

The Design of Online Advertising Markets

Benjamin Edelman—Draft

Online advertising is big business, already reaching some \$40 billion per year. For advertisers, online advertising offers the triple promises of reaching just the right consumers, at fair prices, with robust measurement of the effects of online ad campaigns. For web site publishers, advertising offers an opportunity way to make money from their sites—an important consideration since few consumers appear willing to provide money payments for the sites and services they use. For users, in principle online ads can be affirmatively useful in finding new products or suppliers; in any event, online ads are often easily overlooked (compared with, say, the temporal and auditory interruption of television advertisements).

Because the market for online advertising is both new and fast-changing, participants experiment with all manner of variations. Should an advertiser's payment reflect the number of times an ad was shown, the number of times it was clicked, the number of sales that resulted, or the dollar value of those sales? Should ads be text, images, video, or something else entirely? Should measurement be performed by an ad network, an advertiser, or some intermediary? Market participants have chosen all these options at various points, and prevailing views have changed repeatedly. Online advertising therefore presents a natural environment in which to evaluate alternatives for these and other design choices.

In this piece, I review the basics of online advertising, then turn to design decisions as to ad pricing, measurement, incentives, and fraud.

Defining the Product: Payment Structure and Purchasing Incentives

The fundamental product in online advertising markets is a lead—a customer who might make a purchase from a given advertiser, or otherwise respond to an advertiser's offer. An advertiser typically prefers to reach customers especially likely to buy its product or service, and observable customer

characteristics indicate varying degrees of interest in an advertiser's offer. For example, consider an advertiser selling motorcycles. The advertiser could attempt to reach consumers in particular demographic groups (say, males age 18 to 25), site browsing (reading a motorcycle enthusiast web site), or search terms (searching for "motorcycle deals"). The advertiser's forecast of the likelihood of the user making a purchase would inform the advertiser's willingness to pay to present its offer to that consumer.

Meanwhile, from the perspective of an online publisher, operating a web site or other online resource, advertising is typically an ancillary component to be integrated with, or at least juxtaposed against, a larger offering. If a publisher offers a search function, the publisher could show text ads related to users' search requests. Alternatively, a publisher could place "banner ads" (typically graphical images in industry-standard sizes) adjacent to articles on its site. In principle a publisher could even make individual words on its site into ads leading to advertisers' sites—though with questions about who selects which words link where, and whether and how consumers know they're clicking on ads. A publisher's resource is typically space on its site or service. If the advertiser's site presents too many ads, consumers may reach an unfavorable view of the site.

Online advertising can be measured and sold along any of several metrics. An advertiser could pay a fee each time its ad is shown—a "cost per impression" placement, often known as CPM ("cost per mille" being the price for 1,000 impressions). Alternatively, an advertiser could pay when its ad is clicked—"cost per click" (CPC). Or an advertiser could pay only when a user clicks and subsequently makes a purchase—"cost per action" (CPA). An advertiser could even offer payment proportional to the *amount* of the user's purchase, *ad valorem*, or differing payment scales could apply to the advertiser's various products. In expectation, advertisers and publishers might be indifferent among these payment metrics; with a known click rate, conversion rate, or order size, an advertiser and publisher could agree to use any of these metrics, and fees would be equal in expectation. That said, the metrics have

importantly different implications for parties' incentives, moral hazard, and fraud, as discussed in subsequent sections.

Industry norms associate certain payment metrics with certain advertising formats. Historically, display advertisements were typically priced per impression—a natural approach from the perspective of a publisher who does not know which ads will attract many clicks, and who wants to be able to predict site revenues. That said, selling ads per impression influences participants' behavior: A CPM advertiser wants to attract as many click as possible, even from customers who may ultimately be minimally interested in the advertiser's offer; perhaps some of those marginal customers can be convinced to buy the advertiser's product. CPM advertisers thus have a clear incentive to present banners with overstated claims of relevance or urgency, like those shown in the inset at right. Facing this onslaught of low-value ads, consumers seem to develop "banner blindness": As of 2009, practitioners at iMedia Connection report that for every 1,000 display ads shown to consumers, just 0.2 to 0.3 are clicked. (Stern 2010) Meanwhile, some display ad services have begun to price ads differently—selling ad placements on a per-click basis, encouraging advertisers to design offers that consumers choose to activate.

[Insert deceptive banner exemplars about here]

Ads on search engines typically follow a CPC model—not charging advertisers for their ads to be shown, but charging substantial fees when a user clicks an ad (for some keywords, as much as \$20 or more per click). With CPC pricing, an advertiser seeks to attract only customers reasonably likely to purchase its product or otherwise offer the advertiser some benefit; attracting clicks from uninterested customers means unnecessary marketing expense. On the most favorable view, CPC pricing also invites users to click ads: Knowing that an advertiser was willing to pay to reach users searching for a given keyword, a user may expect that the advertiser's offer will match to the user's request. Indeed, as Overture (later acquired by Yahoo) began offering pay-per-click ads, founder Bill Gross specifically

boasted of the benefits of “us[ing] money as a filter” of which sites to show in search listings. (Hansell 2001)

Affiliate link systems typically follow a conversion-contingent CPA payment model—either paying a publisher only when a user signs up (e.g. a \$15 commission for referring a customer to Netflix), or in proportion to the dollar value of the user’s purchase (e.g. a 6% commission on the user’s purchase from Amazon). To date, few affiliate marketing programs have been willing to pay affiliates for impressions or clicks—seemingly on the view that little-known affiliates, without meaningful vetting or supervision, would have an overwhelming tendency to fake impressions and/or clicks, whereas actual sales are viewed as harder to fake. That said, as detailed in *Advertising Fraud* below, even conversion-based payment methods suffer strategic behavior that inflates advertisers’ costs.

	Display ads	Search ads	Affiliate/links
Pay per impression	standard	unusual	
Pay per click	also used	standard	unusual, but newly-implemented at eBay
Pay per action	used for some campaigns	brief experiment at Google	standard

Search Ads

Auctions & pricing

Historically, online ads were typically sold through posted prices, rate sheets, and person-to-person negotiations—much like ads in print, television, and radio. But auctions and auction-like mechanisms have proven particularly well-suited to online advertising for at least three reasons. For one, there are a multiplicity of items to be sold, including a large number of sites showing ads, as well as multiple ad placements on each such site. With so many items to sell, it would be difficult to announce a price for each or to negotiate the particulars of a placement. Furthermore, values change as market conditions change—making efforts to post or negotiate prices all the more difficult. Finally, the automated online delivery of advertisements seems to complement an automated online sales process;

interconnected systems and servers can accept offers for a given placement, select an ad to be shown, show the ad to the corresponding users, and charge the advertiser accordingly.

The use of auctions and auction-like mechanisms presents a variety of questions of auction design. Should an advertiser be charged its own bid (“first price”) or something less (“second price” or otherwise)? How often may bids be updated, and should an advertiser pay a fee for adjusting its bid? Should advertisers be able to see the bids of competitors seeking the same placements, see how many competitors are interested, or see something less? Should an auction impose a reserve price, below which ads are rejected, or is any payment better than nothing? Ad platforms have reached differing conclusions on all these questions.

The world of sponsored search advertising began in 1998 with pay-per-click text ads developed by Goto.com, later renamed Overture and purchased by Yahoo. Advertisers were suspicious of Overture’s novel approach to pricing: With early fees often reaching \$1 per click or even more, advertisers were concerned that competitors might click their ads, or Overture might charge for clicks that did not actually occur. To attempt to address these concerns, Overture showed advertisers the ads and bids of all competitors—confirming that an advertiser was not alone in its use of Overture’s offering, and that others were onboard too. Showing all bids also helped advertisers adjust to the unfamiliar auction format: With competitors’ bids visible for inspection, an advertiser could better assess the tradeoff between bidding higher (getting more clicks) and bidding lower (reduced price, but lesser exposure).

When a user clicked an advertiser’s ad, Overture charged each advertiser the amount it had bid—a first-price auction. This system was intuitive: If an advertiser reported being willing to pay \$0.70 for a click, why would Overture charge the advertiser anything less? But the game was infinitely repeated, with bid updates allowed frequently. (Initially, it seems, update frequency was limited only by the effort required to log into Overture’s systems and make adjustments. Later, a rule limited updates

to one every fifteen minutes, and a widespread automatic bidding agent adjusted bids every fifteen minutes.) In the Overture first-price auction, each advertiser had an incentive to lower its bid to the minimum increment (\$0.01) above the next-highest advertiser—letting the advertiser retain the same position but pay a reduced price. Edelman and Ostrovsky (2007) show that the resulting instability led to an inefficient allocation of placements—often misordering advertisers, putting a lower-value advertiser above one who valued clicks more highly, and thereby destroying surplus. The resulting instability also reduced total revenue of the mechanism by at least 7% (a conservative bound reflecting the difficulty of estimating advertisers' valuations from historic bid data).

In 2002, Google began to use a mechanism with some characteristics of a second-price auction. Rather than paying its own bid, an advertiser would pay an amount linked to the bid of the next-highest advertiser—reducing the incentive to adjust bids continuously. Moreover, Google adjusted each bid by the estimated likelihood of a user clicking the corresponding ad, thereby selecting the ad with largest expected revenue to Google.

Edelman, Ostrovsky, and Schwarz (2007) (EOS) study this multi-unit second-price mechanism, calling it “generalized second price” or “GSP.” EOS shows that GSP has no dominant-strategy equilibrium, and truth-telling is not an equilibrium. However, the corresponding generalized English auction has a unique equilibrium, and that equilibrium is an ex post equilibrium with bidders' strategies independent of their beliefs about others' types. Moreover, Cary et al. (2007) show that a reasonable myopic bidding strategy converges to the equilibrium identified in EOS. Further overviews of sponsored search appear in Feldman et al. (2008), Lahaie et al. (2007), Liu et al. (2008) and Yao and Mela (2009).

Ad platforms continue to use reserve prices to rule out bids they view as undesirably low. In simulations, Edelman and Schwarz (2010) assess the revenue consequences of an optimally-chosen reserve price. Which bidders face the largest cost increases from a rising reserve price? Edelman and

Schwarz show that, for all advertisers who do not drop out as reserve price increases, the increased reserve price yields an identical dollar-for-dollar increase in total payment.

Most ad platforms offer additional targeting of their ads based on at least the user’s geographic region (“geotargeting”) and day/time (“dayparting”). These targeting functions are typically operated on a binary basis: Either a user request matches the restrictions, and hence is eligible to show the advertiser’s ad, or the advertiser specifies that its ad may not be shown. Microsoft adCenter offers further supplemental targeting based on user self-reports of age and gender at other Microsoft properties (such as Hotmail, MSN, and Windows Live). If a user matches the demographic characteristics an advertiser specifies, the advertiser may opt to increase its bid, potentially increasing its ranking relative to competitors. Thus, in adCenter, an advertiser’s bid is not just a price, keyword, and vector of match conditions, but also additional price adjustments paired with demographic conditions. Despite the additional targeting possible under demographic bid adjustment, uptake of demographic targeting seems to be limited so far.

Transparency of Pricing and Ranking

Ad platforms limit the information available to advertisers, relative to the early list of all advertisers and bids that Overture initially provided. For example, Google has never shown advertisers the bids or identities of competing bidders. Instead, Google provides advertisers a *traffic estimator* tool: An advertiser enters a possible bid, and Google reports the estimated number of clicks it would provide per day as well as the advertiser’s estimated average position in ad listings.

Ranking of advertisers sometimes raises concerns about favoritism or penalties—concerns that tend to focus on Google, given that company’s large market share (discussed further in “Multihoming, Competition...” below). Google states that it ranks advertisers according to both their bids and Google’s various assessments of site characteristics. (Varian 2009) If one site enjoys a more favorable assessment, it can obtain a more prominent placement at considerably lower expense. On one view, a

search engine is a private party entitled to show whatever links it sees fit, in whatever order and prominence it chooses. But some advertisers allege that Google singles out up-and-coming competitors for particularly unfavorable treatment, typically by demanding unreasonably large prices in order to show ads from those would-be competitors.

TradeComet styles itself as a vertical search engine, specifically a potential way for businesses to find the suppliers they require and a potential competitor to Google to the extent that companies use TradeComet, not Google, to find desired resources. In ongoing antitrust litigation, TradeComet claims Google violated the Sherman Act by increasing TradeComet's prices from \$0.05-\$0.10 per click to \$5-\$10 per click, overnight. TradeComet says Google attributed the price increases to "landing page quality." But TradeComet claims Google itself had recently awarded TradeComet "site of the week," and says recognition from others was similarly positive—countering any suggestion that TradeComet was undesirable or low quality.

Foundem (of Bracknell, UK) offers similar allegations. Foundem says Google dramatically reduced the prominence of organic (ordinary, unpaid) links to Foundem's site, which dropped overnight from top-10 to #100 or lower for certain terms in Google, while remaining highly ranked (as high as #1) in Yahoo and Bing searches for the same terms. Foundem also bought advertising placements from Google, but found it faced dramatically increased prices: Foundem says prices spiked from around 5p to £5, a 100-fold increase, overnight. (Foundem 2009)

Foundem attributes its penalties to Google manually cutting Foundem's "quality score" rating. (Foundem 2009) But quality scores are not available to the public, so it is difficult to confirm these allegations except through litigation and discovery. That said, Google policies indicate penalties for sites with "little or no original content." (Google 2009) On one view, many such sites are traps that seek to ensnare users within mazes of advertisements. Yet the Google search service itself offers little or no original content; instead, Google links to content hosted elsewhere. Indeed, a lack of original content is

distinctively characteristic of vertical search sites, like TradeComet and Foundem, that seek to compete with Google. Would-be competitors therefore take these Google exclusions and penalties to be an improper barrier to competition. This aspect of ranking remains a subject of dispute.

Matching Display Ads to Users and Sites

In the realm of search ads, a user's search request provides most of the information required to select suitable advertisements. But in the area of display ads, a user's requests provide significantly less context. Knowing what web page a user is viewing often does not reveal commercial offers the user would be likely to accept.

Matching is made more difficult by preferences of both advertisers and consumers. From an advertiser's perspective, sites are importantly different. Users at some sites may be significantly more likely to accept an advertiser's solicitation. Furthermore, some sites may be viewed as inappropriate for an advertiser's offer, e.g. due to inclusion of offensive, adult, or copyright-infringing material.

Meanwhile, from a user's perspective, ads are also importantly different. Some ads offer products or services users actually want or need. But other ads resort to trickery or deception to attract consumers' attention. See e.g. Edelman 2009a.

To date, most display ad platforms offer relatively limited methods of matching advertisers with sites and users. Typically, platforms begin by excluding placements where the advertiser or site has rejected a counterpart specifically or through various characteristics viewed as undesirable. (For example, Edelman 2009a explores the various characteristics by which Yahoo Right Media allows sites to exclude ads that are deceptive, distracting, or otherwise undesirable.) Then, platforms sort ads from highest expected revenue to lowest, conditioning on the advertiser and/or ad, the site, and sometimes an interaction between advertiser/ad and site. As a user browses a site, the site's chosen ad platform typically begins by showing the ad with highest expected revenue, then onwards to ads expected to yield lower revenue. If the site uses multiple ad platforms, the site typically attempts to pass each ad

placement to the platform expected to pay the most for that placement, and some third-party services aim to assist sites in this effort.

To date, matching rules have been binary, without any notion of pricing or compensating differentials. For example, a publisher typically must either allow or deny a category of ads (e.g. ads that play sounds, deceptive ads), but the publisher ordinarily cannot demand an increased fee for showing disfavored ads. Similarly, an advertiser must either allow or reject placement of its ads on a given site, but ad platforms typically give the advertiser no clear mechanism to demand a lower price for ads placed on a site viewed as less desirable. In this context, it may seem natural to introduce prices for disfavored placements: Prices would increase complexity, but would also reduce deadweight loss by facilitating placements that current rules discard. That said, added payment for placement of unethical or otherwise undesirable ads may be viewed as repugnant. (Roth 2007) Moreover, such payments might have legal consequences. So far, sites have not faced legal liability for showing deceptive ads.¹ However, if sites could be shown to charge extra for deceptive ads, they would reveal themselves to both aware of the problem and, in an important sense, culpable.²

Ad Networks and Syndication

Advertisers typically prefer to buy online advertising in large blocks from known partners, so intermediaries organize multiple sites into *networks*. By helping advertisers buy placements on small to mid-sized sites, networks help fund such sites—fueling the diversity of web content. Furthermore, networks reduce transaction costs by aggregating many small sites into a single line item that an advertiser can buy with a single contract and a single payment.

¹ Google was sued for deceptive advertisements, namely ads for “free” ringtones that actually carried substantial monthly charges. However, Google presented a successful defense grounded in the Communications Decency Act §230, which prohibits treating the provider of an interactive computer service (here, Google) as the publisher of information provided by an independent entity (here, the advertiser who submitted the deceptive ad). See *Goddard v. Google*. N.D.Ca. 2008, Case No. 5:2008cv02738.

² Compare *Gucci America, Inc. v. Frontline Processing Corp.*, 09 Civ. 6925 (HB), holding credit card processing companies liable for contributory trademark infringement when they charged extra fees to “high risk” sellers selling counterfeit merchandise.

Information Disclosure in Ad Networks

Ad networks present a clear question of disclosure of lists of participating sites. When buying online ad placements, advertisers naturally want to know where their ads appear. Some ad networks provide lists of their member sites. But most networks see a strategic downside in providing advertisers with site lists: With a site list, an advertiser could bypass the network—contacting member sites and negotiating direct placements that deny the network compensation for its effort in suggesting the placement. Citing this concern, many networks use a “blind” information structure—selling placements on a bundle of sites, without telling advertisers which sites are included.

It is unclear whether the risk of bypass merits keeping network site lists confidential. For large advertisers running ads on just a few sites, bypassing a network might offer financial benefits sufficient to justify the effort. But such bypasses would require sacrificing networks’ serving, tracking, contracting, and payment functions, which would require considerable effort to replace. Moreover, if networks’ sole concern is bypass, they have other tools at their disposal. For example, affiliate network LinkShare requires that an advertiser commit not to run any affiliate marketing activities through competing networks, while affiliate network Commission Junction prohibits an advertiser from bypassing the network for any relationship initially brokered by the network.

An alternative explanation for blind networks comes from member sites that advertisers would not approve, if an advertiser’s approval were requested. By keeping its member list confidential, a network can avoid advertiser scrutiny of its sites—thereby letting the network include sites of mixed desirability.

Pricing in Ad Networks

When a network bundles placements on multiple web sites, billed to advertisers without itemization as among included sites, a network must allocate payments within the network. If some sites will be paid more than others, what measure will allocate value among sites? Will each impression or each click yield an equal payment? Or are some impressions or clicks more valuable than others?

If a network pays the same price for each impression or each click, it risk underpaying sites where traffic is particularly valuable, i.e. particularly likely to lead to purchases or other desired outcomes. If top sites leave, the network would retain only average to below-average sites—an unraveling that would reduce advertisers' valuation of the network's traffic. Indeed, there is some evidence for such unraveling: The web's top publishers often sell much of their advertising space directly to advertisers; they report that networks offer lower revenue than direct relationships. At the same time, a few premium networks (e.g. Quigo) promise special care in selecting member sites, yielding higher revenues to sites that make the cut.

In response, networks recognize a need to offer different payments to different publishers. For example, Google describes its "smart pricing" as follows: "If our data shows that a click is less likely to turn into business results (e.g. [an] online sale...), we may reduce the price [an advertiser] pay[s] for that click." (Google 2004) That said, it is difficult for networks to condition payments on user behavior at advertisers' sites. For one, such conditioning requires combining multiple data sources, including outcomes of many advertisers' ads on many publishers' sites. Furthermore, advertisers often view post-click outcomes as confidential, lest networks know advertisers' results and raise prices when results are favorable.

Intermediary Counts and the Prospect of Disintermediation

Early intuition on online markets anticipated disintermediation—that online markets would let contracting parties eliminate brokers and middle-men. (Bambury 1998) But disintermediation has not been the dominant outcome in online advertising, especially not in display advertising. Rather, a drop in transaction costs makes it easier and more common to build lengthy relationships not often seen in other contexts. For example, an advertiser's ad might pass through half a dozen brokers en route to a publisher's site—each taking a cut as small as a few percent, such that even these complex relationships may leave adequate surplus to the ultimate buyer and seller. On the other hand, lengthy relationships

reduce accountability when an ad ends up misplaced (e.g. Edelman 2007), while also slowing ad-load times and sometimes yielding lost impressions or error messages.

Measurement, Mismeasurement, and Fraud

Measuring the Value of an Ad Placement

To optimize ad spending, advertisers typically seek to assess the value of an advertisement placement—then buy more of the placements that seem to offer the largest value relative to cost. Simple as it sounds, such measurement often proves difficult. In principle, advertisers can measure the ratio of impressions or clicks to sales, including the gross profit from such sales, thereby calculating the benefit attributable to a given placement. But this measurement calls for an online sales process—a poor fit for those selling through offline channels. Offline sellers can attempt to collect data on ad effectiveness by collecting leads online, e.g. asking would-be car-buyers to submit their contact information for referral to a local dealer. But customers often decline to submit such leads, adding bias or requiring ad hoc manual adjustments.

Most measurement assumes that, without an advertising expenditure, subsequent sales would not have occurred. For example, if a user clicks an ad and then makes a purchase, a typical measurement concludes that the ad “caused” the purchase—asserting that, without the ad, the purchase would not have occurred; and asserting that other advertising efforts did nothing to cause the sale. This assumption tends to reduce the apparent value of display ads, which often offer delayed benefits to advertisers. For example, a user might see an ad on a news site, then begin to consider a possible future purchase of the advertised product. (Fowler 2007) This assumption similarly discounts the value of offline advertising (TV, print, billboard, etc.), which is also hard to tie to specific purchases. Conversely, this assumption tends to increase the apparent value of search ads, which often immediately precede a purchase. For example, a user looking to buy a laptop might search for “laptop” or even “Thinkpad x300 laptop” right before completing the purchase. Yet the user running such a

search might well buy the specified laptop even if no ad were presented. Thus, from the perspective of the advertiser, the relevant comparison may be “pay for the ad and sell the product” versus “don’t pay for the ad, yet still sell the product.” In that context, paying for the ad may be a poor value. Yet most measurement systems nonetheless assume that online advertising directly and solely causes subsequent purchases.

Moreover, all manner of spyware, adware, typosquatting sites, and other interlopers can claim to have referred customers who actually requested a merchant specifically and by name, as detailed in the subsequent section.

Advertising Fraud

Delivered purely electronically, through computer systems without in-person checks or well-developed verifications, online advertising can suffer from a variety of frauds, unjustified charges, and other complications. For example, a site can load many banner ads in invisible windows—then charge advertisers for the resulting “impressions” even though users could not see the resulting ads. (Edelman 2006b) Through spyware or adware installed on users’ computers, or through certain JavaScript within ordinary web pages, sites can fake or simulate pay-per-click ad clicks—imposing costs on advertisers who pay by the click. (Edelman 2006a) Rogue affiliate marketers can invisibly invoke affiliate links so that they receive commission on subsequent purchases from the corresponding affiliate advertisers—managing to overcharge even advertisers who chose what was believed to be a fraud-proof or low-fraud channel. (Edelman 2007)

For most advertisers, measurement efforts are the best defense against improper charges. But sophisticated fraudsters can manipulate the figures most advertisers measure. For example, if a display advertiser is wary of placements with a high ratio of impressions to clicks (too few clicks relative to the number of impressions), the fraudster can fake both clicks and impressions. If a pay-per-click advertiser is measuring the ratio of sales to clicks, the fraudster can design its systems to target users already likely

to make a purchase from a given advertiser—for example, by faking clicks when the user is already at the advertiser’s site. (Edelman 2006a) The resulting costs can be substantial. For example, June 2010 indictments allege that Brian Dunning and Shawn Hogan stole some \$5 million and \$15 million from eBay through eBay’s Commission Junction affiliate program; the indictments allege that these affiliates actually sent eBay worthless traffic, yet eBay’s measurement systems deemed them eBay’s two largest and most productive affiliates.

Incentives, both between firms and within firms, sometimes dull efforts to uncover advertising fraud. Most large advertisers buy online ads through agencies which are paid on a commission basis. Catching fraud would reduce the measured spending and hence reduce the agency’s commission—requiring an investment of time and effort yielding *lower* payment to the agency. Networks’ incentives are also attenuated: In the long run, advertisers will distrust networks with a reputation for fraud. But in the short run, networks can increase revenue by retaining unsavory placements that increase volume.

Furthermore, within-firm incentives invite advertisers’ staff to ignore or tolerate fraud. For many online advertising buyers, the prestige of a position comes in part from the size of the budget under management—limiting the incentive to exclude fraudulent spending which would reduce budgets. Furthermore, some buyers face leveraged incentives that sharply discourage clean-up. For example, some companies pay their affiliate managers based on year-over-year growth of the programs they operate. Ejecting fraud would cut spending and yield a disproportionate drop in compensation. Finally, where a fraud has successfully defrauded a buyer, that person may hesitate to come forward, on the view that admitting the problem would reveal a personal failure. In a forthcoming draft, Edelman attempts to measure some of these effects based on variation in staff and network compensation schemes.

Ad Placement Arbitrage

Industry participants often use the term “arbitrage” to describe buying ad placements from a low-cost source, then showing ads through a network that offers higher payments. If both placements are equally desirable, such arbitrage might equalize prices across markets, improving efficiency and increasing surplus. But if a seller offers lower prices because its placements suffer lower quality, resale of these resources to a high-paying buyer practice does not constitute “arbitrage” as economists use the term. Rather, such resale is more likely to constitute misrepresentation of a low-quality resource as a high-quality resource. (Edelman 2005)

Multihoming, Competition, and Barriers

58% of search advertisers use only Google, not Yahoo or Microsoft adCenter. (Ashlagi et al. 2010) This is arguably puzzling because, from an advertiser’s perspective, competing search ad services seem to be at least orthogonal if not complementary: Some users favor one search engine, while others use another, and an advertiser who foregoes a top ad platform fails to reach those users who rely on the corresponding search engine. Prices cannot explain this puzzle because Google has both the most advertisers and the highest prices. (Edelman 2009b)

Instead, it seems advertisers distinctively favor Google because, despite Google’s higher prices, Google offers access to more users and to a larger volume of searches. Ashlagi (2010) shows that the advertisers that use all of Google, Yahoo, and Microsoft are significantly larger than the advertisers that use just one or two of these platforms. Ashlagi attributes this difference to transaction costs: That advertisers using multiple platforms face extra costs, including signup costs, copying and updating ads, monitoring performance, and adjusting bids.

In principle, advertisers could use automated software systems to copy their campaigns from one ad platform to another—avoiding most costs of transferring and updating ads. Each ad platform provides an application programming interface (API) to let advertisers and tool-makers update and

check ads and bids. That said, Google’s API contract limits how advertisers may use this API—prohibiting tools that copy ads from one platform to another. Edelman (2008) argues that these restrictions are an improper barrier to advertisers seeking to use smaller ad platforms.

Consumer Protection: Disclosures and Deception

Online advertising raises all manner of consumer protection issues. For one, must advertisements be labeled as such? The FTC has called for “clear and conspicuous disclosures” that listings are advertisements, particularly in contexts such as search advertising, where users may reasonably fail to recognize advertisements as such. Through late 2010, most search engines used terms like “sponsored links” to label their advertisements. In an online experiment, Edelman and Gilchrist (2010) show that the more detailed label “paid advertisement” reduces users’ clicks on ads by 25% to 33%, with drops particularly pronounced for users for with low income, low education, and little online experience. Meanwhile, Edelman (2010) critiques Google’s newest advertisement label, “Ads”, pointing out that the new label is so tiny that it substantially fits within an “o” of Google, among other shortfalls.

Some pay-per-click advertisements seek to deceive or defraud users – for example, promising “free ringtones” when in fact the service carries a substantial charge. Edelman (2006c) documents all manner of such schemes. However, *Goddard v. Google*, 640 F. Supp 2d 1193 (N.D.Cal. 2009), finds Google not responsible for deceptive ads it sold and presented to users – even when Google charged for each advertisement, was aware of the untrue statements, and even encouraged the deception through, e.g., a “keyword suggestion tool” that suggested describing ringtones as “free.” This decision reflects an interpretation of the Communications Decency Act §230, which instructs that a web site must not “be treated as the publisher or speaker of any information provided by” anyone else.

Open Questions

The contracts, institutions, and norms of online advertising continue to evolve. Innovation continues even on questions as fundamental as when an advertiser pays—with new payment metrics based on “view-throughs” (a CPM-CPA hybrid requiring an impression followed by a conversion) and “impressions per connection” (a CPM-CPC hybrid charging advertisers for impressions, but providing bonus impressions if click-throughs are sufficiently frequent). These metrics alter incentives for advertisers and publishers, addressing some of the problems with standard approaches but simultaneously creating new concerns. With so much in flux, there remains ample opportunity to identify new metrics that better satisfy participants’ requirements.

Meanwhile, Google’s market share continues to grow—exceeding 90% of search volume in scores of countries. Does Google’s auction mechanism fully determine prices? Or can Google use its increasing popularity to increase prices to advertisers and otherwise enjoy its market power?

The structure of online advertising markets is closely linked to issues of general public concern. For example, despite the rise of online advertising, newspapers receive significantly less revenue for readers reached online rather than in print. But newspapers serve important public functions, so online advertising shortfalls prompts a need to revisit the future of journalism. Funding newspapers through online ads is particularly challenging because it is often unclear what ads are most suitable: What advertiser seeks a placement adjacent to news of war, election, or natural disaster? Some ads could be selected based on a user’s prior activities rather than current browsing, but this approach calls for collecting and retaining ever more information about users’ activities. Balancing these concerns—while satisfying users, advertisers, publishers, and various intermediaries—presents challenging questions at the intersection of economics, computer science, law, and public policy.

References

- Ashlagi, I., B. Edelman, and H.S. Lee (2010). "Competing Ad Auctions: Multi-homing and Participation Costs." Harvard Business School Working Paper No. 10-055.
- Bambury, P. (1998). "A Taxonomy of Internet Commerce." *First Monday*. Volume 3, Number 10 – October 5, 1998.
- Cary, M., A. Das, B. Edelman, I. Giotis, K. Heimerl, A. Karlin, C. Mathieu, and M. Schwarz (2007). "Greedy Bidding Strategies for Keyword Auctions." *Proceedings of the 9th ACM Conference on Electronic Commerce*.
- Commission Junction v. Thunderwood Holdings, Inc., et al.* Sup.Ct. Ca., Case No. 30-2008 00101025. Second Amended Complaint for Damages, Exhibit B. May 16, 2008.
- Edelman, B. (2010). "A Closer Look at Google's Advertisement Labels." November 10, 2010. <http://www.benedelman.org/adlabeling/google-nov2010.html> .
- Edelman, B. (2009a). "False and Deceptive Display Ads at Yahoo's Right Media." January 14, 2009. <http://www.benedelman.org/rightmedia-deception/> .
- Edelman, B. (2009b). "Towards a Bill of Rights for Online Advertisers." September 21, 2009. <http://www.benedelman.org/advertisersrights/> .
- Edelman, B. (2008). "PPC Platform Competition and Google's 'May Not Copy' Restriction." June 27, 2008. <http://www.benedelman.org/news/062708-1.html> .
- Edelman, B. (2007). "Spyware Still Cheating Merchants and Legitimate Affiliates." May 21, 2007. <http://www.benedelman.org/news/052107-1.html> .
- Edelman, B. (2006a). "The Spyware - Click-Fraud Connection -- and Yahoo's Role Revisited." April 4, 2006. <http://www.benedelman.org/news/040406-1.html> .
- Edelman, B. (2006b). "Banner Farms in the Crosshairs." June 12, 2006. <http://www.benedelman.org/news/061206-1.html> .
- Edelman, B. (2006c). "False and Deceptive Pay-Per-Click Ads." October 9, 2006. <http://www.benedelman.org/ppc-scams/> .
- Edelman, B. (2005). "How Yahoo Funds Spyware." August 31, 2005. <http://www.benedelman.org/news/083105-1.html> .
- Edelman, B. and D. Gilchrist (2010). "'Sponsored Links' or 'Advertisements'?: Measuring Labeling Alternatives in Internet Search Engines." HBS Working Paper No. 11-048.
- Edelman, B. and M. Ostrovsky (2007). "Strategic Bidder Behavior in Sponsored Search Auctions. With Michael Ostrovsky." *Decision Support Systems*. v.43(1), February 2007, pp.192-198.
- 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*. v. 97(1), March 2007, pp.242-259.

- Edelman, B. and M. Schwarz (2010). "Optimal Auction Design and Equilibrium Selection in Sponsored Search Auctions." *American Economic Review*. v. 100(2), May 2010.
- Feldman, J. and Muthukrishnan, S. (2008) "Algorithmic Methods for Sponsored Search Advertising." Chapter 4 in Liu, Z. and Xia, C. H., eds. *Performance Modeling and Engineering*, Springer, New York, NY.
- Foundem. "Search Neutrality—Foundem's Google Story." <http://www.searchneutrality.org/foundem-google-story> . August 18, 2009.
- Fowler, J. "Overlap's Impact on Reach, Frequency and Conversions." Atlas Institute. Mimeo. June 5, 2007.
- Google. "Little or No Original Content." Google Webmaster Central. <http://www.google.com/support/webmasters/bin/answer.py?hl=en&answer=66361> . Last Updated: June 10, 2009.
- Google. "Google AdWords News Archive—April 2004." https://www.google.com/intl/en_us/adwords/select/news/sa_mar04.html . April, 2004.
- Hansell, S. 2001. "Paid Placement is Catching On in Web Searches." *New York Times*. June 4, 2001.
- Lahaie, S., Pennock, D. M., Saberi, A. and Vohra, R. V. (2007) "Sponsored Search Auctions." Chapter 28 in Nisan, N., Roughgarden, T., Tardos, E. and Vazirani, V. eds., *Algorithmic Game Theory*, Cambridge University Press, New York, NY.
- Liu, D., Chen, J. and Whinston A. B. (2008) "Current Issues in Keyword Auctions." SSRN Working Paper 1008496.
- Roth, A. "Repugnance as a Constraint on Markets." *Journal of Economic Perspectives*. V.21(3), Summer 2007, pp.37-58.
- Stern, A. "8 Ways to Improve Your Click-Through Rate." iMedia Connection. February 1, 2010. <http://www.imediaconnection.com/content/25781.asp> .
- Varian, H. "Introduction to the Google Ad Auction." <http://www.youtube.com/watch?v=K7l0a2PVhPQ> . March 11, 2009.
- Yao, S. and Mela, C. F. (2009) "A Dynamic Model of Sponsored Search Advertising." SSRN Working Paper 1285775.

Figures

Deceptive banner ads overstate the urgency of clicking through.

