The Pathologies of Online Display Advertising Marketplaces

BENJAMIN EDELMAN
Harvard Business School

Display advertising marketplaces place "banner" ads on all manner of popular sites. While these services are widely used, they suffer significant challenges, including weak user response and low accountability for both advertisers and web site publishers. I survey a few major challenges, flagging possible areas for future research.

Categories and Subject Descriptors: J.1 [**Computer Applications**]: Administrative Data Processing—*Marketing*

General Terms: Measurement, Legal aspects, Security

Additional Key Words and Phrases: Online advertising, banner ads, incentives, intermediaries, ad networks, fraud, impression, pay-per-impression, pay-per-click, conversion

Much has been written about online *search* advertising, where Google enjoys 90%+ market share in numerous countries and a dominant position almost everywhere. That's a remarkable situation, a sea change from Google's late entry in 1998, and a subject of substantial concern for the users, advertisers, and publishers who depend on Google for information, leads, and advertising payments.

But I'd like to look at a different troubled online advertising market – one that suffers quite a different set of problems. *Display* advertising systems place ads – typically, rectangular "banners" – on the majority of popular web sites. Though display ads are widespread, they are also troubled – ignored by many users, priced at levels that fail to adequately support many online publishing businesses, and riddled with complex relationships that hinder accountability.

1. USER PERCEPTION OF BANNERS: AMBIVALENT OR WORSE

A first problem for both display advertisers and web site publishers is "banner blindness" – users' tendency to ignore banner ads. Indeed, on many sites, click-through rates have dropped as low as 0.10%.

Thanks to a series of efforts by the Internet Advertising Bureau (IAB), most ads follow standard sizes, letting a single set of ads fit in many sites. This standardization brings important benefits – reducing transaction costs by letting advertisers develop a single set of ads for use across the web, and helping users distinguish a site's ads from its editorial content. Yet standardized ads also carry a cost: Users see so many standard-sized ads that eyes automatically overlook ad-sized messages.

Display ads are also ignored because they are perceived to be annoying or irrelevant. Surely many are; we've all seen dancing monkeys and flashing exclamation points. Importantly, every irrelevant ad poses a negative externality – discouraging users from even looking at other ads. To some extent a publisher can internalize

Author's address: bedelman@hbs.edu

this externality by excluding irrelevant ads. But there's a tension: Every excluded ad means foregone revenue, and the benefit in increased user engagement is delayed, uncertain, and hard to measure. Plus, irrelevant ads on one site discourage users from examining ads elsewhere – a classic externality the first site would be unlikely to consider. (For an analogous result in sponsored search, see Ghosh and Mahdian [2008].)

Fraudulent ads present a further externality. Some banners trick or defraud users – promising "free ringtones" that actually carry a charge, claiming a user has "already won" when in fact there's no prize at all, or sneaking browser add-ons into users' computers. (Edelman [2006a] presents specific examples.) These scams further undermine display advertising – giving users good reason to ignore banner ads.

What would a solution entail? Some sites experiment with ads in unusual sizes and formats – "takeovers" that cover or push aside a site's editorial content, or "interstitials" that give a user an ad when the user requests an article. These make sites a bit more money – but at the expense of user satisfaction and convenience. If there's a sensible middle ground, few sites have found it yet.

2. ACCOUNTABILITY IN ONLINE ADVERTISING MARKETPLACES

Display advertisers, publishers, and networks report concerns about all manner of ad placements. Where will an advertiser's offers appear? What ads will a publisher's site show? Often no one quite knows. Meanwhile, as to costs, all parties seem to be unhappy: Advertisers think prices are too high, while publishers argue revenues are unsustainably low.

I believe these problems stem, at least in part, from a common cause: Multiple intermediaries insert themselves between advertisers and publishers. In the best case, an advertiser buys a placement through a network which has contracted to show ads on a publisher's site. But in practice, relationships are often considerably more complex. Network A could buy from B, from C, from D, from the publisher – an "intermediary chain."

Long intermediary chains add cost: Each network takes a fee as it repackages the traffic it brokers. Often networks defend their fees as well-earned: They bought traffic more cheaply than their clients had otherwise found it. But this view assumes all traffic carries equal quality. In fact some traffic is far better than others – showing ads to users more inclined to purchase, showing ads in more prominent areas on publishers' sites. Conversely, other traffic can be strikingly damaged: "below the fold" ads that a user only sees after scrolling, ads that a publisher shows only for an instant (before substituting an alternative), ads in popups, spyware or adware, and invisible ads that users can't see because publishers find ways to cover ads with other material or load ads in tiny subwindows. (Edelman [2006] offers examples.) Advertisers typically struggle to detect these shenanigans: For example, banner ad advertising is often designed to spur long-term brand recognition rather than immediate purchases, while other banners seek to encourage offline purchases that can't easily be tied to a specific ad placement. But with numerous possibilities for malfeasance, an ad placement that seems to be a great deal may prove to be too good to be true. So convoluted traffic brokering often adds less value than

advertisers are led to believe.

Furthermore, long chains hinder accountability. Even when an advertiser finds an unwanted source of traffic, the advertiser often cannot block it: The advertiser could refuse to buy the traffic A buys from B from C from D. But what of the traffic A buys from E from F from D? Through a different chain, the advertiser still ends up with the unwanted D traffic. (See, e.g., Edelman [2007].) Meanwhile publishers have a corresponding challenge: A publisher may flag ads deemed offensive or deceptive, but when multiple intermediaries serve ads onto the publisher's site, the ad is likely to reoccur despite the publisher's best efforts to exclude it.

What would a solution entail? One alternative cuts out many intermediaries in favor of direct relationships. In particular, an advertiser could buy placements only from networks that deal directly from known web sites, and not from networks that buy from other networks and brokers. A few top-tier networks use this approach, but costs are high: networks must negotiate relationships with each individual site. Furthermore, this approach tends to disfavor small sites – harming both advertisers (which lose access to legitimate but lesser-known sites) and the overall diversity of content on the web (to the detriment of small sites and users alike).

If one ad network emerged as leader, it might naturally simplify relationships: All advertisers would buy from that network, all publishers would sell to that network, and chain-type relationships would disappear. But in fact the display advertising ecosystem is awash in small networks – hundreds or more thanks to low barriers to entry plus advertisers' and sites' willingness to work through multiple networks. Moreover, whatever the benefits for simplifying relationships and increasing accountability, increased market concentration would invite higher prices to advertisers, lower payments to publishers, and all manner of unsavory non-price terms.

Ad exchanges present a possible way to simplify relationships without the costs of market concentration. If many networks traded traffic on a common exchange, each could sell traffic to any other with low transaction costs thanks to standardized technical configuration and automatic record-keeping. Furthermore, lengthy chains could disappear: If network A wants traffic from network D, A could contact D directly – without any need to pass through other intermediaries.

Despite these benefits, ad exchanges have struggled to offer the expected benefits. For one, traffic quality has proven hard to verify. For example, fraudsters found they could sell worthless traffic through multiple exchange participants simultaneously. A savvy buyer might find one of the exchange participants brokering this traffic, but finding them all was no easier on an exchange than in prior chain-style relationships. Meanwhile, exchanges seem to encourage heated price competition: Without clear proof that some traffic is worth more than others, buyers favor the lowest-price suppliers – even as those suppliers often turn out to engage in all manner of shenanigans. Meanwhile, as in non-exchange environments, publishers continue to risk showing ads they consider offensive, deceptive, or otherwise undesirable. An active exchange can attempt to block these problems, but diverse participants – from different regions, backgrounds, and sectors – have strikingly different visions of what ads are offensive and even what traffic is fraudulent. (Edelman [2008] offers problems drawn from Right Media's experience.) The most challenging disputes

would require considerable management attention – from an exchange's perspective, all cost with no direct revenue. So exchanges have tended to stay out of these disputes. But in the long run, failure to address such problems taints an exchange's reputation with networks, top advertisers, and top sites. (Edelman [2009])

Ultimately, I believe the answer lies in accountability. If a network places an ad into a location that's fraudulent – say, invisible ads that users can't see at all – the network should refund the advertiser for the resulting wasted spending. Crucially, the network shouldn't just have to pay the few advertisers who happen to figure out what happened; it should be required to make full refunds to everyone affected. The underlying incentives are simple: Only fully internalizing the costs of bad traffic forces a network to properly exclude that traffic from its offerings. Yet current industry norms are understood to let networks keep the proceeds of fraud – charging uninformed advertisers as if nothing was wrong. Remarkably, standard form contracts are often in accord: In fine print, advertisers purportedly agree to pay even for traffic networks know to be invalid. I question whether these contract provisions are enforceable, but in any event sophisticated advertisers should demand better.

3. ROOM FOR RESEARCH

Beyond the areas flagged above, I see abundant room for further work. For one, pricing structures remain strikingly variable: While many advertisers pay each time their banners are viewed, others pay for clicks or even for purchases. Some pay in complicated hybrids – for example, paying if a user sees an ad and later makes a purchase, even if the user never clicks the ad; or submitting both per-impression and per-click bids (as in Goel and Munagala [2009]). These payment systems offer interestingly divergent incentive characteristics – changing parties' behavior and transferring risk. Further work could usefully explore the alternatives, identify optimal uses for each, and perhaps develop new approaches beyond those already in use.

The display advertising ecosystem would also be well-served by further research on apportionment of credit between the various firms that facilitate online advertising (endogenizing the prices in Chatterjee et al. [2003]). If a user sees three banner ads over the course of a week, then finally clicks a fourth, and makes a purchase a week after that, how should the advertiser assess the effectiveness of the networks placing these ads? Most advertisers currently credit the provider of the final click – attributing most value to a search engine, since search engines so often close a sale. But display ads can help a consumer recognize a need or learn that a new class of product is available – efforts that the current approach seems to discount unduly. A sensible alternative would offer large benefits – not just to the display advertising services whose value is currently underappreciated, but also to the many web site operators who seek to use such value to sustain their socially-beneficial publications.

REFERENCES

CHATERJEE, P., HOFFMAN, L., AND NOVAK, T. 2003. Modeling the Clickstream: Implications for Web-Based Advertising Efforts. Marketing Science 22, 520-541.

EDELMAN, B. 2006a. False and Deceptive Pay-Per-Click Ads. benedelman.org/ppc-scams/ .

EDELMAN, B. 2006b. Banner Farms in the Crosshairs. benedelman.org/news/061206-1.html.

EDELMAN, B. 2007. Advertising through Spyware, After Promising to Stop. benedelman.org/news/031407-1.html .

EDELMAN, B. 2008. Ad Classification at Right Media. Harvard Business Publishing Case 909032.

EDELMAN, B. 2009. False and Deceptive Display Ads at Yahoo's Right Media. benedelman.org/rightmedia-deception/ .

GHOSH, A. AND MAHDIAN, M. 2008. Externalities in Online Advertising. In WWW 2008, Beijing, China, August 2008. 161-168.

GOEL, A. AND MUNAGALA, K. 2009. Hybrid Keyword Search Auctions. In WWW 2009, Madrid, Spain, April 2009. 221-230.