Risk, Information, and Incentives in Online Affiliate Marketing *

Ben Edelman Wesley Brandi

June 23, 2014

Abstract

We examine online affiliate marketing programs in which merchants oversee thousands of affiliates they have never met. Some merchants hire outside specialists to set and enforce policies for affiliates, while other merchants ask their ordinary marketing staff to perform these functions. For clear violations of applicable rules, we find that outside specialists are most effective at excluding the responsible affiliates, which we interpret as a benefit of specialization. However, in-house staff are more successful at identifying and excluding affiliates whose practices are viewed as "borderline" (albeit still contrary to merchants' interests), foregoing the efficiencies of specialization in favor of the better incentives of a company's staff. We consider the implications for marketing of online affiliate programs and for online marketing more generally.

1 Introduction

For decades—perhaps centuries—marketers have bemoaned the effectiveness of their advertising campaigns. Paying for advertising up-front and receiving benefits later, advertisers are vulnerable to low-performing or non-performing ad placements. Against this backdrop, *affiliate marketing* appears to offer a refreshing change: In this performance-based approach to online marketing, advertisers pay only when a sale occurs. With robust online tracking that attributes sales to affiliates, advertisers perceive an unprecedented reduction in risk. The Economist (2005) captured advertiser excitement for the apparent alignment of incentives, calling affiliate marketing "the holy grail of online advertising."

As it turns out, affiliate marketing is neither as easy nor as safe as proponents initially anticipated. Most advertisers struggle to find reliable affiliates who deliver new customers in desired quantities, in exchange for reasonable compensation. Meanwhile, despite the promised alignment of incentives, bad affiliates can exploit shortcomings in tracking and attribution to claim commissions

*We thank George Baker, Florian Ederer, Francesca Gino, Robert Glazer, Brian Hall, Zhenyu Lai, Ian Larkin, Tyler Moore, Frank Nagle, Dave Naffziger, Steve Tadelis, NOM seminar participants, and two anonymous referees.

they have not fairly earned. Informed by these problems, affiliate marketing raises longstanding questions of judgment, partnership, and incentives reminiscent of decades of media-buying.

This paper offers two contributions. We begin by presenting affiliate marketing generally, exploring its institutions and participants as well as key risks uncovered to date. Second, we explore advertisers' efforts to address those risks. Specifically, we evaluate advertisers' management structures by measuring relative prevalence of affiliate fraud. By examining the common methods of affiliate program management, we identify the vulnerabilities best addressed by outsourcing marketing management to external specialists, versus the problems better handled by keeping management decisions in-house. We find that outside specialists are most effective at enforcing clear rules, while in-house staff are better at preventing practices viewed as "borderline" under industry norms.

While our results apply most directly to advertisers considering the management structure of their online marketing programs, our analysis also speaks to a broader literature of outsourcing and boundary of the firm. Managers often face a tradeoff between retaining functions in-house (typically with greater supervision and greater control over quality) versus outsourcing to a specialist (who may have greater capability or a cost advantage thanks to scale and experience). In general, these questions make empirical estimation difficult: It is usually challenging to find a context that offers numerous similar insourcing/outsourcing decisions. Furthermore, companies' structures are generally confidential and hence unobservable to researchers. In contrast, we examine an online marketing context where advertisers often reveal their management structures as they recruit marketing affiliates. We enjoy an additional advantage from a novel data set. Ordinarily, both firms and researchers lack top-quality data about opportunistic behavior; those providing low-quality service usually seek to conceal their activities from the principals that pay them. If principals cannot determine quality, researchers usually also struggle to determine what has occurred. In contrast, we develop custom software to examine affiliates' behavior—information often unavailable even to the advertisers and networks who purport to supervise these affiliates.

2 Affiliate marketing and affiliate fraud

Affiliate marketing combines sharp performance incentives with the broader efficiencies of online advertising. In particular, affiliate marketing compensation is usually purely performance-based—

offering perhaps a $5 or 10% advertising fee for each purchase. Under standard rules, an affiliate earns a commission only if 1) a user browses to an affiliate's site, 2) the user clicks the affiliates specially-coded link to the merchant, and 3) the user makes a purchase from the merchant. (Edelman, 2013) These additional requirements importantly differ from better-known methods of online advertising: Most display ads ("banner ads") require an advertiser to pay as soon as a web site serves an ad to the user, and almost all search ads require an advertiser to pay as soon as a user clicks an ad.

Affiliate marketing payment rules are understood to protect advertisers against wasted expenses. Consider payment structures and resulting risks in other online advertising implementations. For example, when buying display advertising, an advertiser might reasonably worry that few users will click its ads: Perhaps the ads are irrelevant to users' interests or are placed in locations where few users notice. Some of these factors are outside the advertiser's control: Standard contracts let advertising networks decide which sites show a given advertiser's ad. In these circumstances, advertisers perceive serious risks that their banner advertising expenditures will be wasted. Similarly, an advertiser buying search ads risks extra expense if uninterested users, competitors, or fraudsters click or purportedly click. (Wilbur and Zhu, 2009) Here too, standard contracts require advertisers to pay even if the advertising leads to few or no purchases. In contrast, affiliate marketing payment is only due if a user makes a *purchase*—aligning advertising expense more closely with an advertiser's revenue and profit.

Affiliate marketing is also distinctive in that most affiliate merchants buy advertising from small marketing affiliates they have never met. Merchants typically accept affiliates with few to no assets, affiliates lacking well-known brand names or established reputations, and affiliates in remote locations. Small, low-asset, distant marketing partners present an obvious risks of unaccountability, but merchants typically consider themselves at least partially protected from rogue affiliates due to the structure of affiliate compensation: So long as a user actually makes a purchase, merchants generally perceive that there is little downside to paying a commission. LinkShare, a leading affiliate network, historically promised advertisers that they would "pay affiliates only when a sale or other qualifying action is completed," an approach which LinkShare touted as "very efficient." (LinkShare, 2009) Similarly, affiliate network Commission Junction notes that "advertisers only pay when a specific action has been completed (e.g. a purchase...)" which, CJ says, makes affiliate

marketing "low risk." (Commission Junction (2014))

Although practitioners initially considered affiliate marketing structurally protected from fraud, there are actually significant risks including the practices we examine in Section 2.2.

2.1 The institutions of affiliate marketing

An affiliate marketing *merchant* is the web site seeking to sell goods or services through online advertising. Affiliate marketing merchants span the gamut of online commerce, from the web's largest sellers, including Amazon, to mom-and-pop specialty sites.

An *affiliate* or *publisher* is a web site that presents links to its visitors. For example, when posting a book to a blog or discussion forum, an affiliate could offer a link to Amazon to facilitate readers' purchases. Similarly, when suggesting a vacation destination, a travel site could link to a page at Expedia offering hotels in that area. In the best case, these affiliate links make the underlying content more useful while also providing payment to the publisher. In practice, some affiliates use the prohibited practices explored in Section 2.2.

A *network* connects merchants and affiliates. Most merchants rely on networks for tracking, administration, and accounting purposes—to record which users clicked which links and made what purchases; to provide a secure web site for affiliates to obtain links and check results; and often to provide efficient consolidated payments to numerous affiliates each month. In principle, merchants could handle these tasks in-house, and some of the web's largest affiliate marketers have done so (including Amazon since the inception of its affiliate program, and more recently eBay and Apple). But most merchants prefer the benefits of specialization. Networks impose some rules about permissible affiliate practices. When a merchant joins a network, the merchant can waive most such rules or add other requirements of its own.

An *affiliate program manager* sets the rules of an affiliate program including how much affiliates will be paid, what behaviors are permitted, and which affiliates to accept or reject. Section 4 explores the various models of affiliate program management.

For our purposes, it is generally not necessary to explore the technology that facilitates affiliate program operations and tracking. The fundamental enabling feature is the browser *cookie*, a technology which lets a web site place data on a user's computer in order to recognize a user upon a further visit. When an affiliate refers a user to a merchant, the merchant or network places a

cookie on the user's computer. Then, if the user later makes a purchase, the cookie will reveal that the purchase followed the affiliate's referral.

2.2 Fraud in affiliate marketing

Our tabulation of affiliate litigation (Edelman, 2012) reveals a dozen disputes large enough to spur legal action. Because the practices at issue satisfy the elements of common law fraud and have been charged as fraud in both civil and criminal litigation, we call these practices *affiliate fraud*.

In most affiliate marketing programs, commission is only paid if a user makes a purchase. Thus, if a rogue affiliate seeks to inflate its charges to a given merchant, the affiliate needs to make the merchant's records indicate that the affiliate has delivered additional sales. In principle, an affiliate could infiltrate a merchant's servers to alter records directly. But an attacker with privileged access to merchants' servers need not stop at affiliate fraud. In practice, affiliate fraud most often focuses on schemes that find users who were already on the verge of making purchases. Through such methods, rogue affiliates claim commission on purchases that were going to occur anyway, even though the affiliate did not genuinely cause or encourage these purchases.

Our data is grounded in four affiliate schemes, which we have found to be the largest areas of affiliate malfeasance:

1. *Adware.* When a user visits a merchant's site on a computer running certain advertising software, the software sees the user's activity and redirects the user through an affiliate's marketing link. If the user subsequently makes a purchase, the affiliate will be credited as the putative cause of that purchase.

2. *Cookie-stuffing.* When a user visits a web page, a section of that page can claim to refer the user to a given merchant. If the user happens to make a purchase from that merchant within a predetermined time thereafter (often seven to 30 days), the affiliate will be credited. In variations, the affiliate can design its page to attract traffic to particular merchants (perhaps by repeatedly mentioning the merchant's name or by promising, truthfully or falsely, to offer coupons for that merchant). The affiliate could also design its cookie-stuffing to be a component of some other web page (a dot on a banner ad, or a section of a comment on a forum or blog).

3. *Typosquatting.* Affiliates register domain names that are misspellings of merchants' domain names. (Moore and Edelman, 2010) When a user misspells a merchant's domain name in the way that the affiliate anticipated, the user will be sent to the affiliate's site, which immediately redirects the user through an affiliate link and onwards to the merchant. If the user makes a purchase, the affiliate will be credited.

4. *Loyalty software.* Affiliates place "loyalty" software on a user's computer to remind the user about possible rebates, points, or other benefits from purchasing through certain merchants. The loyalty software automatically sends a user through an affiliate's link when the user requests a merchant's site directly. Typically, loyalty software becomes installed as part of a bundle when users ask for wholly unrelated software. Often, loyalty software claims affiliate commission even if the user had never registered with the loyalty service and is hence incapable of claiming or receiving benefits. We note that there is some debate about whether "loyalty" software in fact creates customer loyalty—and if so, whether it is to the merchant or to the maker of the loyalty software. But we accept the term since it is widely used by practitioners.

The schemes we examine are a subset of undesirable affiliate behavior. For example, other rogue affiliates engage in trademark bidding by buying search engine advertising for a merchant's name, then sending the resulting users through merchants' affiliate links and claiming affiliate commission on resulting sales. Sophisticated merchants largely disallow this practice because it tends to increase competition in the search engine advertising auction (adding an extra bidder who a merchant must outbid) and because merchants have found that they can buy these same keywords at prices lower than affiliate fees. Typically, merchants impose custom terms and conditions to ban affiliates from engaging in trademark bidding. In principle we could obtain each merchant's stated rules, then check for violations. But some merchants provide selected affiliates with waivers of their general rules, and we do not observe those waivers. Moreover, varying merchant rules would add significant complexity: a given practice might be a violation for one merchant but permissible for another. We therefore focus on the four behaviors listed above, where a merchant's interest is more clear-cut.

2.3 Practitioners' views of affiliate fraud

Affiliate marketing practitioners have differing views on the practices presented in Section 2.2.

In general, practitioners view adware as clearly impermissible, specifically forbidden by most network contracts. They share similar views of cookie-stuffing. We refer to these practices as "clear-cut" violations.

In contrast, practitioners offer varying evaluations of typosquatting. Some affiliate managers view typosquatting traffic as helpful to users and likely to reach users who would otherwise get lost—and not purchase—due to a browser error message or other unhelpful content. For example, if a user misstypes `expendia.com` (s.i.c.), Expedia might be willing to pay a modest commission to obtain that user without delay, thereby avoiding the risk of an error message dulling the user's interest or an advertisement diverting the user to a competitor. In contrast, other affiliate managers view typosquatting as an improper practice that they should not be asked to allow or pay for. They note that typosquatting is contrary to federal law (the Anti-Cybersquatting Consumer Protection Act, 15 USC §1125(d)), that it entails "forcing clicks" in violation of affiliate network rules, and that most web browsers would direct the user to the genuine site without charging a merchant any advertising fee.

Practitioners are even more divided on the issue of loyalty software. Supporters typically argue that users value the points and rebates, that competitors also participate in these loyalty programs, and that users might shift to a competitor if a given merchant leaves a loyalty program. In contrast, critics view loyalty applications as a way to collect affiliate commission on traffic that merchants would otherwise have received without charge. They argue that if a user did not care about the loyalty benefit enough to manually visit the loyalty service's web site, the user would unlikely to shift a purchase to a competitor. Critics of loyalty applications also note that loyalty applications are frequently installed onto users' computers without users requests (Edelman, 2004, 2005a), risking merchants paying commission without the users getting any benefit or being motivated by any rebate or points.

In our view, critics of typosquatting and loyalty software have stronger arguments: Both typosquatting and loyalty software claim commissions on sales that merchants would otherwise receive without charge, which is contrary to merchants' interests. But standard contracts include no explicit prohibition on either typosquatting or loyalty software, whereas standard contracts typically do exclude adware and cookie-stuffing. To the extent that typosquatting and loyalty software are prohibited, the ban comes from more general contract provisions such as a requirement that

a user must "click" a link in order for commission to be earned, whereas no such click occurs in typosquatting or in loyalty software. Despite our firm view that these practices are contrary to merchants' interests, we classify them as "grey area" in recognition of diverging views among relevant practitioners.

2.4 Selected instances of affiliate fraud

In this section, we profile the largest publicly-documented instances of affiliate fraud in order to give a sense of perpetrators, methods, and detection.

The largest and best-known affiliate fraud was perpetrated by Shawn Hogan, founder and CEO of an online advertising network that facilitated the placement of banner ads onto the web sites of independent publishers. According to the indictment in USA v. Hogan (2010), pp.6-7, as well as companion private litigation by eBay, Hogan modified his company's ad network code so that when a user viewed a publisher's site, the "user's computer [would] make a request to eBay's home page merely for the purpose of prompting eBay's servers to serve up [an affiliate tracking] cookie." Then, when these users "thereafter visit[ed] eBay.com and engage[d] in revenue activities... [Hogan] would receive compensation from eBay with respect to those events." During the relevant time period, eBay paid as much as $35 to an affiliate who referred a new user (who within 30 days of the referral, registered and bid on at least one item). eBay also paid an affiliate up to 75% of its revenue (on net, roughly 12% of an item's purchase price) from a referred user's purchases within seven days of the referral. According to the indictment, eBay paid Hogan more than $15 million in 2006 and 2007, making him the highest-paid affiliate in eBay's affiliate program. But the indictment and corresponding private eBay litigation allege that Hogan's referrals were entirely useless—users who would have come to eBay anyway. More generally, the indictment and private litigation claim that Hogan's invisible cookie-stuffing did nothing to cause or increase purchases. Hogan pled guilty in December 2012 and, as of March 2014, is awaiting sentencing.

A similar indictment and private eBay claim were brought against Brian Dunning, also alleging invisible cookie-stuffing. (USA v. Dunning, 2010) Litigation documents reveal that Dunning was, at the time, eBay's second-largest affiliate and received more than $5 million in 2006 and 2007. The defendants' statements reveal collaboration: Dunning indicates that Hogan "offered to help" Dunning by "teaching him" key techniques (Miller, 2007, p.4), while Hogan told FBI agents that

Dunning "ripped off" Hogan's approach to claiming affiliate commissions from eBay (Walbridge, 2007, p.5). Meanwhile, when FBI agents interviewed Dunning, he indicated that he paid a 10% fee to an Andrew Way, an account manager at Commission Junction (at the time, the network tracking affiliate transactions for eBay). Dunning says that Way "provided Dunning with inside information regarding how to take advantage of the affiliate program." (Miller, 2007, p.3) Way's LinkedIn page confirms that he worked at CJ, albeit some months before the events at issue. The litigation docket reveals nothing further about Way's alleged involement, leaving Way's true scope of involvement (if any) in Dunning's activities unclear. Dunning pled guilty in April 2013 and, as of March 2014, is awaiting sentencing.

Indictments and other litigation documents indicate that both Hogan and Dunning took significant steps to conceal their activities from eBay and Commission Junction. The indictments allege that both defendants intentionally avoided stuffing affiliate cookies on computers located in geographic areas that they believed were used by eBay, CJ, and their investigators. Walbridge (2007, p.3) also reports Hogan admitting stuffing cookies only once to each IP address, which prevented an investigator from uncovering the practice via repeated testing.

The next-largest known instance of affiliate fraud resulted from brothers Andrew and Allen Chiu's misuse of an affiliate rebate site, Fatwallet. When a user clicks from Fatwallet to a merchant's service, Fatwallet claims affiliate commission from the merchant, and in turn pays most of that commission to the user. The Chiu brothers found that one retailer, Nordstrom, would pay a commission to Fatwallet even if the order was cancelled (either by Nordstrom or by the buyer). In 2010 and 2011, the Chius placed 4,000 Nordstrom orders worth approximately $23.7 million. They knew that Nordstrom would cancel these orders because the store had previously banned the Chius from its site due to excessive complaints of merchandise purportedly lost in transit. Although Nordstrom canceled the orders and did not charge the Chiu's credit cards, Nordstrom nonetheless paid approximately $2 million of affiliate commission to Fatwallet, which in turn paid approximately $1.1 million to the Chius. The Chius pled guilty and were sentenced to 24 months incarceration as well as repayment of the amount taken. (USA v. Chiu, 2012)

When granting awards to the affiliates who had achieved fastest growth, affiliate network LinkShare in three consecutive instances had to retract awards from recipients who were proven to be engaged in cookie-stuffing. (Fadner, 2004) In each instance, a visitor to the affiliate's site

would receive affiliate cookies even without clicking an affiliate link. While the perpetrators were ultimately removed from LinkShare, no publicly-available documents indicate that refunds were provided to affected merchants.

Because these instances are unusually large, they are preserved in litigation records, news media, and other documents. In contrast, most affiliate frauds yield no such records. Nonetheless, though our data collection methods (detailed in Section 5), we are nonetheless able to identify numerous perpetrators as well as victim merchants.

3 Related literature

Since other advertising formats are significantly more established, one might ask why merchants choose affiliate marketing. The literature offers some insight. Libai et al. (2003) consider the similar context of publishers selling "leads" such as signups from users purportedly interested in a given service. They emphasize the risk of publisher moral hazard (perhaps filling out the form with names from a phone book), suggesting that payment structure can shift incentives to discourage such misbehavior. (For example, the advertiser might pay the publisher only for those customers who actually make purchases.) We develop Libai et al. by evaluating the effectiveness of the resulting relationships including schemes that are measured as productive but do not actually advance advertisers' interests. We also extend Libai et al. by considering the management structure that oversees these relationships.

In the analogous context of sites deciding whether to charge advertisers for ads being displayed versus clicked, Zhu and Wilbur (2011) note the role of advertiser heterogeneity as well as uncertain levels of advertiser effort to attract clicks. If an advertiser pays for every click, it should design its ads to attract only clicks from users who are genuinely interested. In contrast, an advertiser paying for displays might as well invite every click possible, even if an increased proportion of clickers do not make purchases. In the context of affiliate marketing, merchants are generally perceived to be trustworthy, but affiliates are highly heterogeneous. As Zhu and Wilbur note (p.251), cost-per-action affiliate marketing is to cost-per-click ads as cost-per-click is to cost-per-impression, and the Zhu and Wilbur principles flow through accordingly.

A separate stream of research questions the measurability and effectiveness of various online

advertising. In a field experiment, Blake et al. (2013) find search ads offering much lower short-term benefits than conventional estimates suggest, including an absence of benefits from brand-keyword ads. We follow their broad skepticism of the measurability of online advertising, exploring the potential mishaps and merchants' varying abilities both to uncover and to prevent these problems.

Meanwhile, questions of firm boundaries, information, and incentives have arisen in numerous contexts far from online marketing. For example, Baker and Hubbard (2003) consider incentives and contractual incompleteness among truck drivers, noting the role of new technology in shifting market structure by facilitating verification of work done. While we examine quite a different market, we note that affiliate marketing is also grounded in granular tracking—improved information collection broadly similar to the trucking on-board computers that motivate Baker and Hubbard. More broadly, Lafontaine and Slade (2007) survey research on the causes and consequences of firm boundaries across numerous sectors.

4 Affiliate program management and resulting incentives

Having elected to run an affiliate program—usually, for the broad reasons noted in The Economist (2005)—a merchant's decision is how to run the program. For example, an affiliate program manager will need to establish rules and decide which affiliates to accept or reject, as well as which affiliates deserve a bonus. Merchants have found three management structures for affiliate programs:

1. *In-house affiliate management staff.* A merchant can assign or hire an ordinary employee to select and manage affiliates. Discussions with affiliate managers reveal that most such staff are paid on a salaried basis, albeit often with performance objectives. Some receive explicit contingent compensation ("$10,000 bonus if our program grows by 10% next year"). We believe most affiliate managers' performance-based compensation is implicit, with a larger program largely viewed as calling for greater compensation. Of course, affiliate mangers' long-term compensation is also typically tied at least in part to company health, including equity as well as opportunities for advancement. Camaraderie and intrinsic motivation further encourage affiliate managers to consider company objectives.

2. *Specialist affiliate-management companies.* A merchant can retain the services of a vendor that specializes in affiliate marketing management. Practitioners often call these vendors

"outsourced program managers" or OPMs. Industry sources reveal scores of OPMs ranging from sole practitioners to modest-sized firms of at most a few dozen staff. A sole-practitioner OPM might manage three to ten programs, while a large OPM could manage a hundred programs or more. OPM contracts vary widely as well: Some are flat fees (e.g. $3,000 per month to manage a given program), while others are percentage ("20% of spend") and some are hybrids. OPM staffing is also diverse: Some OPMs assign a full-time staff person to a single large merchant. Smaller programs typically share OPM managers: a single OPM staff person may manage a dozen small programs for a dozen different merchants.

3. *Affiliate network provides management services.* Most merchants retain the services of an affiliate network to provide the required *technical* infrastructure, including preparing specially-coded links, tracking which purchases were made through which links, reporting purchases, and streamlining payments to affiliates. Merchants can also turn to affiliate networks for *management* services, including judgment of which affiliates to accept and reject. Merchants' payments to networks are largely proportional to the total commission merchants pay: Networks typically charge percentage fees for their technical and tracking functions. (For example, Commission Junction's public price list historically specified $30 of network fees for every $100 of commissions. Commission Junction (2004)) Practitioners indicate that networks' management fees are also largely a percentage of commissions paid. Thus, when affiliate networks perform management services, their charges are best understood as proportional to merchant spending. Glazer (2013b)

Our discussions with practitioners confirm that some merchants are broadly aware of the diverging incentives resulting from compensation of affiliate program managers. That said, practitioners rarely write about these concerns or appropriate responses, and there is little evidence of merchants discussing these questions. Notable exceptions are Glazer (2013a) and Glazer (2013b).

4.1 Information and incentives for affiliate managers

Affiliate program management structures vary both in the information available to managers and in compensation and resulting incentives.

Management structures differ in their access to information about affiliates' practices. An in-

house affiliate manager has access to whatever data the network chooses to provide plus whatever information the affiliate manager can collect from an affiliate or through independent research. Typically, both sources offer limited insight, particularly as to practices that are concealed and difficult to uncover. In contrast, an OPM can combine data from multiple merchants. For example, an OPM can observe an affiliate's effectiveness or integrity in promoting one of the OPM's merchants, and use that information to evaluate the affiliate's suitability for other merchants that hire the same OPM. Finally, a network enjoys the greatest level of information about affiliates' practices. For one, a network's systems store and tabulate data about each affiliate's actions across the entire network, and network program managers in some instances can access this data via mechanisms unavailable to in-house managers and OPMs. Furthermore, network program managers have closer access to other network staff including the affiliate managers who are affiliates' standard points of contact as well as the "network quality" group that investigates possible violations.

Meanwhile, alternative management structures also imply differing incentives. In-house staff are most likely to have flat compensation, while networks are certain to have an important element of *ad valorem* (proportional) compensation. While OPMs often join networks in using proportional management fees, networks combine both management fees and tracking fees, giving networks greater incentive to take actions that increase merchants' costs. Suppose a network and an OPM both charge 20% fees for management service, while the network charges 30% for tracking service. If an OPM takes an action that increases a merchant's cost by $1, the OPM collects additional revenue of $0.20. But if the network takes that same action, the network enjoys additional revenue of *$0.50*. Since a network incurs minimal marginal cost in providing tracking services, its additional revenue is best understood as pure profit. Thus, networks have a notably stronger incentive to increase merchants' costs, compared to the corresponding incentive for OPMs.

The following table summarizes the information and incentives associated with alternative methods of affiliate management:

	Incentive	Information
in-house	flat or modest performance incentives	limited: networks share only selected data; managers are often generalists
OPM	modest performance incentive	intermediate: can combine data across merchants; staff are specialists
network	significant performance incentive	superior: combine data across merchants; direct access to logs

4.2 Merchants' choice of affiliate management structure

The preceding section offers mixed recommendations to a merchant selecting a management structure for its affiliate program. On one hand, merchants might focus on the importance of obtaining information about affiliates' practices. If information is the most important determinant of program success, then programs managed by affiliate networks should have the best quality thanks to the superior information available to affiliate networks. Merchants with OPM-managed programs should have intermediate quality due to OPMs' ability to combine information across multiple merchants. Merchants with in-house programs should have the lowest affiliate quality in light of the limited information available to them.

Alternatively, one might worry about the incentives of affiliate program managers. If some merchants' programs accept undesirable affiliates due to managers' incentives, then network-managed programs are most vulnerable: Networks charge a fee for each transaction, and these fees provide direct and immediate financial benefits for allowing and retaining rogue affiliates. Indeed, if a network found a given affiliate to be in violation in its promotion of one merchant, the network might be obliged to exclude that affiliate from the entire network. With dozens or hundreds of merchants at issue, such an expulsion would often increase a network's lost revenue by an order of magnitude or more. Networks thus have a particularly acute incentive to avoid detecting violations or declaring affiliates' practices to be violations. In contrast, a merchant running its own program is more likely to run the program to maximize merchant profitability: Even if an individual staff person faces formal performance objectives or informal pressure to expand the affiliate program, these incentives are tempered by the work environment and duty to the employer. Merchants with OPM-managed programs should have intermediate quality due to OPMs' mixed incentives.

In a more nuanced interpretation—what our data supports, as we discuss in Section 6.4—information and incentives interact to provide differing benefits for differing behaviors. As explored

in Section 2.3, affiliate malfeasance includes both practices that are understood to be clear-cut violations of applicable rules, as well as "grey area" practices that are contrary to merchants' interests, yet nonetheless sometimes accepted by practitioners. As to clear violations, the key barrier to taking action is information—figuring out which affiliates are engaged in such practices. A capable affiliate network could use its superior information to be most effective at excluding clear-cut violations of applicable policies. Conversely, an in-house manager would have a comparatively reduced ability to find such violations for lack of required information. Meanwhile, as to "grey area" violations, the crucial question is incentives—correctly determining what is truly in a merchant's best interest. In that regard, in-house managers have an advantage because their objectives are most closely aligned with merchants' goals. In contrast, networks' fees for both management and tracking provide a greater incentive for networks to accept grey area behaviors that are not truly in merchants' interest.

Seeing networks' financial incentive to allow affiliate misbehavior, one might ask why a network-managed program rejects any affiliates at all. But consider the impact if a merchant uncovers a clear-cut violation that the network failed to prevent: This would surely shake the merchant's confidence in the network. (Indeed, after uncovering major affiliate fraud, eBay terminated its seven-year relationship with Commission Junction. Edelman (2012) offers other examples of merchants changing affiliate management structure or closing programs after violations are revealed.) In light of these possible repercussions, networks should hesitate to allow clear-cut violations. Meanwhile, as to grey-area violations, networks anticipate that such risks are much reduced, since a merchant would be less likely to end its use of a network in response to grey-area practices.

5 Data

5.1 Merchant management structure

We begin with data on which merchants use which marketing structures. Each merchant using the largest three US affiliate networks (Commission Junction, Google Affiliate Network, and LinkShare) offers a merchant "detail" page with information about the merchant's general offerings, commission payments to affiliates, and requirements for affiliates. 69% of merchants' pages provide a contact email address for affiliates with questions about a given affiliate program, while 31% of pages

provide no email address whatsoever.

We categorize merchants' posted email addresses to draw inferences about the management structure of each merchant's affiliate program. For example, if the email address is a named individual person at the merchant, we categorize that merchant as managing its own affiliate program. If the email address is a named individual or role account at an OPM, we categorize that merchant as delegating affiliate management tasks to an OPM. If the email address is a named individual or role account at an affiliate network, we categorize the merchant as delegating affiliate management tasks to a network. We are able to categorize 62% of merchants in this way.

Some merchant email addresses are difficult to categorize. For example, a Gmail account could forward mail to one or multiple staff at any combination of merchant, OPM, or network. A Gmail account could also be used to let the merchant more easily switch from one OPM to another or from in-house management to OPM or network, or vice versa. For an ambiguity of this form or for lack of any email address at all, we mark as "unknown" the remaining 38% of merchants.

5.2 Affiliate practices

To evaluate the effectiveness of alternative affiliate management structures, we need data on all manner of affiliate misbehavior. This is a challenging task because perpetrators of affiliate fraud largely seek to avoid being revealed as such, lest their accounts be closed and payments withheld.

This section offers an overview of our data collection. The online appendix presents details.

To uncover affiliate fraud, we run automation to render entire web pages in virtual computers running standard web browsers. In this way, we examine pages just as users see them. (In contrast, ordinary web crawlers load only a page's HTML source code. Such crawlers would typically fail to uncover affiliate fraud using methods beyond pure HTML. For example, most cookie-stuffing uses images, JavaScript, Flash, or a combination of these and other methods.) Our automation also simulates random user interaction with web pages to further mimic standard user activities and to trigger any page or program functions that await user activity. Through this reenactment of users' browsing, our approach attempts to trigger as much affiliate fraud as possible.

We seek to identify and count all the practices listed in Section 2.2. To test adware and loyalty software, we install those programs onto some of our virtual computers, allowing us to mimic the experience of users with those programs installed. Our automation classifies each occurrence with

the type of infraction and the victim merchant, allowing us to estimate the amount of fraud of each type targeting each merchant.

We check for affiliate fraud targeting every merchant using the three largest US affiliate networks as of February 2012 (4,523 merchants in total). We collected all data during February-March 2012. All told, our automation ran more than 2 million page-loads, finding 18,264 distinct observations in which 4,815 rogue affiliates targeted 2,446 merchants. Table 1 presents summary statistics about our data, and Table 2 tabulates the merchants with various numbers of affiliate violations.

When searching intensively for affiliate fraud targeting an individual merchant, our automation spends hundreds or thousands of computer-hours examining the various mechanisms which affiliates might use to target that merchant. But with thousands of merchants to be tested for this project, capacity constraints limited us to briefer searches. Our data is thus best understood as a sample of the affiliate fraud targeting affected merchants.

The online appendix offers additional detail about our data collection systems.

Combining these data sources, we offer a measure of which merchants—with which management structure—suffer how much affiliate misbehavior and of which types.

5.3 Endogeneity concerns

The structure and sequence of merchant decision-making reduce the risk of endogeneity biasing our estimates. Our estimates would be biased if some merchants knew they were at greater risk of fraud, and if those merchants chose particular management structures with an eye to reducing fraud. But our discussions with practitioners indicate that few to no merchants choose management structure in light of merchant-specific information about fraud.

For one, fraud is not typically a primary concern when merchants choose to open an affiliate marketing program. Most merchants view affiliate marketing as a low-risk strategy for the reasons discussed in the introduction. Furthermore, most merchants' choice of management structure seems to reflect a primary focus on capability and cost: Merchants tell us that they most often choose in-house management if they already have suitable expertise on staff or if they deem network management too costly. Conversely, merchants indicate that they most often choose network or OPM management if they lack appropriate expertise and seek accelerated results. In supplemental

results, we attempted to predict merchants' choice of management structure using each merchant's category (in Alexa's taxonomy of web sites) along with controls for size (Alexa traffic rank) and network. A few category dummies were statistically significantly different from zero, but largely marginally so and only to an extent consistent with random chance (e.g. one in twenty category dummies significant at the 5% level). These regressions never predicted more than 2% of variance in merchants' choice of management structure.

Networks' statements confirm the view that the risk of fraud is not a primary impetus for the choice of management structure. Commission Junction offers a flyer and detail page presenting the benefits of its "full program management" offering. The flyer nowhere mentions any benefit of excluding unwanted affiliates, and the detail page mentions this service only in a sub-page reached via an additional click. These marketing materials do not encourage merchants to choose a particular management structure with an eye to excluding rogue affiliates or preventing fraud more generally. If CJ does not view fraud prevention as an important selling point when describing its management service, it is unlikely that this is a major factor influencing merchants' choice of management structure.

Finally, it appears unlikely that merchants possess special information about their merchant-specific risks of fraud at the time they choose a management structure. In general, the fraud we study can affect any merchant, and most merchants are affected roughly in proportion to the size of their affiliate programs. (One important exception is that the web's largest merchants are more vulnerable to untargeted cookie-stuffing, but this problem affects only a handful of exceptionally large merchants.) If a merchant has no information whatsoever about its individual fraud risk at the time when it chooses its management structure, it cannot choose its management structure to reduce fraud, ending the risk of endogeneity. If a merchant has some information (albeit partial or incomplete) about its individual fraud risk, bias might result to the extent that the merchant acts on that information. A merchant's concern about fraud generally, as well as its general desire to choose a management structure robust to fraud, would not bias our results so long as this concern is not correlated with a merchant's knowledge of its distinctive *vulnerability* to fraud.

In principle, endogeneity could also result from the sampling caused by our incomplete search for affiliate fraud (as discussed in Section 5.2). If some types of merchants are systematically targeted by affiliate fraud that is more skillfully concealed, our automation might fail to find those practices

and might conclude, incorrectly, that those merchants are not targeted at all. But within each type of affiliate fraud, most incidents are roughly similar in concealment. We find statistically significant relationships between management structure and prevalence of fraud even within a given type of affiliate fraud (see Table 6), which means that differences across types of fraud are not driving our results.

6 Results

6.1 Summary statistics

Table 1 reports summary statistics for the affiliate practices we observed. Table 2 reports that some merchants suffer much more affiliate fraud than others: For nearly half the merchants we checked, we found no affiliate fraud at all, but for the most-targeted merchants, we found dozens of instances of affiliate fraud.

Table 3 tabulates merchant management structure, both on an overall basis and for specific networks. Merchant-managed programs are most common at all the networks we examine, although we are unable to identify the management structure of approximately 38% of merchants' affiliate programs. (Recall the data limitations discussed in Section 5.1.)

Table 4 reports the average number of affiliate fraud observations we found, by activity type and by network. The listed practices are largely comparable in their prevalence across networks. Google Affiliate Network has the least fraud of each type. Table 4 also totals the number of affiliate IDs engaged in each practice, across all merchants and by network.

6.2 Estimation framework

We now turn to our main results which estimate the effect of management structure on affiliate fraud. We run regressions of the following structure:

$$fraud_{ij} = \alpha + \sum_{k \in \left\{ \begin{smallmatrix} network \\ OPM \\ unknown \end{smallmatrix} \right\}} \beta_k I(management_i = k) + controls_i + \epsilon_i \qquad (1)$$

Here, i indexes merchants. In some specifications, j indexes types of affiliate fraud, required

for separately analyzing various types of affiliate fraud. k indexes types of merchant management structure, with in-house management as the omitted type to which others are compared.

In some specifications, we add controls for merchant characteristics. We control for merchant site popularity via a polynomial in merchant site Alexa traffic rank. (We add higher-order polynomial terms as instructed by Ramsey (1969).) We control for merchant site type using a set of dummy variables for each top-level category in Alexa's taxonomy of web sites. We control for possible differing rates of fraud across affiliate networks by adding a dummy variable for each network. Because affiliates might find it more profitable to defraud the merchants that pay larger commissions, we control for merchant payout per click (Earnings Per Click or "EPC") as reported by affiliate networks' detail pages for each merchant.[1]

We run all regressions using a negative binomial model. The $fraud_{ij}$ variable gives the number of times we observed fraudulent affiliates targeting merchant i. Holding constant the details of a merchant and its management structure, we might think of the merchant accepting each of some large number of affiliate applicants with some constant probability—matching the structure of the negative binomial distribution. (Cameron and Trivedi, 1998) Results are qualitatively similar when we run the estimation using Ordinary Least Squares regression, but OLS is misspecified in that the number of incidents of affiliate fraud (of each type, for each merchant) is a count variable, necesarily nonnegative and better modeled by the negative binomial distribution.

Throughout, our analysis uses a dependent variable of the number of *observations* of affiliate fraud. If we observed a given affiliate engaging in a listed practice multiple times—perhaps buying so much adware traffic that our crawlers observed the affiliate repeatedly even in limited testing, or typosquatting using multiple domains—then that affiliate counts multiple times in the listed variable. This approach captures a portion of variation in affiliate size: an affiliate engaged in a large-scale activity, for example widespread use of adware or numerous typosquatting domains, harms a merchant more than an affiliate whose behavior is more limited. In results not reported here, we also run all analysis at the level of distinct affiliates. Results are qualitatively similar, although some coefficient estimates shift in statistical significance.

[1] Due to a data collection error, we failed to collect contemporaneous EPC data for LinkShare merchants. LinkShare merchant EPC's are therefore absorbed into the LinkShare dummy variable.

6.3 Effect of management structure on affiliate fraud

Table 5 reports regression results summing across all types of affiliate fraud. With and without controls for merchant site popularity, networks suffer more fraud than programs managed by in-house staff. Note the positive coefficient on "Managed by Network" in all specifications of Table 5.

Table 5 finds no statistically significant difference in fraud rates when programs are run by merchants' in-house staff versus by outsourced specialists OPMs. In one specification, the difference is slightly positive, and in another it is slightly negative, but neither is statistically significantly different from zero.

6.4 Interactions between management structure and type of affiliate fraud

Table 6 reports results separated by the type of affiliate fraud observed. Networks have less adware than in-house-managed programs (column 1), denoted by the negative coefficient on Managed by Network. But networks have more typosquatting (column 3), and the difference between network management and in-house management is not significant for cookie-stuffing and loyalty apps (columns 2 and 4).

Tables 6 and 7 offer insight into the mechanism causing in-house programs to suffer, on the whole, less fraud (as found in Section 6.3). Recall that relevant practitioners regard cookie-stuffing and adware as clear violations. Table 6 columns 1 and 2 (aggregated in Table 7 column 1) report that networks are best at detecting these behaviors. Meanwhile, practitioners have not reached a clear consensus on typosquatting and loyalty applications, and Table 6 columns 3 and 4 (aggregated in Table 7 column 2) report that in-house programs are better at excluding those "grey area" practices. Thus, Table 6 indicates that network management best excludes clear-cut violations while in-house management better detects grey area violations. Because grey area violations are considerably more widespread (per Table 4), the overall effect is that in-house-managed programs are more successful than networks at excluding fraud.

Tables 6 and 7 offer a favorable evaluation of OPM efforts. In every specification, OPMs either offer statistically significantly less fraud than in-house management, or an amount statistically indistinguishable from in-house-managed programs. These results seem to confirm the desirability of the OPM approach: OPMs enjoy significant specialization (and hence improved information

compared to what is typically available to merchants), without the stark incentive problems of network management.

6.5 Interpreting coefficient estimates

The estimated coefficients in our regression analyses are modest in magnitude. For example, Table 5 reports that a merchant that shifts from in-house to network management would likely suffer 0.19 to 0.41 more observations of affiliate fraud found via our search methodology, i.e. less than one additional fraudulent affiliate. Though this sounds like a small effect, we believe it is nonetheless economically significant.

For one, our data collection process necessarily uncovers only a small portion of affiliate fraud. As discussed in Section 5.2, our automation expends a limited amount of time searching for practices targeting each individual merchant; with thousands of merchants to evaluate, it is infeasible for our crawlers to find all the fraudulent affiliates targeting all merchants. Based on our long-term examinations of affiliate fraud targeting selected merchants, we estimate that the data analyzed in this paper considers at most one tenth of merchants' affiliate fraud. For example, while our automation found on average 1.1 fraudulent affiliates for each merchant we examined, in separate focused searching, we usually find ten or more fraudulent affiliates in the first year of work. Furthermore, our automation found at most 14 distinct fraudulent affiliates targeting the most-targeted merchant in our data, but in our long-term work with merchants, we have found hundreds of affiliate frauds targeting some merchants. If our crawler's preliminary investigations detected one tenth of each merchant's affiliate fraud, then our estimated affiliate fraud counts should be increased by a factor of ten to estimate the true quantity of affiliate fraud. This adjustment affects interpretation of the magnitude of affiliate fraud but does not alter our estimation of the factors affecting the prevalence of the problem.

Meanwhile, even a single instance of affiliate fraud can be costly. Edelman (2012) reports individual affiliate frauds that reach hundreds of thousands, millions, or even tens of millions of dollars. Fraudulent affiliates are often among a merchant's largest affiliates. Indeed, eBay's losses to Hogan and Dunning totalled $21 million, and they were previously eBay's largest and second-largest affiliates. Similarly, affiliate network LinkShare often grants awards to its fastest-growing affiliates, but in three successive years received proof that winners were engaged in cookie-stuffing.

(Fadner, 2004) Our analysis would greatly benefit from weighting our observations with data about each affiliate's earnings, but networks consider affiliate earning data to be confidential, even when it would help discredit fraudulent affiliates. Affiliate earnings data therefore is not available to us.

7 Conclusion

Seeing all manner of affiliate malfeasance, a merchant might reasonably question whether affiliate marketing is worth pursuing. Despite the problems, we are convinced that affiliate marketing fills a genuine need. For one, affiliate marketing allows a merchant to more confidently advertise via the Internet's many small publishers, even if the merchant would hesitate to buy banner ads or syndicated search ads on little-known sites. Furthermore, affiliate marketing allows little-known publishers to accept greater risk in order to prove their efficacy. If a publisher is confident in the quality of its site and the likely purchases of its visitors, the publisher might reasonably prefer large payments if users make purchases, rather than far smaller payments for ad views or clicks.

For merchants that resolve to pursue affiliate marketing with full knowledge of what can go wrong, our analysis suggests suitable responses to typical vulnerabilities. If the merchant prefers to manage its affiliate program using in-house staff, the merchant may want to encourage its affiliate program manager to take special steps to learn affiliates' practices—perhaps through more detailed inquiries on affiliate intake questionnaires, online discussion forums to share information with counterparts, or extra efforts to attend conferences with other affiliate program managers. Meanwhile, if the merchant chooses to delegate management duties to an outsourced specialist or to network staff, the merchant should be particularly clear in its statements about which practices the program will permit. The merchant should not assume outsourced managers will act in the merchant's interest; quite the contrary, our data suggests that they often will not. Our data also provides some grounds to prefer in-house affiliate program management over network management, and if a merchant seeks the convenience of outsourced management, OPM management may be worth a look. If a merchant nonetheless chooses to proceed with a network, perhaps due to other advantages the network can offer, our analysis suggests that particularly clear and explicit rules can help restrain the network's actions in order to protect the merchant's interests.

Concerned merchants might reasonably look to the legal system—both to recover losses and

to deter infractions in the first place. But legal remedies seem to offer limited protection against affiliate misbehavior. In the cases summarized in Edelman (2012), merchants largely recovered most of the fees they had paid to the rogue affiliates at issue, but there is no suggestion that merchants recovered the large transaction costs such as attorneys fees, technical experts, and distraction of management from their core businesses. Moreover, the disputes that end up in litigation are highly unrepresentative—limited to instances where a merchant realized it was defrauded, was able to find the perpetrator, and anticipated that bringing suit would yield a recovery sufficient to justify the effort. We are aware of literally hundreds of incidents of merchants accepting affiliate fraud as "unavoidable," largely because these factors were not met. Merchants may be correct ex post, but if alternative marketing management structures could reduce such frauds, the size of merchants' losses suggests that such efforts would be cost-effective.

Although we focus on malfeasance in the context of affiliate marketing, similar problems extend to other forms of online advertising. Advertisers buying search engine advertising tend to focus on questions of bidding and targeting, but search syndication networks also place ads in all manner of sites, including sites that are highly undesirable. (Edelman, 2005b, 2009, 2010b) Similarly, display ads risk placements in invisible windows (Edelman, 2010a), in locations covered with other ads (Edelman, 2006), and via automatic reloads (Edelman, 2006), among other infractions. Uncovering and resolving these problems calls for diverse skills as close to computer forensics and law enforcement as to marketing and advertising—a marked change from the simpler contracts and better-understood risks associated with advertising in other media.

Tables

Table 1: Data Overview

Number of merchants examined in our testing	4523
Distinct affiliates observed engaged in the listed practices	4815
Observations of the listed practices	18264
Merchant with most ...	
observations of affiliates engaging in the listed practices	Travelocity (119 observations)
distinct affiliate IDs observed engaging in the listed practices	Logitech (14 affiliate IDs)

Table 2: Affiliate Fraud Observations by Merchant - by Network

	Number of merchants with the specified number of observations of affiliate fraud			
	LinkShare	Commission Junction	Google Affiliate Network	Total
$= 0$	401	1205	471	2077
≥ 1	353	1697	396	2446
≥ 5	158	695	179	1032
≥ 10	96	399	114	609
≥ 20	35	126	33	194
≥ 40	10	33	12	55
≥ 80	3	5	4	12

Table 3: Merchant Management Structure - by Network

	LinkShare	Commission Junction	Google Affiliate Network	Total
Managed by Merchant	401	900	441	1742
Managed by Network	63	198	124	385
Managed by OPM	81	409	182	672
Management Unknown	208	1387	119	1714
... with blank email	138	1228	32	1398
... due to role account	56	122	67	245
Total	754	2902	867	4523

Table 4: Affiliate Fraud Incidence Rate - by Network

	LinkShare	Commission Junction	Google Affiliate Network	Overall
Adware	0.859	0.508	0.355	0.537
	211	680	135	1026
Cookie-stuffing	0.103	0.081	0.042	0.077
	50	186	31	267
Typosquatting	2.869	3.346	3.258	3.250
	451	1790	495	2736
Loyalty apps	0.077	0.216	0.115	0.174
	58	628	100	786
Total	770	3284	761	4815

In each cell, the top value gives the fraud incidence rate per merchant (average number of such frauds per merchant). The bottom value gives the total number of affiliate-fraud incidents observed of that type, where one observation is one affiliate targeting one merchant. If an affiliate targets multiple merchants, those count separately.

Table 5: Effect of Management Structure on Affiliate Fraud - Total

	(1)	(2)	(3)	(4)	(5)
Managed by Network	0.410***	0.220**	0.192*	0.239**	0.250**
	(0.108)	(0.101)	(0.100)	(0.101)	(0.101)
Managed by OPM	-0.106	0.0113	0.113	0.119	0.120
	(0.0882)	(0.0835)	(0.0831)	(0.0836)	(0.0834)
Management Unknown	-0.367***	-0.232***	-0.170***	-0.248***	-0.242***
	(0.0665)	(0.0630)	(0.0628)	(0.0641)	(0.0642)
EPC Dummies					Yes
Network Dummies				Yes	Yes
Category Dummies			Yes	Yes	Yes
Site Popularity Controls		Yes	Yes	Yes	Yes
Constant	Yes	Yes	Yes	Yes	Yes
N	4523	4523	4523	4523	4523

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 6: Effect of Management Structure on Affiliate Fraud - by Type of Affiliate Fraud

	(1) Adware	(2) Cookie-stuffing	(3) Typosquatting	(4) Loyalty apps
Managed by Network	-0.375**	-0.278	0.338***	-0.00653
	(0.174)	(0.286)	(0.124)	(0.0523)
Managed by OPM	0.115	-0.574**	0.153	0.00766
	(0.140)	(0.266)	(0.102)	(0.0419)
Management Unknown	-0.200*	-0.191	-0.224***	-0.0514
	(0.107)	(0.184)	(0.0790)	(0.0328)
EPC Dummies	Yes	Yes	Yes	Yes
Network Dummies	Yes	Yes	Yes	Yes
Category Dummies	Yes	Yes	Yes	Yes
Site Popularity Controls	Yes	Yes	Yes	Yes
Constant	Yes	Yes	Yes	Yes
N	4523	4523	4523	4523

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 7: Effect of Management Structure on Affiliate Fraud - Clear / Grey Area

	(1) Clear Fraud	(2) Clear Fraud	(3) Grey Area	(4) Grey Area
Managed by Network	-0.321*	-0.245	0.295***	0.320***
	(0.164)	(0.164)	(0.111)	(0.110)
Managed by OPM	-0.172	0.0309	0.0436	0.142
	(0.135)	(0.137)	(0.0915)	(0.0911)
Management Unknown	-0.295***	-0.219**	-0.221***	-0.238***
	(0.101)	(0.104)	(0.0690)	(0.0703)
EPC Dummies		Yes		Yes
Network Dummies		Yes		Yes
Category Dummies		Yes		Yes
Site Popularity Controls	Yes	Yes	Yes	Yes
Constant	Yes	Yes	Yes	Yes
N	4523	4523	4523	4523

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

References

George P. Baker and Thomas Hubbard. Make versus buy in trucking: Asset ownership, job design, and information. *American Economic Review*, 93(3):551–572, September 2003.

Thomas Blake, Chris Nosko, and Steven Tadelis. Consumer heterogeneity and paid search effectiveness: a large scale field experiment. *mimeo*, 2013.

Adrian Cameron and Pravin Trivedi. *Regression Analysis of Count Data*. Cambridge University Press, 1998.

Economist. Pay per sale. September 2005.

Benjamin Edelman. Video: Ebates installed through security holes. 2004. URL http://www.benedelman.org/news/121504-1.html.

Benjamin Edelman. Debunking ShopAtHomeSelect. 2005a. URL http://www.benedelman.org/news/081105-1.html.

Benjamin Edelman. How Yahoo funds spyware. 2005b. URL http://www.benedelman.org/news/083105-1.html.

Benjamin Edelman. How Google and its partners inflate measured conversion rates and increase advertisers' costs. 2009. URL http://www.benedelman.org/news/051309-1.html.

Benjamin Edelman. Sony's Crackle: Invisible traffic galore. 2010a. URL http://www.benedelman.org/news/042710-1.html.

Benjamin Edelman. Google still charging advertisers for conversion-inflation traffic from WhenU spyware. 2010b. URL http://www.benedelman.org/news/010510-1.html.

Benjamin Edelman. Affiliate fraud litigation index. 2012. URL http://www.benedelman.org/affiliate-litigation/.

Benjamin Edelman. The design of online advertising markets. In Nir Vulkan, Alvin Roth, and Zvika Neeman, editors, *The Handbook of Market Design*. Oxford University Press, 2013.

Benjmain Edelman. Banner farms in the crosshairs. 2006. URL http://www.benedelman.org/news/061206-1.html.

Ross Fadner. For third time, LinkShare awards, revokes suspected fraudster. *MediaPost Online Media Daily*, October 20 2004.

Robert Glazer. Why cheap affiliate management does not work. *Adotas*, 2013a.

Robert Glazer. Why network-based affiliate management is a conflict of interest. *Acceleration Partners Articles & Insights*, 2013b.

Commission Junction. About CJ Access. URL http://www.cj.com/solutions/adv_access.jsp - archived page was posted from 2004-2006, and is now preserved by Archive.org, 2004.

Commission Junction. Pay for performance: CPA is in our DNA. 2014. URL http://www.cj.com/advertiser/pay-performance.

Francine Lafontaine and Margaret Slade. Vertical integration and firm boundaries: The evidence. *Journal of Economic Literature*, 45:629–685, September 2007.

Barak Libai, Eyal Biyalogorsky, and Eitan Gerstner. Setting referral fees in affiliate marketing. *Journal of Service Research*, 5:303–315, May 2003.

LinkShare. Affiliate Information. page present 2005 - 2009, 2009. URL http://web.archive.org/web/20090207094826/http://linkshare.com/affiliates/affiliates.shtml.

Lisa Miller. Report of service of search warrant for Brian Andrew Dunning. *U.S. District Court for the Northern District of California, CR 10-0494-EJD*, 2007. Document 49.

Tyler Moore and Benjamin Edelman. Measuring the perpetrators and funders of typosquatting. In Radu Sion, editor, *Financial Cryptography*, volume 6052 of *Lecture Notes in Computer Science*, pages 175–191. Springer, 2010.

J.B. Ramsey. Tests for specification errors in classical linear least squares regression analysis. *Journal of the Royal Statistical Society, Series B*, 31:350371, 1969.

United States of America v. Chiu. U.S. District Court for the Western District of Washington, CR12-070-RSM, 2012.

United States of America v. Dunning. U.S. District Court for the Northern District of California, CR 10-0494-EJD, 2010.

United States of America v. Hogan. U.S. District Court for the Northern District of California, CR 10-0495-JF, 2010.

Todd Walbridge. Report of service of search warrant for Shawn Dean Hogan. *U.S. District Court for the Northern District of California, CR 10-0495-JF*, 2007. Document 68-3.

Kenneth Wilbur and Yi Zhu. Click fraud. *Marketing Science*, 2:293–308, 2009.

Yi Zhu and Kenneth Wilbur. Hybrid advertising auctions. *Marketing Science*, 2:249–273, 2011.

Online appendix: data collection

This appendix extends Section 5.2 by providing additional detail about our data collection systems.

We collect data on affiliate fraud via direct observation: Our computer systems simulate users browsing the web in circumstances in which affiliate fraud is reasonably likely to occur. Our systems then monitor whether any affiliate fraud in fact does occur, and if so by which affiliate and targeting which merchant.

We customize our data collection systems to uncover each form of affiliate fraud itemized in Section 2.2. Key adjustments:

1. *Adware.* We install adware onto a virtual computer. Our automation then browses a merchant's site on this virtual computer. The adware observes the browsing and may redirect the user through an affiliate link and set affiliate cookies. We tested several adware programs, all of which are broadly similar in that they a) monitor the web sites and pages that users browse, then b) open popup or popunder windows. Often, the popup or popunder window loads an affiliate link promoting the same merchant the user was browsing in the first place—thereby claiming affiliate commission for the user's purchase, despite not genuinely doing anything to cause or encourage the user's purchase. Of the adware programs we tested, the most widespread and best-known is Zango.

2. *Cookie-stuffing.* We use search engines to obtain web page results related to a merchant—for example, the web pages that arise when searching for the merchant's name or domain plus the words coupons, deals, discount, or savings. On a clean virtual computer (without any adware or other nonstandard software), we load each such search result. If a web page engages in cookie-stuffing, the browser's request for that page will invoke an affiliate link and set affiliate cookies.

3. *Typosquatting.* We check the list of registered domains for domains that are similar to a merchant's domain name—for example, one or two characters away from the domain (insertions, deletions, or transpositions). On a clean virtual computer (without any adware or other nonstandard software), we load the root page of each such domain. If that domain engages in affiliate typosquatting, the browser's request for that page will invoke an affiliate link and set

affiliate cookies.

4. *Loyalty software.* We install loyalty software onto a virtual computer. Our automation then browses a merchant's site on this virtual computer. The loyalty software observes the browsing and may redirect the user through an affiliate link and set affiliate cookies. In this way, the loyalty software claims affiliate commission for the user's purchase, despite not genuinely causing the user's purchase. We tested a number of loyalty programs we tested; a representative example is BeeBucks which a) becomes installed via a bundle as users request other software, b) monitors users' browsing, and c) invokes affiliate links automatically.

We test all Commission Junction, Google Affiliate Network, and LinkShare merchants for each of these four problems, determining the amount of fraud affecting each.

During each test, our automation monitors all network traffic. If an application or site invokes an affiliate link, our automation sees that traffic flowing over the test computer's Internet connection, recognizes the affiliate link, and determines the targeted merchant.

We observe the number of instances of affiliate fraud and the number of distinct affiliate IDs involved. We do not observe the total amount paid to these affiliates; affiliate networks and merchants do not disclose this information to the public.

Since some affiliate fraud is known to target or avoid users in particular geographic locations (including as discussed in Section 2.4), we designed our data collection to be geographically diverse. We ran 150 virtual computers in 80 datacenters in nine countries.

We usually observe even those fraudsters whose attempts have already been discovered by networks or merchants: Most often, an affiliate's links continue to work even if a merchant or network uncovers a fraud and disables the affiliate's account.[2] Our automation can therefore determine the link destination without complication. If a network or merchant instead disables the affiliate's links, our systems do not associate that attempted fraud with a victim merchant, so we discard that attempt, and it does not appear in our data.

[2]Most merchants believe that their best choice, to avoid interrupting transactions or inconveniencing users, is to keep links functional even while withholding payment from the corresponding affiliates. Consider an adware application that covers a merchant's site with a popup loading an affiliate link. If the merchant keeps the link active, the user will continue to receive a redirect to the merchant's site. In contrast, if the merchant disables the link, the user will receive an error message. The merchant would prefer to stop the popups completely, but that requires cooperation from the adware vendor, which is outside the merchant's direct control.