu\right)\right)^{(m)} \mid \omega^1 = ... = \omega^{m-1} = 0\right)$ is the expectation of the m^{th} order statistic for the probability that the advertiser will meet the need of the consumer in group z given that the consumer's need has not yet been met (following Athey and Ellison (2008), $\omega^1,...,\omega^m$ are Bernoulli random variables equal to one with the probability that the consumers need was satisfied by the m^{th} advertiser), and $s_{z,k} \sim G_z$ is the search cost for the k^{th} consumer in group z. Using the fact that G_B first order stochastic dominates G_A and that $\phi_A\left(F\left(\mu\right)\right) \geq \phi_B\left(F\left(\mu\right)\right) \forall F\left(\mu\right)$, we have:

$$\Pr\left(s_{A,k} < \overline{\eta}_{A,m}\right) = G_A\left(\overline{\eta}_{A,m}\right) > G_B\left(\overline{\eta}_{A,m}\right) > G_B\left(\overline{\eta}_{B,m}\right) = \Pr\left(s_{B,k} < \overline{\eta}_{B,m}\right)$$

Thus, $CTR_{A,m} > CTR_{B,m}$. C_A consumers from group A (C_B consumers from group B) search high specificity words with probability α_H (β_H) and low specificity words with probability α_L (β_L) where $\alpha_L + \alpha_H = 1$ ($\beta_L + \beta_H = 1$). Using the retention assumption we have that if $\alpha_H^1 > \beta_H^1$ for position 1, it must be the case for all positions (note: $\alpha_H^2 = \alpha_H^1\left(1-\eta_{A,1}\right) \cdot G_A\left(\overline{\eta}_{A,2}\right)$). Given that $\alpha_H^1 > \beta_H^1$ and defining $CTR_m^H = \dfrac{\alpha_H^m C_A CTR_{A,m} + \beta_H^m C_B CTR_{B,m}}{C_A \alpha_H^m + C_B \beta_H^m}$ and $CTR_m^L = \dfrac{\alpha_L^m C_A CTR_{A,m} + \beta_L^m C_B CTR_{B,m}}{C_A \alpha_L^m + C_B \beta_L^m}$ it is easy to show that the click through rate is larger for high specificity words for a given position (i.e., $CTR_m^H \geq CTR_m^L$) and thus the average click through rate for M positions is higher for high specificity words.

2) The probability of clicking on position T for group z is $CTR_z(T) = \prod_{m=1}^{T-1}(1-\eta_{z,m}) \cdot G_z(\bar{\eta}_{z,T})$, where $\prod_{m=1}^{T-1}(1-\eta_{z,m})$ represents the probability that the consumer has not been satisfied by any of the above advertisers and $G_z(\bar{\eta}_{z,T})$ is the probability that the expected utility for the T^{th} position is less than the known search costs. Given that $1 > \eta_{z,m} \geq \eta_{z,m+1} \quad \forall m$ and that $G_z(\bar{\eta}_{z,m})$ is falling for lower advertisers, $CTR_z(m) = CTR_z(m-1) \cdot (1-\eta_{z,m-1}) \cdot \dfrac{G_z(\bar{\eta}_{z,m})}{G_z(\bar{\eta}_{z,m-1})}$. The last two terms are less than one and by extension, $CTR_z(T)$ and the weighted average of the CTR for group A and B is decreasing in T.

3) The average quality of all the advertisers is defined as $\bar{\mu} = \frac{1}{M}\sum_{m=1}^{M}\mu_m$. Since $\eta_{z,m} = \phi_z(F(\mu_m))$, increasing $\bar{\mu}$ results in increasing probability that consumers will have their need met in any position. The retention assumption states that more consumers are dropping out of the search process from negative expected search costs than having their need met. Thus, ceteris paribus, higher quality-weights for the advertisers will result in higher average $\bar{\mu}$, and larger $\bar{\eta}_{z,m}$ which will increase the CTR due to fewer consumers experiencing negative expected search costs in each position. $\otimes$

Proof of Lemma 2. Using the results from Lemma 1: 1) The conversion rate is defined as the probability of the advertiser meeting the consumers need conditional on clicking. For consumer group z and the m^{th} position, this is defined as: $CR_{z,m} = \eta_{z,m} = \phi_z\left(F\left(\mu_m\right)\right)$. Since $\phi_A\left(F\left(\mu_m\right)\right) \geq \phi_B\left(F\left(\mu_m\right)\right)$, it follows that $CR_{A,m} \geq CR_{B,m}$. Following the methodology in Lemma 1, one can see that if $\alpha_H^1 > \beta_H^1$ and there are C_A consumers from group A and C_B consumers from group B, the position specific conversion rate high specificity words is greater than that for low specificity words ($CR_m^H \geq CR_m^L$) and thus the average conversion rate for high specificity words is greater than that for low specificity words: $CR^H \geq CR^L$.

2) Define the overall CR for position 1 as the weighted average of conversion rates from group A and group B (for notational simplicity the H or L subscript has been suppressed, though this is clearly for only one word specificity): $CR_1 = \left(\alpha_1 C_A \eta_{A,1} + \beta_1 C_B \eta_{B,1}\right) \big/ \left(C_A \alpha_1 + C_B \beta_1\right)$. The proportion of group A (group B) consumers clicking on the second advertiser is $\alpha_2 = \alpha_1 C_A \left(1-\eta_{A,1}\right) G_A\left(\bar{\eta}_{A,2}\right)$ ($\beta_2 = \beta_1 C_B \left(1-\eta_{B,1}\right) G_B\left(\bar{\eta}_{B,2}\right)$) where $\bar{\eta}_{z,m}$ is defined in Lemma 1. Thus, $CR_2 > CR_1$ when:

$$
\begin{aligned}
&\alpha_1 \beta_1 C_A C_B \left\{\left(1-\eta_{A,1}\right) G_A\left(\bar{\eta}_{A,2}\right)\left[\eta_{A,2}-\eta_{B,1}\right] - \left(1-\eta_{B,1}\right) G_B\left(\bar{\eta}_{B,2}\right)\left[\eta_{A,1}-\eta_{B,2}\right]\right\} \\
&\qquad -\alpha_1^2 C_A^2 \left(1-\eta_{A,1}\right) G_A\left(\bar{\eta}_{A,2}\right)\left[\eta_{A,1}-\eta_{A,2}\right] - \beta_1^2 C_B^2 \left(1-\eta_{B,1}\right) G_B\left(\bar{\eta}_{B,2}\right)\left[\eta_{B,1}-\eta_{B,2}\right] > 0
\end{aligned}
$$

The second and third terms are both negative and reflect the fact that the $CR_{z,m}$ falls for group z as the position falls (see Figure 5). For $\phi_z\left(F\left(\mu_m\right)\right)$ distributions with lower variance, such that $\left[\phi_z\left(F\left(\mu_m\right)\right)-\phi_z\left(F\left(\mu_{m+1}\right)\right)\right]\to 0$, the magnitude of these terms fall. The magnitude of the first term determines the sign. While $\left[\eta_{A,2}-\eta_{B,1}\right]<\left[\eta_{A,1}-\eta_{B,2}\right]$, the Retention Assumption allows for the possibility that the first term is positive and could potentially balance out the next two terms. Thus, under the conditions discussed here, it is possible that the conversion rate could increase between the first and second positions and by extension, positions further down the list of advertisers. ⊗

Proof of Proposition 2. Proof of this strategy as a truth-dominant equilibrium strategy requires proof of three items: 1) The strategy always maintains the correct ordering of bidder score (i.e., $q_j b_j^* < q_{j-1} b_{j-1}^*$); 2) There is no profitable deviation to a lower a position; and 3) there is no profitable deviation to a higher position. I prove each of these below:

1) The strategy always maintains the correct ordering of bidder score. Given that: $b_j^* = \left[\left(v_j^{j-1}\cdot CTR_{j-1} - v_j^j \cdot CTR_j\right)\cdot q_j + \overline{P}_j\right] \Big/ \left(q_j \cdot CTR_{j-1}\right)$ and $b_{j+1}^* = \overline{P}_j \Big/ \left(q_{j+1}\cdot CTR_j\right)$ where $\overline{P}_j$ is the effective payment in dollars and quality points that must be paid by the j^{th} advertiser. It follows:

$$q_j b_j^* = \frac{\left(v_j^{j-1}\cdot CTR_{j-1} - v_j^j\cdot CTR_j\right)\cdot q_j + \overline{P}_j}{CTR_{j-1}} > \frac{\overline{P}_j}{CTR_j} = q_{j+1} b_{j+1}^*$$

$$\left(v_j^{j-1}\cdot CTR_{j-1} - v_j^j \cdot CTR_j\right)\cdot q_j \cdot CTR_j > \overline{P}_j \cdot \left(CTR_{j-1} - CTR_j\right)$$

Given the assumption that every advertiser experiences a negative externality by the advertiser above her, we know that $v_j^{j-1} = v_j^j + \delta$. Substituting and rearranging yields:

$$v_j^j \cdot q_j \cdot CTR_j + \frac{\left(\delta \cdot CTR_{j-1} \cdot q_j \cdot CTR_j\right)}{\left(CTR_{j-1} - CTR_j\right)} > v_j^j \cdot q_j \cdot CTR_j > \overline{P}_j$$

The second inequality ($v_j^j \cdot q_j \cdot CTR_j > \overline{P}_j$) must always be true for the VCG auction. That is, no advertiser will pay in dollars or quality-weights more than their value and assigned quality weight (or else they could bid zero and pay zero). Similarly, no advertiser will pay exactly as much as they value the slot (a slightly lower value that does not change the advertiser ordering induces the same payment).

2) There is no profitable deviation to a lower position. Assuming the converse, we compare the profit from truth-telling with the profit from deviating lower $j' > j$:

$$\begin{gathered} v_j^j \cdot CTR_j - \overline{P}_{j+1}/q_j < v_j^{j'} \cdot CTR_{j'} - \overline{P}_{j'+1}/q_j \\ \left(v_j^{j'} \cdot CTR_{j'} - v_j^j \cdot CTR_j\right) - \left(\overline{P}_{j'+1}/q_j - \overline{P}_{j+1}/q_j\right) > 0 \end{gathered}$$

$$\sum_{i=j}^{j'}\left(v_j^{i+1} \cdot CTR_{i+1} - v_j^i \cdot CTR_i\right) - \sum_{i=j}^{j'}\left(\overline{P}_{i+1}/q_j - \overline{P}_i/q_j\right) > 0$$

Noting that $\overline{P}_i = \left(v_{i+1}^i \cdot CTR_i - v_{i+1}^{i+1} \cdot CTR_{i+1}\right) \cdot q_{i+1} + \overline{P}_{i+1}$ yields:

$$1/q_j \sum_{i=j}^{j'}\left(v_{i+1}^i \cdot CTR_i - v_{i+1}^{i+1} \cdot CTR_{i+1}\right) \cdot q_{i+1} - \sum_{i=j}^{j'}\left(v_j^i \cdot CTR_i - v_j^{i+1} \cdot CTR_{i+1}\right) > 0$$

There are two types of advertisers who could experience a negative externality from the advertiser above them: T-type and K-type. T-type advertiser's per click values are independent of position which reduces the above inequality to: $\sum_{i=j}^{j'}\left(\tau_{i+1} \cdot q_{i+1} - \tau_j \cdot q_j\right)\left(CTR_i - CTR_{i+1}\right) > 0$ which is clearly false as $\tau_{i+1} \cdot q_{i+1} < \tau_j \cdot q_j$.

K-type advertiser's per click values depend on position, but in the same way, and thus can be represented as: $v_j^{m+1} = \kappa_j \cdot CV_{m+1} \cdot CTR_{m+1}$ for the j^{th} advertiser and m^{th} position. Again this allows the above inequality to be simplified to: $\sum_{i=j}^{j'} \left(\kappa_{i+1} \cdot q_{i+1} - \kappa_j \cdot q_j\right)\left(CV_i \cdot CTR_i - CV_{i+1} \cdot CTR_{i+1}\right) > 0$. If the K-type advertisers want to move up, it must be the case that: $CV_i \cdot CTR_i > CV_{i+1} \cdot CTR_{i+1}$. And based on the actual ordering, $\kappa_{i+1} \cdot q_{i+1} < \kappa_j \cdot q_j$ which also invalidates the above inequality. Thus, there is no profitable deviation below the truthful position. $j+1^{st}$

3) There is no profitable deviation to a higher position. Assuming the converse, we compare the profit from truth-telling with the profit from deviating higher $j' < j$. Taking some liberties with the above notation, we have:

$$\sum_{i=j'}^{j-1} \left(v_j^i \cdot CTR_i - v_j^{i+1} \cdot CTR_{i+1}\right) - 1/q_j \sum_{i=j'}^{j-1} \left(\bar{P}_i - \bar{P}_{i+1}\right) < 0$$

$$q_j \cdot \sum_{i=j'}^{j-1} \left(v_j^i \cdot CTR_i - v_j^{i+1} \cdot CTR_{i+1}\right) - \sum_{i=j'}^{j-1} \left(v_{i+1}^i \cdot CTR_i - v_{i+1}^{i+1} \cdot CTR_{i+1}\right) \cdot q_{i+1} < 0$$

Using an identical methodology to the proof that it is not optimal to deviate to a lower position for T-type or K-type advertisers, one can see that the above inequality must be weakly positive and thus it is not optimal to deviate to a higher position. The only case where the above inequality equals zero is when $j' = j-1$. To see this: $q_j \cdot \left(v_j^{j-1} \cdot CTR_{j-1} - v_j^j \cdot CTR_j\right) - \left(v_j^{j-1} \cdot CTR_{j-1} - v_j^j \cdot CTR_j\right) \cdot q_j = 0$. This proves that any advertiser is exactly indifferent to swapping scores with the advertiser above her (i.e., each advertiser is locally envy free). $\otimes$

Proof to Corollary to Proposition 2. Following the proof laid out by Edelman et al. (2007), the matching literature makes clear that an advertiser optimal equilibrium must exist in the set of stable matching assignments (Roth and Sotomayor, 1990). Here advertiser optimal is described as profit maximization. Let $\overline{P} = \left(\overline{p}_1, \overline{p}_2, \overline{p}_3, ..., \overline{p}_M\right)$ be the set of payments in the bidder optimal stable assignment. Similarly, let $\overline{P}^V = \left(\overline{p}_1^V, \overline{p}_2^V, \overline{p}_3^V, ..., \overline{p}_M^V\right)$ be the set of payments described under Proposition 2. For the bidder optimal assignment, it must be the case that in the lowest position, $\overline{p}_M \geq CTR_M \cdot q_{M+1} \cdot v_{M+1}$ or the first rejected bidder (i.e., the one in the $M+1^{st}$ position would bid for the M^{th} position). In the Proposition 2 strategy, it is the case that $\overline{p}_M^V = CTR_M \cdot q_{M+1} \cdot v_{M+1}$ and thus, $\overline{p}_M = \overline{p}_M^V$. In the bidder optimal stable assignment, it must be the case that $\overline{p}_{M-1} - \overline{p}_M \geq q_M \cdot \left(v_M^{M-1} \cdot CTR_{M-1} - v_M^M \cdot CTR_M\right)$ or the bidder in the M^{th} position would choose to be in the $M-1^{st}$. From this it follows that:

$$\overline{p}_{M-1} \geq q_M \cdot \left(v_M^{M-1} \cdot CTR_{M-1} - v_M^M \cdot CTR_M\right) + \overline{p}_M \geq q_M \cdot \left(v_M^{M-1} \cdot CTR_{M-1} - v_M^M \cdot CTR_M\right) + \overline{p}_M^V = \overline{p}_{M-1}^V$$

Thus, $\overline{p}_{M-1} = \overline{p}_{M-1}^V$ and following the same methodology we have that $\overline{P} = \overline{P}^V$. Thus, $\overline{P}^V$ is the advertiser-optimal stable assignment and the revenue to the search engine from the $\sum_{m=1}^{M} \overline{p}_m^V$ is the lowest from the class of stable assignments. $\otimes$

Proof to Proposition 3. The B-type advertisers in M_4 report scores only slightly above the highest score of the first rejected advertiser. This prevents the B-type M_4

advertisers from moving into M_3 but allows them to avoid rejection. If any advertiser deviates and bids less, they will drop out of the auction and make zero profit. If any advertiser deviates and bids more, they will win a higher position which is strictly lower ranked due to the fact that values are rising for positions lower in the list. The search engine will randomly assign an ordering to the advertisers in this group.

The M_3 advertisers follow a strategy similar to the Proposition 2 except for the first and last advertisers. As proven in Proposition 2, there are no profitable deviations for these advertisers in the middle of M_3. The lowest advertiser (K_3) bids with respect to the first rejected advertiser who would like to move into M_3. That is, this advertiser bids as if the M_4 group did not exist. Deviating and bidding less would cause R_6 to capture the K_3 position. Then all the advertisers in M_4 will bid more than this bid to force the deviating advertiser out of the auction earning that advertiser zero profit. Deviating and bidding more such that this lowest advertiser wins the next higher position would earn the advertiser zero incremental profit (i.e., as in Proposition 2, the lowest advertiser in this group is indifferent to her position or the position above her at the equilibrium price). The highest advertiser in this group ($K_2 + 1$) must bid the amount that makes her indifferent from advancing to the next higher position. Given that the group of advertisers in M_2 are optimally located in M_2, they will always bid higher than this advertiser's bid. Thus, this highest advertiser bids an amount that makes her indifferent to moving into the lowest position in M_1. Deviations are ruled out by the proof to Proposition 2.

The advertisers in M_1 follow a strategy similar to Proposition 2 except for the lowest advertiser. This advertiser bids an amount that would make her indifferent to the position above her except that she bids with respect to the largest negative externality she imposes on those advertisers below her. This may not be the negative externality imposed on advertiser K_2+1. Deviating and bidding less than this would cause an advertiser from M_2 to move into her position reducing the profit of the K_1 advertiser. Deviating and bidding more would be non-optimal, as shown in the proof to Proposition 2.

The advertisers in M_2 all bid the same amount such that the search engine randomly assigns them to position. If they deviate, they either move down a position guaranteeing lower profit or up a position guaranteeing lower profit.

A few special cases. 1) If there is only 1 advertiser in M_1, this advertiser bids according to the K_1 advertiser. As described in Advertiser Scenario 1, this prevents the M_2 advertisers from bidding away all the profit of the sole M_1 advertiser. 2) If there are no M_2 advertisers, then all advertisers follow the M_3 strategy with the top advertiser bidding value. 3) if there is only 1 advertiser in M_2 (or the max value is equally close to both M_1 and M_3), this advertiser will bid the lower amount $\overline{S}_m^* = \max_S \{M_3\} + \varepsilon$. This advertiser is indifferent to any bid between $\min_S \{M_1\} - \varepsilon$ and $\max_S \{M_3\} + \varepsilon$, but will always break ties by bidding the lower value. This is consistent with the behavioral implications of envy free equilibrium where advertisers are indifferent to bidding above the next higher bidder, but remain at the lower bid level. ⊗

Proof to Corollary to Proposition 3. Consider two groups of identically sized T-type advertisers with identical quality-weights where the sequence of scores for the advertisers in groups 1 (T_1) and group 2 (T_2) are such that $T_1 \geq T_2$ in piecewise fashion. Search engine revenue is defined as:

$$R_{SE} = \sum_{m=2}^{M+1} \left[(m-1) \cdot \tau_m \cdot \left(CTR_{m-1} - CTR_m \right) \right]$$

Comparing search engine revenue under the two advertiser scenarios yields:

$$R_{SE}^1 - R_{SE}^2 = \sum_{m=2}^{M+1} \left[(m-1) \cdot \left(\tau_m^1 - \tau_m^2 \right) \cdot \left(CTR_{m-1} - CTR_m \right) \right] > 0$$

Thus, search engine revenue is strictly larger when all scores are larger and all advertisers have the same quality-weights. Similarly, consider any group of T-type advertisers where one or more T-type is replaced with an advertiser from M_4 (i.e., B-type advertisers) or a single satiated advertiser that does not alter the BOO for T-type advertisers. The scores of these new advertisers will be less than the T-type they replaced by nature of the bidding strategy in Proposition 3 (for single satiated bidder see the 3rd special case). Thus it follows that search engine revenue will also be less. ⊗

Combined Bibliography

Agarwal, Nikhil, Susan Athey, and David Yang. 2009. "Skewed Bidding in Pay Per Action Auctions for Online Advertising." American Economic Review Papers and Proceedings, May 2009.

Aggarwal, Gagan, Ashish Goel, and Rajeev Motwani. 2006. "Truthful auctions for pricing search keywords." In ACM Conference on Electronic Commerce (EC), 2006.

Aggarwal, Gagan, Jon Feldman, S. Muthukrishnan, and Martin Pal. 2008. "Sponsored Search Auctions with Markovian Users." In ACM Conference on Electronic Commerce (EC), 2008.

Anderson, Chris. 2006. "The Long Tail: Why the future of Business is Selling Less to more." Hyperion: New York.

Andreoni, James. 1989. "Giving with Impure Altruism: Applications to Charity and Ricardian Equivalence," *Journal of Political Economy* 97: 1147-1458.

Andreoni, James. 1990. "Impure Altruism, and Donations to Public Goods: A Theory of Warm-Glow Giving," *The Economic Journal* 100: 464-477.

Athey, Susan and Glenn Ellison. 2008. "Position Auctions with Consumer Search." Working paper Harvard and MIT.

Athey, Susan, and Philip Haile. 2002. "Identification of Standard Auction Models." *Econometrica* 70 (6): 2107-40.

Ausubel, Larry and Thayer Morrill. 2008. "The Hungarian Auction." Working paper. University of Maryland.

Ausubel, Lawrence M. and Peter Cramton. 1999. "The Optimality of Being Efficient." Working paper University of Maryland.

Ausubel, Lawrence, Peter Cramton, Nathaniel Higgins, Emel Filiz-Ozbay, Erkut Ozbay, and Andrew J Stocking. 2009. "Common-Value Auctions with Liquidity needs: An Experimental Test of a Troubled Assets Reverse Auctions." Draft Paper, AREC University of Maryland College Park (May).

Ayres, Ian and Peter Cramton. 1996. "Deficit reduction through diversity: How affirmative action at the FCC increased auction competition." Stanford Law Review. 48:761–815.

Bagwell, Kyle & Riordan, Michael H, 1991. "High and Declining Prices Signal Product Quality," *American Economic Review* 81(1), 224-39.

Bajari, Patrick, and Ali Hortacsu. 2003. “The Winner's Curse, Reserve Prices, and Endogenous Entry: Empirical Insights from eBay Auctions.” *RAND Journal of Economics* 34 **(**2): 329-55.

Blattberg, Robert C. and Scott A. Neslin. 1989, “Sales Promotion: The Long and the Short of it,” *Marketing Letters* 1(1), 81-97.

Bulow, Jeremy, and Paul Klemperer. 1996. “Auctions Versus Negotiations.” *American Economic Review* 86 (1): 180-94.

Caillaud, Bernard and Bruno Julien. 2003. “Chicken & egg: competition among intermediation service providers.” RAND Journal of Economics. 34(2): 309-328.

Chen, Yan, Xin Li and Jeffrey MacKie-Mason. 2006. Online Fund-raising Mechanisms: A Field Experiment. *Contributions to Economic Analysis & Policy.* 5(2): Article 4.

Connor, Jeffrey, John Ward, and Brett Bryan. 2007. “Exploring The Cost Effectiveness of Land Conservation Auctions and Payment Policies.” American Agricultural Economics Association Annual Meeting. Portland, Oregon.

COP (Chronicle of Philanthropy). 2008. After the Flood: Online gifts still flow to charities, but the stream slowed in 2007. http://philanthropy.com/premium/articles/v20/i17/17000701.htm

Cornes, Richard and Todd Sandler. 1984. “Easy Riders, Joint Production, and Public Goods,” *Economic Journal* 94: 580-598.

Cramton, P., Shoham, Y., and Steinbert, R. 2006. Combinatorial Auctions. MIT Press: Boston, MA.

Cramton, Peter. 2007a. Columbia’s Forward Energy Market. Working Paper, University of Maryland, August.

Cramton, Peter. 2007b. "Comments on the RGGI Market Design." Submitted to RGGI, Inc. by ISO New England and NYISO, November.

Cramton, Peter. 2008. "A Review of the L-Band Auction." Office of Communications, United Kingdom.

Cramton, Peter. 2009. "Market Design: Auctions and Matching." In John Siegfried (ed.), *Better Living Through Economics*, Harvard University Press, forthcoming.

Crewal, Dhruv, R. Krishnan, Julie Baker, and Norm Borin. 1998.. "The Effect of Store Name, Brand Name and Price Discounts on Consumers' Evaluations and Purchase Intentions", *Journal of Retailing* 74(3): 331-352.

Cummings, Ronald, Charles Holt, and Susan Laury. 2004. "Using Laboratory Experiments for Policymaking: An Example from the Georgia Irrigation Reduction Auction." *Journal of Policy Analysis and Management* 23 (2): 341-63.

Davis, Wendy. 2006. "NBC Expands Web Foothold with iVillage Acquisition." Mediapost archive. Available at: http://www.mediapost.com/publications/index.cfm?fuseaction=Articles.san&s=40634&Nid=18894&p=347224

Delaware Agricultural Land Preservation Foundation (DALPF). 2007. Farmland Preservation in Delaware. http://dda.delaware.gov/aglands/lndpres.shtml November.

Dellarocas, Chrysanthos and Siva Viswanathan. 2008. "The Holy Grail of Advertising? Allocative Efficiency and Revenue Implications of "Pay-per-Action" Advertising in Environments with Quality Uncertainty." In ACM Conference on Electronic Commerce (EC), 2008.

Dixit, Avinash K., and Robert S. Pindyck. 1994. *Investment under Uncertainty.* Princeton, New Jersey: Princeton University Press.

Echevarria, John. 2005. "Regulating Versus Paying Land Owners to Protect the Environment." *Journal of Land, Resources, and Environmental Law* 26 (1): 1-46.

Edelman, B., M. Ostrovsky and M. Schwarz (2007), "Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords," *American Economic Review*, 97(1): 242-259

Falk, Armin. 2007. "Gift Exchange in the Field," *Econometrica* 75(5): 1501-1511.

Folkes, Valerie and Rita D. Wheat. 1995. "Consumers' Price Perceptions of Promoted Products," *Journal of Retailing* 71(3): 317-328.

Gale, David, and Lloyd Shapley. 1962. "College admissions and the stability of marriage. American Mathematical Monthly." 69: 9-15.

Gerstner, Eitan. 1985. "Do Higher Prices Signal Higher Quality?" *Journal of Marketing Research* 22(2): 209-215.

Giving USA. 2008: http://nonprofit.about.com/od/trendsissuesstatistics/a/giving2008.htm

Goel, Ashish, Mohammad Mahdian, Hamid Nazerzadeh, and Amin Saberi. 2008. "Advertisement Allocation for Generalized Second Pricing Schemes." Presented at the 2008 ACM Conference on Electronic Commerce, Chicago, IL.

Goldberger, Arthur S. 1968. "The Interpretation and Estimation of Cobb-Douglas Functions." *Econometrica* 35 (3-4): 464-472.

Goldfarb , Avi and Catherine Tucker. 2008. Search Engine Advertising: Pricing Ads to Context. In ACM Conference on Electronic Commerce (EC), 2008.

Gordy, Michael B. 1998. "Computationally Convenient Distributional Assumptions for Common Value Auctions." *Computational Economics* 12 (1): 61-78.

Gupta, Sunil and Lee G. Cooper. 1992. "The Discounting of Discounts and Promotion Thresholds", *Journal of Consumer Research* 19: 410-411.

Haile, Philip A., Han Hong, and Matthew Shum. 2003. "Nonparametric Tests for Common Values in First-Price Sealed Bid Auctions." NBER Working Papers 10105. National Bureau of Economic Research, Inc.

Harbaugh, W. 1998. "What do Donations buy? A model of philanthropy based on prestige and warm-glow," *Journal of Public Economics*: 269-284.

Heidhus, Paul and Botond Köszegi. 2008. "Competition and Price Variation when Consumers Are Loss Averse," *American Economic Review* 98(4): 1245-1268.

Hendricks, Kenneth, Joris Pinkse, and Robert H. Porter. 2003. "Empirical Implications of Equilibrium Bidding in First-Price, Symmetric, Common Value Auctions." *Review of Economic Studies* 70 (1): 115-45.

Higgins, Nathaniel and Michael Roberts. 2008. "Increasing Competition through Quotas." Working paper. University of Maryland.

Hong, Han, and Matthew Shum. 2002. "Increasing Competition and the Winner's Curse: Evidence from Procurement." *Review of Economic Studies* 69: 871-98.

Horowitz, J.K. and Gayer, T. 2007. Market-based Approaches to Environmental Regulation. From Foundations and Trends in Microeconomics. In press.

Kagel, John H., and Levin, Dan. 2002. "Common Value Auctions and the Winner's Curse: Lessons from the Laboratory." In *Common Value Auctions and the Winner's Curse* ed. J. H. Kagel and D. Levin, Princeton, NJ: Princeton University Press.

Karlan, Dean and John A. List. 2008. "Does Price matter in Charitable Giving? Evidence from a Large-Scale Natural Field Experiment." *American Economic Review* 97(5): 1774-1793.

Kempe, David and Mohammad Mahdian. 2008. "A Cascade Model for Externalities in Sponsored Search." In ACM Conference on Electronic Commerce (EC), 2008.

Kirwan, Barrett, Ruben Lubowski, and Michael J. Roberts. 2005. "How Cost-Effective Are Land Retirement Auctions? Estimating the Difference between Payments and Willingness to Accept in the Conservation Reserve Program." *American Journal of Agricultural Economics* 87 (5): 1239-47.

Klaassen, Abbey. 2007. "Search Marketing Fact Pack 2007: a Supplement to Advertising Age." Published by Advertising Age. Available at: http://adage.com/images/random/datacenter/2007/searchfactpack2007.pdf

Köszegi, Botond and Matthew Rabin. 2006. "A Model of Reference-Dependent Preferences." *Quarterly Journal of Economics*, 121(4): 1133-1166.

Krishna, Vijay. 2002. "Action Theory." Academic Press: New York.

Lahaie, Sebastien, David M Pennock, Amin Saberi and Rakesh V. Vohra. 2007. "Sponsored Search Auctions," in Noam Nisan, Tim Roughgarden, Eva Tardos and Vijay V. Vazirani, eds., Algorithmic Game Theory. Cambridge: Cambridge University Press. Pp 699-716.

Landry, Craig, Andreas Lange, John List, Michael Price, and Nicholas Rupp. 2006. "Toward an Understanding of the Economics of Charity: Evidence from a Field Experiment", *Quarterly Journal of Economics* 121: 747-782.

Landry, Craig, Andreas Lange, John List, Michael Price, and Nicholas Rupp. 2008, "Is a Donor in Hand Better Than Two in the Bush? Evidence from a Natural Field Experiment," NBER Working Paper 14319.

Lange, Andreas, John A. List, and Michael Price. 2007. "Using Lotteries to Finance Public Goods: Theory and Experimental Evidence," *International Economic Review*, 48(3): 901-927.

Latacz-Lohmann, Uwe, and Carel Van der Hamsvoort. 1997. "Auctioning Conservation Contracts: A Theoretical Analysis and an Application." *American Journal of Agricultural Economics* 79 (2): 407-18.

Lynch, Lori, and Sabrina Lovell. 2003. "Combining Spatial and Survey Data to Explain Participation in Agricultural Land Preservation Programs." *Land Economics* 79 (2): 259-76.

Maskin, E. (1992), "Auctions and Privatization," in Horst Siebert, ed., *Privatization: Symposium in Honor of Herbert Giersch*. Tubingen: Mohr (Siebeck), 115–136

McAfee, P. and J. McMillan (1987), "Auctions and Bidding," *Journal of Economic Literature*, 25: 699–738

Messer, Kent D., and William Allen III. 2008. “Applying Optimization and the Analytic Hierarchy Process to Enhance Agricultural Preservation Strategies in the State of Delaware.” Department of Economics: Working Paper No. 2008-18. University of Delaware, Newark, Delaware.

Milgrom, Paul R., and Robert J. Weber. 1982. “A Theory of Auctions and Competitive Bidding.” *Econometrica* 50 (5): 1089-1122.

Milgrom,. Paul and John Roberts. 1986. “Price and Advertising Signals of Product Quality,” *Journal of Political Economy* 94(4): 796-821.

Morgan, John and Martin Sefton. (2000). “Funding Public Goods with Lotteries: Experimental Evidence”. *Review of Economic Studies*, 67: 785-810.

Morgan, John. (2000). “Financing Public Goods by Means of Lotteries”. *Review of Economic Studies*, 67: 761-784.

Myerson, Roger B. (1981), “Optimal Auction Design,” Mathematics of Operations Research, 6, 58-73.

Paarsch, Harry. 1992. “Deciding Between the Common and Private Value Paradigms in Empirical Models of Auctions.” *Journal of Econometrics* 51 (1-2): 191-215.

Pinkse, Joris, and Guofu Tan. 2005. “The Affiliation Effect in First-Price Auctions.” *Econometrica* 73 (1): 263-277.

Plantinga, Andrew, Ruben Lubowski, and Robert Stavins. 2002. “The Effects of Potential Land Development on Agricultural Land Prices.” *Journal of Urban Economics* 52 (3): 561-81.

Quigg, Laura. 1993. “Empirical Testing of Real Option-Pricing Models.” *Journal of Finance* 48 (2): 621-40.

Reichelderfer, Katherine, and William G. Boggess. 1988. “Government Decision Making and Program Performance: The Case of the Conservation Reserve Program.” *American Journal of Agricultural Economics* 70 (1): 1-11.

Rochet, Jean-Charles and Jean Tirole. 2006. “Two-sided markets: a progress report.” RAND Journal of Economics. 37(3): 645-667.

Roth, A.E. 2002. The Economist as Engineer: Game Theory, Experimentation, and Computation as Tools for Design Economics. Econometrica. 70(4): 1341-1378.

Roth, A.E. 2007 "Efficient Kidney Exchange: Coincidence of Wants in Markets with Compatibility-Based Preferences," American Economic Review. 97(3):828-851

Roth, A.E. and Peranson, E. 1999. The Redesign of the Matching Market for American Physicians : Some engineering Aspects of Economic Design. American Economic Review. 89(4) : 748-780.

Roth, Alvin and Marilda A Oliveira Sotomayor. 1990. “Two-sided Matching: A study in Game-Theoretic Modeling and Analysis.” Econometric Society Monographs. Cambridge Press.

Roughgarden, Tim and Mukund Sundararajan. 2007. “Is Efficiency Expensive?” Working paper Presented at the 2007 ACM Conference on Electronic Commerce. Banff, Canada.

Salamon, Lester M., Megan A. Haddock, S. Wojciech Sokolowski, and Helen S. Tice. 2007. “Measuring Civil Society and Volunteering: Initial Findings from Implementation of the UN Handbook on Nonprofit Institutions.” Working Paper No. 23. Baltimore: Johns Hopkins Center for Civil Society Studies.

Sánchez-Azofeifa, Arturo, Alexander Pfaff, Juan Robalino, and Judson Boomhower. 2007. “Costa Rican Payments for Ecological Services Program: Intention, Implementation, and Impact.” *Conservation Biology* 21 (5): 1165-1173.

Schilizzi, Steven, and Uwe Latacz-Lohman. 2007. “Assessing the Performance of Conservation Auctions: An Experimental Study.” *Land Economics* 83 (4): 497-515.

Shang, Jen and Rachel Croson. 2006. “The Impact of Social Comparisons on Nonprofit Fundraising,” Research in Experimental Economics, Vol 11: 143-156.

Smith, Rodney. 1995. “The Conservation Reserve Program as a Least-Cost Land Retirement Mechanism.” *American Journal of Agricultural Economics* **77** (1): 93-105.

Stokes, Richard. 2008. “Search Engine Advertiser Update – Q208.” AdGooroo Advertising insight.

Stoneham, Gary, Vivek Chaudhri, Arthur Ha, and Loris Strappazzon. 2003. “Auctions for conservation contracts: an empirical examination of Victoria’s Bush Tender trial.” *Australian Journal of Agricultural and Resource Economics* 47 (4): 477-500.

Thaler, Richard. 1985, “Mental Accounting and Consumer Choice,” Marketing Science 4(3): 199- 214.

Varian, Hal. 2007. “Position Auctions.” International Journal of Industrial Organization. 25: 1163-1178.

Varian, Hal. 2009. "Online Ad Auctions." American Economic Review. 99(2): 430-434.

Vesterlund, Lise. 2003. "The Informational Value of Sequential Fundraising." *Journal of Public Economics*. 87(3-4): 627-57.

Vukina, Tomislav, X. Zheng, Michele Marra, and Armando Levy. Forthcoming. "Do Farmers Value the Environment? Evidence from the Conservation Reserve Program Auction." *International Journal of Industrial Organization* in press.