Deterring Online Advertising Fraud Through Optimal Payment in Arrears

Benjamin Edelman

Harvard Business School
bedelman@hbs.edu

Abstract. Online advertisers face substantial difficulty in selecting and supervising small advertising partners: Fraud can be well-hidden, and limited reputation systems reduce accountability. But partners are not paid until after their work is complete, and advertisers can extend this delay both to improve detection of improper partner practices and to punish partners who turn out to be rule-breakers. I capture these relationships in a screening model with delayed payments and probabilistic delayed observation of agents' types. I derive conditions in which an advertising principal can set its payment delay to deter rogue agents and to attract solely or primarily good-type agents. Through the savings from excluding rogue agents, the principal can increase its profits while offering increased payments to good-type agents. I estimate that a leading affiliate network could have invoked an optimal payment delay to eliminate 71% of fraud without decreasing profit.

Keywords: online advertising, screening, signaling, contracts, fraud

1 Introduction

When buying online advertising, principals often seek to contract with agents of unknown quality – often thousands of sites on which ads are to be shown, or thousands of affiliates who are to be paid for promotional methods they devise themselves. Ex ante, it is difficult to assess agent quality or to predict which agents will perform unfavorably. Moreover, it is often impractical to extract a penalty from agents ultimately deemed to be nonproductive. These constraints challenge advertisers and ad networks that seek to reduce marketing fraud and to control the presentation of their offers.

Advertisers' evaluation of marketing partners generally mirrors the task of an employer screening prospective employees, as in Spence's defining work on signaling in labor markets [26]. In particular, just as *Spence* employers cannot observe employee productivity, so too are advertisers unable to foresee marketing partners' practices. But online advertisers benefit from two important capabilities beyond *Spence* employers: First, an advertising principal pays its agents "in arrears" – that is, at some time after each agent completes its work. Second, in each period a principal has some positive probability of learning that an agent is engaged in impermissible ("rogue") marketing practices (if in fact the agent is engaged in such practices).

Under conditions derived below, a principal can delay all agents' payments in order to deter rogue agents' participation. Meanwhile, by paying good agents to compensate them for the delay, the principal can make itself and the good agents strictly better off.

In Section 2, I present the relevant characteristics of the online advertising industry. In Section 3, I develop a model of a principal paying agents in arrears, and I derive circumstances in which the principal and good-type agents prefer to delay payments. In Section 4, I apply this model to online advertising markets, and I estimate the benefits a leading affiliate network would have achieved by optimally delaying payment. In Section 5, I compare online advertising to other relevant contexts.

2 Principal-Agent Problems in Internet Advertising

The market for Internet advertising features large advertising principals (ad networks and major advertisers) contracting with a numerous small advertising agents such as web sites, blogs, search syndicators, and other marketing partners. For example, affiliate network LinkShare boasts more than a million affiliates promoting offers from the network's hundreds of merchants [22]. Google contracts with an unknown but large number of independent web sites (at least hundreds of thousands) to include its AdSense ad frames [18].

2.1 Rogue agents in Internet advertising

Some Internet advertising agents claim payments they have not truly earned. Consider a search engine that places ads onto a syndicator's web site. The syndicator can increase its revenue by clicking the ads on its own site – *click fraud*, in that the associated clicks come from the syndicator rather than from bona fide users. In principle a search engine might manage to identify telltale signs of click fraud, e.g. many clicks coming from a single PC. But in practice, perpetrators disguise their efforts, *i.e.* through the use of *botnets* or others systems to submit fake clicks from a large number of computers. *See e.g.* [3], [7].

On one view, a necessary condition for click fraud is that a click is easy to fake: A robot can "click" an ad just as easily as a human can. How better to distinguish bona fide visitors from robots and fakes? One possibility is to measure something more fundamental: Rather than measuring clicks, measure users' actual purchases. Indeed, some advertising intermediaries, *affiliate networks*, promise to charge advertising fees only when 1) a participating affiliate presents a user with a special tracking link to a merchant's web site, 2) the user clicks that link, and 3) the user subsequently makes a purchase from the corresponding merchant. LinkShare, a leading affiliate network, touts its service as requiring a merchant to "pay only when a sale … is completed" [20] – emphasizing reduction in an advertiser's supposed risk.

LinkShare correctly points out that fake clicks (*i.e.* click fraud) do not, in and of themselves, garner payment. But rogue affiliates nonetheless find ways to defraud advertisers. For example, some affiliates perform *cookie-stuffing* to claim commission without a user clicking an affiliate tracking link. Consider a popular merchant from which a large proportion of users make a purchase in a given month – say, Amazon. If a user clicks an affiliate link to Amazon, then makes a purchase from Amazon anytime within the next thirty days, Amazon intends to pay a commission to the corresponding affiliate. But a rogue affiliate could modify a web page, banner advertisement, or email so that merely viewing that page, ad, or email would "click" the affiliate's link to Amazon – thereby crediting the affiliate for any Amazon purchases the user makes within the next thirty days. ([10, 14], [13] - exhibit 70) Through such tactics, an affiliate can trick a merchant into paying a commission when in fact the affiliate did nothing to promote the merchant, and when in fact no commission is due.

In an alternate attack on affiliate marketing tracking systems, an affiliate installs (or pays a partner to use) tracking software on a user's computer – software typically known as spyware (for its intrusive tracking of web site visits) or adware (for its display of pop-up ads). This tracking software monitors what merchant web sites a user visits, then opens affiliate links to the corresponding merchants. For example, if the user browses Dell.com, the software invokes its affiliate link to Dell. If the user then makes a purchase from Dell, Dell would mistakenly conclude the affiliate had referred the transaction, and Dell would pay commission accordingly [9, 11, 15].

2.2 The practical unavailability of the legal system in typical disputes between advertising principals and agents

In general, rogue advertising agents breach their contracts with advertising principals when they fake clicks, stuff cookies, or otherwise overcharge advertising principals. Upon uncovering such a breach of contract, a principal could file suit to demand redress and to prevent future violations. But in practice, the legal system is effectively unavailable in many disputes with advertising agents.

For one, transaction costs (including attorney fees and management time) tend to exceed the amount of harm cause by any single agent. Transaction costs are particularly weighty given the technical complexity of the violations, the absence of physical evidence, and the lack of expertise among investigators, attorneys, and arbiters. Furthermore, rogue agents are dispersed around the world, inviting jurisdictional disputes and increasing litigation costs. (*See e.g.* [28], reporting rule-breaking affiliates on four continents.)

Even when agents can be identified cost-effectively, agents often lack the resources to make principals whole. Some agents abscond with their ill-gotten gains. Others conceal their wealth in stores of value that are difficult for investigators to uncover [19]. Furthermore, bankruptcy laws let some rogue agents shelter assets in homesteads or in other assets that principals cannot seize [24].

Institutional factors further deter some advertising principals from pursuing rogue agents. For example, a principal may be embarrassed to admit to the public, in open court and in the public record, that it was defrauded. (*See e.g.* [5], questioning why Google declined to pursue a click fraud perpetrator.) Embarrassment is particularly pronounced in those circumstances that survive transaction cost analysis: There is special reason to be embarrassed when a perpetrator successfully stole a large amount of money. Revealing a fraud, even for purposes of achieving redress, could undermine confidence in a network or advertiser: Consumers might not want to buy from a merchant they learn has been cheated. (Consumers might worry that if rogue advertising partners defrauded the merchant, perhaps credit card information isn't safe either.) Similarly, advertisers might not want to advertise with a network they learn has cheaters. (If the network admits it has some cheaters, maybe it has more it hasn't yet found.) In other instances, a principal may blame itself: A principal typically could have caught the prohibited activity earlier, and it seems principals often worry that their initial failure to act will weaken legal claims or, in any event, reputation.

2.3 Technical protections against advertising fraud

Even if rogue advertising agents escape legal redress, as suggested in the prior section, advertising principals could attempt to use technical systems to protect themselves from agent fraud. Indeed, by all indications, many advertising principals make substantial efforts to uncover fraud. For example, Google reportedly examines patterns in paid click data in an attempt to identify and negate click fraud [27]. ValidClick supplements pay-per-click links with JavaScript that reports indicia of fraud (e.g. ads purportedly clicked without movement of a user's mouse). ValueClick Commission Junction uses a web crawler from Cyveillance to uncover cookie-stuffing, among other practices [23]. I personally designed an automated system that manipulates spyware-infected virtual computers in search of unexpected advertising links claiming fees not properly earned [12].

Despite these various efforts to catch online advertising fraud, by all indications fraud remains widespread. Click fraud monitoring services estimate that 16% of paid search clicks were actually click fraud [25]. In my hands-on and automated testing, I have uncovered literally thousands of affiliates using spyware, adware, or cookie-stuffing to claim commissions not properly earned. Discussion sites [1] and consulting services [2] confirm the breadth of rogue online advertising agents.

2.4 Defending against agents' multiple and sequential identities

Identity verification further complicates an online advertising principal's supervision of its agents. An advertising principal typically interacts with its agents only through electronic communication systems, making it difficult to prevent an agent from registering under multiple separate identities. Using multiple smaller accounts offers clear benefits to agents who intend to use tactics that principals prohibit: If one account gets caught and cancelled, the agent will retain proceeds associated with its other accounts. ([13] - exhibits 2, 6, 7, 8, and 9)

Furthermore, even if a principal successfully uncovers an agent's improper activities, limited identity verification prevents the principal from reliably severing ties with the agent. The principal may eject the agent from its program, but there is little to stop the agent from reapplying under a new name. Online advertising fraud thus faces the same unavoidable pseudonyms considered in [17].

2.5 Penalizing and deterring rogue affiliates

Because the legal system is largely unavailable to advertising principals, and because limited identity verification hinders principals' efforts even to know who they are dealing with, standard legal remedies offer advertising principals no clear way forward. Yet a principal pays its agents on an ongoing basis, and a principal may structure its contracts as it chooses, subject to agents deciding to focus their efforts elsewhere.

One natural approach would require that each agent post a bond. But advertising agents seem hesitant to pay fees to advertising principals when the entire purpose of the relationship is to facilitate payments flowing in the opposite direction (*i.e.* from the principal to the agent). Furthermore, these fees would tend to penalize newcomers, raising the [17] concerns of hindering growth and flexibility. However, advertising agents may be more inclined to accept delayed payment of their earnings, as suggested in the sections that follow.

3 Delayed Payment: Model

Suppose a principal ordinarily makes payment v when an agent completes (more precisely, appears to have completed) some specified task of gross value V to the principal. I take v to be exogenous, *e.g.* the outside option of agents who could perform similar work elsewhere, in a competitive market beyond this model.

Suppose good-type agents exogenously exist with probability p in the principal's pool of would-be agents. Rogue agents exist with probability $1-p$, and their output is worthless to the principal. Section 3.8 defends the decision to take p to be exogenous.

3.1 Outcome under a simple contract

Suppose a principal pays v for each seemingly-completed task. The principal receives proportion p of good agents who produce V and receive v. The principal also receives $1-p$ rogue agents who provide the principal with 0 value but also receive v. The principal then obtains profit:

$$\pi_{simple} = p(V - v) + (1 - p)(0 - v) = pV - v \tag{1}$$

That is, the principal makes payment v to each agent, but the principal only receives value V from proportion p of agents.

3.2 Delaying payment: good agents' demands and principal's costs

Suppose the principal imposes a delay in payment to agents. Agents' payments are set by a competitive outside market: If the principal merely delays payment, without offering any corresponding bonus, all good agents will leave the principal for its competitors. To retain good agents in the face of delayed payment, the principal must compensate agents for the delay, *e.g.* via bonus payments.

The principal and agents differ in their relative time preferences. The principal's deposits yield r, the market risk-free real interest rate. Good-type agents discount their future payments from the principal by a higher discount rate, $r + s$. The difference, $s > 0$, is good agents' relative impatience – because they worry the principal will not pay them as promised, or because they lack access to low-cost capital.

Suppose the principal elects to pay its agents with a delay given by proportion q of a year (*e.g.* $q = 0.5$ signifies a 6-month delay). With such a payment delay, good agents will require a larger payment w to accept the principal's offer:

$$w = v(1 + (r + s)q) \tag{2}$$

Here, $r + s$ is the annual bonus percentage required for good agents to accept the delay.[1]

The principal's gross additional cost in making such payments is:

$$w - v = v(1 + (r + s)q) - v = vq(r + s) \tag{3}$$

But in the interim, the principal could invest the amount v for duration q at rate of return r, yielding revenue vqr. Thus the principal's net additional cost of delayed payments is:

$$w - v - vqr = vqs \tag{4}$$

3.3 Delaying payment: probability of detection

Let $\tilde{T}$, a random variable, be the time until a given rogue agent is revealed as such. Let d be the mean time to detection, *i.e.* $E[\tilde{T}] = d$.

Suppose the principal detects rogue agents with a delay that follows an exponential distribution.[2] Let the principal wait time q before paying a given agent. Then the probability that the principal learns the agent is rogue before the principal pays is given by the cumulative distribution function of the exponential distribution:

$$F_T(q) = 1 - e^{-q/d} \tag{5}$$

3.4 Outcome under the delayed-payment contract: agents' profits

Suppose a rogue agent's profit margin in serving the principal is m. (Section 3.7 considers outcomes when rogue agents' margins vary in an interval.) Then the rogue agent incurs cost of $c = (1 - m)v$ in producing one unit for the principal.

---

[1] For simplicity, I ignore compounding of interest.

[2] Other distributions of detection time generally yield similar results. However, the exponential distribution is a particularly natural choice due to its uniform hazard rate: The exponential distribution implies that, in each period, a principal catches a constant proportion of those rogue agents not yet revealed to be rogue.

Let a rogue agent have outside option 0. Rogue agents are therefore deterred from serving a principal if the expected profit from such service is less than 0.[3] Substituting:

$$
\begin{aligned}
\text{[expected revenues]} - \text{[costs]} &< 0 \\
v(1 - F_T(q)) - c &< 0 \\
v(1 - F_T(q)) - v(1 - m) &< 0 \\
1 - F_T(q) &< 1 - m \\
F_T(q) &> m
\end{aligned}
\tag{6}
$$

This is the *rogue-type non-participation constraint* – the condition that must be satisfied to prevent rogue agents from participating. The left side gives the probability that a rogue agent is caught by the principal within time q, *i.e.* that the rogue agent does not receive the payment. The right side is the agent's margin (as a proportion of the principal's payment). If the agent gets caught more often (in percent) than its margin (in percent), the agent will lose money in expectation and will be deterred from participating.

If rogue agents are detected with an exponential delay, constraint (6) becomes:

$$e^{-q/d} < 1 - m \tag{7}$$

Rearranging yields the range of q that deters rogue agents:

$$q > -d \ ln(1 - m) \tag{8}$$

3.5 Outcome under the delayed-payment contract: principal's profit

Suppose the principal can set a q such that only good-type agents choose to work for the principal. The principal then achieves a profit of:

$$\pi_{good-only} = p\big(V - v(1 + (r + s)q) + vqr\big) = p(V - v(1 + sq)) \tag{9}$$

The principal prefers $\pi_{good-only}$ over π_{simple} from (1) if:

$$
\begin{aligned}
\pi_{good-only} &> \pi_{simple} \\
p\big(V - v(1 + sq)\big) &> pV - v \\
q &< \tfrac{1-p}{sp}
\end{aligned}
\tag{10}
$$

This is the *principal profit constraint* – the condition allowing a principal to pay good agents the required bonus for the delay, while simultaneously increasing principal profit.

If the principal succeeds in deterring all rogue agents by imposing a payment delay of length q, the principal's profit increases as follows:

$$
\begin{aligned}
\Delta\pi = \pi_{good-only} - \pi_{simple} &= p(V - v(1 + sq)) - (pV - v) \\
&= -pv - pvsq + v \\
&= v(1 - p - psq)
\end{aligned}
\tag{11}
$$

[3] By implication, an agent can serve – and a rogue agent can defraud – many principals simultaneously. That is, accepting a relationship with one principal does not require an agent to forego relationships with others. So an agent will accept any relationship that offers positive profit.

$\Delta\pi$ is decreasing in q: All else equal, the principal prefers a shorter payment delay.

3.6 Incentive-compatible choice of delay

To retain good agents while increasing profit, a principal must satisfy (6) and (10) simultaneously. In particular, a principal needs a delay q that is large enough to deter rogue agents, yet small enough not to increase the principal's costs excessively.

In principle, there need not be a value of q that simultaneously satisfies the requirements of both the principal and the good agents. For example, if the probability of detection were very close to 0, rogue agents would know they have little chance of being caught, no matter the payment delay. Conversely, if good agents were overly impatient, it might be too costly for merchants to satisfy good agents while deterring rogue agents.

But if the q ranges do overlap, the principal can satisfy both inequalities. Graphically, the principal seeks a q that falls between the dashed lines in Figure 2.

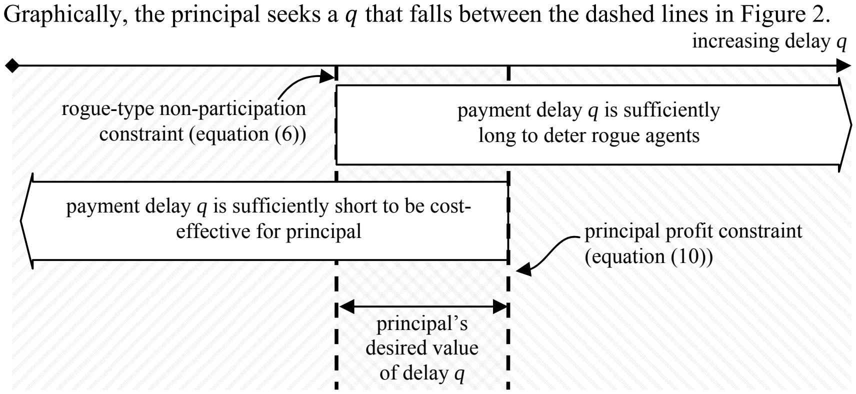

Figure 1: Incentive-Compatible Choice of Delay

In Section 4.1, I calibrate the model to estimate the permissible ranges of q, yielding estimates that suggest the inequalities do overlap.

3.7 Variations in rogue agents' profit margins

Suppose a principal faces a variety of rogue agents with varying profit margins m_i, rather than the single m proposed in Section 3.4. The principal then sets q to deter as many rogue agents as possible while satisfying its profit constraint (10) and retaining good agents.

Suppose all profit margins from 0 to 1 are equally likely. (That is, profit margins following the standard uniform distribution.) A given choice of q will then deter all rogue agents whose margin m_i satisfies the rogue-type non-participation constraint:

$$1 - F_T(q) < 1 - m_i \qquad (12)$$

Using the cumulative distribution function of the uniform distribution, $F_M(m) = P(M < m) = m$, a given choice of q deters the following proportion of rogue agents:

$$P(1 - F_T(q) < 1 - m_i) = P(m_i > F_T(q)) = F_T(q) \tag{13}$$

Then the principal achieves the following profit:

$$\pi_{some-rogue} = pV - pv(1 + sq) - (1-p)v(1 + sq)(1 - F_T(q)) \tag{14}$$

The final term reflects the principal's loss from paying commissions to those rogue agents whose high profit margins allow them to remain despite payment delay q.

To optimally set q, the principal uses the first-order condition of (14):

$$\frac{d\pi_{some-rogue}}{dq} = -pvs - (1-p)v(s - (1 + sq)f_T(q) - sF_t(q)) = 0 \tag{15}$$

The principal prefers $\pi_{some-rogue}$ from (14) to π_{simple} from (1) if $\pi_{some-rogue} > \pi_{simple}$ for q selected to satisfy (15). If so, the principal offers the delayed-payment contract specified in Section 3.2. If not, the principal offers only the simple contract of Section 3.1.

If rogue agents are detected with exponential delay, (14) and (15) become

$$\pi_{some-rogue} = pV - pv(1 + sq) - (1-p)v(1 + sq)e^{-\frac{q}{d}} \tag{16}$$

$$\frac{d\pi_{some-rogue}}{dq} = -pvs - (1-p)ve^{-q/d}\left(s - \frac{1 + sq}{d}\right) = 0 \tag{17}$$

With knowledge (or estimates) of detection speed d, good-type prevalence p, and good agents' impatience s, a principal can evaluate (17) to find the payment delay q that maximizes the principal's profit. I present this approach in the following section.

3.8 The exogeneity of *p*

Section 3 takes p, the prevalence of good agents, to be exogenous. In principle, p could vary as advertising principals change their anti-fraud tactics. Nonetheless, I view the fixed model of p as appropriate under the circumstances. In particular, experience suggests that few advertising agents shift from fraud to non-fraud, or vice versa. Rather, industry experience indicates that agents are either fraudsters or legitimate, but do not often change back and forth. Thus, the key moral hazard worry, i.e. that an otherwise-good agent would see a principal's compensation scheme and turn to fraud in response, appears less urgent.

My decision to model two types of agents – good and rogue – is also consistent with the literature. For example, [16] similarly presents a mdoel of "good types" and "bad types."

4 Application to Internet Advertising

4.1 Calibrating the model

To calibrate the model in Section 3, I received data from a major US advertising network. The network specializes in relationships between advertisers and small publishers ("affiliates"), paying publishers in proportion to their sales. Publisher infractions include those described in Section 2, as well as additional infractions such as falsely or deceptively describing the merchants' products or pricing.

The network's 2006-2007 detections of publisher infraction yield an estimate of good-type prevalence $p = 0.86$. Among 2006 active affiliates who were ultimately terminated for cause, the mean time to termination, d, was 0.59 years (217 days). (2006 is the last full year for which such data is available.) (Compare [4], estimating a range of plausible detection rates for other browser-based attacks.)

For a worst-case bound on an affiliate's cost of capital, consider an affiliate whose funds come from a consumer credit card with annual real interest rate of 20%. In contrast, the affiliate network might earn a 2% real return in a low-risk investment. Then $s = 0.18$, $r = 0.02$.

A typical rogue affiliate might have a profit margin $m = 0.5$. This value reflects that the rogue affiliate's efforts require limited out-of-pocket expenditures, as in the examples in Section 2.1. Substituting into (8):

$$q > -0.59\, ln(1 - 0.5) = 0.41 \tag{18}$$

If $q > 0.41$, then rogue affiliates will earn negative profits and will cease to participate.

Meanwhile, from Section 3.5, the principal prefers to pay with delay q if that delay increases profit while retaining good affiliates. Substituting from (10), increasing principal profit requires:

$$q < \frac{1 - 0.86}{(0.18)(0.86)} = 0.90 \tag{19}$$

For such a principal, the gain from excluding all rogue affiliates is so large that the principal would be willing to pay nearly a year of interest (at a rate given by the difference between the principal's discount rate and the agent's discount rate) in order to exclude all rogue affiliates.

Combining (18) and (19), any q in the range $0.41 < q < 0.90$ will deter rogue affiliates while increasing the principal's profit.

4.2 Variation in rogue agents' profit margins

Rogue agents incur a variety of costs in attempting to defraud advertising principals. For example, agents typically buy traffic (e.g. from banners, pay-per-click search campaigns, spyware, or adware). Agents also face an imputed cost from the value of their own time in planning and coordinating their tactics.

The preceding section estimates that delayed payments can profitably deter rogue agents if rogue agents all have margin $m = 0.5$. But what if some agents' margins are larger or smaller than that value? Figure 2 shows the relationship between payment delay and agent margin. For a variety of agent profit margins $\boldsymbol{m}$, the plot shows the range of delay q that lets an advertising principal profitably delay payment, consistent with the other parameters estimated in the preceding section. Within the double-hatched area, the principal's profit increases from the use of a delayed-payment contract, and the principal successfully deters rogue agents from participating. If all rogue agents have profit margins below 0.83 (the value of m where (7) and (10) cross), the advertising principal can deter participation of all rogue agents, obtain the increased profit derived in Section 3.5, and pay good-type agents the increased fee described in (2).

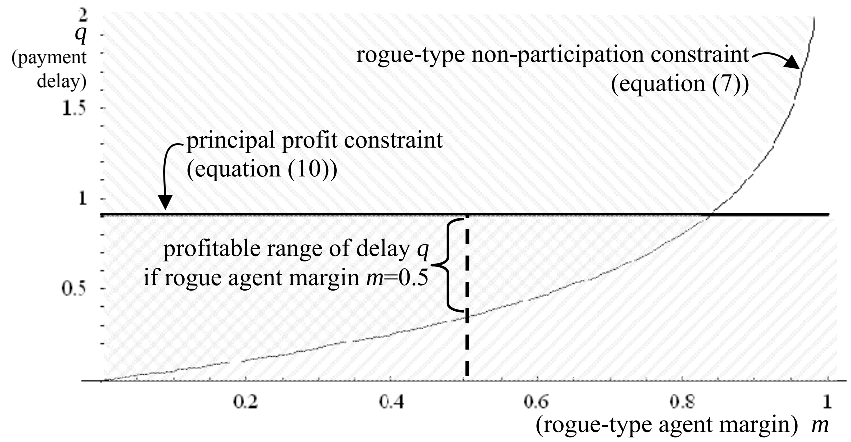

Figure 2: Profitable Delay as Rogue Agents' Margins Vary

If some rogue agents have margins that exceed the value of m where (8) and (10) cross, the advertising principal must turn to the approach presented in Section 3.7. Suppose rogue agents' profit margins follow the distribution posited in 3.7, with all values between 0 and 1 equally likely. Equation (16) reports how the principal's profit varies in its choice of delay. Figure 3 plots the principal's change in profit as payment delay varies.

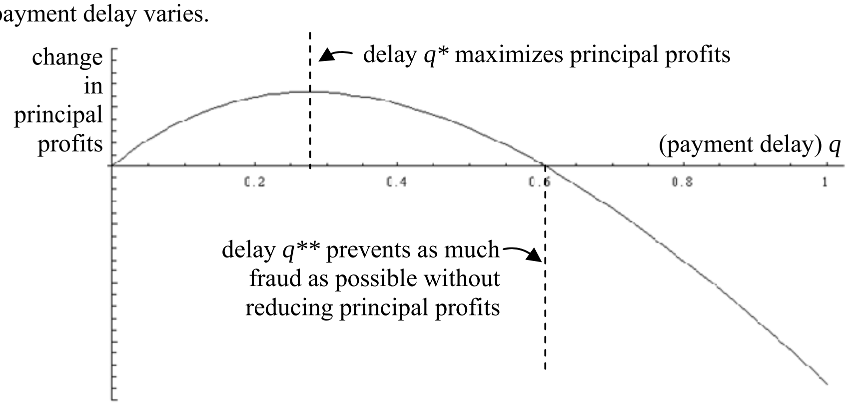

Figure 3: Effect of Payment Delay on Principal Profits

Consistent with (17), Figure 3 confirms that the principal's maximum profit occurs with delay $q^* = 0.28$ (*i.e.* 15 weeks). At this payment delay, (13) indicates that the principal will deter 44% of rogue agents. Alternatively, the principal could choose a payment delay $q^{**} = 0.61$ (*i.e.* 32 weeks) – foregoing any profit increase from deterring rogue agents, but deterring more rogue agents (namely, 71%).

As the principal further increases its payment delay, it deters participation by additional higher-margin rogue agents. But deterring the highest-margin agents requires that the principal lose good-type agents or accept a reduction in profit (relative to profit under the simple contract in (1)). In particular, if the principal delayed payment long enough to deter the highest-margin rogue agents' participation, the principal would face increasing costs in compensating good-type agents for the delay, and the principal would be unable to pay those costs from the proceeds of excluding rogue agents.

4.3 Implementation in practice

In general, an advertising principal might not know all the parameter values set out above. But the preceding analysis suggests that a substantial payment delay could be profitable under reasonable market conditions.

Despite the benefits of delaying payment, many advertising industry participants seem to think affiliates should be paid *more* frequently. Consider LinkShare's 2007 move to pay affiliates as often as once per week [11], a move made possible by the transition from printed checks to electronic funds transfers. LinkShare claims to offer "the most publisher-friendly payment plan of the major affiliate networks" – presenting weekly payments as a boon to affiliates. Indeed, both good and rogue affiliates prefer to be paid quickly, all else equal. But by paying its affiliates more often, a network limits its ability to punish affiliates ultimately found to be violating its rules or defrauding merchants. Although good affiliates appreciate being paid quickly, the preceding estimation suggests an interested affiliate network could offer an *increased* payment that good affiliates would value even more than *rapid* payment.

Table 1 reports payment delays of selected marketing programs, ad networks, and affiliate networks. Payment delays range from one week (LinkShare as well as large affiliates of Clickbooth and CPA Empire) to 15 days after the end of each month (ordinary Clickbooth affiliates) to 30 days after the end of each month (Google AdSense) to 60 days after the end of each month (Yahoo's Right Media Network). The web appendix to this paper [8] expands Table 1 to include additional details of applicable rules, as well as citations and links to governing agreements.

In implementing delayed payments, an affiliate network would face the problem that good affiliates' profit margins vary substantially. For example, *content affiliates* place affiliate links within their own material (*e.g.* articles or blogs) – yielding high gross margins because these distribution methods present few direct costs and, in any event, few marginal out-of-pocket costs. Conversely, *search affiliates* buy ad placements from search engines and sell the resulting traffic to merchants via affiliate networks – yielding low net margins due to search engine fees and due to competition from other search affiliates with a similar approach.

A payment delay that satisfies most good-type affiliates might nonetheless prove unworkable for search affiliates due to their lower profit margins. But affiliate networks and merchants could review requests for faster payments on a case-by-case basis – using appropriate indicia of legitimacy (*e.g.* reputation, audit results, *HTTP*

Marketing Program	***Payment Frequency & Delay***
Amazon Associates	Monthly payment, paid approximately 60 days after the end of each month
Clickbooth	Monthly payment, paid on the 15th day of the next month. Weekly payment, paid 7 days after the end of the week (for affiliates earning >$5000/week)
Commission Junction	Monthly payment, paid on the 20th day of the next month
CPA Empire	Monthly payment, paid 20 days after the end of the month Weekly payments (for affiliates earning >$1000/week)
Google AdSense	Monthly payment, paid approximately 30 days after the end of each month
Hydra Network	Monthly payment, paid within 15 days of the end of the month
LinkShare	Weekly payment
NeverblueAds	Monthly payment, paid within 30 days of the end of each month
PrimaryAds	Monthly payment, paid approximately 30 days after the end of each month. "Aggressive payment terms" ("receive [a] commission check every week") for "high-volume affiliates"
Regnow	Monthly, delay unspecified. Weekly payment "may" be available if an affiliate pays an additional fee and meets an unspecified additional qualification threshold
Right Media Network	Monthly payment, paid within 60 days of the end of each month
Yahoo Publisher Network	Monthly payment, paid 3-4 weeks after the end of the month

Table 1. Payment Delays for Selected Marketing Programs

Referrer headers showing traffic truly coming from search engines) to confirm the claims of affiliates seeking faster payment. By limiting fast payment to affiliates that survive heightened verification, affiliate networks could reduce fraud while avoiding burdensome investigations of all their affiliates. At present, Table 1 indicates that few marketing systems invoke such subtle analysis: Regnow indicates that fast payment "may" be available to affiliates who meet unspecified additional qualifications, but no other network's public statements report additional substantive requirements for accelerated payment. While some networks (e.g. Clickbooth, CPA Empire, PrimaryAds) offer accelerated payment to large affiliates, substantial earnings in and of themselves are not a clear indicator of trustworthiness.

Facing the prospect of substantially delayed payments, advertising agents might be concerned about the creditworthiness of their advertising principals. Beginning in 2001, a series of affiliate merchants entered bankruptcy, and in some instances affiliates did not receive the commissions they had earned. But existing institutions can help assure that affiliates are paid as expected. For example, affiliate network Commission Junction now requires that merchants tender prepayments sufficient to cover their anticipated monthly advertising expenses [6], and Commission Junction holds these funds to assure affiliates' subsequent payments. Under a delayed payment regime, merchants would continue to pay networks as usual, on the current schedule – substantially protecting affiliates from lost payments if a merchant became insolvent. Merchants particularly determined to demonstrate their creditworthiness could turn to a formal escrow service or other mechanism to accept and hold affiliates' accrued earnings.

Improving detection technology offers important benefits beyond delayed payment. In particular, improved detections are particularly important if the model in Section 3.3 misstates the probability of detecting a rogue affiliate, *i.e.* if some rogue affiliates have exceptionally effective technologies for avoiding detection no matter how long networks search. But improving enforcement is costly – spiders and crawlers for automated enforcement, human review teams for manual investigations, and managers and attorneys to make final decisions. Delayed payment offer a more expedient alternative – a useful stopgap strategy for use when primary enforcement systems prove inadequate.

5 Other Applications and Future Work

Online advertising markets are one of many markets where agents may be effectively unreachable through the legal system. But in other such contexts, institutions and norms develop to deter misbehavior. For example, apartment tenants generally prepay a security deposit plus first and month's rent. Because tenants have prepaid these fees, landlords are well protected from typical damage – without having to incur litigation costs if damage occurs. Similarly, *neafarios* require payment in advance for their immigration services, protecting them from clients disappearing and failing to pay the promised fee. Conversely, a contingent fee agreement protects a client from the risk of low attorney effort by delaying payment until a better measure of effort (namely, success) becomes available.

Each of these payment rules addresses a market-specific information asymmetry. Although online advertising features similar risk of agent misbehavior, online advertising contracts presently lack any similar institution by which payment structure can enforce good practices. Online advertising would still suffer somewhat from the context-specific unavailability of a bond or other prepayment from the judgment-proof agent. But appropriate selection of a payment delay can achieve the valuable benefits offered by contingent payments in other markets.

I have offered an initial model of agent behavior – with agents moving from one advertising principal to another, but never shifting from rogue to good or vice versa. Future work might appropriately extend my approach to consider agents who respond to changing incentives by modifying their behavior as to a given advertiser, *i.e.* a model in which agents are subject to both moral hazard and adverse selection.

Acknowledgments

I thank George Baker, Eric Budish, Peter Coles, Fuhito Kojima, Jeff Molander, Tyler Moore, David Parkes, Al Roth, and five anonymous referees. I gratefully acknowledge data from an advertising network that prefers not to be referenced by name. This research was supported by the Division of Research and Faculty Development at Harvard Business School.

References

1. ABestWeb, 2008. "Parasiteware." http://forum.abestweb.com .
2. Affiliate Fair Play, 2008. http://www.affiliatefairplay.com .
3. Barth, Adam, et al., 2007. "Detecting Fraudulent Clicks from BotNets 2.0." Mimeo.
4. Boneh, Dan, et al., 2008. "Crimeware in the Browser." Crimeware. Addison-Wesley.
5. BusinessWeek, 2006. "Click Fraud: The Business of Cyberstealing." http://www.businessweek.com/technology/content/dec2006/tc20061204_923336.htm .
6. Commission Junction, 2008. "Generate Advertising Order Form." http://help.cj.com/en/ja25)@a/Generate_Advertising_Order_Form_.htm .
7. Daswani, Neil and Michael Stoppelman, 2007. "The Anatomy of Clickbot.A." Proceedings of the First Conference on Hot Topics in Understanding Botnets.
8. Edelman, Benjamin, 2008. "Delaying Payment to Deter Online Advertising Fraud." http://www.benedelman.org/paymentdelay/ .
9. Edelman, Benjamin, 2008. "Auditing Spyware Advertising Fraud: Wasted Spending at VistaPrint." http://www.benedelman.org/news/093008-1.html .
10. Edelman, Benjamin, 2008. "CPA Advertising Fraud: Forced Clicks and Invisible Windows." http://www.benedelman.org/news/100708-1.html .
11. Edelman, Benjamin, 2007. "Spyware Still Cheating Merchants and Legitimate Affiliates." http://www.benedelman.org/news/052107-1.html .
12. Edelman, Benjamin, 2007. "Introducing the Automatic Spyware Advertising Tester." http://www.benedelman.org/news/052107-2.html .
13. Edelman, Benjamin, 2006. "People of the State of New York v. Direct Revenue, LLC. – Documents and Analysis." http://www.benedelman.org/spyware/nyag-dr/ .
14. Edelman, Benjamin, 2005. "Cookie-Stuffing Targeting Major Affiliate Merchants." http://www.benedelman.org/cookiestuffing/ .
15. Edelman, Benjamin, 2004. "The Effect of 180solutions on Affiliate Commissions and Merchants." http://www.benedelman.org/spyware/180-affiliates/ .
16. Ely, Jeffrey and Jusso Valimaki, 2002. "Bad Reputation." *NAJ Economics* 4:2.
17. Friedman, Eric and Paul Resnick, 2001. "The Social Cost of Cheap Pseudonyms." *Journal of Economics and Management Strategy* 10(2): 173-199.
18. Google, 2008. "Content Network." https://adwords.google.com/select/afc.html?sourceid=awo&subid=en-us-et-awhp_related .
19. Leyden, John, 2006. "AOL Seeks Spammer's Buried Gold." *Register*. August 17, 2006.
20. LinkShare, 2008. "Affiliate Information." http://www.linkshare.com/affiliates/affiliates.shtml.
21. LinkShare, 2007. "LinkShare Announcements." http://www.linkshare.com/rc/announcements.html . September 24, 2007.
22. LinkShare, 2002. "Industry Leaders Convene to Examine Trends and Share Key Success Strategies at LinkShare Symposium." http://www.linkshare.com/press/convene.html .
23. Livingston, Brian, 2005. "Commission Junction Hires Web Detectives." Datamation. http://itmanagement.earthweb.com/columns/executive_tech/article.php/3558911 .
24. Scanlan, Emma, 2005. "The Fight to Save America's Inbox: State Legislation and Litigation in the Wake of CAN-SPAM." 2 Shidler J. L. Com. & Tech. 12.
25. Scholz, Karl, 2008. "Industry Click Fraud Rate Hovers at 16 Percent for Third Quarter 2008." ClickForensics. October 23, 2008.
26. Spence, Michael, 1973. "Job Market Signaling." *Quarterly Journal of Economics* 87 (3): 355-374.
27. Tuzhilin, Alexander, 2006. "The Lane's Gifts v. Google Report." http://googleblog.blogspot.com/pdf/Tuzhilin_Report.pdf .
28. Zango, 2005. "180solutions Sues Former Affiliates for Illegal Software Installations." http://www.zango.com/Destination/Corporate/ReadArticle.aspx?id=29 .