

The Effect of Fraud Investigation Cost on Pay-Per-Click Advertising

Min Chen *

School of Business, Richard Stockton College of New Jersey, Galloway, NJ 08205
min.chen@stockton.edu

Varghese S. Jacob, Suresh Radhakrishnan, Young U. Ryu
School of Management, University of Texas at Dallas, Richardson, TX 75080
{vjacob, sradhokr, ryoung}@utdallas.edu

Abstract

Click fraud is a critical problem in pay-per-click (PPC) advertising industry wherein advertisers pay a service provider (SP) when their ad is clicked. While both SPs and advertisers employ technologies to identify fraudulent clicks, prior work shows that they cannot be induced to make further improvements to their respective technologies and suggests that PPC industry would benefit from using other mechanisms, such as reputation and third party investigation, to deter click fraud. In this paper, we consider the use of third-party investigation to address this problem. In particular, we model the problem of identifying click fraud in a three stage process where the SP first classifies clicks as fraudulent or not using a *detection* technology, and then the advertiser does the same for those clicks classified as valid using a *verification* technology and if there are disagreements, a third party inspects them using an *investigation* technology and his conclusion is considered binding. The SP and advertiser can choose a high- or low-precision technology to identify click fraud, but the choices are not observable to the other, leading to a double moral hazard problem. On the other hand, the quality of the third-party's investigation technology is observable and known to both the SP and the advertiser. We consider two payment schemes for the third party's investigation cost, the SP-pay scheme (SS), where SP pays for third party investigation, and the advertiser-pay scheme (AS), where the advertiser pays the investigation cost. In this paper, we examine the question of who has the incentive to pay the investigation cost (i.e., finance the third party) when a third party handles the investigation of contestable clicks? We show that when the cost of SP's detection technology is large, the SP would choose not to pay for the investigation cost even though he may receive a higher payment from the advertiser. Nevertheless, the SP would be willing to pay for third party investigation when the cost of his detection technology is small. This happens because if the advertiser pays for the investigation in this case, then the responsibility of payments to the third party will help mitigate the advertiser's incentive problem. This leads to a large decrease in PPC and overall less rent to the SP. Therefore, the SP may still choose to pay for third party investigation for a higher payment from the advertiser even though this would incur a direct cost to the SP.

Key Words: Click Fraud, Online Advertising, Game Theory, Double Moral Hazard, Incentives, Payment Schemes

*The author gratefully acknowledges the travel support from the Volkswagen Foundation.

1. Introduction

The past decade has witnessed extraordinary growth in the online advertising market. As the most popular payment model that accounts for more than 50% of the entire market (Hamner 2009)¹, the pay-per-click model is widely accepted by both service providers (SP) and advertisers and is the primary source of income for many SPs.² One primary advantage of this model is its simplicity and accountability: advertisers pay the SP only when someone clicks on the advertisement. However, because the pay-per-click model relies on the assumption that a person clicking on an ad has an interest in the advertised product or service, it is vulnerable to *click fraud*, a practice of imitating a legitimate user to click on an ad to generate a charge per click without having an actual interest in the target of the ad (Liu et al. 2009).³ Click Forensics (2010a) estimates the average click fraud rate to be 18.6% for the second quarter of 2010.

Currently, both SPs and advertisers are taking actions against click fraud. Tuzhilin (2006) reports that SPs commonly employ online filters relying on click patterns to identify invalid/fraudulent clicks and charge the advertisers only for valid clicks. Many advertisers choose to examine the clicks classified as valid by the SP to verify that they are indeed valid and flag contestable clicks.⁴ Currently, the contestable clicks are resolved by the SP with the help of advanced offline tools or human experts.⁵ The advertiser pays for clicks that are identified as valid by both the advertiser and the SP.

Although both SPs and advertisers are concerned about click fraud, they do not disclose the methodologies used to identify invalid clicks for fear of the information falling into the wrong hands; as such, the qualities of the technologies in identifying invalid clicks are not directly observable and contractible.

¹ Other payment models are pay-per-impression (PPM) model, where advertisers are charged based on the number of times (impressions) that advertisements are shown to customers, and pay-per-action (PPA) model, where advertisers pay a commission for every referred customer who performs a desired action, such as filling out a form, creating an account or making a purchase.

² SPs derive revenues from two main sources: (1) SPs' own Web sites, e.g., Google's AdWords program, and (2) third-party Web site publishers (known as "content network") that work with SPs in a revenue-sharing arrangement: they agree to serve ads supplied by SPs on their Web sites and receive part of the revenues generated by clicks in return, e.g., Google's AdSense program.

³ Corresponding to SPs' two main sources of revenues, there are two primary types of click fraud: (1) competitive click fraud, where advertisers click rivals' ads to drive up their costs or deplete their ad budgets, and (2) inflationary click fraud, where unscrupulous third parties publishers click ads to inflate their revenues (Wilbur and Zhu 2009). As to be explained later, our analysis does not assume any particular type of click fraud attack and hence our insights hold for both types of click fraud.

⁴ Many firms use click fraud auditing firms to help verify the clicks classified by the SP, see Click Forensics (2010b) and Megna (2008a) for examples.

⁵ Google and Yahoo, for example, accept and investigate electronic traffic quality reports from Click Forensics on behalf of its advertisers (Olsen 2008, Perez 2008).

Since the SPs are better off with more clicks and advertisers are worse off if many of the clicks are fraudulent/invalid, the incentives for click fraud identification are not aligned. It follows that the SP and the advertiser can choose a technology of low rather than high precision to identify invalid clicks.

The click fraud problem has drawn increasing attentions in both academics and industry. The results from several recent research studies (e.g., Wilbur and Zhu 2009, Chen et al. 2012) suggest that the pay-per-click industry would benefit from using a neutral third party to audit SPs' click fraud detection algorithms. This is also in line with suggestions from many industry practitioners (see Arrington 2007, Hedger 2008, Megna 2008b, for examples).

Our study is the first one that formally and analytically examines the case where a neutral third-party is introduced into the current pay-per-click market. In this setting, a click can be either valid or invalid, the SP and the advertiser have common knowledge of the probability that a click is valid but cannot observe the true nature of the click. Once a click occurs, three stages of identifying click fraud proceed in the following order: First, the SP classifies clicks using a click fraud *detection* system. Second, the clicks classified as valid by the SP's detection technology are verified by the advertiser using a *verification* technology. Finally, a neutral third-party auditor inspects the clicks using an *investigation* technology if there is a disagreement between the classification results from the first two stages. The advertiser pays for the clicks that are jointly identified as valid in all of the three stages.

The classifications by the detection, verification, and investigation technologies are imperfect: they do not identify all invalid clicks and may incorrectly classify some valid clicks as invalid.⁶ We use the term "precision" to measure the quality of the technologies, where the precision of a technology is defined as the probability that a valid click is correctly identified by the technology. The SP and advertiser can choose either a high-precision or a low-precision technology, where a high-precision technology is more likely to classify clicks correctly than a low-precision technology; as such, high-precision technolo-

⁶ Click identification methods can be classified into three categories (Kshetri 2010): (1) the anomaly-based approach, which considers invalid clicks to be those deviating significantly from normal behaviors; (2) the rule-based approach, which uses heuristics to classify clicks based on specific rules defined by human experts; (3) the classifier-based approach, which employs data mining classifier to detect invalid clicks based on the past data about valid and invalid clicking activities. Like the classification methods in other domains, because of the nature of these approaches, clicks may not be classified into valid or invalid with certainty, hence, the classification technologies are imperfect. We use the term "precision" to refer to their performance.

gies are also more costly. On the other hand, the quality/precision of the neutral third-party's investigation technology is observable and known to both the SP and the advertiser. The SP and the advertiser choose the high- or low-precision technology based on the payment-per-click (PPC) to maximize their own expected profit. The choice of the high- or low-precision technology is not jointly observable and contractible which leads to a double moral hazard problem. We model the problem in a principal-agent setting with the SP as the principal and the advertiser as the agent.

When disagreements occur, the neutral auditor inspects clicks using the investigation technology and the outcome from the investigation is considered binding to both parties. As such, the immediate question is who should pay for the investigation. In this paper, we consider two payment schemes for the investigation cost: (1) the SP-pay scheme (SS), where the SP is held responsible for investigating the advertiser's disputes, and (2) the advertiser-pay scheme (AS), where the advertiser pays for the investigation cost. The current pay-per-click market is dominated by few giant SPs such as Google and Yahoo. As such, we compare the SP's profits across the two schemes to examine the question of which is the SP's preferred scheme when a third party handles the inspection of contestable clicks.

The probability that a click is paid for, i.e., a click is identified by all the three technologies as valid, is higher with high-precision than low-precision detection technology and is lower with high-precision verification technology. Similarly, the probability that a click is investigated by the third party is lower with high-precision technologies. Accordingly, both the SP and the advertiser have benefits of choosing the current high-precision technology over the low precision technology, and they balance their choice with the cost of the technology and benefits from paying for investigation.

We show that either the incentive problem of inducing advertiser's high-precision verification technology or the incentive problem of inducing SP's high-precision detection technology can be more severe in both schemes. To keep the intuition straightforward, the SP's incentive problem is more severe if the cost of the high-precision detection technology is sufficiently high. Otherwise, the advertiser's incentive problem is more severe. Thus, if the cost of detection technology is (not) sufficiently high, the PPC induces the SP's (advertiser's) choice of the high-precision detection (verification) technology, and the veri-

fication (detection) technology can be automatically induced. Similar results are also found in other double-moral hazard principal agent models (see Chen et al. 2012, Hwang et al. 2006, Jayanth et al. 2011).

Depending on the cost of SP's detection technology, we need to compare the two investigation payment schemes in three possible cases: (a) the SP's incentive problem is more severe in both schemes, i.e., cost of detection technology is sufficiently large, (b) the SP's incentive problem is less severe in SS but more severe in AS, i.e., cost of detection technology is moderate, and (c) advertiser's incentive problem is more severe in both schemes, i.e., cost of detection technology is not large.

We show that when the cost of detection technology is not large, surprisingly, the SP would receive higher profits in SS. In other words, the SP is better off paying for the investigation cost himself. This occurs because opting for AS has two effects. First, the SP can save on investigation cost by having the advertiser pay for it. This is a *direct effect* that increases SP's profits. Second, when advertiser pays for the investigation, he will have more incentives to choose the high-precision verification technology, because it helps reduce the investigation probability and hence lower the investigation cost. In other words, when the advertiser's incentive problem is more severe, opting for AS will also cause an *indirect effect* that mitigates the advertiser's incentive problem and decreases the PPC made to SP. As such, even though the SP does not pay for investigation in AS, the indirect effect that decreases the PPC is sufficiently large to offset the savings in investigation cost, leading to a decrease in SP's profit. As a result, the SP would rather choose to pay for the investigation for a larger PPC in return.

The rest of the paper is organized as follows: section 2 reviews related literature, section 3 describes the key elements of the model, section 4 analyzes the model and provides the results, and section 5 discusses managerial insights and concludes the paper.

2. Related Literature

Our paper is closely related to two streams of literature: one that examines the click fraud problem, the other that examines the double moral hazard problem.

2.1. Literature on Click Fraud Problem

The issue of click fraud has drawn increasing attention in the recent years. Wilbur and Zhu (2009) analyze the effects of click fraud on the online advertising industry. They show that when advertisers know the level of click fraud, they will lower their bids to the point that click fraud has no impact on total advertising expenditures. Nevertheless, when the level of click fraud is uncertain, the SPs' revenues will rise when the keyword auction is less competitive and fall when the keyword auction is more competitive. Their result that SPs are sometimes helped and sometimes hurt by click fraud reinforces the need for a neutral third party to audit SPs' click fraud detection algorithms. Chen et al. (2012) examine the click fraud problem in a similar setting, with the exception that the investigation is handled by SP and the quality of investigation is not observable to the advertiser. They find that when the investigation cost is sufficiently high, even though both players may be better off with improvements in both the detection and verification technologies, they will not unilaterally do so: a classic prisoner's dilemma result. Their results also suggest that a third-party investigation may be needed to help effect improvements to detection and verification technologies. However, the reasons for such implication are quite different from that from Wilbur and Zhu (2009): specifically, Chen et al. (2012) find that interaction among the detection, verification and investigation technologies makes it difficult for the PPC to induce both parties to effect improvements in their technologies.

Other works focus on the use of alternative payment methods to mitigate the click fraud problem. Mahdian and Tomak (2009) propose the PPA model. They examine the challenges in the design of incentive-compatible PPA mechanisms and suggest approaches to tackle some of them. The PPA model, however, creates an incentive for advertisers not to report all conversions occurred on their Web sites and hence produces a new fraudulent activity known as *action* fraud. Goodman (2005) proposes a pricing scheme that sells advertisers a particular percentage of all impressions rather than user clicks. Nevertheless, the per-percentage of impressions model, apart from being difficult to implement, deviates from the developed industry standard that sells clicks, and risks a negative backlash in the marketplace (Immorlica et al. 2005).

In contrast to these studies we consider the interaction between the technologies employed for classi-

fyng clicks. We examine the case where a third party handles inspection of contestable clicks to provide insights into whether the use of third party inspection helps mitigate the click fraud problem.

2.2. Literature on Double Moral Hazard Problem

Double moral hazard problems have been examined extensively in economics and have been applied in a variety of settings. Balachandran and Radhakrishnan (2005) examine a double moral hazard case where both the supplier's and the buyer's qualities are unobservable. They consider not only whether contract payments and penalties based on either information from incoming inspection or information from external failures induces first-best quality, but also whether the penalty satisfies the fairness criterion. Hwang et al. (2006) examine a two-tier supply chain between a supplier and a buyer, with the supplier's quality and the buyer's inspection effort being unobservable. They compare the inspection regime with the certification regime and show that inspection leads to additional agency cost due to the presence of agency problems, which provides a rationale for the shift to the certification regime even though the direct cost of inspection is low.

We extend these studies to the click fraud setting by examining the interactions between sequential actions by the two parties. To our knowledge, there have been no studies that examine such sequential actions in double moral hazard settings.

3. Model Description

We examine a double moral hazard problem between a risk-neutral advertiser and a risk-neutral service provider (SP), as shown in Figure 1. The advertiser's sponsored link receives x clicks enabled by the SP. Without loss of generality, we set $x=1$. A click could either be valid (non-fraudulent) or invalid (fraudulent). The information on whether a click is fraudulent is not observable to the advertiser or the SP. Both the advertiser and the SP assess a probability of α that the click is valid, i.e., with probability $(1-\alpha)$ the advertiser and the SP expect the click to be fraudulent.

[Insert Figure 1 here]

The SP uses a detection technology that classifies a valid click as valid with probability q^S and an

invalid click as invalid with probability $\phi^S q^S$, where $\phi^S \in (0, 1/q^S)$.⁷ Correspondingly, the detection technology's type I error rate of classifying a valid click as invalid is $(1 - q^S)$, and the type II error rate of classifying an invalid click as valid is $(1 - \phi^S q^S)$. We refer to q^S as the precision of the detection technology because increasing q^S would result in a decrease in both type I and type II errors.⁸ The report of the classification by the detection technology is denoted $r^S = v, f$ for valid and invalid/fraudulent clicks, respectively. The SP can choose a detection technology with either a high precision or low precision, i.e., $q^S \in \{q_H^S, q_L^S\}$ with $q_H^S > q_L^S$. The corresponding cost of the high-precision and low-precision detection technology is $C^S(q^S) \in \{C_H^S, C_L^S\}$ with $C_H^S > C_L^S$. The choice of the detection technology is not observable and thus, is subject to moral hazard. In addition, there is a unit cost m^S for each click. We assume that the unit cost is the same irrespective of whether the high or the low precision detection technology is chosen by the SP.

The clicks that are classified as valid by the SP are verified by the advertiser.⁹ The advertiser's verification technology classifies a truly valid click that is classified as valid by the SP, as valid with probability q^A . Also, the verification technology classifies a truly invalid click that is classified as valid by the SP, as invalid with probability $\phi^A q^A$, where $\phi^A \in (0, 1/q^A)$. The precision of the verification technology is characterized by q^A . The report of the classification by the verification technology is denoted $r^A = v, f$, for valid and invalid/fraudulent clicks, respectively. The advertiser can choose a verification technology with either high or low precision, i.e., $q^A \in \{q_H^A, q_L^A\}$, where q_H^A (q_L^A) denotes the precision of high- (low-) precision technology with corresponding verification cost $C^A \in \{C_H^A, C_L^A\}$, where $q_H^A > q_L^A$ and $C_H^A > C_L^A$. The advertiser's choice of verification technology is not observable and thus, is subject to moral hazard as

⁷ The SP typically uses large amounts of click stream data from many advertisers and for many keywords to develop algorithms to identify click fraud (Greenberg 2008). Besides advertisers' log data, SPs collect other data that can be used to assist in click fraud detection in a variety of ways, including JavaScript page tags (ClickFacts 2010), URL redirect service and real-time API (Click Forensics 2010c) and proprietary server-side data collection method (Clicklab 2006).

⁸ Some research on intrusion detection systems use receiver operating characteristics (ROC) curve to characterize the performance of the systems. A ROC curve is a graphical plot of the true positive rate against false positive rate. In our model, $\phi^S q^S$ is the true positive rate and $(1 - q^S)$ is the false positive rate. In terms of ROC curves, an increase in q^S , the precision of detection technology, implies that the detection system is of a better quality because it is able to classify more fraudulent clicks ($\phi^S q^S$) with less errors in misclassifying valid clicks into fraud ($1 - q^S$).

⁹ While some information, e.g., the IP addresses of clicks, is available to both the advertiser and SP, the advertisers have page-visit data after a visitor is directed to his Web site from the SP (Greenberg 2008). This is used by advertisers to provide the advertiser with more precise information to classify clicks.

well. In addition, there is a unit cost m^A for each click that the SP classifies as valid and the advertiser verifies. We assume that the unit cost is the same irrespective of whether the high- or the low-precision verification technology is chosen by the advertiser. The probability that a click is verified by the advertiser is denoted by $I_i^A = \Pr(r_i^S = v)$ where the subscript i denotes that the report is conditioned on q_i^S the detection technology. Without loss of generality we assume that $C_L^S = C_L^A = 0$.

The disagreements in the SP's and advertiser's classification of invalid clicks are resolved by a neutral third party carefully investigating such disagreements. The investigation technology used by this third party classifies a contestable click that is truly valid as valid with probability q^I , and it classifies a contestable click that is truly invalid as invalid with probability $\phi^I q^I$, where $\phi^I \in (0, 1/q^I)$. The report of the classification by the investigation technology is denoted $r^I = v, f$ for valid and invalid/fraudulent clicks, respectively. The precision of the investigation technology used by the third party, i.e., q^I , is common knowledge to both players. In other words, unlike the detection precision and verification precision which are not observable or contractible, the investigation precision is known to both the SP and the advertiser. The third party charges a unit cost m^I for each click investigated, i.e., when a click is classified as valid by the SP and as invalid by the advertiser.

The advertiser pays the SP θ for each click classified jointly by all the three stages as valid. This is equivalent to the advertiser paying θ for each click classified by the SP as valid, and then obtaining a refund of θ for clicks classified by the verification and investigation technologies as invalid.

The sequence of events unfolds as follows. First, the advertiser and the SP agree on payment per click (PPC) θ . Second, the SP chooses the detection technology and the advertiser chooses the verification technology. Finally, the clicks occur, the technologies classify the clicks and the payment occurs when the clicks are classified as valid by the detection and verification technologies or the detection and the investigation technologies. We assume that the classification obtained from the technologies cannot be manipulated by the SP or the advertiser. In effect, the choice of the technologies, i.e., the program/code and the team who are dedicated to examine/monitor the clicks, constitutes the processes that spew out the classifi-

cation. The events are depicted in Figure 1.

We denote by T_{ij} the probability that a click is classified by either the detection and verification technologies as valid or is inspected by the investigation technology as valid, where the subscripts $i, j \in \{H, L\}$ indicate the detection and verification technologies chosen. In particular, T_{ij} is the probability that a click is paid for (with \Pr denoting probability) and is given by:

$$\begin{aligned} T_{ij} &= \Pr[\text{a click is paid} | q_i^S, q_j^A, q^I] \\ &= \Pr[r_i^S = v, r_{ij}^A = v] + \Pr[r_i^S = v, r_{ij}^A = f, r^I = v] \\ &= \alpha q_i^S [q^I + q_j^A (1 - q^I)] + (1 - \alpha) (1 - \phi^S q_i^S) (1 - \phi^A q_j^A \phi^I q^I). \end{aligned}$$

The advertiser expects to receive a benefit of γ from a true valid click: the expected benefit of γ includes not only the revenues obtained from the customer who has clicked through to his site, but also the probability that the valid customer's visit to the advertiser's site results in a sale. Thus given that not all truly valid customers may buy the advertiser's product/services, the outcome of the sales cannot be used to provide information on truly valid clicks. The expected profit for the advertiser (U) for $i, j \in \{H, L\}$ under investigation payment scheme $k \in \{A, S\}$ is given by

$$U_{ij}^k = \gamma \alpha - \theta T_{ij} - C_j^A - m^A I_i^A - D_1 m^I I_{ij}.$$

Note that the indicator variable $D_1 = 1$ if $k = A$ and $D_1 = 0$ if $k = S$. U_{ij}^k is the advertiser's expected benefit from the click, less (a) the expected payment to the SP, (b) the cost of verification technology, (c) the expected cost of verifying the click classified by the detection technology as valid, and (d) the expected cost per click of investigation if the advertiser is responsible for the investigation cost, i.e., $k = A$ and $D_1 = 1$. We assume that the expected benefit the advertiser receives from a click is high enough such that he has incentives to participate in PPC advertising. The probability that a click needs to be inspected by the third party's investigation technology (i.e., the SP classifies it as valid but the advertiser disagrees), is given by:

$$I_{ij} = \alpha q_i^S (1 - q_j^A) + (1 - \alpha) (1 - \phi^S q_i^S) \phi^A q_j^A$$

The expected profit for the SP (V) for $i, j \in \{H, L\}$ and $k \in \{A, S\}$ is given by

$$V_{ij}^k = \theta T_{ij} - C_i^S - m^S - D_2 m^I I_{ij}.$$

Note that the indicator variable $D_2 = 1$ if $k = S$ and $D_2 = 0$ otherwise. V_{ij}^k is the SP's expected click payment from the advertiser less (a) the cost of the detection technology, (b) the SP's expected cost per click of detection, and (c) the SP's expected investigation cost if the SP is responsible for the investigation cost, i.e., $k = S$ and $D_2 = 1$. The detection and verification technologies use automated programs, so that unit costs of conducting detection and verification are likely to be very small. Hence, without loss of generality, we let $m^S = m^A = 0$ in our analysis.

We make some assumptions to ensure that the representation of the problem conforms to the standard double moral hazard settings. First, we assume that the probability that a click is paid for and the probability that a click is investigated satisfy AI below.

AI. (a) $T_{Hj} > T_{Lj}$, (b) $T_{iL} > T_{iH}$, (c) $I_{HH} < I_{LH}$, and (d) $I_{HH} < I_{HL}$ for $i, j \in \{H, L\}$ and $k \in \{A, S\}$.

Assumption AI(a) states that compared to the low-precision detection technology, the high-precision detection technology results in an increase in probability that a click is paid for, i.e., $[dT/dq^S] > 0$. This will provide the SP the motivation to choose high-precision detection technology. Assumption AI(b) states that compared to the high-precision verification technology, the low-precision verification technology increases the probability that the click is paid for $[dT/dq^A] < 0$. This will provide the motivation for the advertiser to choose the high-precision verification technology over the low-precision verification technology. Similarly, assumptions AI(c) and AI(d) state that compared to the low-precision detection technology, choosing the high-precision detection or verification technologies results in a decrease in probability that a click is investigated, i.e., $[dI/dq^S] < 0$ and $[dI/dq^A] < 0$. Overall, AI is similar to the assumption in the standard agency problems where high action is more productive than low action.

$$\mathbf{AII.} \left[\frac{\alpha}{(1-\alpha)} \right] > \left[\frac{\phi^A (1 - \phi^S q_H^S) (1 + \phi^I)}{q_H^S} \right]$$

Assumption II states that the proportion of valid clicks needs to be sufficiently large. Intuitively, ad-

vertisers will not have incentives to have his ads hosted by the SP when a large portion of clicks are fraudulent. The estimated click fraud rate reported typically ranges from 15% to 25%, suggesting that this assumption is likely to be valid.

Table 1 provides a glossary of notations used in this paper. We now proceed to analyze the model.

[Insert Table 1 here]

3.1. The SP-pay Scheme (SS)

In the SP-pay scheme (SS), the SP pays the neutral third-party for investigating the clicks disputed by the advertiser. The SP's and advertiser's profit functions are obtained by setting $k = S$, $D_1 = 0$ and $D_2 = 1$ in equations (1) and (2), respectively. We examine the case where the PPC induces the choice of high-precision detection and verification technologies from the SP and the advertiser, respectively. The SP's problem in SS is provided in Program 1.

Program 1

$$\begin{aligned}
 & \max_{\theta} V_{HH}^S && \text{(OBJ1)} \\
 \text{subject to } & V_{HH}^S \geq \bar{V} = 0 && \text{(PCS1)} \\
 & U_{HH}^S \geq \bar{U} = 0 && \text{(PCA1)} \\
 & V_{HH}^S \geq V_{LH}^S && \text{(ICS1)} \\
 & U_{HH}^S \geq U_{HL}^S && \text{(ICA1)}
 \end{aligned}$$

Program 1 is a standard double moral hazard model (see Demski et al., 2004; Hwang et al., 2006). The SP's expected profit function in SS is given in (OBJ1). The participation constraints (PCS1) and (PCA1) ensure that both parties receive at least the reservation profit which is set to zero. Constraint (ICS1) is the incentive compatibility constraint for the choice of high-precision detection by the SP. Constraint (ICA1) and is the incentive compatibility constraint with respect to the advertiser's verification technology. We make an assumption on the unit cost of investigation.

AIII. $m^l < C_H^S T_{LH} / (T_{HH} I_{LH} - T_{LH} I_{HH})$.

Assumption AIII states the cost of investigating a click is not too large. This ensures that the incentive problem of inducing SP's high-precision detection technology is sufficiently severe and that the solution to Program 1 is not determined by the individual rationality constraint. Thus, the solution to Program 1 when AIII is violated is to pay the SP the expected cost of employing the high-precision detection technology and the expected cost of third party investigation. Of course, this solution will not have any agency cost and thus incentive design problem is moot. Hence, we do not consider this solution.

We characterize the solution to the Program 1 in the proposition below. Our analysis focuses on the parameter region where the advertiser and the SP are induced to choose their respective high-precision technologies. We denote the PPC by θ^{kl} , the advertiser's profits by U^{kl} , and the SP's profits by V^{kl} , where $k \in \{S, A\}$ represents the investigation payment scheme used and $l \in \{S, A\}$ indicates whose incentive constraint is binding in the solution. All proofs are presented in the Appendix.

Proposition 1

Under the SS, if the cost of high-precision detection technology is sufficiently large (small), then the incentive problem of inducing the high-precision detection technology is more (less) severe than the incentive problem of inducing the high-precision verification technology. Technically, if condition C1 is satisfied, constraint (ICSI) is binding, and if condition C1 is not satisfied, constraint (ICAI) is binding; and in particular,

1. *if C1 is satisfied, the solution is*

$$\theta^{SS} = [C_H^S - m^l(I_{LH} - I_{HH})]/(T_{HH} - T_{LH}), U^{SS} = \gamma\alpha - \theta^{SS}T_{HH} - C_H^A,$$

$$V^{SS} = \theta^{SS}T_{HH} - C_H^S - m^l I_{HH}.$$

2. *If C1 is not satisfied, the solution is*

$$\theta^{SA} = C_H^A/(T_{HL} - T_{HH}), U^{SA} = \gamma\alpha - \theta^{SA}T_{HH} - C_H^A,$$

$$V^{SA} = \theta^{SA}T_{HH} - C_H^S - m^l I_{HH}.$$

where condition C1 is given by

$$[[C_H^S - m^l(I_{LH} - I_{HH})]/C_H^A] > [\{T_{HH} - T_{LH}\}/\{T_{HL} - T_{HH}\}].$$

Proposition 1 characterizes the solution to the Program 1. It shows that either the advertiser’s incentive problem or the SP’s incentive problem can dictate the PPC, as determined by the condition C1. This is standard in binary action double moral hazard problems, i.e., one parties’ agency problem is more severe than the other (see Jayanth et al. 2011; Arya et al. 2007). Particularly, if the PPC is obtained from the SP’s incentive compatibility constraint we refer to this as the SP’s incentive problem being more severe than the advertiser’s incentive problem. On the other hand, if the PPC is obtained from the advertiser’s incentive constraint we refer to this as the advertiser’s incentive problem being more severe.

There are two other drivers that determine the severity of the incentive problem. The first is the cost part: if the cost of the high-precision detection technology is sufficiently large, then the SP’s incentive problem is likely to be more severe. The second is the effectiveness of technologies. *The effectiveness of a technology* is defined as the increase in the probability that a click is paid for when SP (advertiser) chooses the high- (low-) instead of the low- (high-) precision technologies. If the effectiveness of the detection (verification) technology, i.e., $[T_{HH} - T_{LH}]$ ($[T_{HL} - T_{HH}]$), is sufficiently large, then the advertiser’s (SP’s) incentive problem is likely to be more severe.

To keep the intuition straightforward, we focus our explanation on the cost part: if the cost of high-precision detection technology is sufficiently large then C1 is satisfied and the incentive problem of inducing the SP’s high-precision detection technology is more severe; and vice versa. We now proceed to examine the advertiser-pay scheme.

3.2. The Advertiser-pay Scheme (AS)

In the advertiser-pay scheme (AS), the advertiser pays the third-party for investigating a click if it is classified as valid by the detection technology but classified as invalid by the verification technology. The SP’s and advertiser’s profit functions are obtained by setting $k = A$, $D_1 = 1$ and $D_2 = 0$ in equations (1) and (2), respectively. We again examine the case such that the PPC induces the choice of high-precision technologies from the two parties. The SP’s problem in AS is provided in Program 2.

Program 2

$$\max_{\theta} V_{ij}^A \quad (\text{OBJ2})$$

$$\text{subject to } V_{HH}^A \geq \bar{V} = 0 \quad (\text{PCS2})$$

$$U_{HH}^A \geq \bar{U} = 0 \quad (\text{PCA2})$$

$$V_{HH}^A \geq V_{LH}^A \quad (\text{ICS2})$$

$$U_{HH}^A \geq U_{HL}^A \quad (\text{ICA2})$$

Program 2 is similar to Program 1 with the exception that the profit functions for the two parties are different from those in Program 1, because it is the advertiser instead of the SP that pays for the investigation cost in AS. The SP's expected profit function in AS is given in (OBJ2). The participation constraints (PCA2) and (PCS2) ensure that the SP and the advertiser receive at least their reservation profits. Constraint (ICS2) is the incentive compatibility constraint that incentivizes SP's high-precision detection technology. Constraint (ICA2) is the incentive compatibility constraint that induces the advertiser's high-precision verification technology. The solution to Program 2 is examined in the following Proposition.

Proposition 2

Under the AS, if the cost of high-precision detection technology is sufficiently large (small), then the incentive problem of inducing the high-precision detection technology is more (less) severe than the incentive problem of inducing the high-precision verification technology. Technically, if condition C2 is satisfied, constraint (ICS2) is binding, and if condition C2 is not satisfied, constraint (ICA2) is binding; and in particular,

1. *if C2 is satisfied, the solution is*

$$\theta^{AS} = C_H^S / (T_{HH} - T_{LH}), U^{AS} = \gamma\alpha - \theta^{AS}T_{HH} - C_H^A - m^I I_{HH},$$

$$V^{AS} = \theta^{AS}T_{HH} - C_H^S.$$

2. *If C2 is not satisfied, the solution is*

$$\theta^{AA} = [C_H^A - m^I (I_{HL} - I_{HH})] / (T_{HL} - T_{HH}), U^{AA} = \gamma\alpha - \theta^{AA}T_{HH} - C_H^A - m^I I_{HH},$$

$$V^{AA} = \theta^{AA}T_{HH} - C_H^S.$$

where condition C2 is given by

$$[C_H^S / [C_H^A - m^I(I_{HL} - I_{HH})]] > [\{T_{HH} - T_{LH}\} / \{T_{HL} - T_{HH}\}].$$

Proposition 2 characterizes the solution to the Program 2, as determined by the condition C2. Similar to Proposition 1, both the cost of technologies and their effectiveness can determine the severity of the SP's and advertiser's incentive problems. Particularly, if the cost of the high-precision detection technology is sufficiently large (small), then the SP's (advertiser's) incentive problem is likely to be more severe and C2 is likely to be (not) satisfied. On the other hand, if the effectiveness of the detection technology is sufficiently large, then the advertiser's incentive problem is likely to be more severe and C2 is more likely to be not satisfied. To keep the intuition straightforward, we again focus our explanation on the cost part. In simple terms, if the cost of high-precision detection technology is sufficiently large then C2 is satisfied and the incentive problem of inducing the SP's high-precision detection technology is more severe; and vice versa. The advertiser should consider this incentive problem when considering his auction bids: if the bids are not incentive compatible in the sense represented in Program 1 or Program 2, then the SP may not have an incentive to detect frauds.

Proposition 1 and Proposition 2 are standard solutions and provide a benchmark of the characteristics of the PPC that incentivizes the SP and the advertiser to choose the high-precision technologies in the two schemes. If the PPC is not incentive compatible then even the choice of high-precision technology is infeasible and thus the question of improvements to high-precision technologies is a moot question. We now proceed to gain insights into which payment scheme will be chosen by the SP and whether the SP and the advertiser will have incentives to improve their respective high-precision technologies unilaterally.

4. Comparison of the Investigation Payment Schemes

We compare the SP's profits across the SS and the AS to provide insights into the factors that drive SP's choices of schemes. As shown in Proposition 1 and Proposition 2, there can be two solutions for each of the two schemes, depending on whether the conditions C1 and C2 are satisfied. This leads to four possible comparisons: the SP's incentive problem is (a) more severe in both schemes (C1 and C2 are satisfied), (b)

more severe in SS and less severe in AS (C1 is satisfied but C2 is not satisfied), (c) less severe in SS and more severe in AS (C1 is not satisfied but C2 is satisfied), and (d) less severe in both schemes (C1 and C2 are not satisfied). Before proceeding to the comparison, we present the following observation which shows that one of the four possible comparisons is not feasible.

Observation 1

The advertiser's incentive problem is more severe in SS if it is more severe in AS. Technically, C1 is not satisfied if C2 is not satisfied.

Observation 1 shows that the fact that the advertiser's incentive problem is more severe in AS, i.e., C2 is not satisfied, implies that the advertiser's incentive problem is also more severe in SS, i.e., C1 is not satisfied. This observation helps eliminate case (b) from consideration because the case where the SP's incentive problem is more severe in SS but less severe in AS, i.e., C1 is satisfied but C2 is not satisfied, is not feasible.

We now proceed to compare schemes to provide insights into the conditions under which the SP prefers one scheme over the other.

4.1. When SP's incentive problem is more severe in both schemes

We first compare the two schemes when the cost of detection technology is sufficiently large, i.e., SP's incentive problem is more severe and the PPC is determined by the SP's incentive compatibility constraint, as characterized by Proposition 1.1 and 2.1, respectively. Technically, this occurs when both C1 and C2 are satisfied. The result is summarized in the following proposition.

Proposition 3

When SP's incentive problem is more severe in both schemes, i.e., the cost of detection technology (C_H^S) is sufficiently large, (a) the PPC (θ) is higher in AS, (b) the SP's profit (V) is higher in AS, (c) the advertiser's profit (U) is higher in SS. Technically, (a) $\theta^{SS} < \theta^{AS}$, (b) $V^{SS} < V^{AS}$, (c) $U^{SS} > U^{AS}$.

Proposition 3 compares the schemes when the incentive problem of inducing SP's high-precision de-

tection is more severe in both schemes, i.e., the cost of detection technology (C_H^S) is sufficiently large and both C1 and C2 are satisfied. It shows that in this case, the SP always receives a higher PPC and higher profits when he opts for AS and has the advertiser pay for investigation cost.

It is important to note that in this case, two effects occur when the SP opts for SS instead of AS. First, the SP incurs an investigation cost ($m^I I_{HH}$) and this is a *direct effect* that leads to a decrease in SP's profits. Second, opting for SS will alleviate the SP's incentive problem because the high-precision detection technology is more productive and reduces the investigation probability for the SP, i.e., $I_{HH} < I_{LH}$ (assumption AI(c)). *The productivity of a technology* is defined as the decrease in the probability that a click is investigated (I_{ij}) when the high- instead of the low- precision technology is chosen. As such, the SP has motivations to choose the high-precision detection technology to reduce the expected investigation cost by $m^I(I_{LH} - I_{HH})$ in SS and hence the PPC is lower in SS. The second effect is an *indirect effect* on SP's and advertiser's profits because it reduces the PPC needed to incentivize the SP's high-precision detection technology, which in turn decreases (increases) the SP's (advertiser's) profits.

Proposition 3 shows that the SP prefers AS over SS when the cost of detection technology is sufficiently high as he receives a higher profit in AS. This occurs because both direct and indirect effects from opting for SS would decrease the SP's profit. On the other hand, the advertiser receives a lower profit in AS because of the increase in PPC and additional investigation cost when the SP chooses AS. Thus, both parties prefer not to pay for investigation cost when the cost of SP's detection is sufficiently high.

4.2. When SP's incentive problem is more severe in AS but less severe in SS

We then proceed to compare SS against AS when SP's incentive problem is more severe in AS but less severe in SS. This occurs when the cost of detection technology (C_H^S) is moderate. Technically, this means that condition C1 is not satisfied, i.e., the advertiser's incentive constraint (ICA1) binds in SS, but C2 is satisfied, i.e., the SP's incentive constraint (ICS2) binds in AS.

Proposition 4

(i) When SP's incentive problem is more severe in AS but less severe in SS, i.e., the cost of detection technology (C_H^S) is moderate, (a) the PPC (θ) is higher in AS if the cost of detection technology (C_H^S) is moderately large; (b) the PPC (θ) is higher in SS if the cost of detection technology (C_H^S) is moderately small. Technically, (a) $\theta^{SA} < \theta^{AS}$ if $X < C_H^S < X + m^l(I_{LH} - I_{HH})$; (b) $\theta^{SA} > \theta^{AS}$ if $[1 - m^l(I_{LH} - I_{HH})/C_H^A]X < C_H^S < X$, where $X = \frac{C_H^A(T_{HH} - T_{LH})}{(T_{HL} - T_{HH})}$.

(ii) When SP's incentive problem is more severe in AS but less severe in SS, i.e., the cost of detection technology (C_H^S) is moderate, the SP's (advertiser's) profits is higher (lower) in SS if and only if the cost of detection (C_H^S) is moderately small and the cost of investigation (m^l) is sufficiently small. Technically, $V^{SA} > V^{AS}$ and $U^{SA} < U^{AS}$ if and only if $[1 - m^l(I_{LH} - I_{HH})/C_H^A]X < C_H^S < X$ and $m^l < T_{HH}[\theta^{SA} - \theta^{AS}]/I_{HH}$, where $X = \frac{C_H^A(T_{HH} - T_{LH})}{(T_{HL} - T_{HH})}$.

Proposition 4 compare the two schemes when SP's incentive problem is more severe in AS but less severe in SS, i.e., the cost of detection technology (C_H^S) is moderate. In this case, the PPC induces advertiser's high-precision detection in SS and SP's high-precision detection in AS. Proposition 4(i) shows that the PPC is higher in AS (SS) if the cost of detection technology is moderately large (small). Note that the PPC in SS and AS are characterized by Proposition 1.2 and 2.1, respectively. There is no indirect effect from opting for one of the two schemes because the PPC is not affected by the investigation cost. In other words, the PPC is determined only by the cost of the technologies and their effectiveness. In particular, if the cost of SP's detection technology is relatively large, the incentive problem of inducing the SP's high-precision detection in AS is likely to be more severe and hence the PPC is higher in AS. Similarly, if the cost of SP's detection technology is relatively small as compared to the advertiser's verification technology, the incentive problem of inducing advertiser's high-precision verification technology in SS is more severe and the PPC is higher in SS.

Proposition 4(ii) shows that both parties would prefer to pay for investigation themselves when the

SP's high detection cost is moderately small and the cost of investigation is sufficiently small. In other words, the SP will choose SS and be willing to pay for third party investigation. This finding appears to be surprising because the SP incurs an investigation cost in the SS and this cost could decrease the SP's profits. This occurs because the SP would receive a higher PPC by opting for SS when the detection cost is moderately small (Proposition 4(i)) and the increase in PPC is high enough to offset the investigation cost incurred in SS if this additional cost is sufficiently small. Similarly, the advertiser would also prefer to pay for investigation because the benefit from decreased PPC in AS is sufficiently high to result in a net increase in advertiser's profits. On the other hand, if the SP's detection cost is moderately large, the PPC is higher in AS (Proposition 4(i)), and the increase in the PPC together with the savings on investigation cost makes the SP better off choosing the AS.

4.3. When advertiser's incentive problem is more severe in both schemes

Lastly, we compare the two schemes when the advertiser's incentive problem is more severe in both schemes. This occurs when the cost of detection technology is sufficiently small, i.e., neither C1 nor C2 is satisfied. As characterized by Proposition 1.2 and 2.2, the incentive problem of inducing advertiser's high-precision verification technology is more severe, and the PPC is obtained by setting advertiser's incentive compatibility constraints, i.e., (ICA1) and (ICA2), binding in the two schemes. The following proposition presents the result of the comparison.

Proposition 5

When advertiser's incentive problem is more severe in both schemes, i.e., the cost of detection technology (C_H^S) is sufficiently small, (a) the PPC (θ) is higher in SS, (b) the SP's profits (V) is higher in SS, (c) the advertiser's profits (U) is higher in AS. Technically, (a) $\theta^{SA} > \theta^{AA}$, (b) $V^{SA} > V^{AA}$ (c) $U^{SA} < U^{AA}$.

Proposition 5 shows that when the incentive problem of inducing the advertiser's high-precision verification is more severe, i.e., the cost of SP's detection technology is small, the PPC is always higher in SS. This is because when advertiser pays for investigation (AS), choosing the high-precision verification technology helps reduce the investigation probability, i.e., $I_{HH} < I_{HL}$ (assumption AI(d)). As such, in ad-

dition to reducing the payment probability (T_{ij}), choosing the high-precision verification technology also helps decrease the expected investigation cost by $m^I(I_{HL} - I_{HH})$ in AS. This is an indirect effect that provides the advertiser more motivations to choose the high-precision verification. In other words, the advertiser's benefit from choosing high-precision verification is higher in AS and hence less incentive payment (PPC) is needed.

Proposition 5 shows that both parties would receive higher profits from paying for investigation themselves when the cost of SP's detection technology is small, i.e., the incentive problem of inducing advertiser's high-precision verification is more severe. Surprisingly, the SP is willing to pay for the investigation cost in this case. This occurs because opting for AS will cause an indirect effect that mitigates the advertiser's incentive problem and decreases the PPC made to SP. As such, even though the SP does not need to pay for investigation in AS, the indirect effect that decreases the PPC is sufficiently large to offset the savings in investigation cost, leading to a decrease in SP's profit. As a result, the SP would rather choose to pay for the investigation because he could receive a much larger PPC in return.

5. Managerial Implications and Concluding Remarks

We examined a model of click fraud identification in a three stage process: (1) the SP classifies clicks using a detection technology; (2) the advertiser does the same to those classified as valid by the SP's detection technology using a verification technology; (3) a neutral third party auditor examines the disagreements using an investigation technology. The SP is paid for a click only when it is classified jointly as valid by all three stages. The choice of technologies by the advertiser and the SP are not jointly observable and cannot be used for contracting, but the quality of the third party's investigation technology is observable to both the SP and the advertiser. We modeled this problem in a double moral hazard, principal-agent setting with the SP as the principal and the advertiser as the agent.

The third party charges for inspecting the contestable clicks and two payment schemes for investigation cost are considered: the first is SP-pay scheme (SS), where SP pays for investigation; the second is advertiser-pay scheme (AS), where the investigation is paid by the advertiser. We examine the question of who has the incentive to pay the investigation cost when a third party handles the investigation of contest-

able clicks? In other words, who would have the incentives to finance the third party for inspecting contestable clicks?

We show that either the SP's incentive problem or the advertiser's incentive problem could be more severe in both schemes. In simple terms, if the cost of the high-precision detection technology is sufficiently high, then the SP's incentive problem is more severe and the PPC is obtained by setting the SP's incentive constraint binding. Otherwise, the advertiser's incentive problem is more severe and the PPC is determined by the advertiser's incentive constraint.

Given that the cost of SP's detection technology is likely to be small due to use of inexpensive and automated online filters, we show that the SP would rather prefer to pay for investigation himself, because if he does so, he would receive higher profits. This occurs because opting for AS has two effects. First, the SP can save on investigation cost. This is a *direct effect* which increases SP's profit. Second, when advertiser pays for third party investigation, he will have more incentives to choose the high-precision verification technology, because it helps reduce the investigation probability and hence lower the investigation cost. In other words, when the advertiser's incentive problem is more severe, opting for AS will also cause an *indirect effect* that mitigates the advertiser's incentive problem and decreases the PPC made to SP. The indirect effect that decreases the PPC is sufficiently large to offset the savings in investigation cost, leading to a decrease in SP's profit. As a result, the SP would rather choose to pay for the investigation as he could receive a larger PPC in return. Therefore, we showed that the SP would still be willing to pay for third party for investigation even though this would incur a direct cost to the SP.

Our study can be extended in several directions. For example, we only discussed advertiser-pay and SP-pay schemes in this paper. In the future research, we could also consider a "loser-pay" scheme in which the party who loses the dispute pays for investigation cost. In other words, the SP pays for the investigation cost if a contestable click is classified as invalid by the third party; otherwise, the advertiser pays for investigation cost. As another example, we only examined the choices of the payment schemes for the case where a third party is introduced to inspect the contestable clicks, but did not discuss in detail the consequent impacts to both parties. In the future study, we plan to analyze whether the responsibility

of payments to the third party could help mitigate the incentive problems in the click fraud setting.

References

- Arrington M. 2007. Anchor Intelligence to Audit Click Fraud. <http://techcrunch.com/2007/12/09/anchor-intelligence-to-audit-click-fraud/>
- Arya, A., J. Glover, S. Radhakrishnan. 2007. *The Controllability Principle in Responsibility Accounting: Another Look, Essays in Accounting Theory in Honour of Joel S. Demski*. Springer, New York, 183–198.
- Balachandran, K. R., S. Radhakrishnan. 2005. Quality Implication of Warranties in a Supply Chain. *Management Science* **51**(8) 1266–1277.
- Chen, M., V.S. Jacob, S. Radhakrishnan, Y.U. Ryu. 2012. Can Payment-per-click Induce Improvements in Click Fraud Identification Technologies? Working Paper.
- ClickFacts. 2010. ClickFacts Ad Network Suite for Advertising Networks and Publishers. <http://www.clickfacts.com/ad-buddy.html>
- Clicklab. 2006. Clicklab Click Fraud Detection Frequently Asked Questions. <http://www.clicklab.com/Clicklab-click-fraud-detection-FAQ.pdf>
- Click Forensics. 2010a. Click Fraud Rate Rises Slightly in Q2 2010 to 18.6 Percent. <http://www.clickforensics.com/newsroom/press-releases/165-click-fraud-rate-rises-slightly-in-q2-2010-to-186-percent.html>.
- Click Forensics. 2010b. Current Click Forensics for Advertisers. <http://www.clickforensics.com/products/for-advertisers-and-agencies/advertiser-advertisers.html>.
- Click Forensics. 2010c. Collecting Click Data. <http://www.clickforensics.com/whatwedo/how-we-do-it/getting-data.html>.
- Demski, J. S., H. Frimor, D. E. M. Sappington. 2004. Efficient Manipulation in a Repeated Setting. *Journal of Accounting Research* **42**(1) 31–49.
- Goodman, J. 2005. Pay-per-Percentage of Impressions: An Advertising Method that is Highly Robust to Fraud. *Workshop on Sponsored Search Auctions*, Vancouver, BC, Canada.
- Greenberg, A. 2008. Yahoo! Cozies Up To Its Click-Fraud Critics. http://www.forbes.com/2008/03/17/click-fraud-yahoo-tech-security-cx_ag_0317click.html.
- Hammer, S. 2009. Pay-Per-Click Web Advertisers Combat Costly Fraud. *The New York Times*. <http://www.nytimes.com/2009/05/13/business/media/13adco.html>.
- Hedger, J. 2008. Click Fraud and the Need for Third Party Verification. <http://www.searchenginejournal.com/third-party-click-fraud-verification/6647/>
- Hwang, I., S. Radhakrishnan, L. Su. 2006. Vendor certification and appraisal: Implications for supplier quality. *Management Science* **52**(10) 1472-1482.
- Immorlica, N., K. Jain, M. Mahdian, K. Talwar. 2005. Click Fraud Resistant Methods for Learning Click-through Rates. In *Lecture Notes in Computer Science*. Springer-Verlag, New York, 34–45.
- Jayanth, R., V. Jacob and S. Radhakrishnan. 2011. Vendor/Client Interaction for Requirements Assessment in Software Development: Implications for Feedback Process. *Information Systems Research* **22**(2).
- Kshetri, N. 2010. The Economics of Click Fraud. *IEEE Security and Privacy* **8**(3) 45-53.

- Liu, D., J. Chen, and A. B. Whinston. 2009. Current Issues in Keyword Auctions. In *Business Computing (Handbooks in Information Systems)*. Emerald Group Publishing Limited. Volume 3, 69-97.
- Mahdian, M., K. Tomak. 2009. Pay-Per-Action Model for On-line Advertising. *International Journal of Electronic Commerce* **13**(2) 113–128.
- Megna, M. . 2008a. Click Forensics Submitting Click Fraud Reports to Google. <http://www.ecommerce-guide.com/news/news/article.php/3777551>.
- Megna, M. 2008b. eBiz News: Study Says Search Firms Profit from PPC Fraud. <http://www.ecommerce-guide.com/article.php/3738776/eBiz-News-Study-Says-Search-Firms-Profit-from-PPC-Fraud.htm>
- Olsen, S. 2008. Click Forensics, Yahoo Take on Click-fraud Cases. http://news.cnet.com/8301-1023_3-9986373-93.html.
- Perez, J. C. 2008. Google allies with click-fraud-detection firm Click Forensics. http://www.computerworld.com/s/article/print/9116932/Google_allies_with_click_fraud_detection_firm_Click_Forensics?taxonomyId=17&taxonomyName=Security
- Tuzhilin, A. 2006. The Lane's Gifts v. Google Report. [http://googleblog.blogspot.com/pdf/Tuzhilin Report.pdf](http://googleblog.blogspot.com/pdf/Tuzhilin_Report.pdf).
- Wilbur, K. C., Y. Zhu. 2009. Click Fraud. *Marketing Science* **28**(2) 293–308.

Appendix

Proof of Proposition 1

From (PCS1) note that $\theta > 0$. The solution to Program 1 is given by θ such that either (PCS1), (ICS1), or (ICA1) is binding. Rearranging constraints (PCS1) and (ICS1) requires $\theta \geq (C_H^S + m^l I_{HH})/T_{HH} = A1$ and $\theta \geq [C_H^S - m^l(I_{LH} - I_{HH})]/(T_{HH} - T_{LH}) = A2$, respectively. Using Assumption AI and AIII, it follows that $A2 \geq A1$. Hence, whenever (ICS1) is satisfied, (PCS1) will also be satisfied, and (PCS1) is not binding. Thus, the solution to Program 1 is determined by either (ICS1) or (ICA1). Solving for θ using (ICS1) yields θ^{SS} and using (ICA1) yields θ^{SA} . Thus, the solution to Program 1 is given by $\max\{\theta^{SS}, \theta^{SA}\}$. Using condition C1, it can be verified that $\theta^{SS} > \theta^{SA}$ when C1 is satisfied, and vice versa. The other expressions are derived by substituting optimums θ^{SS} and θ^{SA} in $U(\cdot)$ and $V(\cdot)$, respectively.

Proof of Proposition 2

From (PCS2) note that $\theta > 0$. The solution to Program 2 is given by θ such that either (PCS2), (ICS2), or (ICA2) is binding. Rearranging constraints (PCS2) and (ICS2) requires $\theta \geq C_H^S/T_{HH} = A3$ and $\theta \geq C_H^S/(T_{HH} - T_{LH}) = A4$, respectively. Using assumption AI, i.e., $T_{HH} > T_{LH}$, it follows that $A4 \geq A3$. Thus, whenever (ICS2) is satisfied, (PCS2) will also be satisfied, and hence (PCS2) is not binding. Therefore, the solution to Program 2 is determined by either (ICS2) or (ICA2). Solving for θ using (ICS2) yields θ^{AS} and using (ICA2) yields θ^{AA} . Thus, the solution to Program 2 is given by $\max\{\theta^{AS}, \theta^{AA}\}$. Using condition C2, it can be verified that $\theta^{AS} > \theta^{AA}$ when C2 is satisfied, and vice versa. The other expressions are derived by substituting optimums θ^{AS} and θ^{AA} in $U(\cdot)$ and $V(\cdot)$, respectively.

Proof of Observation 1

Rearrange C1 and C2 yields $C_H^S > C_H^A [\{T_{HH} - T_{LH}\}/\{T_{HL} - T_{HH}\}] + m^l(I_{LH} - I_{HH})$ and $C_H^S > [C_H^A - m^l(I_{HL} - I_{HH})][\{T_{HH} - T_{LH}\}/\{T_{HL} - T_{HH}\}]$, respectively. Using assumption AI, it follows that $C_H^A [\{T_{HH} - T_{LH}\}/\{T_{HL} - T_{HH}\}] + m^l(I_{LH} - I_{HH}) > C_H^A [\{T_{HH} - T_{LH}\}/\{T_{HL} - T_{HH}\}] > [C_H^A - m^l(I_{HL} - I_{HH})][\{T_{HH} - T_{LH}\}/\{T_{HL} - T_{HH}\}]$. Hence, C2 is satisfied whenever C1 is satisfied, and similarly, C1 is not satisfied whenever C2 is not satisfied.

Proof of Proposition 3

When the cost of detection technology (C_H^S) is sufficiently large, the incentive problem of inducing SP's high-precision detection technology is more severe, and the PPC in the two schemes are characterized in Proposition 1.1 and 2.1, respectively.

(a) Take the difference between θ^{SS} and θ^{AS} yields:

$$\begin{aligned}\theta^{SS} - \theta^{AS} &= [C_H^S - m^I(I_{LH} - I_{HH})]/(T_{HH} - T_{LH}) - C_H^S/(T_{HH} - T_{LH}) \\ &= -m^I(I_{LH} - I_{HH})/(T_{HH} - T_{LH}) < 0, \text{ where the last inequality follows from assumption AI.}\end{aligned}$$

(b) Take the difference between V^{SS} and V^{AS} yields:

$$\begin{aligned}V^{SS} - V^{AS} &= [\theta^{SS}T_{HH} - C_H^S - m^I I_{HH}] - [\theta^{AS}T_{HH} - C_H^S] = T_{HH}[\theta^{SS} - \theta^{AS}] - m^I I_{HH} < 0, \\ \text{where the last inequality follows from Proposition 3(a) that } \theta^{SS} < \theta^{AS}.\end{aligned}$$

(c) Take the difference between U^{SS} and U^{AS} yields:

$$\begin{aligned}U^{SS} - U^{AS} &= [\gamma\alpha - \theta^{SS}T_{HH} - C_H^A] - [\gamma\alpha - \theta^{AS}T_{HH} - C_H^A - m^I I_{HH}] \\ &= -T_{HH}[\theta^{SS} - \theta^{AS}] + m^I I_{HH} > 0.\end{aligned}$$

Proof of Proposition 4

When the cost of detection technology (C_H^S) is moderate, the incentive problem of inducing SP's high-precision detection technology is less severe in SS but more severe in AS. When the SP's incentive problem is less severe in SS, C1 is not satisfied, and we can rewrite the condition as:

$$C_H^S < C_H^A \left[\frac{\{T_{HH} - T_{LH}\}}{\{T_{HL} - T_{HH}\}} \right] + m^I(I_{LH} - I_{HH}).$$

When the SP's incentive problem is more severe in AS, C2 is satisfied, we can rewrite the condition as:

$$C_H^S > [C_H^A - m^I(I_{LH} - I_{HH})] \left[\frac{\{T_{HH} - T_{LH}\}}{\{T_{HL} - T_{HH}\}} \right].$$

In this case, the PPC in the two schemes are characterized in Proposition 1.2 and 2.1, respectively.

i. Take the difference between θ^{SA} and θ^{AS} yields:

$$\theta^{SA} - \theta^{AS} = C_H^A/(T_{HL} - T_{HH}) - C_H^S/(T_{HH} - T_{LH})$$

Hence, (a) $\theta^{SA} - \theta^{AS} < 0$, if $X < C_H^S < X + m^I(I_{LH} - I_{HH})$, and (b) $\theta^{SA} - \theta^{AS} > 0$, if $[1 - m^I(I_{LH} - I_{HH})/C_H^A]X < C_H^S < X$, where $X = C_H^A \{T_{HH} - T_{LH}\}/\{T_{HL} - T_{HH}\}$.

ii. Take the difference between V^{SA} and V^{AS} yields:

$$\begin{aligned}V^{SA} - V^{AS} &= [\theta^{SA}T_{HH} - C_H^S - m^I I_{HH}] - [\theta^{AS}T_{HH} - C_H^S] = T_{HH}[\theta^{SA} - \theta^{AS}] - m^I I_{HH} = \\ &= -(U^{SA} - U^{AS}).\end{aligned}$$

Clearly, $V^{SA} < V^{AS}$ and $U^{SA} > U^{AS}$, if $\theta^{SA} < \theta^{AS}$, which happens when $X < C_H^S < X + m^I(I_{LH} - I_{HH})$, where $X = C_H^A \{T_{HH} - T_{LH}\}/\{T_{HL} - T_{HH}\}$.

When $\theta^{SA} > \theta^{AS}$, which happens when $[1 - m^I(I_{LH} - I_{HH})/C_H^A]X < C_H^S < X$, we have $V^{SA} > V^{AS}$ and $U^{SA} < U^{AS}$, if $m^I < T_{HH}[\theta^{SA} - \theta^{AS}]/I_{HH}$, where $X = C_H^A \{T_{HH} - T_{LH}\}/\{T_{HL} - T_{HH}\}$.

Proof of Proposition 5

When the cost of detection technology (C_H^S) is sufficiently small, the incentive problem of inducing advertiser's high-precision verification technology is more severe, and the PPC in the two schemes are characterized in Proposition 1.2 and 2.2, respectively.

(a) Take the difference between θ^{SA} and θ^{AA} yields:

$$\begin{aligned}\theta^{SA} - \theta^{AA} &= C_H^A / (T_{HL} - T_{HH}) - [C_H^A - m^I (I_{HL} - I_{HH})] / (T_{HL} - T_{HH}) \\ &= m^I (I_{HL} - I_{HH}) / (T_{HL} - T_{HH}) > 0, \text{ where the last inequality follows from assumption AI.}\end{aligned}$$

(b) Take the difference between and V^{SA} and V^{AA} yields:

$$\begin{aligned}V^{SA} - V^{AA} &= [\theta^{SA} T_{HH} - C_H^S - m^I I_{HH}] - [\theta^{AA} T_{HH} - C_H^S] = T_{HH} (\theta^{SA} - \theta^{AA}) - m^I I_{HH} \\ &= T_{HH} m^I (I_{HL} - I_{HH}) / (T_{HL} - T_{HH}) - m^I I_{HH} = m^I (I_{HL} T_{HH} - I_{HH} T_{HL}) / (T_{HL} - T_{HH}) > 0, \\ &\text{where the last inequality follows because using AII it can be shown that } I_{HL} T_{HH} - I_{HH} T_{HL} > 0 \\ &\text{when } q^I \leq 1.\end{aligned}$$

(c) Take the difference between and U^{SA} and U^{AA} yields:

$$\begin{aligned}U^{SA} - U^{AA} &= [\gamma\alpha - \theta^{SA} T_{HH} - C_H^A] - [\gamma\alpha - \theta^{AA} T_{HH} - C_H^A - m^I I_{HH}] \\ &= -T_{HH} (\theta^{SA} - \theta^{AA}) + m^I I_{HH} = -(V^{SA} - V^{AA}) < 0.\end{aligned}$$

Table 1: Summary of Notations

Notation	Definition
α	The probability that a click is valid, i.e., the valid click rate
γ	The advertiser's expected profit from a valid click
θ	The advertiser's PPC to the SP
q_i^S	The precision of the SP's detection technology $i \in \{H, L\}$
q_j^A	The precision of the advertiser's verification technology $j \in \{H, L\}$
q^I	The precision of the third party's investigation technology
ϕ^S	The ratio of true positive to true negative for the SP's detection technology
ϕ^A	The ratio of true positive to true negative for the advertiser's verification technology
ϕ^I	The ratio of true positive to true negative for the neutral auditor's investigation technology
C_i^S	The SP's detection cost when choosing detection technology $i \in \{H, L\}$
C_j^A	The advertiser's verification cost when choosing verification technology $j \in \{H, L\}$
m^S	The unit cost of classifying a click by SP's detection technology
m^A	The unit cost of classifying a click by advertiser's verification technology
m^I	The unit cost of classifying a click by neutral auditor's investigation technology
I_{ij}	The probability that a click is investigated by the neutral third party
T_{ij}	The probability that a click is paid to the SP in scheme

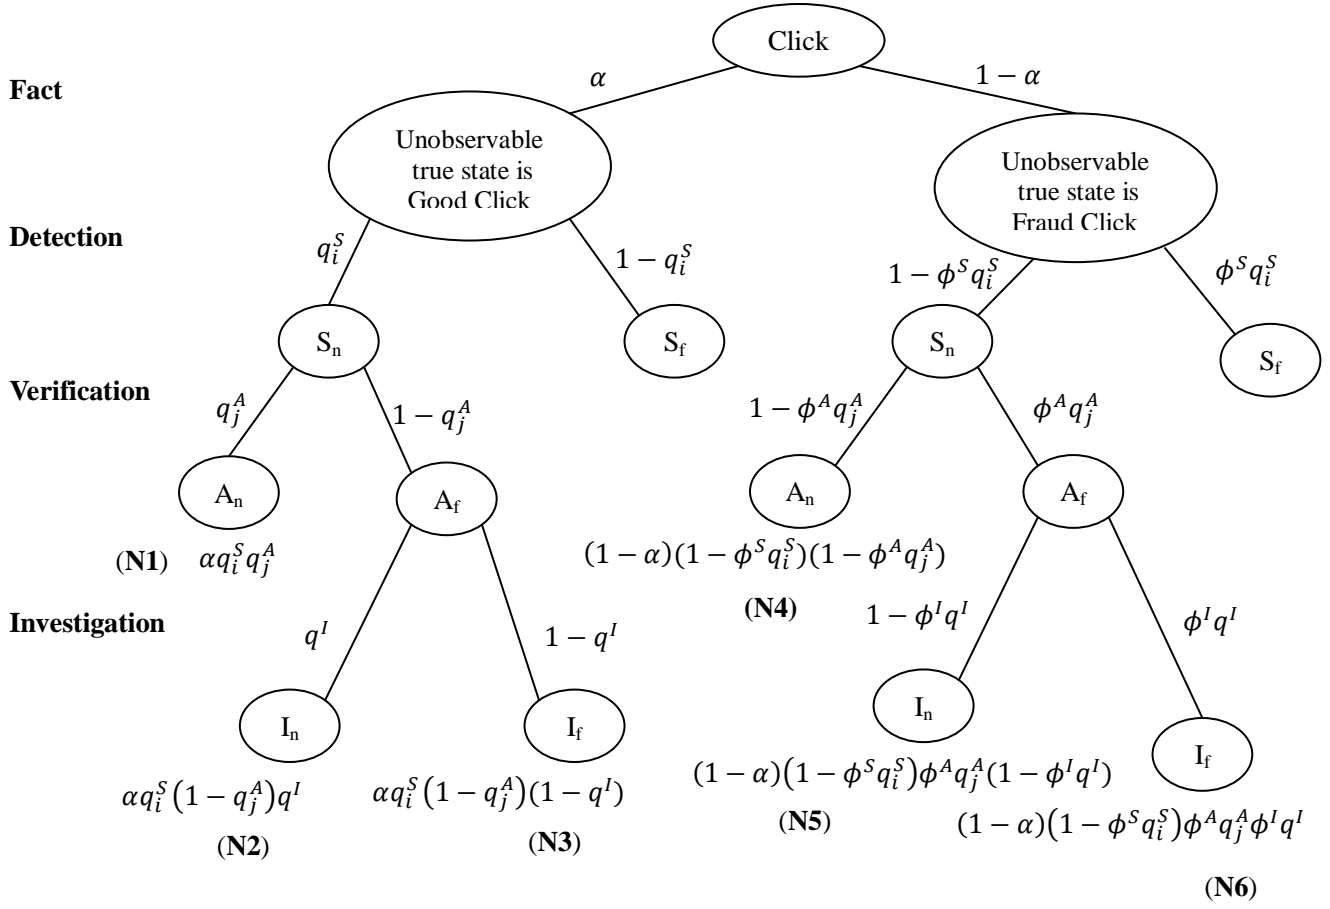


Figure 1: Game Tree for the Click Fraud Problem