

The Dark Underbelly of Online Advertising

by Benjamin Edelman

November 17, 2009

Benjamin Edelman (bedelman@hbs.edu) is an assistant professor at Harvard Business School in Boston. A shorter version of this piece, titled “How to Combat Online Ad Fraud,” appears in the December 2009 issue of *Harvard Business Review*.

The internet is sold to advertisers as a highly measurable medium that is the most efficient way to target exactly the right customers. But online advertising is also easily subverted — letting fraudsters claim advertising fees for work they did not actually do. The trickiest frauds deceive advertisers so effectively that measurements of ad effectiveness report the fraudsters as exceptionally productive and high quality, rather than revealing that their traffic was actually worthless. This is a quiet scandal. In a time of tightening ad budgets, losses to advertising fraud come straight from the bottom line—but savings can be equally dramatic.

Here’s a look behind the veil—an explanation of ad practices that have cheated even the web’s largest advertisers. Advertising scams take plenty of victims, both witting and not, but I offer strategies to help determined marketers protect themselves.

Invisible Banner Ads and CPM Attacks

Inqwire’s Rotating and Stacked Banners

Banner ads are most often sold CPM (“cost per mille”), with publishers charging a predetermined price for every thousand

Duration v. Rich Media Club
IPR2023-00953, 11,443,329
RMC EX2061

During 2006, Sausalito, California, ad broker Hula Direct promoted a set of sites with a remarkable volume of ads. For example, as shown in the screenshot, Hula's Inqwire site showed multiple banners stacked on top of each other. (Notice the Universal Orlando Resort ad covering most of a car rental ad at screen left.) Whenever a CPM ad was placed into the Hula site, the advertiser paid full price for the corresponding impression—even if only a corner of its ad was actually visible.

<!--



-->

To further increase revenues,

Hula showed many ads for just seconds each. For example, the code below confirms Hula's efforts to replace ("refresh") an ad after the ad was visible for only nine seconds. Whenever a CPM ad was placed in an ad reloader, the advertiser paid full price even though its ad was visible only briefly.

```
<meta http-equiv="Refresh"
content="9 url=728x90.asp?
jscode=...>"
```

Hula was ultimately called to account for these practices. Ad network Red McCombs Media

times an ad is purportedly shown. For example, an advertiser might agree to pay a web site \$10 for every thousand views of the advertiser's ad. On legitimate web sites, this works well: Advertisers and sites can each assess the value of the site's users and negotiate an appropriate fee. The problem is that CPM payments make it all too easy for fraudsters to game the system.

Fraudsters can increase the number of impressions—and hence their own revenues—by claiming to have shown ads that were not actually, or not meaningfully, displayed to users. In one early attack, fraudsters rapidly switched ("rotated") ads, showing each ad for just a few seconds. An ad visible for a few seconds offers an advertiser little benefit, because it gives the user little chance to notice the ad, not to mention click. Yet under standard CPM rules, the advertiser would be paying just as much for this three-second display as for a proper, extended display. The sidebar "Inqwire's Rotating and

refused to pay a \$200,000 bill from Hula, and Hula sued in response. At the same time, I presented Hula's practices in a detailed article on my web site. Hula sent me a cease-and-desist letter and threatened suit, but I had the proof to back up my claims. Soon after, Hula abandoned these techniques.

Stacked Banners" presents a site that used this rotation technique.

Other fraudsters, termed "banner farms," found ways to inflate impressions by showing many ads at once. For example, some sites stack multiple banner ads above

one another, so that users see only the corners of many ads. The site in the Inqwire sidebar used stacked banners to further increase its revenue. In the most noxious banner frauds, sites load dozens of ads in invisible windows, where users cannot see the ads at all. The sidebar "Mytoursinfo's Invisible Banner Ads" dissects a site that loaded ads invisibly.

Banner farms offer advertisers little real value: When an ad is visible only briefly or in part, the ad is less likely to get the user's attention and less likely to spur a purchase or even increased brand awareness. But available data can mislead advertisers into thinking their banner farm placements are effective. Consider sellers focused on brand building—seeking to show banner ads with little expectation of receiving many clicks because, for example, the seller's products are largely sold offline. These advertisers might not be alarmed by the few clicks occurring at banner farms; these advertisers are *expecting* few clicks. Meanwhile, in their search for wide distribution of their ads, these advertisers are pleased to find sites willing to show banners at a low cost per impression. Looking only at banner farms' low prices, these advertisers often conclude that banner farms are a great deal — not realizing that their spending is going to waste.

The savviest of fraudsters create roadblocks to avoid detection. For example, the Mytoursinfo site shown in the sidebar splits its traffic between dozens of different ad-loader alter egos—the Pointtrip and Fluentcall sites listed in the sidebar inset, among dozens more. To sever ties with Mytoursinfo, an advertiser would need to identify and ban all of Mytoursinfo's subsidiary sites. Furthermore, advertisers typically buy banner ads through ad networks, and ad networks tend to trade and reroute traffic as

they try to optimize ad campaigns. So even after an ad network cut its direct link to Mytoursinfo, other networks might continue to pass along the tainted traffic. With dozens of networks involved, there were hundreds of possible relationships and sequences to be banned. All told, Mytoursinfo continued operation for at least three months before finally shutting down.

Mytoursinfo's Invisible Banner Ads



The image shows the Mytoursinfo site as it appeared to

ordinary users in February 2009. The site looks innocuous—perhaps a bit dull—but, at first glance, hardly fraudulent. However, a visit to Mytoursinfo also created 41 invisible windows using the code excerpted in the inset. First, Mytoursinfo loaded external JavaScript files called “stat.js” and “counter.js” (yellow). Far from tracking statistics or recording counts, these files each created dozens of inline frames (“IFRAMES” in the parlance of web page HTML code) (green), and each of these frames loaded a further page (blue) which in turn embedded several banner ads. With each of dozens of inline frames proceeding to load several banner ads, a single view of the Mytoursinfo site typically loaded more than a hundred banner ads at once. Because the inline frames were invisible (red), advertisers

Banner farms are possible only because advertisers are paying in proportion to the number of times their ads are (purportedly) shown. Might some other method of measurement create better incentives?

Defrauding Pay-Per-Click Advertisers

The paid search revolution brought a new norm for internet advertising: Paying not merely when an ad is shown, but only when the ad is clicked. Under a pay-per-click model, banner farms are no longer profitable: Few to no users *click* on fast-disappearing, stacked, or invisible ads, so a fraudster would gain little through such trickery. Nonetheless, a new generation of fraud was on the way.

For one, even clicks can be faked with surprising ease. A 2004 *Times of India* article presents a New

Delhi company hiring staff to click on ads—yielding

were charged for hundreds of ad-views even though users could not see the associated ads.

```
GET / HTTP/1.1
Host: www.mytoursinfo.com


HTTP/1.1 200 OK
<html>
<script src="/js/counter.js"
type="text/javascript"> ...
<script src="/js/stat.js"
type="text/javascript"></script>

GET /js/stat.js HTTP/1.1

HTTP/1.1 200 OK
document.write("<iframe width=0 height=0
src='http://www.pointtrip.com/florida_tour
.html'>");
document.write("<iframe width=0 height=0
src='http://www.fluentcall.com/pda_phones.
html'>");
document.write("<iframe width=0 height=0
src='http://www.webhotshop.com/shopping.ht
m'>");
document.write("<iframe width=0 height=0
src='http://www.freebiespack.com/freebies
insider.htm'>");
document.write("<iframe width=0 height=0
src='http://www.onlinemoneytrading.net/for
ex_trading.htm'>");
...
```

SmileyCentral Interceding on Users' Requests for Specific, Named Web Sites


A standard Internet Explorer 6.0 browser features the layout shown in the top screenshot. To request a web site, the user clicks in the large box at top left, enters the site's address, and presses Enter. This approach is remarkably standard; it has been used by dozens of browsers dating back to the 1990s.



SmileyCentral changes Internet Explorer by inserting a search box where the browser's address bar belongs.

revenues to the sites where those ads appear, but simultaneously inflating advertisers' costs with no bona fide advertising benefit. More generally, it's easy to write software to fake a click—to make a user's computer request a paid search ad, as if the user had clicked the ad link, when the user did no such thing. Tens of millions of computers are infected with botnets that can fake ad clicks, and these infections have proven remarkably difficult to eradicate.

To defend against click fraud, some advertisers monitor the sales that seem to result from their pay-per-click campaigns. If an advertiser saw thousands of clicks with no sales, perhaps it would lower its bids or cease buying pay-per-click advertising at all. So search engines have been quick to claim that click fraud does not harm advertisers. (For example, in 2006 Google commissioned a report from NYU professor Alexander Tuzhilin, reaching that conclusion.) But I'm less convinced. For example, an algebraic



If a user with SmileyCentral clicks in the top-left box, enters a site address, and presses Enter, the user does not reach the site he or she specifically requested. Rather, SmileyCentral takes the user to a list of ads triggered by that site's address. Of course the top ad typically promotes the site the user requested, so a user clicking such an ad will ultimately end up where the user wanted to go. But due to SmileyCentral's intervention, the site has to pay a fee to reach a user it would otherwise have received for free.



model from Kenneth C. Wilbur of the Fuqua School of Business shows how advertisers lose if they delay even slightly in adjusting their bids or if they are uncertain about the amount of click fraud. Moreover, in industries with offline sales or long sales cycles, advertisers are ill-equipped to measure ad effectiveness, so an advertiser might not know when (or to what extent) it is suffering click fraud.

At least as important as click fraud, *syndication fraud* charges paid search advertisers for traffic they should have gotten anyway, for free. Suppose a user installs the

SmileyCentral toolbar from IAC (the advertising powerhouse that also owns Ask.com and Match.com and, until recently, owned Expedia). Typically, a user seeking a particular web site clicks in the browser's upper-left address bar, types the address, and presses Enter or Go. (See the sidebar "SmileyCentral Interceding on Users' Requests for Specific, Named Web Sites.") But SmileyCentral breaks that convention by pushing a user's address bar to the right, inserting a search box where users naturally expect the address bar to appear. Then, when a user goes to the top-left box and enters a domain name, SmileyCentral runs a *search* instead. When the user clicks the desired web site on the search page, chances are it's a sponsored link, so the destination site pays a fee to Google, which in turn pays IAC. Advertisers agree to pay high Google prices because they value Google's legendary ability to find suitable buyers for advertisers' products. But syndication fraud leads advertisers to pay Google for traffic that the advertisers would have received anyway, without charge.

Crucially, syndication fraud inflates advertisers' measured conversion rates—making advertisers think their campaigns are working better than they really are.

Affiliate Cookie-Stuffing

Some online scammers have developed strategies for attacking “cost-per-action” (CPA) and affiliate programs that were once believed to be fraud-proof. The basic CPA idea is simple enough: Pay advertising commissions only when a sale occurs—not merely when a user purportedly sees an ad or purportedly clicks. Unfortunately, affiliates have learned how to claim CPA commissions they did nothing to earn.

For example, in August 2008, I uncovered the “AlleBrands” banner ad shown in the sidebar “The AlleBrands Cookie-Stuffing Scam.” The ad doesn’t look like much, and its generic design hardly invites users to click. But this fraudster’s strategy doesn’t require clicks. Rather, as soon as a computer loads the ad, code within the ad *fakes* clicks on three CPA links—invisibly setting “tracking cookies” in the user’s computer falsely claiming that AlleBrands referred the user to security software offerings from McAfee, Microsoft, and Symantec. If the user happens to make a purchase from one of these vendors within the next 30 to 45 days, AlleBrands is credited as the purported cause of the sale and rewarded with a commission of up to 25%—even though AlleBrands did nothing to promote any of these companies. AlleBrands’ approach is a scam—charging advertisers for sales they would have gotten anyway, for free. I happened to have contacts at all three targeted companies, so I was able to notify them, and they promptly terminated AlleBrands. But similar attacks target other CPA advertisers.

With tricky software on a user’s computer, attacks can be even more flagrant. Consider a fraudster that wants to claim CPA payments from Dell, say, for purportedly recommending and selling Dell computers. Since CPA payments are issued only when a user makes a purchase, the fraudster needs to find users who are highly likely to buy from Dell. Spyware and adware offer an easy solution: These programs track users’ browsing, so they can provide users who are already at Dell’s site—users who are reasonably likely to buy from Dell in the coming hours or days.

partnership with the spyware and adware vendors, the fraudster can open its affiliate link to Dell while a user browses the Dell site—thereby claiming commission on the purchases users were already about to submit. I first reported these tactics in summer 2004, and they have only grown since then; all told, I have found more than two thousand different CPA affiliates using these fraudulent techniques. I even wrote a software robot that catches these affiliates automatically.

The ALLEBrands Cookie-Stuffing Scam



In August-September 2008, the ad shown here was running on numerous web sites. It doesn't look like much, and it seems it didn't get much user attention. But within the ad, invisible subwindows placed tracking codes onto users' computers, claiming commission for users' subsequent purchases from top computer security vendors.

Here's the code that did the ad's dirty work:

```
<a
href='http://allebrands.com'>
<img
src='images/allebrands.JPG'>
</a>
<iframe src
='http://click.linksynergy.com/fs-
bin/ click?
id=Ov83T/v4Fsg&offerid=144797-10000067&type=3&subid=0'
width = '0' height = '0'>
<iframe src
='http://www.microsoftaffiliates.net/t-
aspx?

```

The Roots of the Problem

A hundred years ago, newspapers realized they could increase advertising revenues by claiming to distribute more copies than were requested by actual subscribers—sometimes not even printing the extras, or printing them only to dump them in the trash. In due course the publishing industry largely ended these “circulation fraud” tactics through circulation bureaus and harsh penalties to the reputations of publications that inflated their circulation. Why haven't similar sanctions worked online?

For one, auditing advertising spending is unpalatable to most marketing professionals. Marketers are trained to win customers—not to run gumshoe-style investigations. And

Duration v. Rich Media Club
IPR2023-00953, 11,443,329
RMC EX2061

<pre> kbid=9066&p=http%3a%2f%2fcontent.affiliates.net %2fWLTtoolbar.aspx%2f&m=27&cid= width='0' height='0'> <iframe src ='http://send.onenetworkdirect.com/1/4/ CD98773' width ='0' height = '0'> oneNetwork affiliate network)] </pre>	typically feature a significant creative component (like designing ads and building promotions), whereas overseeing ad spending requires tedious verification and follow-up—a poor fit for many marketers.
--	---

Marketers often turn to their legal departments for the formalities of partner relationships and for resolution of any disputes. But legal staff usually lack adequate knowledge of the institutions of online advertising and the unusual vulnerabilities and risks. Furthermore, even the worst online advertising frauds operate pursuant to some form of contract, so they typically appear legitimate in a superficial legal review.

An optimal investigation of advertising fraud draws on an unusual diversity of skills—not just ad industry savvy, but also a lawyer’s interest in parsing contracts, an economist’s attention to resulting incentives, an engineer’s technical analysis, and a detective’s nose-to-the-ground attention to detail. Where to find such a guru? No wonder companies fall short!

Moreover, the incentives of the ad-buying process tend to hinder efforts at accountability. Any attempt by an ad network or agency to find fraud would entail research and investigation expenses, and if deception were discovered, the network or agency would *lose* its markup on the fraud it found, while the advertisers would reap the savings.

Ad networks worsen the problem of unwanted ad traffic through the legal terms they typically impose on advertisers. Consider Google’s AdWords Terms and Conditions. Google requires advertisers to accept the remarkably one-sided agreement excerpted in the sidebar “Google AdWords Terms and Conditions (excerpted).” Remarkably, Google doesn’t even promise refunds when ads go astray. Instead, Google insists that “any refunds for suspected invalid impressions or clicks are within Google’s sole discretion.” These harsh provisions may not actually be

enforceable—particularly when imposed on small advertisers in no position either to investigate fraud or to negotiate better terms. But, whether or not the terms are enforceable, advertisers tend to *treat* them as enforceable and to believe they cannot sue an ad network for improper charges. Nor is it any solution to suggest that Google’s advertisers detect fraud on their own. After all, as the operator of the underlying platform, Google enjoys far better data to more effectively identify anomalies; any individual advertiser receives a tiny fraction of the data automatically within Google’s possession.

Google AdWords Terms and Conditions (excerpted)

“To the fullest extent permitted by law, Google disclaims all warranties, express or implied, including without limitation for noninfringement, satisfactory quality, merchantability and fitness for any purpose. To the fullest extent permitted by law, Google disclaims all guarantees regarding positioning, levels, quality, or timing of: (i) costs per click; (ii) click through rates; (iii) availability and delivery of any impressions, Creative, or Targets on any Partner Property, Google Property, or section thereof; (iv) clicks; (v) conversions or other results for any ads or Targets; (vi) the accuracy of Partner data (e.g. reach, size of audience, demographics or other purported characteristics of audience); and (vii) the adjacency or placement of ads within a Program. Customer understands that third parties may generate impressions or clicks on Customer’s ads for prohibited or improper purposes, and Customer

The internet’s breadth further compounds online advertising fraud. With millions of web sites showing ads, it is hard for advertisers to track sites’ reputations. Moreover, even after a fraudster gets caught under one name, it can easily change to another address and continue business as usual. Computer scientists Eric J. Friedman and Paul Resnick identified this problem in their 2001 “The Social Cost of Cheap Pseudonyms,” and sequential identities plague online advertising in spades. All too often, I catch a fraudster and reveal its practices, only to find the scam reappear weeks later by the same offenders, now using a new web address and a new name.

Finally, the measurability of online advertising seems to hinder many

accepts the risk of any such impressions and clicks.”
–Google AdWords Terms and Conditions

efforts at accountability. If a fraudster’s efforts were exceptionally simplistic—perhaps claiming to have shown a banner ad millions of times, but

without a single click resulting—measurement and oversight would reveal the scam. But suppose a fraudster’s numbers seem to check out—fake clicks match the fake impressions, for example. Examining these records, an auditor would have little basis to second-guess the ad partners. Quite the contrary, fraud sites might look *good*—low-cost and seemingly effective, if measurements are taken at face value. In my work with online advertisers, this problem is increasingly widespread: When management by the numbers replaces a genuine understanding of partners’ practices, fraudsters escape detection by manipulating reported values to their advantage. And with automated, standardized measurement systems, manipulation has become easier than ever.

A Way Forward

Preventing online advertising fraud begins with recognizing the problem and assigning responsibility. Advertisers need plans to audit their spending and confirm that they are getting what they’re promised. In traditional purchasing contexts, few companies would spend millions of dollars through thousands of partners without rigorous oversight of who is doing what. Similar institutions—procurement managers, quality control, and reputation—are long overdue for online advertising.

Advertisers must also recognize the structural challenges that make online advertising fraud especially pernicious. Consider an advertiser facing hundreds of frauds costing thousands of dollars each. With six or more digits of spending on the line, it is no solution to accept these losses as a cost of doing business. Yet with hundreds of separate incidents—from hundreds of separate perpetrators—it is equally unattractive to file suit, for litigation costs would swamp likely recoveries. Instead, effective enforcement requires a response coordinated across many advertisers. In principle, class action litigation could be used, but with so many key intermediaries cloaking themselves in

language of the law, this is a tall order. Neither is public enforcement realistic: The FTC and state attorneys general typically stand aside in business-to-business disputes, on the premise that their limited resources are better directed to cases where consumers, not companies, are the direct victims. On this point, perhaps the online advertising industry also shoulders some responsibility: Sectors like financial services have long histories of referring fraud to law enforcement, but such referrals are nearly unprecedented in the world of advertising.

Despite these challenges, I'm cautiously optimistic about prospects for deterring online advertising fraud. For one, advertisers can make better use of the fact that they are ultimately paying the bills. A savvy advertiser might resolve to pay its partners more slowly. That way, when the advertiser uncovers a fraud, it is more likely that the fraudster has not yet been paid for the traffic at issue. In "Deterring Online Advertising Fraud Through Optimal Payment in Arrears," I offer an algebraic model that develops this approach—offering a mechanism to compute an advertiser's optimal delay in light of its fraud rate and its partners' impatience.

Advertisers can also benefit from redoubled efforts to verify and confirm partners' identities. If an ad network seeks to put an advertiser's ads on thousands of sites but will not tell the advertiser what sites are included, that should be an immediate cause for concern. At the very least, an advertiser could sensibly demand a heightened level of assurance or a guarantee when an ad network seeks to keep its partner list confidential. Meanwhile, for an advertiser evaluating prospective partners, increased diligence is in order. If a prospective partner is located in a far-flung location with a limited judicial system, an advertiser might wonder whether that partner can appropriately promote the world's top brands. No matter the partner's capabilities, a shrewd advertiser should also consider the high costs of any subsequent dispute with a partner so far away. Conversely, if a company claims to be based in the United States but cannot offer a U.S. mailing address or always logs into an advertiser's servers from computers abroad, that is ample grounds for a second look. With pseudonymity so easy in a world of web-based contracting, advertisers should constantly ask themselves whether partners are who they claim to be.

Ad partners' misleading traffic statistics require similar skepticism. At their peril do advertisers assume that a count of "views," "clicks," or even "purchases" means what it claims. Ad buyers should ask at every turn: How is that measured? How do we know it's real? Online marketing certainly lends itself to quantitative analysis, and advertisers cannot ignore the numbers. But even purchases can be faked, as I've shown. So advertisers ought not proceed on the assumption that every partner has actually done the work that reports seem to indicate. An ad buyer's instinct and professional judgment can be just as valuable as automatic reports from a measurement system that can be gamed.

Large advertisers, in particular, would benefit from clarifying responsibility for resolving advertising fraud. Advertisers should determine, in advance, who will be looking for fraud (the advertiser, the agency, or the network), how fraud will be investigated, and what actions will be taken if it is found. These steps can help advertisers reduce waste and, ultimately, improve the effectiveness of their ad campaigns.

References

Ben Edelman. "Deterring Online Advertising Fraud Through Optimal Payment in Arrears." Proceedings of the Thirteenth International Conference on Financial Cryptography and Data Security. Springer-Verlag Lecture Notes in Computer Science. 2009.

Eric J. Friedman and Paul Resnick. "The Social Cost of Cheap Pseudonyms." Journal of Economics and Management Strategy 10(2): 173-199. 2001.

Google AdWords Terms and Conditions.
<https://adwords.google.com/select/tsandcsfinder>.

Alexander Tuzhilin. "The Lane's Gifts v. Google Report,"
http://googleblog.blogspot.com/pdf/Tuzhilin_Report.pdf. July 2006.

N Vidyasagar. "India's Secret Army of Online Ad 'Clickers.'" The Times of India. May 3, 2004.

Kenneth C. Wilbur and Yi Zhu. "Click Fraud." Marketing Science. 27(3): 293-308. 2009.

Benjamin Edelman is an associate professor at Harvard Business School and an adviser to various companies that compete against major platforms.

Recommended For You

PODCAST


How Do I Handle Negative Feedback from My Team?



How Netflix Expanded to 190 Countries in 7 Years



10 Common Job Interview Questions and How to Answer Them



38 Smart Questions to Ask in a Job Interview

