Should Ad Networks Bother Fighting Click Fraud? (Yes, They Should.)

Bobji Mungamuru
Stanford University
bobji@i.stanford.edu

Stephen Weis
Google
sweis@google.com

Hector Garcia-Molina
Stanford University
hector@cs.stanford.edu

Abstract

Suppose an ad network decides that a given click-through is invalid (or "fraudulent"). The implication is that the ad network will not bill the advertiser for that click-through. On the other hand, if the ad network decides that the click-through was valid, they could charge full price. Therefore, arguably, the ad network is "leaving money on the table" by marking click-throughs invalid. As such, should ad networks even bother fighting click fraud? We analyze a simple economic model of the online advertising market and conclude that the answer is, unequivocally, "yes".

1 Introduction

Advertising fraud, particularly click fraud, is a growing concern to the online advertising industry. Broadly, the online advertising market is comprised of three types of parties: content publishers, advertising networks, and advertisers. Users engage in the market indirectly by clicking on advertisements on publishers' content pages. At first glance, the incentives regarding fighting fraud may seem somewhat perverse. If an advertiser is billed for clicks that are fraudulent, the advertising network's revenues might increase. As such, is it even in an advertising network's interest to fight fraud at all? Perhaps it makes more sense for an advertising network to just "let it happen" [1]. If not, can an advertising network actually gain a market advantage by aggressively combating fraud? In this paper, we address these and other questions by studying the economic incentives related to combating fraud, and how these economic incentives might translate into behavior.

An economic analysis of online advertising fraud (or, "ad fraud") is appropriate because, unlike many other online security threats, ad fraud is primarily motivated by financial gain. Successfully committing ad fraud yields immediate monetary gains for attackers at the expense of the victims (although it is not immediately clear who the actual victims of ad fraud are). The threat of fraud to the advertising business model and the technical challenge of detecting fraud have been topics of great concern in the industry (e.g., [2, 4]). There have been many informal conjectures in online forums and the media attempting to answer the questions we have posed above (e.g., [7, 8, 9]). The arguments, while sometimes intuitive, typically are not backed by a sound economic analysis. Thus, the conclusions arrived at differ widely. We attempt to fill this gap by performing just such an analysis[1].

Our goal, then, is to construct and analyze a simplified economic model that hones in on the market effects of fighting fraud. However, conducting such an analysis is difficult because faithful models of the market can

[1] A summary of our results (5 pages), omitting proofs and mathematical details, appears in [6].

quickly become very complex. A complete specification of the players' types, decision variables and signals would be intractable. For example, a publisher's type includes, among other things, the volume of traffic they receive, the quality of their content, and their user demographics and interests. Advertisers can be differentiated by the size of their advertising budgets, their valuation of traffic that they receive through online ads, the quality of their campaign, and their relevance to particular demographics. Ad networks differ in their ability to detect ad fraud, as well as the quality and relevance of their ad serving mechanisms. The challenge is to retain those aspects that are relevant to fighting ad fraud.

This paper will focus solely on *click fraud* in pay-per-click advertising systems. Click fraud refers to the act of clicking on advertisements, either by a human or a computer, in an attempt to gain value (e.g., deplete a competitor's budget or earn commissions) without having any actual interest in the advertiser's website. Click fraud is probably the most prevalent form of online advertising fraud today [3, 4, 5]. There are other forms of ad fraud[2] that will not be addressed here. We do not analyze the reasons why fraud occurs (see [2] for such an analysis). Instead, given that fraud does occur in practice, we ask whether advertising networks have an incentive to aggressively combat it.

1.1 Pay-Per-Click Advertising

We model the *pay-per-click (PPC) advertising market* as a game between three classes of players: publishers, advertising networks and advertisers. *Publishers* create online content and display advertisements alongside their content. *Advertisers* design advertisements (or, "ads"), as well as bid on *queries* that summarize what their target market might be interested in. *Advertising networks* (or, "ad networks") act as intermediaries between publishers and advertisers by first judging which queries best describe each publisher's content, and then delivering ads to the publisher from the advertisers that have bid on those queries. For example, an ad network might deduce that the query "foreign automobile" is relevant to an online article about cars, and serve an ad for used car inspection reports.

When a user views the publisher's content and clicks on an ad related to a given query, they are redirected to the advertiser's site – we say that a *click-through* (or, "click") has occurred on that ad. The advertiser then pays a small amount to the ad network that delivered the ad. A fraction of this amount is in turn paid out to the publisher who displayed the ad. The exact amounts paid out to each party depend on several factors including the advertiser's bid, the auction mechanism being used and the revenue sharing agreement between the ad network and publisher. Advertisers are willing to pay for click-throughs because some of those clicks may turn into *conversions*[3] (or, "customer acquisitions" or "sales"). The publishers and ad networks, of course, hope that users will click on ads because of the payment they (i.e., the publishers and ad networks) would receive from the advertiser. Figure 1 gives a graphical illustration of the "lifecycle" of a click-through.

In some cases, a publisher and an ad network are owned by the same business entity. For example, major search engines often display ads next to their own search results. Similarly, a publisher and an advertiser can be owned by the same entity. Online newspapers are a common example. In our model, even if a publisher and an ad network are owned by the same entity, they will nevertheless both act independently. Consequently, the model may predict some behaviors that, while economically rational, are unlikely to occur in practice. For

[2] See [2] for a detailed discussion of the various types of ad fraud.

[3] The definition of a conversion depends on the agreement between the advertiser and the ad network, varying from an online purchase to joining a mailing list. In general, a conversion is some agreed-upon action taken by a user.

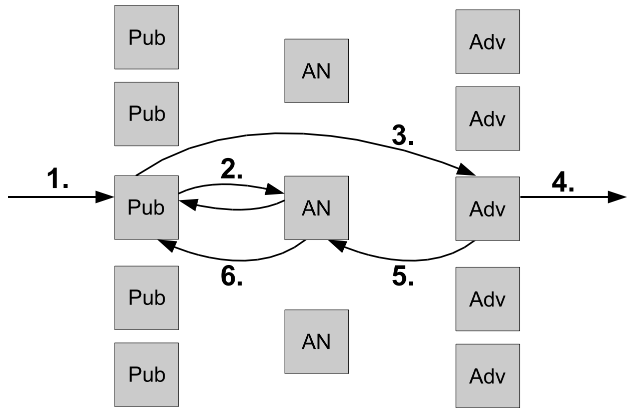

Figure 1: Lifecycle of a click-through in PPC advertising: 1) A user visits a publisher's site. 2) The publisher requests ads from an ad network. 3) If the user clicks on an ad, they are redirected to the advertiser's site. 4) The click-through sometimes becomes a conversion. 5) If the click-through is marked valid, the advertiser pays the ad network. 6) A fraction of the advertiser's payment is shared with the publisher.

example, a real-world entity that owns both a publisher and an ad network is unlikely (for strategic reasons) to display ads from a rival ad network on its properties, even if it might yield an immediate economic advantage[4].

1.2 Invalid Clicks

Apart from ad delivery, advertising networks serve a second important function, namely, trying to detect invalid clicks. *Invalid clicks* can be (informally) defined as click-throughs that have zero probability of leading to a conversion[5]. Invalid clicks include fraudulent click-throughs as well as *unintentional clicks*. For example, if a user unintentionally double-clicks on an ad, only one of the two clicks has a chance at becoming a conversion, so the other click is considered invalid. Going forward, we will speak of valid and invalid clicks, rather than "legitimate" and "fraudulent" clicks. In practice, advertisers are never billed for clicks that ad networks detect as invalid, although the user is still forwarded to the advertiser's site.

The algorithms used by ad networks to detect invalid clicks are prone to error. In particular, their algorithms may produce *false negatives* by identifying invalid clicks as valid, and *false positives* by identifying valid clicks as invalid. A false negative implies that an advertiser has been unfairly billed for a click that could not lead to a conversion. A false positive, on the other hand, is a valid click that the advertiser has received for free. We assume that if an ad network detects an invalid click, it will then actually mark the click invalid (and therefore not bill the advertiser for it). The reader may wonder if an ad network might detect a click as invalid, but then choose to mark it valid. However, for our purposes, doing so is equivalent to simply not detecting the click as invalid, since the user is forwarded to the advertiser's site irrespective of the ad network's decision. Therefore, without loss of generality, we will use the phrases "detect as invalid" and "mark as invalid" synonymously.

In our model, we make an important distinction between how effective and how aggressive an ad network is at detecting invalid clicks. *Effectiveness* refers to the quality of the algorithms used to detect invalid clicks, whereas *aggressiveness* refers to the manner in which the algorithms are applied. For example, suppose an ad network uses an algorithm that takes as input a single click, and outputs a score between 0 and 1, where a

[4]On the other hand, recent anecdotal evidence [10] suggests that such an event might indeed occur in practice.
[5]It is still a topic of some debate what the exact definition of an invalid click should be.

higher score indicates a higher confidence that the click is invalid. The ad network might then run the algorithm on each click, and apply a threshold to the output score to decide whether to mark the click invalid e.g., they might mark clicks invalid if the score is 0.5 or higher. In this example, effectiveness refers to the accuracy of the scoring algorithm – a good algorithm should output a high score when invalid clicks are input and a low score when valid clicks are input. Aggressiveness, on the other hand, is related to the threshold value chosen – a more aggressive ad network would choose a lower threshold, even though doing so may result in more false positives.

Each ad network in our model uses an algorithm whose effectiveness is known (i.e., it is an inherent property of each ad network). The aggressiveness, on the other hand, is a decision variable that is subject to change over time. Our goal in this paper is to study how aggressive a rational ad network would be in equilibrium:

- Is it in an ad network's interest to aggressively detect invalid clicks?
- If an ad network is more effective at detecting invalid clicks, do they gain an edge in the market?

1.3 Decision Variables

Each class of players in our model is faced with a different decision. Publishers must decide which ad networks to sell their click-throughs[6] on. Advertisers must decide what value to place on clicks coming from each ad network[7]. Advertisers' valuations are based mainly on their desired *return on investment* (ROI). In this context, ROI refers to the revenue generated by converted click-throughs divided by the cost of these click-throughs. Finally, ad networks are faced with the decision of how aggressively to detect invalid clicks.

Both the publishers and the advertisers do business with the ad networks. However, there is a key conceptual difference between the decisions they face. Publishers, essentially, are choosing between the ad networks. They allocate their finite inventory of click-throughs across the ad networks. Advertisers, on the other hand, do not choose between ad networks. They are willing to buy an unlimited number of click-throughs from any and all of the ad networks, as long as the ROI on each ad network is sufficiently high. Only publishers have users – advertisers and ad networks do not. That is why we distinguish between publishers, advertisers and ad networks, even when they are owned by the same business entity.

1.4 Results

Our results can be summarized as follows:

1. In equilibrium, all publishers (i.e., even the "spammers") will prefer to display ads from the ad network that promises advertisers the "highest quality" traffic. Why? Because higher-quality click traffic leads to increased ROI for advertisers, which in turn leads to higher bids and more money for publishers. Stated differently, the ad network that promises the highest quality traffic will not only attract the high-quality publishers, but also the spammers.

[6] In practice, publishers sell ad impressions (or "page views" or "eyeballs"), not click-throughs. However, under certain assumptions, it is equivalent to think of click-throughs as the items that are bought and sold in the market. In Section 4, we discuss these assumptions in detail.

[7] Clicks are sold to advertisers using an auction process, so the advertisers' valuations are subsequently translated into bids

2. All other factors remaining equal, the ad network that is most effective at detecting invalid clicks is able to provide advertisers the highest quality traffic. Why? Because filtering effectively ensures that advertisers are billed for fewer invalid (and hence, non-converting) clicks.

3. However, the ad network that is most effective at detecting invalid clicks can only deliver high-quality traffic if it is sufficiently aggressive. In particular, it is suboptimal to simply mark all clicks valid. As its lead in effectiveness narrows, it is forced to be increasingly aggressive. As ad networks become more aggressive, the high-quality publishers benefit disproportionately.

The first conclusion – that all publishers will choose the same ad network – is perhaps the most surprising, and certainly is the most difficult to prove. It is easy to believe that high-quality publishers will be attracted to the ad network with the highest quality traffic, due to the increased bids on that network. However, we predict that even the low-quality ones will prefer this network. The key phenomenon behind this conclusion is that the algorithms used by ad networks are prone to error – no ad network is able to detect 100% of the invalid clicks coming from a publisher (unless, of course, every single click is marked invalid). Even though an effective, aggressive ad network will detect most invalid clicks, a few will always "slip through the cracks". The advertisers' bids will be high enough, in equilibrium, that the few invalid clicks that are mistaken for high-quality traffic generate sufficient revenues to attract even the low-quality publishers.

The rest of the paper is structured as follows. In Section 2 we analyze a simplified single-period model of the online advertising market (a sequential game), and demonstrate our results in this simple model. Although the model in Section 2 is simple, it captures the essential intuitions behind our results. In Section 3, we present a more realistic (and more complex) multi-period model of the market, as a sequence of interactions between publishers, ad networks and advertisers over a long horizon. Much of the intuition and the proof techniques from the single-period model turn out to be directly applicable in the multi-period setting. Under some additional assumptions, we establish essentially the same result in the multi-period setting as we did for the single-period game. We provide a thorough discussion in Section 4 of the impact of various assumptions in our model, and how they relate to practice, and conclude in Section 5. The proofs of our results are presented in the appendices.

2 Single-Period Model

In this section, we model the online advertising market as a single-period, two-step sequential game between publishers, ad networks and advertisers. The two steps of our game are as follows:

1. In the *first step*, ad networks decide how aggressively to filter for invalid clicks.

2. In the *second step* (or, the *subgame*), the publishers and advertisers react to the ad networks' decisions. Publishers decide which ad networks to sell their clicks on, and advertisers decide how much they are willing to pay for clicks from each ad network.

After the second step, profits are realized: a) publishers sell clicks (i.e., display ads) on their chosen ad networks, b) ad networks mark a subset of these clicks invalid and c) advertisers pay the ad networks for the clicks that are marked valid. The game in its entirety (i.e., the first step followed by the second) is referred to as the

Symbol	Description
I, J, K	Number of publishers, ad networks and advertisers (respectively)
i, j, k	Index over publishers, ad networks and advertisers (respectively)
h	Fraction of revenue paid by ad networks to publishers
β	Fraction of clicks that become conversions
V_i	Volume of clicks on publisher i's site
r_i	Fraction of publisher i's clicks that are valid
α_j	Effectiveness of ad network j
y_k	Advertiser k's revenue per conversion
R_k	Advertiser k's target ROI
$c_{i,j}$	Fraction of publisher i's clicks sent to ad network j
x_j	Aggressiveness (false positive rate) of ad network j
$v_{k,j}$	Advertiser k valuation of ad network j's clicks
$N_{i,j}$	Fraction of i's clicks marked valid by j
$\pi_{i,j,k}$	Publisher i's revenue from clicks sold to advertiser k via ad network j
π_i	Publisher i's total revenue
$Y_{k,j}$	Advertiser k's revenue from ad network j's clicks
$Z_{k,j}$	Number of clicks advertiser k's is billed for by ad network j
$R_{k,j}$	Advertiser k's ROI on ad network j
η_j	Ad network j's total revenue
M_j	Ad network j's multiplier
j^*	Ad network that is chosen by publishers in a given equilibrium
$\bar{\pi}_{i,j}$	Publisher i's profit in an equilibrium where ad network j is chosen
$\bar{r}$	Average publisher quality
x^*	Aggressiveness required for the most effective ad network to win the market

Table 1: List of variables introduced in Section 2.

supergame. Our goal is to predict how aggressive ad networks would be in equilibrium, and how the market of publishers and advertisers would react. Our single-period model is simplified in four important ways:

1. It is a single-period game.
2. Ad networks act first, followed by the publishers and advertisers.
3. There is no private information.
4. There are no auctions for click-throughs.

In Section 3, we relax the above simplifications to arrive at a more realistic model of the market. However, our main results and intuitions will still apply.

2.1 Model

Table 1 is a summary of the notation used in Section 2. Consider a market for click-throughs on a single query. The market is comprised of I publishers, J ad networks, and K advertisers. In practice, we have $I \gg K \gg J$. Click-throughs originate on publishers' pages, and are forwarded on to advertisers via the ad networks.

Each publisher i receives V_i click-throughs on its websites, of which only a fraction r_i are valid. That is, publisher i receives $r_i V_i$ valid clicks and $(1 - r_i)V_i$ invalid clicks. We assume that x_j is the fraction of valid clicks that ad network j mistakenly marks invalid[8] (i.e., false positives), and that $x_j^{\alpha_j}$ is the fraction of invalid clicks that j correctly marks invalid (i.e., true positives). A simple calculation then shows that

$$N_{i,j} \equiv (1 - x_j)r_i + (1 - x_j^{\alpha_j})(1 - r_i) \tag{1}$$

is the fraction of publisher i's clicks marked valid by j. As we discuss later, x_j measures ad network j's aggressiveness at detecting invalid clicks, whereas α_j measures j's effectiveness.

Let $c_{i,j}$ be the fraction of publisher i's click-throughs that are on advertisements from ad network j. Then,

$$N_{i,j}V_i c_{i,j} \tag{2}$$

is the total number of i's clicks that ad network j marks valid. Note that all clicks are forwarded onto advertisers, not just those marked valid. Marking clicks valid or invalid only affects how much advertisers are billed and how much publishers and ad networks are paid.

For simplicity, let us assume each advertiser k receives an equal fraction of the clicks from each ad network[9]. Equivalently, ad networks choose an advertiser uniformly at random on each ad impression. Thus, each advertiser k receives a fraction $\frac{1}{K}$ of each ad network j's clicks (recall that there are K advertisers in total). We denote by $v_{k,j}$ the amount k pays j per click[10]. A fraction h of each dollar of revenue is in turn paid out by each ad network to the publisher from which the click originated. Therefore,

$$\pi_{i,j,k} \equiv \frac{1}{K} N_{i,j}V_i c_{i,j} v_{k,j} h \tag{3}$$

is the amount ad network j pays publisher i for clicks that are sent to advertiser k.

Now, ad network j receives $\sum_i r_i V_i c_{i,j}$ valid clicks in total, of which a fraction $\frac{1}{K}$ is sent to advertiser k. Assuming that each valid click becomes a conversion with probability β, and that advertiser k earns y_k dollars from each conversion, advertiser k earns

$$Y_{k,j} \equiv \left(\sum_i r_i V_i c_{i,j} \right) \frac{1}{K} \beta y_k \tag{4}$$

dollars of total revenue from clicks coming from ad network j.

From (2), ad network j marks $N_{i,j}V_i c_{i,j}$ of publisher i's clicks valid. Therefore,

$$Z_{k,j} \equiv \frac{1}{K} \sum_i N_{i,j} V_i c_{i,j} \tag{5}$$

is the total number of clicks that ad network j marks valid and bills advertiser k for. Advertiser k pays $v_{k,j}Z_{k,j}$ dollars in total to ad network j for these clicks. Advertiser k's return on investment (ROI) on ad network j is,

[8] Only to simplify the exposition, we assume throughout this paper that $x_j < 1\ \forall j$ i.e., no ad network gives away 100% of its clicks for free. It is possible to prove that $x_j = 1$ will never occur in equilibrium, but for brevity we will omit the proof. Similarly, in Section 3, we assume $x_{j,t} < 1\ \forall (j,t)$.

[9] As we will see, the manner in which an ad network allocates clicks across advertisers does not impact our results. Thus, in Section 3, we will replace this simplifying assumption by a more realistic one (i.e., auctions), without changing our conclusions.

[10] We are assuming implicitly that the amount paid per click does not vary across clicks. We will drop this assumption in Section 3, once we introduce auctions into the multi-period model.

therefore:

$$R_{k,j} \equiv \frac{Y_{k,j}}{v_{k,j} Z_{k,j}} \tag{6}$$

Finally, ad network j's profit is simply the total revenue received from advertisers, less the payments made to publishers – a fraction h of each dollar of revenue from the advertisers is paid out, and the other $1 - h$ is retained. Therefore, using (3), ad network j's total profit, η_j, is simply:

$$\eta_j \equiv \frac{1}{K}(1-h) \sum_k \sum_i N_{i,j} V_i c_{i,j} v_{k,j} \tag{7}$$

2.2 Player Objectives

In our single-period model, publisher i's type is the ordered pair (r_i, V_i), which describes i's click-through traffic. The allocations $\{c_{i,j}\ \forall j\}$, on the other hand, are i's decision variables. Publisher i's objective is to choose $\{c_{i,j}\ \forall j\}$ so that its total profit across all ad networks is maximized. Define π_i to be publisher i's total profit:

$$\pi_i \equiv \sum_j \sum_k \pi_{i,j,k} \tag{8}$$

Publisher i chooses $\{c_{i,j}\ \forall j\}$ by solving the following optimization problem:

$$\max_{\{c_{i,j}\}} \ \pi_i \quad \text{s.t.} \quad \sum_j c_{i,j} = 1 \quad \text{and} \quad c_{i,j} \geq 0 \tag{9}$$

Advertiser k's decision variables are the *valuations* $\{v_{k,j}\ \forall j\}$ i.e., the amount that k pays ad network j per click. Advertiser k's type is the ordered pair (y_k, R_k), where R_k is k's target *return on investment (ROI)*. In our simplified single-period model, advertiser k's objective is not to maximize profits – instead, we assume for simplicity that k's sole objective is to achieve an ROI of exactly R_k on every single ad network[11]. That is, k selects $v_{k,j}$ such that:

$$R_{k,j} = R_k \ \forall j \tag{10}$$

Intuitively, R_k represents the returns k can realize by advertising offline or through other media. Recall that advertisers enter bids on all ad networks.

Finally, ad network j's type is its effectiveness α_j, whereas j's decision variable is the level of aggressiveness x_j. Ad network j chooses x_j so that their profit η_j is maximized.

2.3 Numerical Example

To gain some intuition for the single-period model, we present a small numerical example. Note that this example is not an equilibrium scenario.

Consider a market with just $I = 2$ publishers, $J = 2$ ad networks and $K = 1$ advertiser. Assume $r_1 = 0.9$, $r_2 = 0.1$, and $V_1 = V_2 = 100$ i.e., the publishers receive the same volume of clicks but publisher 1's traffic is of much higher quality. Let $\alpha_1 = 0.322$ and $\alpha_2 = 0.555$ i.e., ad network 1 is more effective than ad network 2 at detecting invalid clicks.

[11]In Section 3, we present a more realistic (but slightly more complicated) model where advertisers choose $v_{k,j}$ to maximize profits, under the constraint that their ROI is at least R_k.

Now, suppose $x_1 = 0.5$ and $x_2 = 0.4$ i.e., ad network 1 is more aggressively detecting invalid clicks than ad network 2. Using equation (1), we can compute $N_{i,j}$ for $i = 1, 2$ and $j = 1, 2$. The results are shown in Table 2. From Table 2, we see that ad network 2 would mark more clicks valid than ad network 1, for both publishers

	Ad Network 1	Ad Network 2
Publisher 1	47%	58%
Publisher 2	23%	42%

Table 2: Fraction $N_{i,j}$ of publisher i's clicks marked valid by ad network j.

1 and 2. Note that even though only 10% of publisher 2's clicks are valid, ad network 2 marks 42% of them valid. Thus, ad network 2 will be billing advertisers for many invalid clicks that do not generate revenue.

Suppose that $v_{1,1} = 15$ and $v_{1,2} = 10$ i.e., the advertiser is willing to pay \$15 per click on ad network 1 and \$10 per click on ad network 2. The advertiser is willing to pay more per click on ad network 1 than 2 because ad network 1 is giving more clicks away for free (i.e., marking fewer clicks valid), leading to a higher ROI. Again, this is just a numerical example – we are not claiming that these are equilibrium valuations. Assume $h = 0.4$ i.e., both ad networks pay out 40% of their revenue to publishers.

Define $\pi_{i,j} \equiv \pi_{i,j,1} \mid c_{i,j} = 1$. From equation (3), $\pi_{i,j} = N_{i,j} V_i v_{1,j} h$. Since $K = 1$ in our example, we have $\pi_{i,j} = \pi_i \mid c_{i,j} = 1$ i.e., $\pi_{i,j}$ is the total profit to publisher i if it decides to sell all of its clicks to ad network j exclusively. Table 3 shows $\pi_{i,j}$ for $i = 1, 2$ and $j = 1, 2$, assuming $v_{1,1} = 15$ and $v_{1,2} = 10$. We see that

	Ad network $j = 1$	Ad network $j = 2$
Advertiser's valuation, $v_{1,j}$	\$15.00	\$10.00
Publisher 1's profit, $\pi_{1,j}$	\$282	\$232
Publisher 2's profit, $\pi_{2,j}$	\$138	\$168

Table 3: Publisher profits $\pi_{i,j}$ assuming $v_{1,1} = 15.00$ and $v_{1,2} = 10.00$.

publisher 1 would earn \$282 from ad network 1 for all of its traffic compared to only \$232 from ad network 2. Publisher 2 would earn \$168 from ad network 2 compared to only \$138 from ad network 1. Therefore, based on the valuations $v_{1,1} = 15$ and $v_{1,2} = 10$, publisher 1 would send traffic to ad network 1 and publisher 2 would send traffic to ad network 2.

The important features of this numerical example are as follows. Ad network 2 marks more clicks valid than ad network 1, for both publishers 1 and 2. However, the advertiser is willing to pay more for ad network 1's clicks than ad network 2's clicks, since per-click ROI is higher on ad network 1. Based on total profits, publisher 1 would prefer to send its traffic to ad network 1 (i.e., since $\pi_{1,1} > \pi_{1,2}$), whereas publisher 2 would prefer to send its traffic to ad network 2 (i.e., since $\pi_{2,2} > \pi_{2,1}$).

2.3.1 A Subgame Equilibrium

We can carry this numerical example a bit further to predict what would happen in equilibrium.

Suppose publisher 1 does indeed send all of its clicks to ad network 1 and publisher 2 sends its clicks to ad network 2. The publishers' profits computed in Table 3 assumed $v_{1,1} = 15$ and $v_{1,2} = 10$. However, the advertiser had not yet accounted for the difference in traffic quality between the ad networks. Ad network 1's clicks (which come from publisher 1) are much higher quality than ad network 2's clicks (which come from publisher 2). The advertiser, therefore, will adjust $v_{1,1}$ and $v_{1,2}$ to reflect this quality difference.

Using equations (4), (5), (6) and (10), the advertiser decides that it is now willing to pay $v_{1,1} = 19.15$ per click to ad network 1, and only $v_{1,2} = 2.38$ per click to ad network 2. Using these new valuations, we re-calculate the publishers' profits. The results are shown in Table 4. From Table 4, we see that both publisher 1 and 2 make higher profits with ad network 1 than 2, so they both would now send all their traffic to ad network 1.

	Ad Network 1	Ad Network 2
Advertiser's valuation, $v_{1,j}$	\$19.15	\$2.38
Publisher 1's profit, $\pi_{1,j}$	\$360	\$55
Publisher 2's profit, $\pi_{2,j}$	\$176	\$40

Table 4: Publisher profits $\pi_{i,j}$ assuming $v_{1,1} = 19.15$ and $v_{1,2} = 2.38$.

However, publisher 2's decision to use ad network 1 instead of 2 will lower ad network 1's traffic quality. Thus, the advertiser will again re-evaluate what it is willing to pay per click on each network. It now sets $v_{1,1} = 14.29$ and $v_{1,2} = 5$. Re-calculating the values in Table 4, we get Table 5. Observe that under the updated valuations, both publishers' choices remain unchanged – they both would still display ad network 1's ads. As such, the advertiser no longer needs to re-compute its valuations. Therefore, $v_{1,1} = 14.29$ and $v_{1,2} = 5$, with both publishers choosing ad network 1 (i.e., $c_{1,1} = c_{2,1} = 1$) is an *equilibrium* for the subgame when $x_1 = 0.5$ and $x_2 = 0.4$.

	Ad Network 1	Ad Network 2
Advertiser's valuation, $v_{1,j}$	\$14.29	\$5.00
Publisher 1's profit, $\pi_{1,j}$	\$269	\$116
Publisher 2's profit, $\pi_{2,j}$	\$131	\$84

Table 5: Publisher profits $\pi_{i,j}$ assuming $v_{1,1} = 14.29$ and $v_{1,2} = 5.00$.

The interesting feature of this equilibrium is that both publishers agree on their choice of ad network. As we show next, this effect is not at all specific to this numerical example – it applies irrespective of the numbers of publishers, ad networks and advertisers, and irrespective of the distribution of their types.

2.4 Equilibria

We will now characterize the equilibria of our single-period sequential game. Recall that the ad networks act first i.e., in the first step, each ad network j must choose x_j without knowing the publishers' and advertisers' decisions. In the subgame[12], publishers and advertisers simultaneously decide on their allocations and valuations, respectively. Publishers and advertisers know $\{x_j \ \forall j\}$ (i.e., the outcome of the first step) before making their decisions in the second step. In this sense, our single-period sequential game captures how market participants react to an ad network's choice of aggressiveness level.

Assumption 1. *Publishers are not homogeneous with respect to quality of clicks:*

$$\exists m, n \ s.t. \ r_m \neq r_n \tag{11}$$

[12] Recall that we refer to the second step as the *subgame*, while the complete sequential game including both the first and second steps is referred to as the *supergame*.

Assumption 1 will hold in any non-trivial problem instance. Define ad network j's *multiplier*, M_j, as the ratio of the total number of converted clicks to the total number of clicks j marked valid, assuming $c_{i,j} = 1 \ \forall i$:

$$M_j \equiv \frac{\beta \sum_i r_i V_i}{\sum_i N_{i,j} V_i} \tag{12}$$

Under Assumption 1, in equilibrium, all publishers will choose the same ad network in the subgame:

Theorem 1. *Suppose Assumption 1 holds. Then, for any first-step outcome, there exist pure-strategy Nash equilibria (NE) in the subgame. Moreover, in any pure-strategy subgame NE, there exists a $j^* \in 1, \ldots, J$ such that*

$$\begin{aligned} c_{i,j^*} &= 1 \quad \forall i && (13) \\ v_{k,j^*} &= \frac{y_k}{R_k} M_{j^*} \quad \forall k, && (14) \end{aligned}$$

whereas for all $j \neq j^*$,

$$\begin{aligned} c_{i,j} &= 0 \quad \forall i && (15) \\ v_{k,j} &\leq \frac{y_k}{R_k} M_j \quad \forall k. && (16) \end{aligned}$$

We say that ad network j^* has been "chosen" by publishers and advertisers in that equilibrium.

Proof. See Appendix A.1. □

Theorem 1 is a very strong statement. It holds for any number of publishers, ad networks and advertisers, and irrespective of the distribution of player types. Equation (13) says that all publishers will choose the same ad network j^* in the subgame – unlike the example in Section 2.3, publishers will never disagree (in equilibrium) on which ad network they choose. Or, from the ad networks' perspective, a single network will win 100% market share, and the rest will get no business at all from the publishers. From equation (14), advertiser k's equilibrium valuation v_{k,j^*} will be a product of two factors: an advertiser-specific quantity (i.e., $\frac{y_k}{R_k}$) and an ad-network-specific quantity (i.e., M_{j^*}). For ad networks $j \neq j^*$, advertiser k's valuation is given by the inequality (16).

Theorem 1 does not claim that there is a unique equilibrium. Many equilibria exist, and they differ both in which ad network is chosen and the advertiser valuations on those ad networks that aren't chosen. In fact, for any ad network j' and any first-step outcome $\{x_j \ \forall j\}$, there exists an equilibrium for the subgame in which $j^* = j'$ in Theorem 1. In words, irrespective of what the ad networks do in the first step, any ad network can be chosen by publishers in the second step. Theorem 1 only says that all publishers will agree on their choice[13]. Thus, Theorem 1 alone cannot predict which ad network will be chosen, or how aggressive ad networks will be.

To make a stronger prediction than Theorem 1, we need the notion of average publisher quality. Define the *average publisher quality*, $\bar{r}$, as follows:

$$\bar{r} \equiv \frac{\sum_i r_i V_i}{\sum_i V_i} \tag{17}$$

[13] In the numerical example in Section 2.3.1, we studied an equilibrium in which ad network 1 was chosen by the publishers.

We say that a publisher i is *high-quality* if $r_i \geq \bar{r}$. Conversely, a publisher i is *low-quality* if $r_i < \bar{r}$.

Although any ad network j can be chosen in equilibrium, a given publisher will find some equilibria to be more profitable than others. For a fixed first-step outcome $\{x_j \ \forall j\}$, define $\bar{\pi}_{i,j}$ as the profit to publisher i in an equilibrium in which ad network j is chosen by all publishers. We say that publisher i "prefers" ad network j if $\bar{\pi}_{i,j} \geq \bar{\pi}_{i,j'} \ \forall j' \neq j$. In words, publisher i will prefer ad network j if, amongst the set of all equilibria, publisher i's revenues are highest in equilibria where ad network j is chosen.

The following lemma says that all high-quality publishers will agree on which equilibria they prefer:

Lemma 1. *If publisher i is high-quality, it will prefer ad network j if and only if*

$$\frac{1 - x_j}{1 - x_j^{\alpha_j}} \geq \frac{1 - x_n}{1 - x_n^{\alpha_n}} \quad \forall n \neq j \tag{18}$$

Proof. See Appendix A.2. □

That is, all high-quality publishers prefer the ad network j that delivers the highest ratio of true negatives (i.e., $1 - x_j$) to false negatives (i.e., $1 - x_j^{\alpha_j}$). As such, we can interpret the ratio $\frac{1-x_j}{1-x_j^{\alpha_j}}$ as a measure of "ad-network quality". Lemma 1 simply says that high-quality publishers prefer high-quality ad networks.

Assumption 2. *The ad network chosen in the subgame is the one preferred by high-quality publishers.*

Under assumptions 1 and 2, publishers will choose the ad network that is most effective at detecting invalid clicks, as long as the ad network is sufficiently aggressive:

Theorem 2. *Suppose Assumptions 1 and 2 hold, $J \geq 2$, and that ad network 1 is the most effective at detecting invalid clicks i.e., $\alpha_1 < \alpha_j \ \forall j \neq 1$. Then,*

1. *There exists $x^* \in (0, 1)$ such that if $x_1 > x^*$, then $c_{i,1} = 1 \ \forall i$, irrespective of what the other ad networks do (i.e., irrespective of $\{x_j \ \forall j \neq 1\}$). Therefore, it is a dominant strategy for ad network 1 to choose $x_1 > x^*$ in the first step.*

2. *Suppose ad network 2 is the second-most effective at detecting invalid clicks i.e., $\alpha_2 \leq \alpha_j \ \forall j \geq 2$. Then, as $\alpha_2 - \alpha_1 \to 0$, we get $x^* \to 1$.*

3. *Ad network 1's total revenues are the same for all $x^* < x_1 < 1$. As $x_1 \to 1$, high-quality publishers earn an increasingly large share of the total publisher revenues.*

Proof. See Appendix A.3. □

The intuition behind Theorem 2 is as follows. By choosing an x_1 that is sufficiently large, ad network 1 is able to deliver the highest-quality traffic (Lemma 1) to its advertisers. Since high-quality publishers decide which equilibrium is played (Assumption 2), ad network 1's high-quality traffic is enough to win the market over. In particular, it is suboptimal to simply ignore invalid clicks (i.e., set $x_j = 0$) – high-quality publishers would react to an insufficiently aggressive ad network by choosing a more aggressive competitor instead. If ad network 1's detection algorithms are only slightly better than its competitors', it is forced to mark a large fraction of clicks invalid in order to win over the market. High-quality publishers prefer ad networks that filter aggressively, since they would disproportionately benefit compared to low-quality publishers.

Figure 2 gives some graphical intuition for Theorems 1 and 2. Consider a market with only $J = 2$ ad networks, where $\alpha_1 = 0.20$ and $\alpha_2 = 0.25$. The set of possible outcomes (x_1, x_2) for the first-step is depicted in Figure 2(a). According to Theorem 1, depending on the first-step outcome, all publishers will choose either ad network 1 or 2 in the subgame. Using Assumption 2 and the inequality (18), "Region 1" in Figure 2(a) is the set of points (x_1, x_2) where ad network 1 is chosen, whereas "Region 2" is the set of outcomes where ad network 2 is chosen. The boundary between the two regions is the set of points where (18) holds with equality. For example, if $x_1 = 0.4$ and $x_2 = 0.6$, publishers will prefer ad network 1, which means the point $(0.4, 0.6)$ lies in Region 1.

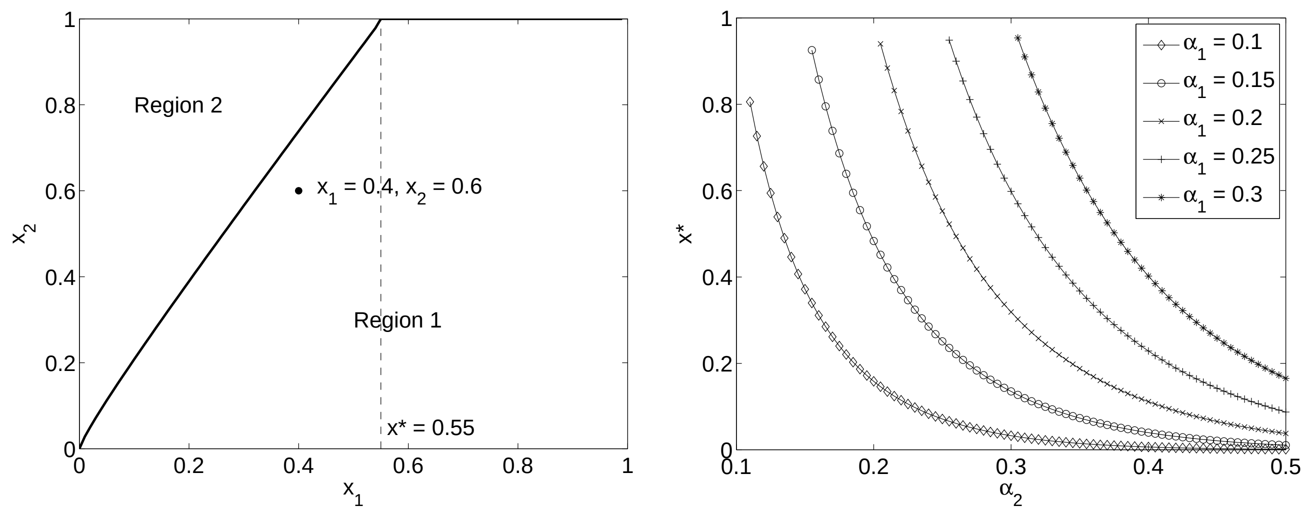

(a) Decision regions for $\alpha_1 = 0.20$, $\alpha_2 = 0.25$.

(b) Dependence of x^* on α_1 and α_2.

Figure 2: Graphical interpretations of Theorems 1 and 2.

Observe that for any $x_1 \leq 0.55$, there is an $x_2 \in [0, 1]$ that lies in Region 2. On the other hand, if $x_1 > 0.55$, no such x_2 exists. Therefore, $x^* = 0.55$ (see Theorem 2). If ad network 1 chooses $x_1 > x^*$, then there is no x_2 that ad network 2 can choose that will cause publishers to prefer it over ad network 1. It can be shown that ad network 1's total profit is the same for all points in Region 1. As such, it is a dominant strategy for ad network 1 to choose any $x_1 \in (x^*, 1)$.

Figure 2(b) illustrates the dependence of x^* on α_2, for various values of α_1. Each curve is obtained by letting $x_n \to 1$ in the right-hand side of (18), with $j = 1$ and $n = 2$. Then, x^* is the solution of the resulting polynomial in x_1, solved for various (α_1, α_2). On each curve, as $\alpha_2 \to \alpha_1$, note that $x^* \to 1$. As the gap between α_1 and α_2 increases, x^* decreases i.e., ad network 1 can get away with being less aggressive. Note that the plots in Figure 2 do not depend on the numbers of publishers and advertisers, or their types.

Observe that an analogue to Theorem 2 would not hold for low-quality publishers i.e., if low-quality publishers were to decide which equilibrium is played, the inequality in (18) would be reversed. However any ad network can achieve the minimum by setting $x_j = 0$. That is, only an effective, aggressive ad network can satisfy high-quality publishers, but any ad network can appease the low-quality publishers by simply not filtering at all.

3 Multi-Period Model

The main intuition of the single-period model in Section 2 was that, if all other factors remain equal, the ad network that is most effective at detecting invalid clicks would win over 100% of publishers, provided it is sufficiently aggressive. In particular, the ad network that could deliver the highest ratio of true negatives to false negatives wins.

Recall that the single-period model made the following simplifying assumptions:

1. It is a single-period game.
2. Ad networks act first, followed by the publishers and advertisers.
3. There is no private information.
4. There are no auctions for click-throughs.

In this section, we will relax all of these assumptions and obtain a new, more realistic multi-period model of the market. Fortunately, the results, intuitions and proof techniques developed for the single-period model are mostly still applicable in this new setting. Rather than presenting the model in complete detail, we will focus on the differences between the multi-period and single-period models.

3.1 Model

Consider once again a market comprised I publishers, J ad networks and K advertisers. We model the online advertising market as a multi-period dynamic game, played over a sequence of discrete time periods, indexed by $t \in 1, \ldots, \infty$. We model each period t as a two-step sequential game, which we refer to as the *stage game*. The two steps in the stage game are as follows:

1. At the beginning of period t, publishers decide which ad networks to sell their clicks on in period t, and advertisers decide how much they are willing to pay for clicks from each ad network in period t. Publishers and advertisers make their decision without knowing how aggressive each ad network will be.
2. At the end of period t, ad networks decide how aggressively to detect invalid clicks from their period t traffic.

Each period t represents a "billing cycle" for the ad networks (e.g., one month). At the beginning of each billing cycle, advertisers design ads and enter bids on queries, and publishers publish code snippets on their sites to request ads from their chosen ad networks. In the middle of the cycle, users visit websites, click on ads and sometimes buy products. At the end of each billing cycle, ad networks examine their logs and try to detect which click-throughs were invalid, advertisers pay for ones marked valid, and publishers receive their share of the revenue. Observe that the two-step stage game in the multi-period model is very similar to the two-step sequential game in Section 2. However, there is one key difference: in the multi-period model, the publishers and advertisers act first, whereas in Section 2 the ad networks acted first. The multi-period model is more realistic, in view of the billing cycle interpretation of each period t.

Table 6 is a summary of the new notation we will introduce in Section 3.

Symbol	Description
$F_{y,R}$	Joint probability distribution over advertisers' type
$c_{i,j,t}$	Fraction of publisher i's clicks sent to ad network j in period t
$x_{j,t}$	Aggressiveness (false positive rate) of ad network j in period t
$v_{k,j,t}$	Advertiser k valuation of ad network j's clicks in period t
$c_{i,j,t}^{(k)}$	Advertiser k's forecast of $c_{i,j,t}$
$v_{k,j,t}^{(i)}$	Publisher i's forecast of $v_{k,j,t}$
$x_{j,t}^{(i)}$	Publisher i's and advertiser k's forecast of $x_{j,t}$
$\hat{x}_{j,t}$	Publishers' and advertisers' common forecast of $x_{j,t}$
$N_{i,j,t}$	Fraction of i's clicks marked valid by j in period t
$\pi_{i,j,t}$	Publisher i's revenue from clicks sold via ad network j in period t
$\pi_{i,t}$	Publisher i's total revenue in period t
$\xi_{k,j,t}$	Fraction of ad network j's clicks send to advertiser k in period t
$Y_{k,j,t}$	Advertiser k's revenue from ad network j's clicks in period t
$Z_{k,j,t}$	Number of clicks advertiser k's is billed for by ad network j in period t
$R_{k,j,t}$	Advertiser k's ROI on ad network j in period t
$\theta_{j,t}$	Ad network j's expected revenue per click in period t
$\eta_{j,t}$	Ad network j's total revenue in period t

Table 6: List of variables introduced in Section 3.

Publisher i's type is (r_i, V_i), just like in the single-period model. It receives V_i clicks per period, of which a fraction r_i are valid. Ad network j's period-t aggressiveness is $x_{j,t}$, and j's effectiveness (i.e., its type) is α_j. The fraction of i's clicks marked valid by j in period t is, therefore:

$$N_{i,j,t} \equiv (1 - x_{j,t})r_i + (1 - x_{j,t}^{\alpha_j})(1 - r_i) \tag{19}$$

Publisher i sends a fraction $c_{i,j,t}$ of its clicks to ad network j in period t. The total number of i's clicks marked valid by j is then $N_{i,j,t}V_i c_{i,j,t}$.

3.2 Advertisers

Advertiser k's type is still (y_k, R_k). However, in the multi-period model, y_k and R_k are private information to advertiser k. In particular, y_k and R_k are not known to ad networks or publishers. As we discuss in Section 3.3, ad networks use auctions as a mechanism to extract this information from each advertiser k. On the other hand, the joint probability distribution (denoted $F_{y,R}$) of y and R over the population of advertisers is public information. Clearly, modeling the advertisers' conversion revenues and ROI targets as private information is more realistic.

Let $\xi_{k,j,t}$ be the fraction[14] of ad network j's clicks sent to advertiser k in period t. Ad network j receives $\sum_i r_i V_i c_{i,j,t}$ clicks in total in period t. Assuming a conversion rate of β, advertiser k earns

$$Y_{k,j,t} \equiv \left(\sum_i r_i V_i c_{i,j,t} \right) \xi_{k,j,t} \beta y_k \tag{20}$$

[14]We will never need to compute the exact value of $\xi_{k,j,t}$ since it "cancels out" in any calculation that is important for our purposes (e.g., equation (22)).

dollars of revenue in period t from clicks coming from ad network j. The total number of clicks that k is billed for by j in period t is:

$$Z_{k,j,t} \equiv \left(\sum_i N_{i,j,t} V_i c_{i,j,t} \right) \xi_{k,j,t} \tag{21}$$

Suppose $v_{k,j,t}$ is the amount k is willing to pay j per click in period t. Then, k's period-t ROI would be:

$$R_{k,j,t} \equiv \frac{Y_{k,j,t}}{v_{k,j,t} Z_{k,j,t}} \tag{22}$$

3.3 Auctions

In the single-period model, $v_{k,j}$ was the actual dollar amount that k paid j per click. Since advertiser k's objective was simply to achieve an ROI each period of exactly R_k – no more, no less – it had no reason to falsely report its valuation $v_{k,j}$ to the ad networks. Each advertiser also received a fixed fraction $\frac{1}{K}$ of each ad network's clicks, so there was no connection between $v_{k,j}$ and the number of clicks k was allocated by j.

In the multi-period model, we are in a more realistic setting where advertisers aim to maximize profits (see (28)). If advertisers knew they would receive a $\frac{1}{K}$ fraction of clicks irrespective of $v_{k,j,t}$, then every advertiser would simply pick $v_{k,j,t} = 0$ every period – profits are clearly maximized if clicks are free, and ROI is infinite! Advertisers would have no incentive to report how much they are actually willing to pay for clicks. Ad network j might try charging k a fixed amount $\frac{y_k}{R_k} M_j$ per click (as in Theorem 1). However, doing so is impossible because y_k and R_k are private information, and not known to j. The solution is for ad networks to create a link between the valuation $v_{k,j,t}$ and the allocation $\xi_{k,j,t}$. Advertisers who are willing to pay more should receive a larger fraction of clicks. However, as it stands, ad network j does not know $v_{k,j,t}$.

To encourage advertisers to report their valuations, we assume that each ad network runs an auction to determine which advertiser receives each click. On each click, a fixed number advertisers are selected uniformly-at-random from the pool of advertisers who have entered a bid on a given query[15]. The click is then awarded to the highest bidder, and the amount billed is related to the randomly-selected advertisers' bids. The benefit of using auctions is that advertisers have an incentive to report $v_{k,j,t}$. The downside for ad networks is that, in most cases, advertiser k ends up paying less than $v_{k,j,t}$ per click.

Let $\theta_{j,t}$ be ad network j's expected revenue per click[16] in period t i.e., if j marks Z clicks valid in period t, then its period-t gross revenue (before paying publishers) would be $Z\theta_{j,t}$. Our focus in this paper is not on the specifics of the auction mechanism. For our purposes, we only assume the following:

1. The expected revenue per click is some function of advertisers' valuations i.e., $\theta_{j,t} = g_j(\{v_{k,j,t}\})$
2. All ad networks use the same auction mechanism i.e., $g_j = g \ \forall j$
3. The auction mechanism is *linear* in the following sense: if every single advertiser's valuation is scaled by a constant factor γ, then the expected revenue from the auction is also scaled by γ i.e., if $v_{k,j,t} \leftarrow \gamma v_{k,j,t} \ \forall (k,j)$, then $\theta_{j,t} \leftarrow \gamma \theta_{j,t}$

[15]Auction participants are chosen randomly for each click-through to ensure that the top bidder does not receive 100% of the clicks for a given query. This practice is referred to as "throttling".

[16]The revenue from any given click is a random variable because the bidders are randomly selected.

The first and third assumptions are quite weak, and hold for most auction mechanisms used in practice. We make the second assumption (i.e., $g_j = g \ \forall j$) because our focus here is on click fraud, rather than comparing the efficiency of ad networks' auction mechanisms. In particular, we do not assume that the auction mechanism is truthful, or any other property about the mapping from advertisers' valuations to bids.

From these assumptions alone, we can show that if advertisers compute $v_{k,j,t}$ optimally, then:

$$\theta_{j,t} = \kappa a_{j,t} \tag{23}$$

where κ is a constant across all ad networks, and $a_{j,t}$ is ad network j's *adjustment factor* in period t:

$$a_{j,t} \equiv \frac{\left(\sum_i r_i V_i c_{i,j,t}\right) \beta}{\sum_i N_{i,j,t} V_i c_{i,j,t}} \tag{24}$$

We omit the proof here. Intuitively, $a_{j,t}$ is the ratio of the total number of converted clicks on ad network j to the total number of clicks j marks valid, in period t. Changing the (common) auction mechanism only affects the value of κ. In Section 4, we discuss the conditions under which we can think of click-throughs, rather than impressions, as the items that are auctioned off.

3.4 Player Objectives

Publisher i's expected revenue from ad network j is:

$$\pi_{i,j,t} \equiv \left(N_{i,j,t} V_i c_{i,j,t}\right) h \theta_{j,t} \tag{25}$$

Publisher i's total period-t expected revenues are:

$$\pi_{i,t} \equiv \sum_j \pi_{i,j,t} \tag{26}$$

Therefore, in each period t, i chooses $\{c_{i,j,t} \ \forall j\}$ by solving the following optimization problem:

$$\max_{\{c_{i,j,t}\}} \pi_{i,t} \quad \text{s.t.} \quad \sum_j c_{i,j,t} = 1 \quad \text{and} \quad c_{i,j,t} \geq 0 \tag{27}$$

Advertiser k selects its valuation $v_{k,j,t}$ so that its revenues from ad network j are maximized, subject to a lower bound R_k on ROI:

$$\max_{v_{k,j,t}} Y_{k,j,t} \quad \text{s.t.} \quad R_{k,j,t} \geq R_k \tag{28}$$

It can be shown that if advertiser k chooses $v_{k,j,t}$ for each j by solving (28), its combined profits from both online and offline advertising would be maximized. Moreover, it can be shown that the constraint is binding at the optimum:

$$R_{k,j,t} = R_k \ \forall (j,t) \tag{29}$$

Intuitively, R_k represents the rate of return k can achieve by advertising through channels other than PPC. For example, if k can achieve an ROI higher than R_k by advertising in, say, print newspapers, then its advertising dollars are better spent on print than online.

Let $\eta_{j,t}$ be ad network j's total expected revenue in period-t:

$$\eta_{j,t} = (1-h)\left(\sum_i N_{i,j,t} V_i c_{i,j,t}\right)\theta_{j,t} \tag{30}$$

Each ad network j has a *discount factor*, $\delta_j \in [0,1)$, which is used to compute its total expected long-term revenue, η_j:

$$\eta_j = \sum_t (\delta_j)^t \, \eta_{j,t} \tag{31}$$

The discount factor δ_j is part of j's type. Its objective is to choose $\{x_{j,t}\ \forall (j,t)\}$ such that η_j is maximized. Note that in our multi-period model, publishers and advertisers are "greedy", in the sense that their objective is to maximize period-t profits. Ad networks, on the other hand, aim to maximize long-term profits. We make this simplifying assumption in order to focus on the interactions between ad networks, rather than possible collusions between publishers and/or advertisers.

Finally, publishers and advertisers base their decisions partially on *forecasts* of what other players will do in period t (how the forecasts are actually computed is not important here). At the start of each period, publisher i computes forecasts $\{v^{(i)}_{k,j,t}\ \forall (k,j)\}$ of each advertiser k's bid on each ad network j, as well as forecasts $\{x^{(i)}_{j,t}\ \forall j\}$ of the ad networks' actions. Publisher i's period-t then chooses allocations that are the best responses to $\{v^{(i)}_{k,j,t}\ \forall (k,j)\}$ and $\{x^{(i)}_{j,t}\ \forall j\}$ i.e., i solves (27) assuming its forecasts are accurate. Similarly, advertiser k computes forecasts $\{c^{(k)}_{i,j,t}\ \forall (i,j)\}$ of each publisher i's allocation to each ad network j, as well as forecasts $\{x^{(k)}_{j,t}\ \forall j\}$ of the ad networks' actions. Advertiser k's period-t valuations are then best responses to $\{c^{(k)}_{i,j,t}\ \forall (i,j)\}$ and $\{x^{(k)}_{j,t}\ \forall j\}$ i.e., k solves (28) under the assumption that its forecasts are accurate.

3.5 Equilibria

Thus far in Section 3, we have relaxed some key assumptions made in the single-period model in Section 2. We can now state generalized versions of Theorems 1 and 2 for the multi-period case. We omit the proofs since they are, in many respects, similar to the single-period model.

Assumption 3. *Publishers and advertisers agree on their forecast of ad networks' actions:*

$$x^{(i)}_{j,t} = x^{(k)}_{j,t} \equiv \hat{x}_{j,t} \quad \forall (i,j,k,t) \tag{32}$$

Assumption 4. *Publishers and advertisers are able to forecast each others actions accurately:*

$$c^{(k)}_{i,j,t} = c_{i,j,t} \quad \forall (i,j,k,t) \tag{33}$$

$$v^{(i)}_{k,j,t} = v_{k,j,t} \quad \forall (i,j,k,t) \tag{34}$$

Informally, Assumption 3 says that nobody has "insider information" regarding the ad networks. However, it does not say that the forecasts are accurate i.e., it certainly is possible that $\hat{x}_{j,t} \neq x_{j,t}$. Given Assumption

3, Assumption 4 is equivalent to simply assuming that a "Nash equilibrium" is played each period between publishers and advertisers in the first step of each stage game.

Under Assumptions 1, 3 and 4, we obtain an analogue of Theorem 1. That is, in each period, all publishers will choose the same ad network in equilibrium:

Theorem 3. *Suppose Assumptions 1, 3 and 4 hold. Then, there exist subgame perfect Nash equilibria (SPNE) for the multi-period game with the following property – In every period t, there exists a $j^* \equiv j^*(t)$ such that*

$$c_{i,j^*,t} = 1 \quad \forall i \tag{35}$$

$$v_{k,j^*,t} = \frac{y_k}{R_k} M_{j^*} \quad \forall k, \tag{36}$$

whereas for all $j \neq j^*$,

$$c_{i,j,t} = 0 \quad \forall i \tag{37}$$

$$v_{k,j,t} \leq \frac{y_k}{R_k} M_j \quad \forall k. \tag{38}$$

We say that ad network j^* has been "chosen" by publishers and advertisers in period t.

Proof. Omitted. □

Observe that j^* can vary between periods i.e., it is possible that $c_{i,j,t} = 1$ and $c_{i,j,s} = 0$ for $t \neq s$. In words, publishers may choose a different ad network each period. However, Theorem 3 assures us that all publishers will agree on their choice each period i.e., $c_{i,j,t} = c_{m,j,t}$ $\forall (i, m)$.

Assumption 5. *Publishers and advertisers are able to forecast ad networks' actions accurately:*

$$\hat{x}_{j,t} = x_{j,t} \quad \forall (j, t) \tag{39}$$

Stated differently, Assumption 5 says that ad networks behave "predictably". For example, Assumption 5 would hold in a "steady state" where $x_{j,t} = x_{j,t-1}$, so that $\hat{x}_{j,t} = x_{j,t-1} = x_{j,t}$.

Assumption 6. *The ad network chosen in each period t is the one preferred by high-quality publishers.*

Assumption 6 is simply an analogue of Assumption 2. Under assumptions 1, 3, 4, 5 and 6 we obtain an analogue to Theorem 2 for the multi-period model. In equilibrium, publishers will always choose the ad network that is most effective at detecting invalid clicks, as long as the ad network is sufficiently aggressive:

Theorem 4. *Suppose Assumptions 1, 3, 4, 5 and 6 hold, $J \geq 2$, $\delta_1 > 0$, and that ad network 1 is the most effective at detecting invalid clicks i.e., $\alpha_1 < \alpha_j$ $\forall j \neq 1$. Then,*

1. *There exists $x^* \in (0, 1)$ such that if $x_{1,t} > x^*$, then $c_{i,1,t} = 1$ $\forall (i, t)$, irrespective of what the other ad networks do (i.e., irrespective of $\{x_{j,t}$ $\forall (j, t), j \neq 1\}$). Therefore, it is a dominant strategy for ad network 1 to choose $x_{1,t} > x^*$ in every period t.*

2. *Suppose ad network 2 is the second-most effective at detecting invalid clicks i.e.,* $\alpha_2 \leq \alpha_j \ \forall j \geq 2$*. Then, as* $\alpha_2 - \alpha_1 \to 0$*, we get* $x^* \to 1$*.*

3. *Ad network 1's total revenues are the same for all* $x^* < x_{1,t} < 1$*. As* $x_{1,t} \to 1$*, high-quality publishers earn an increasingly large share of the total period-t publisher revenues.*

Proof. Omitted. □

4 Discussion

The single- and multi-period models presented in Sections 2 and 3 were somewhat complex, involving a large number of players and a rich decision space. To make our models analytically tractable and hone in on the incentives surrounding click fraud, we made a number of assumptions and modeling decisions, which we have not yet fully discussed. We devote this section to a discussion, albeit brief, of some of these issues.

4.1 Modeling Decisions

Single query market. Our models consider the market for a single search query. In real markets involving an entire collection of queries, we envision an entire "collection" of independently functioning markets. For example, ad network 1 might earn 100% market share for the query "digital camera", while ad network 2 might earn 100% share for the query "mortgage". We have not considered compound queries, or other interactions between queries.

Pricing model. We did not consider *cost-per-mille* (CPM) or *cost-per-acquisition* (CPA) pricing models in our analysis. In these pricing models, fraud occurs in different ways. We refer the reader to [2] to learn about other forms of online advertising fraud.

No direct deals. We did not explicitly model "direct deals" between publishers and advertisers. These direct agreements were more prevalent in the earlier days of the Internet, before the emergence of the larger ad networks and ad aggregators. Suppose a publisher with many users decides not to deal with any ad networks, and opts for direct deals with advertisers only. Advertisers would certainly be interested in advertising on this publisher's site, because of its large user base. In effect, this publisher would be starting up its own ad network.

Innocent publishers. We have assumed that publishers are "innocent" – they are unaware whether a given click is valid, and are therefore unable to send all their valid clicks to one ad network and all their invalid clicks to another. The innocent-publishers assumption is without loss of generality, since a publisher that can distinguish between valid and invalid traffic can simply be modeled as two distinct publishers – one that receives valid traffic, and one that generates a stream of fraudulent click-throughs. Similarly, publishers are assumed to be profit maximizers who do not collude with each other. Those that do collude can simply be combined into a single publisher.

Frictionless markets. Our models assume that switching between ad networks is frictionless. That is, there is no cost to publishers or advertisers for changing over to a different ad network.

Offline filtering. Our games only capture "offline" or "batch" detection of invalid clicks. In practice, detection algorithms that run offline (e.g., over a click log, at the end of a billing cycle) are much more effective than online ones, since offline algorithms can leverage temporal and statistical information, whereas the latter is typically limited to techniques such as IP address filtering or blacklists.

No investing in technology. Ad networks cannot invest in improving their invalid click detection algorithms. For example, an ad network j cannot hire more engineers in order to lower their α_j. Theorems 2 and 4 imply that if such investment is feasible, it is absolutely in an ad network's interest to invest heavily – we have seen that the ad network whose is most effective at detecting invalid clicks has a dramatic market advantage.

4.1.1 Impressions vs. Clicks

In our model, publisher i has an inventory of V_i clicks-throughs to allocate across the ad networks. On each click-through sent to ad network j, D advertisers are chosen at random. An auction is run between these D advertisers to see which one will receive the click. The advertiser then pays the ad network for the click, if it is marked valid.

In reality, however, the PPC workflow is somewhat different. Publisher i has P_i impressions (or, "page views") to allocate across the ad networks. On each ad impression, an auction is run between D_i randomly chosen advertisers, and the impression is sold to the A_i advertisers with the highest valuations (i.e., there are A_i ad slots on the webpage). Then, at most one advertiser (but usually none) will have its ad impression become a click-through (although some click-throughs will be invalid). Payment is collected from the advertiser only if the ad impression is clicked on.

Under the following conditions, it is equivalent to model clicks (rather than impressions) as the object being bought and sold:

- D_i and A_i don't vary across impressions, for a fixed publisher.
- D_i and A_i don't vary across ad networks, for a fixed publisher.
- Selection of the D_i advertisers is independent and identical across impressions, for a fixed publisher.
- All impressions for a fixed publisher have equal probability of being clicked on.
- For any impression that has been clicked on, the probability of being valid is exactly r_i.
- The click-through rate for a given advertiser does not vary across ad networks.

The first four conditions imply that, for a fixed publisher, the revenue generated from each impression (possibly zero) is an independent identically-distributed random variable. The final two conditions imply that there is a fixed volume V_i of click-throughs, of which a fraction r_i are valid. Thinking of click-throughs as the item being bought and sold allows us, for example, to sidestep the issue of click-through rates and impression spam.

4.1.2 ROI Targeting

The decision to model advertisers as agents that maximize profits subject to a lower bound on ROI merits special attention. Advertisers enter bids on each ad network in order to attain their target ROI on every ad network. An advertiser would be willing to pay for an unlimited number of clicks as long as the return on their investment in these clicks is sufficiently high. Therefore, in principle, advertisers have infinite budgets.

It is, of course, possible (and seemingly more natural) to model advertisers as allocating a finite advertising budget across the ad networks. However, in practice, ROI-targeting is a more common behaviour amongst advertisers. It is actually quite rare that advertising budgets are ever exhausted. Advertisers typically are willing to pay for more click-throughs if they are available, since doing so simply means they are making more money from clicks becoming conversions.

In fact, in can be shown via a simple economic argument that in a budget-allocation model, if an advertiser's budget for click-throughs is being exhausted, it is rational for the advertiser to simply lower their bid. Doing so would result in strictly more click-throughs for the exact same cost.

4.2 Simplifying Assumptions

The results in this paper rely on a number of simplifying assumptions. These assumptions can be divided into two groups: "essential" and "non-essential". Essential assumptions are those without which our results would not hold (our would hold only approximately). Viewed differently, deviations in the real world from our models' predictions can be explained by the extents to which our essential assumptions are violated. Non-essential assumptions are those that we can relax easily, and have our key results still hold.

4.2.1 Essential Assumptions

Equal revenue shares. All ad networks are assumed to offer the same revenue share, h, to their publishers. It is intuitive why this assumption is essential – if an ad network does not have the most effective invalid click detection algorithm, one way to draw publishers over to its network is to simply offer them a larger fraction of the revenue. It can be shown analytically that such a tactic can, in fact, attract publishers (albeit low-quality ones) to a lower-quality ad network. In such cases, *market segmentation* amongst the publishers is observed, contrary to the predictions of Theorems 1 and 3.

No discounts. Similarly, we disallow the practice of offering advertisers discounts on particular publishers' traffic based on their conversion rates[17]. In practice, such discounts would alter the bids submitted by advertisers, and would therefore impact the attractiveness of various ad networks to market participants. Since we assume that conversion rates are constant across all publishers, advertisers and ad networks, such a tactic would not be ineffective in our model.

Assumption 2. The real-world applicability of Assumption 2 merits comment. We argue that the assumption is quite realistic. Large publishers tend to have very much influence over the market, in practice. Larger, higher-quality publishers also have more money at stake, and consequently may behave "more rationally".

4.2.2 Non-Essential Assumptions

ROC curve. Our models assumed, for simplicity, that if an ad network j's aggressiveness is x_j and its effectiveness is α_j, then its true positive rate is $f(x_j, \alpha_j) \equiv x_j^{\alpha_j}\ \forall j$. However, our results apply for any function

[17]This practice is referred to as "smart pricing" by Google and "quality-based pricing" by Yahoo!.

$f(x_j, \alpha_j)$ that is non-decreasing over $x_j \in [0,1]$ and non-increasing over $\alpha_j \in [0,1]$. In the most general case, an ad network's effectiveness would be described by the entire curve $f_j(x_j)$. The function f_j is known as a *receiver operating characteristic* (ROC).

No sticky publishers. There are no "sticky" publishers in our model. A sticky publisher is an "irrational" publisher who decides to display ads from a particular ad network, even though it may not be the most profitable one. A common example is a search engine company that is both a publisher and an ad network. Another example is a publisher that colludes with a particular ad network. If publisher i is sticky in its choice of ad network j, we can simply fix $c_{i,j} = 1$ – our results remain qualitatively unchanged.

Fixed click-through rates. To reduce the number of parameters in our model, we assumed that all ad impressions have equal probability of being clicked on. Moreover, click-through rates do not vary across ad networks i.e., V_i for each publisher i is fixed irrespective of whose ads are displayed. We also assume that V_i is also not time-varying. That is, spammers don't target different websites over time.

Time-invariant types. We assumed in the multi-period model that the players' types are time-invariant. That is, $r_i, V_i, \alpha_j, \delta_j, y_k$ and R_k do not vary across time periods (e.g., $r_{i,t}$. However, the theorems in Section 3 can easily be generalized to account for the time-varying case. Essentially, we apply Theorem 1 on a period-by-period basis, and recognize that playing an Nash equilibrium in every stage game is a subgame perfect Nash equilibrium for the supergame.

5 Conclusion

We have presented an economic model of the pay-per-click advertising market, focusing on the effects of detecting invalid clicks (and click fraud in particular). Our results suggest that, indeed, letting fraud go unchecked (i.e., choosing $x_j = 0$) is suboptimal for ad networks – a network that sets $x_j = 0$ would lose business to a more aggressive competitor. In particular, under reasonable assumptions on how the market functions, the ad network that can filter most effectively has a significant competitive advantage.

In practice, no ad network is earning 100% market share. On the other hand, publishers in the real world do often choose the most profitable ad network, and do switch ad networks when revenue prospects seem higher. So, to the extent that our assumptions hold true and publishers and advertisers act rationally, we conjecture that our predictions hold true in practice. Accounting for differences between ad networks in revenue sharing, ad targeting and "quality-based pricing" may also explain deviations from our predictions, and would be a promising extension to our model.

References

[1] Bogatin, D. Google CEO on click fraud: 'Let it happen' is perfect economic solution. *ZDNet* (July 2006).

[2] Daswani, N., Mysen, C., Rao, V., Weis, S., Gharachorloo, K., and Ghosemajumder, S. *Crimeware*. Addison-Wesley Professional, February 2008, ch. Online Advertising Fraud. To Appear.

[3] Daswani, N., and Stoppelman, M. The anatomy of clickbot.A. In *Hot Topics in Understanding Botnets (HotBots)* (April 2007), Usenix.

[4] Gandhi, M., Jakobsson, M., and Ratkiewicz, J. Badvertisements: Stealthy click-fraud with unwitting accessories. *Journal of Digital Forensic Practice 1*, 2 (2006), 131–142.

[5] Immorlica, N., Jain, K., Mahdian, M., and Talwar, K. Click fraud resistant methods for learning click-through rates. In *Internet and Network Econonomics* (November 2005), vol. 3828 of *Lecture Notes in Computer Science*, Springer, pp. 34–45.

[6] Mungamuru, B., and Weis, S. Competition and fraud in online advertising markets. In *Financial Cryptography and Data Security* (January 2008), vol. 5143 of *Lecture Notes in Computer Science*, Springer, pp. 187–191.

[7] Penenberg, A. L. Click fraud threatens web. *Wired News* (October 2004).

[8] Savitz, E. Eric Schmidt: Click fraud is self correcting. Mark Cuban: No, it isn't. *Barron's* (July 2006).

[9] Schneier, B. Google's click-fraud crackdown. *Wired News* (July 2006).

[10] Shankland, S. Yahoo to try using Google for search ads. *CNET News* (July 2006).

A Proofs

A.1 Proof of Theorem 1

A.1.1 Proof Outline

There are three main steps involved in proving Theorem 1:

1. For a fixed first-step outcome, derive the *best-response functions* for publishers and advertisers. (Sections A.1.3 - A.1.4)

2. Show that the best-response functions have a fixed point, implying the existence of a pure-strategy equilibrium in the subgame for any first-step outcome. (Section A.1.5)

3. Show that for any fixed point of the best-response functions, equations (13) - (16) hold. (Section A.1.6)

A.1.2 Notation

We will use the symbol $\mathbf{C}$ to denote a matrix whose (i, j)-element is publisher i's allocation $c_{i,j}$. Similarly, $\mathbf{V}$ is a matrix whose (k, j)-element is advertiser k's valuation $v_{k,j}$, and $\mathbf{x}$ is a vector whose element j is x_j. We also define $\mathbf{c}_i$ to be the length-J vector whose element j is $c_{i,j}$:

$$\mathbf{c}_i \equiv [c_{i,1}\ c_{i,2}\ \cdots\ c_{i,J}]^T \tag{40}$$

Define $S_1 \equiv [0,1]^{IJ} \times \mathbb{R}_+^{KJ} \times [0,1)^J$ and $S_2 \equiv [0,1]^{IJ} \times \mathbb{R}_+^{KJ}$. The symbol "$\times$" denotes a cross-product of sets. We refer to S_1 and S_2 as *decision spaces*, and elements of S_1 and S_2 as *decision profiles*. Each decision profile

$(\mathbf{C}, \mathbf{V}, \mathbf{x}) \in S_1$ is a possible outcome of the supergame (i.e., first- and second-steps), whereas each decision profile $(\mathbf{C}, \mathbf{V}) \in S_2$ is a possible outcome of the subgame (i.e., the second step). Each $\mathbf{x} \in [0, 1)^J$ is a possible outcome of the first step.

We assume throughout the proof that $x_j \in [0, 1)\ \forall j$ i.e., that $x_j = 1$ is not allowed. In words, no ad network marks 100% of its clicks invalid. This assumption allows us to sidestep some "corner cases" in our analysis. However, assuming $x_j < 1$ is by no means necessary for our results, since it can be shown that $x_j = 1$ never occurs in equilibrium.

Table 7 is a list of notation we introduce for the proof of Theorem 1.

Symbol	Description
$\mathbf{C}$	I-by-J matrix whose (i, j)-element is $c_{i,j}$
$\mathbf{V}$	K-by-J matrix whose (k, j)-element is $v_{k,j}$
$\mathbf{x}$	Length-J vector whose element j is x_j
$\mathbf{c}_i$	Length-J vector whose element j is $c_{i,j}$
S_1	Set of possible outcomes $(\mathbf{C}, \mathbf{V})$ for the subgame
S_2	Set of possible outcomes $(\mathbf{C}, \mathbf{V}, \mathbf{x})$ for the supergame
$BR^{Adv}_{k,j}$	Advertiser k's best response on ad network j
$BR^{Pub}_{i,j}$	Publisher i's best response on ad network j
BR	Best response function for all publishers and advertisers
a_j	Advertisers' adjustment factor for ad network j
$X_{i,j}$	Publisher i's expected revenue for sending a single click to ad network j
X^*_i	Publisher i's maximal per-click expected revenue
$\Phi(\{X_{i,j}\})$	Set of solutions $\mathbf{c}'_i$ to equations (48) - (50)
$(\mathbf{C}^*, \mathbf{V}^*)$	A fixed point of the best-response function, BR

Table 7: Notation used in Appendix A.

A.1.3 Advertiser k's Best Response

Fix a first-stage outcome, $\mathbf{x} \in [0, 1)^J$.

Let $BR^{Adv}_{k,j} : S_2 \to 2^{\mathbb{R}_+}$ be advertiser k's *best-response function* on ad network j. That is, for any decision profile $(\mathbf{C}, \mathbf{V}) \in S_2$, each element $v'_{k,j}$ of the set $BR^{Adv}_{k,j}(\mathbf{C}, \mathbf{V}; \mathbf{x})$ is an optimal choice of $v_{k,j}$ for advertiser k on ad network j, assuming $\mathbf{x}$ was played in the first stage and other players' choices in the subgame are given by $(\mathbf{C}, \mathbf{V})$. $BR^{Adv}_{k,j}$ is a set-valued function because k's optimal valuation $v'_{k,j}$ on ad network j need not be unique.

We will now derive an expression for $BR^{Adv}_{k,j}$. Fix a decision profile $(\mathbf{C}, \mathbf{V}) \in S_2$ and a particular ad network j. There are two possible cases: 1) $c_{i,j} \neq 0$ for some i, and 2) $c_{i,j} = 0\ \forall i$.

1. $c_{i,j} \neq 0$ for some i.

   Advertiser k chooses $v_{k,j}$ so that $R_{k,j} = R_k$. From (4), (5) and (6):

$$R_{k,j} = \frac{Y_{k,j}}{v_{k,j} Z_{k,j}} = \frac{\left(\sum_i r_i V_i c_{i,j}\right) \frac{1}{K} \beta y_k}{v_{k,j} \frac{1}{K} \sum_i N_{i,j} V_i c_{i,j}} \tag{41}$$

   Notice that the factor $\frac{1}{K}$ cancels out from the numerator and denominator in (41). The factor $\xi_{k,j,t}$ defined in Section 3 cancels out in a similar manner. As such, our theorems are independent of the schemes used by ad networks to distribute clicks.$\mathbf{c}_i$ Setting $R_{k,j} = R_k$ in (41) and solving for the optimal $v_{k,j}$, we get:

$$v'_{k,j} \equiv \frac{\left(\sum_i r_i V_i c_{i,j}\right) \beta y_k}{R_k \sum_i N_{i,j} V_i c_{i,j}} = \frac{y_k}{R_k} a_j \tag{42}$$

where we have defined the *adjustment factor*, a_j, as:

$$a_j \equiv \frac{\left(\sum_i r_i V_i c_{i,j}\right) \beta}{\sum_i N_{i,j} V_i c_{i,j}} \tag{43}$$

Intuitively, (42) says that advertiser k's optimal valuation $v'_{k,j}$, in response to $(\mathbf{C}, \mathbf{V}; \mathbf{x})$, is equal to their "nominal valuation" $\frac{y_k}{R_k}$, scaled by the per-ad network adjustment factor a_j. Observe that if $c_{i,j} = 1\ \forall i$, then $a_j = M_j$, where M_j was defined in (12).

2. $c_{i,j} = 0\ \forall i$.

   The ROI $R_{k,j}$ is undefined since publishers do not send any clicks to ad network j, and so j doesn't have any clicks to sell to advertiser k. The optimal valuation is not unique – in fact, any non-negative value of $v_{k,j}$ is reasonable. Therefore, when $c_{i,j} = 0\ \forall i$, we define $BR^{Adv}_{k,j}(\mathbf{C}, \mathbf{V}; \mathbf{x}) = [0, \infty)$.

Combining these two cases, we get:

$$v'_{k,j} \in BR^{Adv}_{k,j}(\mathbf{C}, \mathbf{V}; \mathbf{x}) = \begin{cases} [0, \infty) & \text{if } c_{i,j} = 0\ \forall i \\ \frac{y_k}{R_k} a_j & \text{if } c_{i,j} > 0 \text{ for some } i \end{cases} \tag{44}$$

A.1.4 Publisher i's Best Response

Fix a first-stage outcome, $\mathbf{x} \in [0, 1)^J$.

Let $BR^{Pub}_i : S_2 \to 2^{[0,1]^J}$ be publisher i's *best-response function* on ad network j. That is, for any given decision profile $(\mathbf{C}, \mathbf{V}, \mathbf{x}) \in S_1$, each element $\mathbf{c}'_i$ of the set $BR^{Pub}_{k,j}(\mathbf{C}, \mathbf{V}; \mathbf{x})$ is an optimal vector of allocations $\mathbf{c}_i$ for publisher i, assuming $\mathbf{x}$ was played in the first stage and other players' choices in the subgame are given by $(\mathbf{C}, \mathbf{V})$. BR^{Pub}_i is a set-valued function because i's optimal allocation $c'_{i,j}$ on each ad network j need not be unique.

We will now derive an expression for BR^{Pub}_i. Fix a decision profile $(\mathbf{C}, \mathbf{V}, \mathbf{x}) \in S_1$. From (9), publisher i chooses $\mathbf{c}_i$ to maximize its total profits, as follows:

$$\max_{\mathbf{c}_i} \sum_j \sum_k \pi_{i,j,k} \quad \text{s.t.} \quad \sum_j c_{i,j} = 1 \quad \text{and} \quad c_{i,j} \geq 0\ \forall j \tag{45}$$

From (3), $\sum_k \pi_{i,j,k}$ is linear in $c_{i,j}$:

$$\sum_k \pi_{i,j,k} = N_{i,j} V_i c_{i,j} \frac{1}{K} h \sum_k v_{k,j} \tag{46}$$

Therefore, (45) is very easy to solve – publisher i simply sends all its clicks to the single ad network (or multiple ad networks, if there is a tie) that will generate the highest expected revenue per click.

Unfortunately, formally expressing the solution to (45) is somewhat cumbersome notationally. Define $X_{i,j}$

as follows:

$$X_{i,j} \equiv N_{i,j} \frac{1}{K} h \sum_k v_{k,j} \tag{47}$$

$X_{i,j}$ is a key quantity in the proof of Theorem 1 – it is the criterion used by publisher i to decide what fraction of clicks to allocate to j. Intuitively, $X_{i,j}$ is i's expected revenue for sending a single click to j. $X_{i,j}$ is equal to $\sum_k \pi_{i,j,k}$ when $c_{i,j} = 1$ and $V_i = 1$.

Let $X_i^* \equiv \max_j X_{i,j}$. Any $\mathbf{c}_i$ that satisfies:

$$\begin{aligned} c_{i,j} &\geq 0 \text{ if } X_{i,j} = X_i^* & (48)\\ c_{i,j} &= 0 \text{ if } X_{i,j} < X_i^* & (49)\\ \sum_j c_{i,j} &= 1 & (50) \end{aligned}$$

is a best response to $(\mathbf{C}, \mathbf{V}, \mathbf{x})$ for publisher i.

Let $\Phi(\{X_{i,j}\})$ denote the set of vectors $\mathbf{c}_i$ that satisfy equations (48) - (50). Using this notation, we can write the solution to (45) as:

$$\mathbf{c}'_i \in BR_i^{Pub}(\mathbf{C}, \mathbf{V}; \mathbf{x}) = \Phi(\{X_{i,j}\}) \tag{51}$$

A.1.5 Existence of Fixed Points

In this section, we show that the best-response functions always have a fixed point. Define $BR : S_2 \to 2^{S_2}$ as follows:

$$\begin{aligned} BR(\mathbf{C}, \mathbf{V}; \mathbf{x}) \equiv\ & BR_1^{Pub}(\mathbf{C}, \mathbf{V}; \mathbf{x}) \times BR_2^{Pub}(\mathbf{C}, \mathbf{V}; \mathbf{x}) \times \cdots \times BR_I^{Pub}(\mathbf{C}, \mathbf{V}; \mathbf{x}) \times \\ & BR_{1,1}^{Adv}(\mathbf{C}, \mathbf{V}; \mathbf{x}) \times BR_{1,2}^{Adv}(\mathbf{C}, \mathbf{V}; \mathbf{x}) \times \cdots \times BR_{1,J}^{Adv}(\mathbf{C}, \mathbf{V}; \mathbf{x}) \times \\ & \cdots \times \\ & BR_{K,1}^{Adv}(\mathbf{C}, \mathbf{V}; \mathbf{x}) \times BR_{K,2}^{Adv}(\mathbf{C}, \mathbf{V}; \mathbf{x}) \times \cdots \times BR_{K,J}^{Adv}(\mathbf{C}, \mathbf{V}; \mathbf{x}) \end{aligned} \tag{52}$$

Each element $(\mathbf{C}', \mathbf{V}') \in BR(\mathbf{C}, \mathbf{V}; \mathbf{x})$ is a pair of matrices of optimal allocations $c'_{i,j}$ and valuations $v'_{k,j}$ in response to the decision profile $(\mathbf{C}, \mathbf{V}) \in S_2$, assuming $\mathbf{x}$ was played in the first step.

We claim that for any first-step outcome $\mathbf{x}$, $BR(\mathbf{C}, \mathbf{V}; \mathbf{x})$ has at least one fixed point i.e., for any $\mathbf{x} \in [0, 1)^J$ there exists $(\mathbf{C}^*, \mathbf{V}^*)$ such that:

$$(\mathbf{C}^*, \mathbf{V}^*) \in BR(\mathbf{C}^*, \mathbf{V}^*; \mathbf{x}) \tag{53}$$

Our claim can be proved by construction. Pick any integer $n \in 1 \ldots J$. From (44) and (48) - (50), it is easy to see that the following is always a fixed point of BR:

$$c_{i,j} = \begin{cases} 1 & j = n \\ 0 & j \neq n \end{cases} \qquad v_{k,j} = \begin{cases} \frac{y_k}{R_k} M_j & j = n \\ 0 & j \neq n \end{cases} \tag{54}$$

A fixed point of the form (54) exists for every $n \in 1 \ldots J$. The fact that this construction works for any n is why any ad network can be chosen in equilibrium.

A pure-strategy Nash equilibrium in the subgame is, by definition, a fixed point of BR. Therefore, for any first-stage outcome $\mathbf{x} \in [0,1)^J$, there exist pure-strategy Nash equilibria in the subgame.

A.1.6 Analysis of Fixed Points

In this section, we will show that for any fixed point of BR, equations (13) - (16) must hold, thereby proving Theorem 1.

Fix an $\mathbf{x}$ and let $(\mathbf{C}^*, \mathbf{V}^*)$ be any fixed point of BR i.e., $(\mathbf{C}^*, \mathbf{V}^*) \in BR(\mathbf{C}^*, \mathbf{V}^*; \mathbf{x})$. Let element (i,j) of $\mathbf{C}^*$ be $c_{i,j}$, and let element (k,j) of $\mathbf{V}^*$ be $v_{k,j}$. Going forward, we will only be interested in analyzing fixed points of BR. Unless specified otherwise, an allocation $c_{i,j}$ and valuation $v_{k,j}$ will always correspond to a fixed point $(\mathbf{C}^*, \mathbf{V}^*)$ of BR.

To prove Theorem 1, we only need to show that for any $(\mathbf{C}^*, \mathbf{V}^*)$, equations (13) and (15) will hold for any fixed point i.e., that $\exists j^*$ such that $c_{i,j^*} = 1\ \forall i$ and $c_{i,j} = 0\ \forall i, j \neq j^*$. If $c_{i,j^*} = 1\ \forall i$ and $c_{i,j} = 0\ \forall i, j \neq j^*$, (48) - (50) would imply that $X_{i,j^*} \geq X_{i,j}\ \forall i, j \neq j^*$. Equations (14) and (16) would then follow directly.

On our way to proving that equations (13) and (15) hold for any fixed point, we will state and prove Lemmas 2, 3 and 4. But first, we observe from equations (48), (49) and (50) that:

- If $c_{i,j} > 0$ then $X_{i,j} \geq X_{i,n}\ \forall n \neq j$
  - If publisher i allocates any clicks to ad network j, then i's per click revenue is highest on j (there might be a tie, though)
- If $X_{i,j} > X_{i,n}\ \forall n \neq j$ then $c_{i,j} = 1$
  - If publisher i's per click revenue on ad network j is strictly higher on j than any other ad network, i will sends all its clicks to j

We will use these observations repeatedly in the proofs of Lemmas 2, 3 and 4.

The assumption that $(\mathbf{C}^*, \mathbf{V}^*)$ is a fixed point of BR means that publishers and advertisers are both playing their best responses. Therefore, we can substitute (42) into (47). We get:

$$X_{i,j} = N_{i,j} \frac{1}{K} h \sum_k v_{k,j} \tag{55}$$

$$= N_{i,j} \frac{1}{K} h \frac{\beta \sum_i r_i V_i c_{i,j}}{\sum_i N_{i,j} V_i c_{i,j}} \sum_k \frac{y_k}{R_k} \tag{56}$$

$$= N_{i,j} \frac{1}{K} h \kappa \frac{\beta \sum_i r_i V_i c_{i,j}}{\sum_i N_{i,j} V_i c_{i,j}} \tag{57}$$

where we have defined $\kappa \equiv \sum_k \frac{y_k}{R_k}$.

Lemma 2. *If $\exists (i,j,m,n)$ such that $n \neq j$, $c_{i,j} > 0$, and $X_{m,j} > X_{m,n}$, then $c_{p,n} = 0\ \forall p$. It is sufficient but not necessary that $m = i$.*

In words, Lemma 2 says that if ad network j receives any traffic at all, and that at least one publisher strictly prefers ad network j over some other ad network n, then that ad network n will receive no traffic at all from any publisher.

Proof. WOLOG, assume $i = 1, j = 1, m = 1, n = 2, p = 2$ in the statement of Lemma 2. We will show that if $c_{1,1} > 0$ and $X_{1,1} > X_{1,2}$, then $c_{2,2} = 0$. Since the choice of (i, j, m, n, p) was arbitrary, we will obtain the desired result.

Assume that $c_{1,1} > 0$ and $X_{1,1} > X_{1,2}$. Suppose $c_{2,2} > 0$ – we will show that $c_{2,2} > 0$ leads to a contradiction, which implies $c_{2,2} = 0$.

By assumption:

$$\begin{aligned} X_{1,1} &> X_{1,2} \\ N_{1,1}\frac{1}{K}h\kappa\frac{\beta\left(\sum_i r_i V_i c_{i,1}\right)}{\sum_i N_{i,1}V_i c_{i,1}} &> N_{1,2}\frac{1}{K}h\kappa\frac{\beta\left(\sum_i r_i V_i c_{i,2}\right)}{\sum_i N_{i,2}V_i c_{i,2}} \\ N_{1,1}\frac{\sum_i r_i V_i c_{i,1}}{\sum_i N_{i,1}V_i c_{i,1}} &> N_{1,2}\frac{\sum_i r_i V_i c_{i,2}}{\sum_i N_{i,2}V_i c_{i,2}} \end{aligned} \tag{58}$$

Define, for convenience:

$$A_1 \equiv \sum_i r_i V_i c_{i,1} \tag{59}$$

$$A_2 \equiv \sum_i r_i V_i c_{i,2} \tag{60}$$

$$B_1 \equiv \sum_i (1 - r_i) V_i c_{i,1} \tag{61}$$

$$B_2 \equiv \sum_i (1 - r_i) V_i c_{i,2} \tag{62}$$

Since $c_{1,1} > 0$ and $c_{2,2} > 0$ by assumption, we know A_1, A_2, B_1, and B_2 are all strictly positive. Using equation (1) along with this notation, we observe that:

$$\sum_i N_{i,1}V_i c_{i,1} = (1 - x_1)A_1 + (1 - x_1^{\alpha_1})B_1 \tag{63}$$

$$\sum_i N_{i,2}V_i c_{i,2} = (1 - x_2)A_2 + (1 - x_2^{\alpha_2})B_2 \tag{64}$$

Therefore, the inequality (58) can be written as:

$$\frac{\left((1 - x_1)r_1 + (1 - x_1^{\alpha_1})(1 - r_1)\right)A_1}{(1 - x_1)A_1 + (1 - x_1^{\alpha_1})B_1} > \frac{\left((1 - x_2)r_1 + (1 - x_2^{\alpha_2})(1 - r_1)\right)A_2}{(1 - x_2)A_2 + (1 - x_2^{\alpha_2})B_2} \tag{65}$$

The numerators and denominators in (65) are all strictly positive. Therefore, rearranging (65) gives:

$$\begin{aligned} (1 - x_1)(1 - x_2^{\alpha_2})(r_1 A_1 B_2 - (1 - r_1)A_1 A_2) &+ \\ (1 - x_1^{\alpha_1})(1 - x_2)((1 - r_1)A_1 A_2 - r_1 A_2 B_1) &+ \\ (1 - x_1^{\alpha_1})(1 - x_2^{\alpha_2})((1 - r_1)A_1 B_2 - (1 - r_1)A_2 B_1) &> 0 \end{aligned} \tag{66}$$

Multiplying (66) through by $V_1 c_{1,1}$ gives:

$$\begin{aligned}(1-x_1)(1-x_2^{\alpha_2})(r_1 V_1 c_{1,1} A_1 B_2 - (1-r_1) V_1 c_{1,1} A_1 A_2) &\quad + \\ (1-x_1^{\alpha_1})(1-x_2)((1-r_1) V_1 c_{1,1} A_1 A_2 - r_1 V_1 c_{1,1} A_2 B_1) &\quad + \\ (1-x_1^{\alpha_1})(1-x_2^{\alpha_2})((1-r_1) V_1 c_{1,1} A_1 B_2 - (1-r_1) V_1 c_{1,1} A_2 B_1) &\quad > \quad 0 \end{aligned} \tag{67}$$

Let $\Delta_1 \equiv \{i \mid X_{i,1} \geq X_{i,2}\}$ i.e., the set of publishers that weakly prefer ad network 1 over ad network 2 – the assumption that $X_{1,1} > X_{1,2}$ means Δ_1 is non-empty. Note that $c_{i,1} = 0 \; \forall i \notin \Delta_1$, so that:

$$A_1 = \sum_i r_i V_i c_{i,1} = \sum_{i \in \Delta_1} r_i V_i c_{i,1} \tag{68}$$

$$B_1 = \sum_i (1-r_i) V_i c_{i,1} = \sum_{i \in \Delta_1} (1-r_i) V_i c_{i,1} \tag{69}$$

We can write an expression similar to (67) for each $i \in \Delta_1$:

$$\begin{aligned}(1-x_1)(1-x_2^{\alpha_2})(r_i V_i c_{i,1} A_1 B_2 - (1-r_i) V_i c_{i,1} A_1 A_2) &\quad + \\ (1-x_1^{\alpha_1})(1-x_2)((1-r_i) V_i c_{i,1} A_1 A_2 - r_i V_i c_{i,1} A_2 B_1) &\quad + \\ (1-x_1^{\alpha_1})(1-x_2^{\alpha_2})((1-r_i) V_i c_{i,1} A_1 B_2 - (1-r_i) V_i c_{i,1} A_2 B_1) &\quad \geq \quad 0 \end{aligned} \tag{70}$$

Therefore, the inequality will still hold when we sum over $i \in \Delta_1$:

$$\begin{aligned}\sum_{i \in \Delta_1} (1-x_1)(1-x_2^{\alpha_2})(r_i V_i c_{i,1} A_1 B_2 - (1-r_i) V_i c_{i,1} A_1 A_2) &\quad + \\ (1-x_1^{\alpha_1})(1-x_2)((1-r_i) V_i c_{i,1} A_1 A_2 - r_i V_i c_{i,1} A_2 B_1) &\quad + \\ (1-x_1^{\alpha_1})(1-x_2^{\alpha_2})((1-r_i) V_i c_{i,1} A_1 B_2 - (1-r_i) V_i c_{i,1} A_2 B_1) &\quad > \quad 0 \end{aligned} \tag{71}$$

Note that the inequality is strict since it is strict for $i = 1$ in (67). Moving the summation in (71) inside and simplifying gives:

$$\begin{aligned} 0 \;&<\; (1-x_1)(1-x_2^{\alpha_2})(A_1 A_1 B_2 - B_1 A_1 A_2) \\ &+\; (1-x_1^{\alpha_1})(1-x_2)(B_1 A_1 A_2 - A_1 A_2 B_1) \\ &+\; (1-x_1^{\alpha_1})(1-x_2^{\alpha_2})(B_1 A_1 B_2 - B_1 A_2 B_1) \\ &=\; (1-x_2^{\alpha_2})((1-x_1) A_1 + (1-x_1^{\alpha_1}) B_1)(A_1 B_2 - B_1 A_2) \end{aligned} \tag{72}$$

We know that the following quantity is strictly positive:

$$(1-x_2^{\alpha_2})((1-x_1) A_1 + (1-x_1^{\alpha_1}) B_1) \tag{73}$$

Therefore, so we conclude that:

$$(A_1 B_2 - B_1 A_2) > 0 \tag{74}$$

Now, since $c_{2,2} > 0$ by assumption, we know that $X_{2,2} \geq X_{2,1}$. Let $\Delta_2 \equiv \{i \mid X_{i,2} \geq X_{i,1}\}$ i.e., the set of publishers that weakly prefer ad network 2 over ad network 1 – the assumption that $c_{2,2} > 0$ means Δ_2 is non-empty. Note that $c_{i,2} = 0 \ \forall i \notin \Delta_2$, so:

$$\begin{aligned} A_2 &= \sum_i r_i V_i c_{i,2} = \sum_{i \in \Delta_2} r_i V_i c_{i,2} && (75) \\ B_2 &= \sum_i (1-r_i) V_i c_{i,2} = \sum_{i \in \Delta_2} (1-r_i) V_i c_{i,2} && (76) \end{aligned}$$

Using exactly analogous reasoning as above, we can write an expression similar to (70) for each $i \in \Delta_2$:

$$\begin{aligned} (1-x_1)(1-x_2^{\alpha_2})(r_i V_i c_{i,2} A_1 B_2 - (1-r_i) V_i c_{i,2} A_1 A_2) &\quad + \\ (1-x_1^{\alpha_1})(1-x_2)((1-r_i) V_i c_{i,2} A_1 A_2 - r_i V_i c_{i,2} A_2 B_1) &\quad + \\ (1-x_1^{\alpha_1})(1-x_2^{\alpha_2})((1-r_i) V_i c_{i,2} A_1 B_2 - (1-r_i) V_i c_{i,2} A_2 B_1) &\quad \leq \quad 0 \end{aligned} \tag{77}$$

Note that the direction of the inequality in (77) is opposite that of (70), since the publishers weakly prefer ad network 2 rather than 1.

Summing over $i \in \Delta_2$ now gives:

$$\begin{aligned} 0 \quad &\geq \quad (1-x_1)(1-x_2^{\alpha_2})(A_2 A_1 B_2 - B_2 A_1 A_2) \\ &+ \quad (1-x_1^{\alpha_1})(1-x_2)(B_2 A_1 A_2 - A_2 A_2 B_1) \\ &+ \quad (1-x_1^{\alpha_1})(1-x_2^{\alpha_2})(B_2 A_1 B_2 - B_2 A_2 B_1) \\ &= \quad (1-x_1^{\alpha_1})((1-x_2) A_2 + (1-x_2^{\alpha_2}) B_2), (49) and (A_1 B_2 - B_1 A_2) \end{aligned} \tag{78}$$

But (78) implies:

$$(A_1 B_2 - B_1 A_2) \leq 0 \tag{79}$$

which is a contradiction, due to (74). We conclude that $c_{2,2} > 0$ is impossible, which means $c_{2,2} = 0$.

□

Repeated application of Lemma 2 will yield Lemma 4, which is at the heart of proving Theorem 1. We will also need Lemma 3 in the proof of Lemma 4:

Lemma 3. *$\exists (j, n)$ such that $X_{i,j} = X_{i,n} \ \forall i$ if and only if $r_i = r \ \forall i$ for some constant r.*

In words, Lemma 3 says that all publishers will be indifferent between ad networks j and n if and only if all publishers have exactly the same quality, r.

Proof. If $X_{i,j} = X_{i,n} \ \forall i$, then it is easy to show that $\forall \ i$:

$$\frac{N_{i,j}}{N_{i,n}} = \frac{(x_j^{\alpha_j} - x_j) r_i + 1 - x_j^{\alpha_j}}{x_n^{\alpha_n} - x_n) r_i + 1 - x_n^{\alpha_n}} = \psi \tag{80}$$

where ψ is some constant. Therefore:

$$r_i = \frac{\psi(1 - x_n^{\alpha_n}) - (1 - x_j^{\alpha_j})}{(x_j^{\alpha_j} - x_j) - \psi(x_n^{\alpha_n} - x_n)} \equiv r \tag{81}$$

Conversely, if $r_i = r\ \forall i$, it is easy to show that $X_{i,j} = X_{i,n}\ \forall (j,n)$.

□

By repeated application of Lemma 2, we arrive at the next lemma:

Lemma 4. *If* $\exists (i,j)$ *such that* $c_{i,j} = 1$*, then* $c_{m,n} = 0\ \forall m \neq i, n \neq j$*.*

In words, Lemma 4 says that if publisher i sends ad network j all of its traffic, then no other ad network will receive any traffic from any publishers.

Proof. As discussed earlier, since $(\mathbf{C}^*, \mathbf{V}^*)$ is a fixed point of BR, we know from equations (48) - (50) that:

- If $c_{i,j} > 0$ then $X_{i,j} \geq X_{i,n}\ \forall n \neq j$.
- If $X_{i,j} > X_{i,n}\ \forall n \neq j$ then $c_{i,j} = 1$.

WOLOG, assume $i = 1, j = 1$. We will show that if $c_{1,1} = 1$, then $c_{m,n} = 0\ \forall m \neq 1, n \neq 1$. Since the choice of (i,j) was arbitrary, we will obtain the desired result.

Suppose $c_{1,1} = 1$. There are two possible cases: 1) $X_{1,1} > X_{1,n}\ \forall n \neq 1$; and 2) $X_{1,1} = X_{1,n}$ for some $n \neq 1$.

1. $X_{1,1} > X_{1,n}\ \forall n \neq 1$.

   Applying Lemma 2 with $i = j = m = 1$, we get for each $n \neq 1$ that $c_{p,n} = 0\ \forall p$.

2. $X_{1,1} = X_{1,n}$ for some $n \neq 1$.

   Let $\Delta \equiv \{n \mid n \neq 1,\ X_{1,n} = X_{1,1}\}$. Since $c_{1,1} > 0$, we know $X_{1,1} \geq X_{1,n}\ \forall n$. Therefore, $X_{1,1} > X_{1,n}\ \forall n \notin \Delta$. Applying Lemma 2 with $i = j = m = 1$ to each $n \notin \Delta$, we get $c_{p,n} = 0\ \forall p, n \notin \Delta$.

   Since $c_{1,1} > 0$ and $X_{1,1} = X_{1,n}\ \forall n \in \Delta$, Lemma 2 implies that there do not exist $m \neq 1, n \in \Delta$ such that $c_{m,n} > 0$ and $X_{m,n} > X_{m,1}$. Observe that if $c_{m,n} = 0$ and $X_{m,n} > X_{m,1}$ for some $m \neq 1, n \in \Delta$, then there must be some other $q \notin \{n, 1\}$ such that $c_{m,q} > 0$ and $X_{m,q} \geq X_{m,n} > X_{m,1}$, since $c_{m,n} = c_{m,1} = 0$. Therefore, $X_{m,n} \leq X_{m,1}\ \forall m \neq 1, n \in \Delta$.

   Fix an $n \in \Delta$. By Lemma 2, if $\exists m$ such that $X_{m,1} > X_{m,n}$, then $c_{p,n} = 0\ \forall p$. Therefore, if $\exists m$ such that $n \in \Delta$ and $c_{m,n} > 0$, then $X_{i,1} = X_{i,n}\ \forall i$. However, from Lemma 3, this would imply $r_i = r\ \forall i$ and some constant r, violating Assumption 1. Therefore $\forall m, n \in \Delta$ we have $c_{m,n} = 0$. We conclude that $c_{m,n} = 0\ \forall m, n \neq 1$.

In both Case 1 and Case 2, we get $c_{m,n} = 0\ \forall m \neq 1, n \neq 1$, which is the desired result.

□

We can now prove Theorem 1, which says $\exists j^*$ such that $c_{i,j^*} = 1\ \forall i$ and $c_{i,j} = 0\ \forall i, j \neq j^*$.

Proof. The proof involves three simple steps:

1. $\exists (i, j^*)$ such that $c_{i,j^*} = 1$.

   Otherwise, $\exists (j,n)$ such that $X_{m,j} = X_{m,n}\ \forall m$. From Lemma 3 and Assumption 1, this is impossible.

2. If $c_{i,j^*} = 1$ then $c_{m,n} = 0 \; \forall m, n \neq j^*$.

   See Lemma 4.

3. Therefore, we conclude that $c_{i,j^*} = 1 \; \forall i$.

$\square$

A.2 Proof of Lemma 1

Suppose $\mathbf{x}$ is played in the first stage and $(\mathbf{C}^*, \mathbf{V}^*)$ is a NE in the subgame (notation defined in Table 7). From Theorem 1, we know that $c_{i,j} = 1 \; \forall i$ for some $j \in 1, \ldots, J$.

Suppose publisher i (weakly) prefers an equilibrium where $c_{i,j} = 1$ over any other equilibrium:

$$\bar{\pi}_{i,j} \geq \bar{\pi}_{i,n} \; \forall n \neq j \tag{82}$$

Recall that $\bar{\pi}_{i,j}$ is publisher i's profit in any equilibrium when $c_{i,j} = 1 \; \forall i$ i.e., if all publishers choose ad network j.

By definition, $\bar{\pi}_{i,j} = \sum_j \sum_k \pi_{i,j,k}$ with $c_{i,j} \equiv 1 \; \forall i$ i.e., from (3):

$$\bar{\pi}_{i,j} \equiv N_{i,j} V_i \frac{1}{K} h \sum_k v_{k,j} \tag{83}$$

where, using $c_{i,j} = 1 \; \forall i$ in (42), we have:

$$v_{k,j} = \frac{y_k}{R_k} \frac{\beta \sum_i r_i V_i}{\sum_i N_{i,j} r_i V_i} \tag{84}$$

Using (83) and (84) in (82) and simplifying, we get:

$$\left(\frac{1 - x_j}{1 - x_j^{\alpha_j}} - \frac{1 - x_n}{1 - x_n^{\alpha_n}} \right) (r_i - \bar{r}) \geq 0 \;\; \forall n \neq j \tag{85}$$

WOLOG, we have assumed in (84) and (85) that $x_j < 1$ since $\bar{\pi}_{i,j} = 0$ whenever $x_j = 1$.

Under assumption 2, the high-quality publishers decide which equilibrium is chosen. Therefore, in equilibrium, the inequality (85) holds for $r_i > \bar{r}$. We conclude that $c_{i,j} = 1 \; \forall i$ in equilibrium if and only if:

$$\frac{1 - x_j}{1 - x_j^{\alpha_j}} \geq \frac{1 - x_n}{1 - x_n^{\alpha_n}} \;\; \forall n \neq j \tag{86}$$

A.3 Proof of Theorem 2

From Lemma 1 and Assumption 2, (86) will hold for the ad network j that is chosen in equilibrium.

We define the function $f : [0, 1] \times (0, 1] \to \mathbb{R}$ as follows:

$$f(x, \alpha) \equiv \begin{cases} \frac{1-x}{1-x^\alpha} & x < 1 \\ \frac{1}{\alpha} & x = 1 \end{cases} \tag{87}$$

Thus, (86) says that $c_{i,j} = 1\ \forall i$ if and only if $f(x_j, \alpha_j) \geq f(x_n, \alpha_n)\ \forall n \neq j$. The function $f(x, \alpha)$ has the following properties:

1. $f(x, \alpha)$ is continuous on $x \in [0, 1)$.
2. By definition, $f(x, \alpha)$ is also continuous in the limit as $x \to 1$.

$$\lim_{x \to 1} f(x, \alpha) = \lim_{x \to 1} \frac{1 - x}{1 - x^{\alpha}} = \lim_{x \to 1} \frac{\frac{\partial}{\partial x}(1 - x)}{\frac{\partial}{\partial x}(1 - x^{\alpha})} = \frac{1}{\alpha} \equiv f(1, \alpha) \tag{88}$$

3. $\frac{\partial}{\partial x} f(x, \alpha) = \left(\frac{1}{1-x^{\alpha}}\right)^2 \left(x^{\alpha}(1 - \alpha) + (\alpha x^{\alpha-1} - 1)\right) > 0$.

   $\frac{\partial}{\partial x} f(x, \alpha)$ is strictly positive because a) $\lim_{x \to 0} \frac{\partial}{\partial x} f(x, \alpha) = \infty$ b) $\lim_{x \to 1} \frac{\partial}{\partial x} f(x, \alpha) = \frac{1-\alpha}{2\alpha} > 0$ and c) $\frac{\partial^2}{\partial x^2} f(x, \alpha) < 0$. Showing b) and c) is straightforward, but requires some arithmetic.
4. $\frac{\partial}{\partial \alpha} f(x, \alpha) = \frac{1-x}{(1-x^{\alpha})^2} x^{\alpha} \ln x < 0$.

Suppose ad network 1 is (strictly) the most effective at detecting invalid clicks, and ad network 2 is (weakly) the second-most effective i.e., $\alpha_1 < \alpha_2 \leq \alpha_j\ \forall j > 2$. Since $\frac{\partial}{\partial x} f(x, \alpha) > 0$, we get:

$$1 \leq f(x_2, \alpha_2) \leq f(1, \alpha_2) = \frac{1}{\alpha_2} < \frac{1}{\alpha_1} = f(1, \alpha_1) \tag{89}$$

$f(x_j, \alpha_1)$ is continuous on $x_j \in [0, 1]$ with $f(0, \alpha_1) = 1$ and $f(1, \alpha_1) = \frac{1}{\alpha_1}$. So, by the intermediate value theorem, $\exists x^*$ such that $0 < x^* < 1$ and $f(x^*, \alpha_1) = \frac{1}{\alpha_2}$. Since $\alpha_2 \geq \alpha_j\ \forall j > 2$, a similar statement also applies for any other ad network $j > 2$.

Therefore, as long as $x_1 > x^*$, ad network 1 can guarantee that $f(x_1, \alpha_1) \geq f(x_2, \alpha_2)$, and consequently that $c_{i,1} = 1\ \forall i$. It is therefore a dominant strategy for ad network 1 to choose $x_1 > x^*$ in the first-stage game.

Finally, it is easy to see that $x^* \to 1$ as $\alpha_2 - \alpha_2 \to 0$. For $x \geq x^*$:

$$\frac{1}{\alpha_1} \geq f(x, \alpha_1) \geq f(x^*, \alpha_1) = \frac{1}{\alpha_2} \tag{90}$$

As $\alpha_2 - \alpha_1 \to 0$, we also have $\frac{1}{\alpha_2} \to \frac{1}{\alpha_1}$. Therefore, by the squeeze theorem:

$$f(x^*, \alpha_1) \to \frac{1}{\alpha_1} = \lim_{x \to 1} f(x, \alpha_1) \tag{91}$$

which, by continuity, implies that $x^* \to 1$.

Observe that:

$$\frac{h}{1 - h} \eta_j = \sum_i \bar{\pi}_{i,j} = \frac{\beta h}{K} \left(\sum_i r_i V_i\right) \sum_k \frac{y_k}{R_k} \tag{92}$$

which is independent of j. That is, ad network 1's total profit the total profit across all publishers is the same irrespective of which equilibrium is chosen and how aggressive the ad networks are. However, it can be shown from (83) that:

$$\frac{\partial^2 \bar{\pi}_{i,j}}{\partial x_j \partial r_i} > 0 \tag{93}$$

which means as $x^* \to 1$, higher-quality publishers get a larger fraction of the total profits.