

BENJAMIN EDELMAN and WESLEY BRANDI*

The authors 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, whereas other merchants ask their marketing staff to perform these functions. For clear violations of applicable rules, the authors find that outside specialists are the most effective at excluding the responsible affiliates, which can be interpreted 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), forgoing the efficiencies of specialization in favor of the better incentives of a company’s staff. The authors consider the implications for marketing of online affiliate programs and for online marketing more generally.

Keywords: affiliate marketing, fraud, marketing management, incentives, outsourcing

Online Supplement: <http://dx.doi.org/10.1509/jmr.13.0472>

Risk, Information, and Incentives in Online Affiliate Marketing

For decades—perhaps centuries—marketers have bemoaned the effectiveness of their advertising campaigns. When paying for advertising up front and receiving benefits later, advertisers are vulnerable to low-performing or nonperforming ad placements. Against this backdrop, affiliate marketing seems 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 can attribute sales to affiliates, advertisers often perceive an unprecedented reduction in risk. *The Economist* (2005) captured advertisers’ excitement for the apparent alignment of incentives, calling affiliate marketing “the holy grail of online advertising.”

As it turns out, however, affiliate marketing is neither as easy nor as safe as proponents initially anticipated. Most

advertisers struggle to find reliable affiliates that 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 they have not fairly earned. Informed by these problems, affiliate marketing raises long-standing questions of judgment, partnership, and incentives reminiscent of decades of media-buying.

This article offers two contributions. We begin by presenting affiliate marketing in a general sense; we explore its institutions and participants as well as key risks uncovered to date. Then, 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, whereas 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 the broader literature

*Benjamin Edelman is Associate Professor, Harvard Business School, Harvard University (e-mail: bedelman@hbs.edu). Wesley Brandi is Chief Executive Officer, iPensatori (e-mail: wesleyb@ipensatori.com). The authors 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 *JMR* anonymous reviewers for comments and suggestions on earlier drafts of this article.

on outsourcing and boundaries of the firm. Managers often face a trade-off between keeping functions in-house (typically with greater supervision and greater control over quality) versus outsourcing to a specialist (that may have greater capabilities or a cost advantage due 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, thus, unobservable to researchers. In contrast, we examine an online marketing context in which advertisers often reveal their management structures as they recruit marketing affiliates. We enjoy the additional advantage of a novel data set. Typically, both firms and researchers lack top-quality data on opportunistic behavior because those providing low-quality services usually attempt 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 developed custom software to examine affiliates' behavior—information often unavailable even to the advertisers and networks that purport to supervise these affiliates.

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 affiliate's 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 website serves an ad to the user, and almost all search ads require an advertiser to pay as soon as a user clicks on 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 on its ads: Perhaps the ads are irrelevant to users' interests or are placed in locations where few users notice them. Some of these factors are outside the advertiser's control; for example, 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: as long as a user actually makes a purchase, merchants usually 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 that LinkShare touted as "very efficient" (LinkShare 2009). Similarly, the affiliate network Commission Junction notes that "advertisers only pay when a specific action has been completed (e.g., a purchase...)" which, Commission Junction notes, makes affiliate marketing "low risk" (Commission Junction 2014). Although practitioners initially considered affiliate marketing structurally protected from fraud, there are nonetheless some significant risks, including the practices we examine subsequently.

The Institutions of Affiliate Marketing

An affiliate marketing "merchant" is the website aiming 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.com, to mom-and-pop specialty sites.

An "affiliate" or "publisher" is a website 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.com 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, however, some affiliates use prohibited practices, which we explore in the next section.

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 website 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.com since the inception of its affiliate program and, more recently, eBay and Apple). However, 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. In the section titled "Affiliate Program Management and Resulting Incentives," we explore the various models of affiliate program management in greater detail.

For our purposes, it is not necessary to explore the technology that facilitates affiliate program operations and tracking. The fundamental enabling feature is the browser "cookie," a technology that enables a website to place data on a user's computer to recognize a user on a later 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.

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 paid only if a user makes a purchase. Thus, if a rogue affiliate tries 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. However, 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 are 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 7–30 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 onward 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 thus incapable of claiming or receiving benefits. We note that there is some debate about whether loyalty software creates customer loyalty—and, if so, whether it is to the merchant or to the maker of the loyalty software. Nonetheless, we accept the term because 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 and 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 that 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 and then check for violations. However, 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 aforementioned behaviors listed, for which a merchant's interest is more clear-cut.

Practitioners' Views of Affiliate Fraud

Affiliate marketing practitioners have differing views on the practices presented in the previous section. In general, practitioners view adware as clearly impermissible, and it is specifically forbidden by most network contracts. They also 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 mistypes "expedia.com" (sic), 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 U.S.C. §1125(d)); that it entails "forcing clicks," which is 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 website, the user would be unlikely to shift a purchase to a competitor. Critics 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. However, standard contracts include no explicit prohibition on either typosquatting or loyalty software but 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 on a link 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 "gray area" in recognition of diverging views among relevant practitioners.

Selected Instances of Affiliate Fraud

In this section, we profile the largest publicly documented instances of affiliate fraud to communicate a sense of the perpetrators, their methods, and their detection. The largest and best-known affiliate fraud was perpetrated by Shawn Hogan, founder of an online advertising network that facilitated the placement of banner ads onto the websites of independent publishers. According to the indictment in *United States of America 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 that 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, approximately 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. However, the indictment and corresponding private eBay litigation allege that Hogan's referrals were entirely useless, consisting of 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 was sentenced to five months in federal prison, three months' probation, and a fine of \$25,000 in May 2014.

A similar indictment and private eBay claim were brought against Brian Dunning, also alleging invisible cookie stuffing (*United States of America 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 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 Commission Junction, albeit some months before the events at issue. The litigation docket reveals nothing further about Way's alleged involvement, leaving Way's true scope of involvement (if any) in Dunning's activities unclear. Dunning pled guilty in April 2013 and was sentenced to 15 months in federal prison in August 2014.

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, Commission Junction, and their investigators. Walbridge (2007, p. 3) also reports that Hogan admitted to stuffing cookies only once to each IP address, which prevented investigators from uncovering the practice through 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 (*United States of America v. Chiu* 2012).

When granting awards to the affiliates that achieved fastest growth, affiliate network LinkShare in three consecutive instances had to retract awards from recipients proved 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. Although 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, through our data collection methods (detailed subsequently in the "Data" section), we are able to identify numerous perpetrators as well as victim merchants.

RELATED LITERATURE

Because other advertising formats are significantly more established, one might ask why merchants would choose

affiliate marketing. The literature offers some insight. Libai, Bialogorsky, and Gerstner (2003) consider the similar context of publishers selling “leads” such as sign-ups 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 customers who actually make purchases.) We build on Libai, Bialogorsky, and Gerstner 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 their research 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, Nosko, and Tadelis (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 outside 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. Although we examine a different market, we note that affiliate marketing is also grounded in granular tracking—that is, improved information collection broadly similar to the trucking on-board computers that motivated Baker and Hubbard’s research. More broadly, Lafontaine and Slade (2007) survey research on the causes and consequences of firm boundaries across numerous sectors.

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 then *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 in-house 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 that most affiliate managers’ performance-based compensation is implicit, with a larger program viewed as calling for greater compensation. Of course, affiliate managers’ 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” (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 3 to 10 programs, while a large OPM could manage 100 programs or more. Contracts for OPMs vary widely as well: some stipulate flat fees (e.g., \$3,000 per month to manage a given program), while others offer a percentage (“20% of spend”), and some are hybrids. Staffing is also diverse for OPMs: some OPMs assign a full-time staff person to a single large merchant. Smaller programs typically share OPM managers: a single OPM staff person might 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 include Glazer (2013a, b).

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 regarding 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 these data through 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, whereas networks are certain to have an important element of *ad valorem* (proportional) compensation. Whereas 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, but the network also 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 \$.20. However, if the network takes that same action, the network receives additional revenue of \$.50. Because a network incurs minimal marginal cost to provide tracking services, its additional revenue is best understood as pure profit. Thus, networks have a notably stronger incentive to increase merchants' costs compared with the corresponding incentive for OPMs.

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

	<i>Incentive</i>	<i>Information</i>
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

Merchants' Choice of Affiliate Management Structure

The preceding section offers mixed recommendations to a merchant selecting a management structure for its affiliate program. On the one hand, merchants might focus on the importance of obtaining information about affiliates' practices. If information is the most important determinant of program success, 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.

On the other hand, one might worry about the incentives of affiliate program managers. If some merchants' programs accept undesirable affiliates because of 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 because of OPMs' mixed incentives.

In a more nuanced interpretation—what our data support, as we discuss subsequently (see the section "Interactions Between Management Structure and Type of Affiliate Fraud")—information and incentives interact to provide differing benefits for differing behaviors. As we discussed previously, affiliate malfeasance includes both practices that are understood to be clear-cut violations of applicable rules and gray-area practices that are contrary to merchants' interests and yet nonetheless sometimes accepted by practitioners. As to clear violations, the key barrier to taking action is information—that is, 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 gray-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 gray-area behaviors that are not truly in merchants' interest.

In light of networks' financial incentive to allow affiliate misbehavior, one might ask why a network-managed program rejects any affiliates at all. 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. Regarding gray-area violations, networks anticipate that such risks are consider-

ably reduced because a merchant would be less likely to end its use of a network in response to gray-area practices.

DATA

Merchant Management Structure

We begin with data regarding which merchants use which marketing structures. Each merchant using the largest three U.S. 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. Sixty-nine percent of merchants’ pages provide a contact e-mail address for affiliates with questions about a given affiliate program, while 31% of pages provide no e-mail address whatsoever.

We categorize merchants’ posted e-mail addresses to draw inferences about the management structure of each merchant’s affiliate program. For example, if the e-mail address is a named individual person at the merchant, we categorize that merchant as managing its own affiliate program. If the e-mail 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 e-mail 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 e-mail 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.

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 try to avoid being revealed as such, lest their accounts be closed and payments withheld.

This section offers an overview of our data collection. The Web 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 aim to identify and count all the practices we noted in the “Fraud in Affiliate Marketing” section. To test adware and loyalty software, we installed those programs onto

some of our virtual computers, thus enabling us to mimic the users’ experience with those programs. Our automation classifies each occurrence with the type of infraction and the victim merchant, enabling 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 U.S. 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 that 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. Thus, our data are best understood as a sample of the affiliate fraud targeting affected merchants. The Web Appendix offers additional details 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.

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 to reduce the instances of

Table 1
DATA OVERVIEW

	Totals
Number of merchants examined in our testing	4,523
Distinct affiliates observed engaged in the listed practices	4,815
Observations of the listed practices	18,264
<i>Merchant with most ...</i>	
...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

Observations of Affiliate Fraud	Merchants			Total
	LinkShare	Commission Junction	Google Affiliate Network	
0	401	1,205	471	2,077
≥1	353	1,697	396	2,446
≥5	158	695	179	1,032
≥10	96	399	114	609
≥20	35	126	33	194
≥40	10	33	12	55
≥80	3	5	4	12

fraud. However, our discussions with practitioners indicate that few to no merchants choose the management structure in light of merchant-specific information about fraud.

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 previously. Furthermore, most merchants' choice of management structure seems to reflect a primary focus on capability and cost: merchants report 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 websites) 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., 1 in 20 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. Nowhere does the flyer mention any benefit of excluding unwanted affiliates, and the detail page mentions this service only in a subpage reached through 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 Commission Junction does not view fraud prevention as an important selling point when describing its management service, it is unlikely that it is a major factor influencing merchants' choice of management structure.

Finally, it seems 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 about its individual fraud risk at the time when it chooses its management structure, it cannot choose its management structure to reduce fraud, thus 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 in general, as well as its general desire to choose a management structure robust to fraud, would not bias our results as 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 we discuss in the "Affiliate Practices" section). 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. However, within each type of affiliate fraud, most incidents are approximately similar in concealment. We find statistically significant relationships between management structure and prevalence of fraud even within a given type of affiliate fraud (as we discuss subsequently), which means that differences across types of fraud are not driving our results.

RESULTS

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 we discussed in the "Merchant Management Structure" section.)

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 preva-

Table 3
MERCHANT MANAGEMENT STRUCTURE, BY NETWORK

	LinkShare	Commission Junction	Google Affiliate Network	Total
Managed by merchant	401	900	441	1,742
Managed by network	63	198	124	385
Managed by OPM	81	409	182	672
Management unknown	208	1,387	119	1,714
...with blank email	138	1,228	32	1,398
...due to role account	56	122	67	245
Total	754	2,902	867	4,523

Table 4
AFFILIATE FRAUD INCIDENCE RATE, BY NETWORK

	LinkShare	Commission Junction	Google Affiliate Network	Overall
Adware	.859	.508	.355	.537
	211	680	135	1026
Cookie stuffing	.103	.081	.042	.077
	50	186	31	267
Typosquatting	2.869	3.346	3.258	3.250
	451	1,790	495	2,736
Loyalty apps	.077	.216	.115	.174
	58	628	100	786
Total	770	3,284	761	4,815

Notes: 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, in which one observation is one affiliate targeting one merchant. If an affiliate targets multiple merchants, those incidents count separately.

lence 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.

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:

$$(1) \text{fraud}_{ij} = \alpha + \sum_{k \in \left\{ \begin{array}{l} \text{network} \\ \text{OPM} \\ \text{unknown} \end{array} \right\}} \beta_k I(\text{management}_i = k) + \text{controls} + \varepsilon_i.$$

Here, i indexes merchants. In some specifications, j indexes types of affiliate fraud, which is required for separately analyzing various types of affiliate fraud. Finally, 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 through 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 websites. 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 [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). The results are qualitatively similar when we run the estimation using ordinary least squares regression, but the number of incidents of affiliate fraud (of each type, for each merchant) is a count variable, necessarily 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, the 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 analyses at the level of distinct affiliates. The results are qualitatively similar, although some coefficient estimates shift in statistical significance.

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 shows 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.

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

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

Table 5
EFFECT OF MANAGEMENT STRUCTURE ON AFFILIATE FRAUD, TOTAL

	(1)	(2)	(3)	(4)	(5)
Managed by network	.410*** (.108)	.220** (.101)	.192* (.100)	.239** (.101)	.250** (.101)
Managed by OPM	-.106 (.0882)	.0113 (.0835)	.113 (.0831)	.119 (.0836)	.120 (.0834)
Management unknown	-.367*** (.0665)	-.232*** (.0630)	-.170*** (.0628)	-.248*** (.0641)	-.242*** (.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	4,523	4,523	4,523	4,523	4,523

* $p < .10$.

** $p < .05$.

*** $p < .01$.

Notes: Standard errors are in parentheses.

Table 6
EFFECT OF MANAGEMENT STRUCTURE ON AFFILIATE FRAUD, BY TYPE OF AFFILIATE FRAUD

	(1) <i>Adware</i>	(2) <i>Cookie Stuffing</i>	(3) <i>Typosquatting</i>	(4) <i>Loyalty Apps</i>
Managed by network	-.375** (.174)	-.278 (.286)	.338*** (.124)	-.00653 (.0523)
Managed by OPM	.115 (.140)	-.574** (.266)	.153 (.102)	.00766 (.0419)
Management unknown	-.200* (.107)	-.191 (.184)	-.224*** (.0790)	-.0514 (.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	4,523	4,523	4,523	4,523

* $p < .10$.

** $p < .05$.

*** $p < .01$.

Notes: Standard errors are in parentheses.

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 we found in the “Effect of Management Structure on Affiliate Fraud” section). Recall that relevant practitioners regard cookie stuffing and adware as clear violations. Columns 1 and 2 of Table 6 (aggregated in column 1 of Table 7) report that networks are best at detecting these behaviors. Meanwhile, practitioners have not reached a clear consensus on typosquatting and loyalty applications, and columns 3 and 4 of Table 6 (aggregated in column 2 of Table 7) report that in-house programs are better at excluding those gray-area practices. Thus, Table 6 indicates that network management best excludes clear-cut violations, whereas in-house management better detects gray-area violations. Because gray-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 offer either 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 thus improved information compared with what is typically available to merchants), without the stark incentive problems of network management.

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 .19 to .41 more observations of affiliate fraud found using our search methodology (i.e., less than one additional fraudulent affiliate). Although this sounds like a small effect, we believe that it is nonetheless economically significant.

First, our data collection process necessarily uncovers only a small portion of affiliate fraud. As we discussed in

Table 7
EFFECT OF MANAGEMENT STRUCTURE ON AFFILIATE FRAUD, CLEAR/GRAY AREA

	(1) <i>Clear Fraud</i>	(2) <i>Clear Fraud</i>	(3) <i>Gray Area</i>	(4) <i>Gray Area</i>
Managed by network	-.321* (.164)	-.245 (.164)	.295*** (.111)	.320*** (.110)
Managed by OPM	-.172 (.135)	.0309 (.137)	.0436 (.0915)	.142 (.0911)
Management unknown	-.295*** (.101)	-.219** (.104)	-.221*** (.0690)	-.238*** (.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	4,523	4,523	4,523	4,523

* $p < .10$.

** $p < .05$.

*** $p < .01$.

Notes: Standard errors are in parentheses.

the “Affiliate Practices” section, 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. Drawing on our long-term examinations of affiliate fraud targeting selected merchants, we estimate that the data analyzed in this article reflect, at most, one-tenth of merchants’ affiliate fraud. For example, whereas our automation found an average of 1.1 fraudulent affiliates for each merchant examined, in separate focused searching, we usually found 10 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 instances of affiliate fraud targeting some merchants. If our crawler’s preliminary investigations detected one-tenth of each merchant’s affiliate fraud, our estimated affiliate fraud counts should be increased by a factor of ten to estimate the true extent 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 instances of affiliate fraud 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 totaled \$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 it received proof that winners were engaged in cookie stuffing (Fadner 2004). Our analysis would greatly benefit from weighting our observations with data on each affiliate’s earnings, but networks consider affiliate earning data confidential, even when it would help discredit fraudulent affiliates. Therefore, affiliate earnings data are not available to us.

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. First, affiliate marketing enables a merchant to more confidently advertise with 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 enables little-known publishers to accept greater risk 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, it 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 informa-

tion 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 that outsourced managers will act in the merchant’s interest; on the contrary, our data suggest that they often will not. Our data also provide 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 considering. 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 to protect the merchant’s interests.

Concerned merchants might reasonably rely on the legal system both to recover losses and to deter infractions in the first place. However, legal remedies seem to offer limited protection against affiliate misbehavior. In the cases that Edelman (2012) summarizes, 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 fees to attorneys and technical experts, not to mention the distraction of management from their core businesses. Moreover, disputes that end up in litigation are highly unrepresentative—limited to instances in which 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 kinds of sites, including those that are highly undesirable (Edelman 2005b, 2009, 2010a). Similarly, display ads risk placements in invisible windows (Edelman 2010b), in locations covered with other ads (Edelman 2006), and through 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.

REFERENCES

- Baker, George P. and Thomas Hubbard (2003), “Make Versus Buy in Trucking: Asset Ownership, Job Design, and Information,” *American Economic Review*, 93 (3), 551–72.
- Blake, Thomas, Chris Nosko, and Steven Tadelis (2013), “Consumer Heterogeneity and Paid Search Effectiveness: A Large Scale Field Experiment,” mimeo, eBay Research Labs.
- Cameron, Adrian and Pravin Trivedi (1998), *Regression Analysis of Count Data*. Cambridge, UK: Cambridge University Press.

- Commission Junction (2004), "About CJ Access," [http://www.cj.com/solutions/adv_access.jsp], archived page was posted from 2004–2006, and is now preserved by Archive.org [http://web.archive.org/web/*/http://www.cj.com/solutions/adv_access.jsp].
- (2014), "Pay for Performance: CPA Is in Our DNA," (accessed October 30, 2014), [available at http://www.cj.com/advertiser/pay-performance].
- The Economist* (2005), "Pay per Sale," (September 29), (accessed October 30, 2014), [available at http://www.economist.com/node/4462811].
- Edelman, Benjamin (2004), "Ebates Installed Through Security Holes," video, (December 15), [available at http://www.benedelman.org/news/121504-1.html].
- (2005a), "Debunking ShopAtHomeSelect," (October 14), [available at http://www.benedelman.org/news/081105-1.html].
- (2005b), "How Yahoo Funds Spyware," (September 5), [available at http://www.benedelman.org/news/083105-1.html].
- (2006), "Banner Farms in the Crosshairs," (June 23), [available at http://www.benedelman.org/news/061206-1.html].
- (2009), "How Google and Its Partners Inflate Measured Conversion Rates and Increase Advertisers' Costs," (May 13), [available at http://www.benedelman.org/news/051309-1.html].
- (2010a), "Google Still Charging Advertisers for Conversion-Inflation Traffic from WhenU Spyware," (January 7), [available at http://www.benedelman.org/news/010510-1.html].
- (2010b), "Sony's Crackle: Invisible Traffic Galore," (April 27), [available at http://www.benedelman.org/news/042710-1.html].
- (2012), "Affiliate Fraud Litigation Index," [available at http://www.benedelman.org/affiliate-litigation/].
- (2013), "The Design of Online Advertising Markets," in *The Handbook of Market Design*, Nir Vulkan, Alvin Roth, and Zvika Neeman, eds. Oxford, UK: Oxford University Press.
- Fadner, Ross (2004), "For Third Time, LinkShare Awards, Revokes Suspected Fraudster," MediaPost, (October 20), (accessed October 30, 2014), [available at http://www.mediapost.com/publications/article/19661/for-third-time-linkshare-awards-revokes-suspecte.html?edition].
- Glazer, Robert (2013a), "Why Cheap Affiliate Management Does Not Work," Adotas, (May 23), (accessed October 30, 2014), [available at http://www.adotas.com/2013/05/why-cheap-affiliate-management-does-not-work/].
- (2013b), "Why Network-Based Affiliate Management Is a Conflict of Interest," Acceleration Partners, (accessed October 30, 2014), [available at http://www.accelerationpartners.com/blog/Why-Network-Based-Affiliate-Management-Is-A-Conflict-Of-Interest].
- Lafontaine, Francine and Margaret Slade (2007), "Vertical Integration and Firm Boundaries: The Evidence," *Journal of Economic Literature*, 45 (September), 629–85.
- Libai, Barak, Eyal Biyalogorsky, and Eitan Gerstner (2003), "Setting Referral Fees in Affiliate Marketing," *Journal of Service Research*, 5 (May), 303–315.
- LinkShare (2009), "Affiliate Information," [http://web.archive.org/web/20090207094826/http://linkshare.com/affiliates/affiliates.shtml], page present 2005–2009.
- Miller, Lisa (2007), "Report of Service of Search Warrant for Brian Andrew Dunning," U.S. District Court for the Northern District of California, CR 10-0494-EJD, Document 49.
- Moore, Tyler and Benjamin Edelman (2010), "Measuring the Perpetrators and Funders of Typosquatting," in *Financial Cryptography: Lecture Notes in Computer Science*, Vol. 6052, Radu Sion, ed. Berlin: Springer, 175–91.
- Ramsey, J.B. (1969), "Tests for Specification Errors in Classical Linear Least Squares Regression Analysis," *Journal of the Royal Statistical Society, Series B*, 31 (2), 350–71.
- United States of America v. Chiu* (2012), U.S. District Court for the Western District of Washington, CR12-070-RSM.
- United States of America v. Dunning* (2010), U.S. District Court for the Northern District of California, CR 10-0494-EJD.
- United States of America v. Hogan* (2010), U.S. District Court for the Northern District of California, CR 10-0495-JF.
- Walbridge, Todd (2007), "Report of Service of Search Warrant for Shawn Dean Hogan," U.S. District Court for the Northern District of California, CR 10-0495-JF, Document 68-3.
- Wilbur, Kenneth and Yi Zhu (2009), "Click Fraud," *Marketing Science*, 28 (2), 293–308.
- Zhu, Yi and Kenneth Wilbur (2011), "Hybrid Advertising Auctions," *Marketing Science*, 30 (2), 249–73.